Specialization by design: the unequal geographic effects of modular product design

Job Market Paper – Click Here For Latest Version

Vishan Gandhi Nigam (MIT)*

December 22, 2025

Abstract

I show that modular design – a revolution in how firms organize innovation – concentrates industrial production in large countries. Modular products follow common rules called design platforms, and thus can share inputs while remaining customized for local needs. Combining six new datasets on global automotive design and trade, within-firm event studies of platform rollouts and mergers, and a model with scale economies in shared input production, I find that design platforms reshape global trade in two phases. First, platform-sharing across destinations *increases trade* and enables *within-firm specialization*; for instance, poor countries export engines for affordable car segments. Second, platform-sharing across product segments creates *winner-take-all supply chains* in which a firm's largest markets produce most inputs. In both phases, design platforms expand the scope of home-market effects, concentrating input production in large markets for shared platforms rather than individual products. By quantifying the model, I show modular design has unequal effects across countries and shapes the returns to industrial policy: in particular, universal platforms (expected by 2030 for EVs) double American and Chinese input production shares, reduce production by over 80% in most smaller countries, and imply that U.S. input tariffs on China (but not on third countries) have larger reshoring effects.

*Author contact: vishan@mit.edu. I am deeply grateful to my advisors David Atkin, Abhijit Banerjee, Arnaud Costinot and Scott Stern for their invaluable guidance and support, to Amit Khandelwal, Jose Vasquez, Isabela Manelici, Treb Allen, Carliss Baldwin, Dan Gross, Bob Gibbons, Simon Johnson, Hong Luo, Tishara Garg, Vincent Rollet, Ananya Kotia, Ararat Gocmen, Sukrit Puri, and Anya Marchenko for extensive feedback, to participants in the MIT Trade, Development, and Organizational Economics seminars, the MIT Priority Technologies Group, and the Dartmouth Globalization Seminar for comments, to the NSF Graduate Research Fellowships Program, George and Obie Shultz Fund, and Jerry A. Hausman Fellowship for financial support, and to Kamna Ralhan at the US NHTSA for sharing data.

> The appearance of module-level [design] has an important economic side effect... the module can be traded. It may then come to have its own markets... becoming a separate and distinct product in the larger economic system.
>
> Carliss Baldwin and Kim Clark (2000), *Design Rules Vol. 1*

1 Introduction

Why are the world's key technologies, such as engines, batteries, electronics, and now AI models, increasingly produced in China and the United States? Recent debates emphasize policy determinants, including industrial policies (e.g., *Made in China 2025* and the US *CHIPS Act*) and local government support for new construction. Far less attention has been paid to ongoing technological changes that might favor production in large economies. This paper studies one such change – a transformation in how firms organize innovation known as *modular product design* – finding it has large and unequal implications for the geography of global industrial activity.

The core innovation of modular design is that engineers no longer design each product in isolation. Instead, firms develop basic design rules called *platforms*, invent standardized inputs called *modules* for each platform, and finally mix-and-match modules like LEGO bricks to create new products. By sharing a design platform, differentiated products – customized for consumer tastes and national regulations – can still share value-added inputs. First widely used by IBM in 1965 for its mainframe computers, modular design has diffused to goods as wide-ranging as airplanes, semiconductors, furniture, phones, and automobiles.

In this paper, I show that modular design concentrates value-added input production in large economies. Combining new data from the global car industry with event studies and gravity regressions, I find that modularity relocates supply chains in two phases. First, as products (car models) assembled in different destinations share design platforms, *input trade increases* and production shifts to large final markets for each product segment (car size). For instance, India exports engines for small car segments, while the USA exports engines for large SUVs. Second, as design platforms are shared across segments, *production shifts to large overall final markets* for each firm, often China and the United States. Modular design thus first creates specialization and trade within each firm, then renders supply chains winner-take-all.

Why would a change in how firms organize R&D – not in production processes, trade costs, or industrial policies – reshape how countries specialize? Using a model of trade in shared inputs, I show that *shared design platforms expand the scope of home-market effects* (Krugman, 1980): to exploit economies of scale and save trade costs, firms concentrate input production in large final markets for each platform rather than for each product. Modularity is thus a novel form of technological change, reshaping where intermediate goods are demanded rather than how they are produced.

My results suggest that modularity may inhibit widespread participation in global supply chains, especially for poorer and less-populated countries without large domestic markets. In model counterfactuals, I find that these effects are economically significant and will be magnified in

the future. Specifically, anticipated future increases in platform-sharing due to the EV transition double the combined automotive input production shares of China and the United States, from 37 to 75 percent, while reducing production shares by a median of 80% in other countries.

Measurement My analysis proceeds in four steps. First, to study the global production implications of modularity, I require design and supplier data disaggregated by product. In most industries, these data are trade secrets, and thus rarely made public or standardized across firms.

To overcome these issues, I collect and manually harmonize six product-level datasets covering the global automotive industry from 2005-2024. These data exist because substantial design and production details by car model are made public, both for cultural reasons and to facilitate sale and repair by independent dealers and repair shops. Specifically, I obtain modular design platforms and adoption dates by car model from global product portfolio timelines collected by *Marklines Automotive*. To complement these realized design choices, I collect archival product design plans and dates from *Wards Automotive*. To understand how these design choices affect car input choices, I use product-level parts lists from the *Hollander Interchange* automotive junkyards database. To measure input demand, I combine this design information with separate data on product assembly volumes that *Marklines* obtains from car registrations. Finally, to obtain input supply (i.e. sourcing choices) for each car model and year, I use global data on supply chain relationships from the *Marklines Who Supplies Whom* database, as well as newly-digitized individual manufacturer filings under the *American Automotive Labeling Act* (AALA).

Using the data, I first show how auto firms adopt modularity – in other words, how each firm uses shared design platforms. Adoption occurs in two phases. The first is modularity across *destinations* (regulatory regimes), e.g., between American and Indian cars, which was widely adopted by 2006. The second is modularity across product *segments* (physical characteristics), e.g., between small cars and large SUVs, which increased from 2007-2024 and will be fully adopted as universal EV platforms arrive by 2030.

Combining the product timeline and junkyard parts databases, I then document how firms use shared platforms: engineers standardize value-added inputs while keeping final cars differentiated for consumers. In pairwise event studies, car models that adopt shared design platforms use the same technology-intensive inputs (engines, powertrains, and electronics) but do not share parts of the car that are visible to consumers (car bodies and interiors).

Model The second step of my analysis is theoretical. Having described how firms adopt modular design, I examine how modularity might affect the geography of production. This is challenging because modularization is a nonstandard technological change: rather than changing how inputs are produced, design platforms change the set of final goods that share inputs. I make progress via an insight from the management literature: as argued by Baldwin and Clark (2000), *design platforms function as input markets* with module producers as sellers and final products as buyers. I thus develop a two-country model in which many products ("cars") share each design platform.[1]

[1]For instance, if Ford small and large cars use separate platforms, then there is input trade in two "sectors".

The key force in the model is that production of platform-specific inputs ("engines") features economies of scale, reflecting both static increasing returns and learning-by-doing. To exploit these economies of scale while avoiding trade costs, firms concentrate production in large markets for each *platform* (Krugman, 1980). In other words, shared design platforms expand the scope of economies of scale, and thus of countries' market size advantages, from the product level to the platform level. Increases in modularity (i.e. in platform-sharing) relocate production by creating and destroying these platform market size advantages.

The model predicts that the two phases of modularity – platform-sharing first across *destinations* within a product segment, then across *segments* – reshape countries' market size advantages in distinct ways. Without shared platforms, inputs are product-specific, so countries have market size advantages in inputs for locally-assembled products. With platform-sharing across *destinations*, inputs are segment-specific, so countries have market size advantages in locally-assembled segments. With platform-sharing across destinations and *segments*, inputs are only firm-specific, so a firm's largest assembly locations have market size advantages in inputs for all of its products.

These predictions imply that – holding fixed factor prices, productivities, trade costs, and industrial policies – modularization first creates export opportunities for many countries at the expense of local suppliers, then destroys most such opportunities by concentrating production in fewer locations. For instance, within *Ford*'s supply chain, India might shift from producing engines for India-built cars, to producing (and exporting) engines for affordable small cars, to importing all engines from China and the USA.

Empirics My third step is to test empirically for these shifts in international specialization. In doing so, I face an identification challenge: products that share platforms have similar characteristics, and thus may source from common locations for reasons unrelated to design. I thus develop an empirical strategy that exploits *within-product* variation in modularity.

Specifically, I estimate the product-level sourcing effects of modularity in an event-study framework. To do so, I construct a panel of design platform and sourcing choices from 2007-2024 for all North America-built car models. The events are years in which a car model first shares a platform with models in other destinations or product segments. My approach exploits two key features of automotive design: products are regularly redesigned to use new platforms, and redesigns are staggered within firms due to limited engineer time.[2] To guide the analysis, I derive three testable implications of the model for product-level sourcing (e.g., for changes in the engine origin country for the American *Ford Focus*), which I describe below.

I first show that *modularity across destinations increases trade.* Adoption of a platform used in other regions increases product-level import shares by 11 p.p. (45%), which is 55% of the US-wide 1998-2017 increase in import content. This effect is consistent with the first phase of the model:

[2]Because redesign timing could be endogenous to trends in production and trade costs, I control for firmwide sourcing trends via firm-year fixed effects, use planned adoption dates from archival redesign plans, and verify no effect of planned-but-delayed events or placebo redesigns that involve no change in modularity.

inputs previously used only in American-built cars are now used in cars assembled abroad, so production of some shared inputs concentrates there.

If supply chains are offshored, then to where? My second key finding is that *modularity across destinations shifts input production towards the shared platform's largest market.* In a triple-difference design, engine and transmission sourcing disproportionately rises from countries with large input demand (predicted from pre-adoption volumes) on the shared platform, consistent with the presence of platform market size effects. To directly test this mechanism, I use separate variation in platform market size due to 2008 recession-induced firm mergers, which differentially exposed car models within each firm to input demand from other locations. I find that firms increase sourcing from origins with larger merger-induced increases in platform market size (with elasticity 0.11), and that proposed-but-unimplemented mergers from the same period have no effect.

My third finding is that *modularity across segments weakens cross-country specialization in locally-assembled segments*. In event studies, small car inputs are sourced 20 p.p. less from countries that disproportionately assemble small cars, and thus 11 p.p. less from developing countries, after adopting design platforms (and thus inputs) that are also used by larger cars. In gravity regressions estimated using the *Marklines* global supplier relationship data, I show this loss of segment-specific export opportunities generalizes to the full automotive industry: the strength of cross-country specialization in inputs for locally-assembled segments weakens by 56% when platforms cover multiple segments. This result is consistent with the second phase of modularity: as product segments (e.g., car sizes) share inputs, high assembly volume in a product segment no longer implies high relative demand for that segment's inputs.

In summary, modularization first increases trade and shifts input production to large assembly locations for each segment (e.g., India makes inputs for small cars), then eliminates these production opportunities for many countries by weakening segment specialization. Both shifts occur because, as predicted by the model, countries specialize in platforms – not in products – and therefore production concentrates in each platform's largest markets.

Quantification Having shown that modularity concentrates input sourcing for a particular vehicle in its platform's largest and most-productive markets, my fourth and final step is to use a quantitative version of the model to study which *countries* will gain and lose input production as platform-sharing increases over time. To do so, I extend the Dekle et al. (2007) exact-hat counterfactuals method to allow for changes in design technology. I estimate an input scale elasticity of 6% using variation from the event studies, use final product and input demand elasticities from existing literature, and calibrate the model to match supplier relationships, assembly volumes, and design platform choices from *Marklines*.

I first find that *modularity to date concentrates automotive supply chains in several large assemblers with distinct preferences*. Relative to a no-modularity scenario, present-day adoption increases aggregate input production shares for only 6 of 28 countries. Korea, Germany, and Japan see the largest increases, partly because they are large assemblers for certain firms (Hyundai, Volkswagen, and

Toyota) and segments (compact and midsize cars), and partly because platforms are not shared across all car segments: the world's two largest assemblers (China and the USA) build full-size SUVs and trucks that use different platforms at present, and thus demand different inputs.

In contrast with modularity to date, *firmwide universal design platforms – expected by 2030 for electric vehicles – geographically concentrate input production in China and the United States*. Universal platforms increase production shares for the USA and China by 101% and 79% respectively, and for Mexico and Germany by smaller amounts. Shares fall in all other countries, including by over 80% in a majority of locations. American and Chinese shares rise because only firmwide input demand matters, with its composition – for instance, that poor countries mainly assemble small cars – becoming irrelevant. This scenario is rapidly approaching: *BYD* uses one platform (the *e-Platform 3.0*) for all cars, and *Ford* will adopt a universal platform by 2027.

Finally, I find that by altering the strength and locations of production concentration, modularity affects the returns to industrial policy. First, *with universal platforms (but not at baseline), a trade war decreases the US share of global input production*. This result reflects that high trade costs weaken platform market size advantages, which are more important for the US with universal platforms than at baseline. Second, studying unilateral US input tariffs such as the recent *Section 232* levies, I show that *universal platforms magnify the "reshoring" effects of China-specific tariffs, but not of uniform tariffs on all countries.* This is because, under universal platforms, production relocating from China concentrates in the other large-scale producer (the US), further reinforcing its relative scale advantage.

Literature This work contributes to several literatures. The first is on modular design. An established management literature examines how modular architectures shape firm organization and competitiveness (Henderson and Clark, 1990; Baldwin and Clark, 2000; MacDuffie, 2013; Jacobides et al., 2018; Thun et al., 2025), [3] and a small engineering literature studies the optimal adoption of design platforms (de Weck et al., 2003; Case et al., 2023). In contrast, modularity is only beginning to be explored by economists, via two recent theory papers (Matouschek et al. (2025) and Agrawal et al. (2024)) that study the within-firm communication and AI adoption effects respectively. Theoretically, my contribution is to show that modular design – a change inside R&D offices – can affect industrial production and trade. Empirically, this paper is the first to harmonize product-level design data over time for any industry, and to use natural experiments to study the causal effects of design.

Second, a small trade and macroeconomics literature studies how shared inputs affect firm size, scope, and productivity (Argente et al., 2025; Ding, 2023; Atalay et al., 2014; Boehm et al., 2022; Bassi et al., 2023; Autor et al., 2020) as well as intrafirm trade (Alfaro et al., 2025). I depart by studying the production location effects of input-sharing, and by showing that new technologies increase input-sharing over time. As a result, I also relate to work on the implications of new

[3] This literature includes several papers (Sako and Murray, 1999; Helper et al., 2003; Sako, 2005; Sturgeon, 2002; Cusumano and Nobeoka, 1998) under the *MIT International Motor Vehicle Project (IMVP)*, which provided valuable institutional context for this project.

technologies for international specialization, which typically emphasizes skill- and capital-biased technical changes due to offshoring and automation (Grossman and Rossi-Hansberg, 2008; Kikuchi, 2025; Rodrik, 2015; Diao et al., 2024) rather than scale-enabling innovations such as modularity.

Third, my paper relates to an established literature on the trade and industrial development effects of market size (Linder, 1961; Krugman, 1979, 1980; Murphy et al., 1989; Alesina and Spolaore, 1997; Alesina et al., 2005; Hanson and Xiang, 2004; Fajgelbaum et al., 2011; Dingel, 2016; Costinot et al., 2019; Goldberg and Reed, 2023; Dingel et al., 2023). In most work, exogenous population size, density and demographics determine market size. I show that in intermediate goods markets, firm design decisions can determine relative market sizes and thus reshape industrialization opportunities. Specifically, modularity eliminates two low-hanging development paths – backward linkages from downstream assembly (Hirschman, 1958; Alfaro-Ureña et al., 2022) and export-led industrialization (Balassa, 1978) via specialization in low-quality product segments – because inputs are no longer specific to countries or segments. Without this differentiation, global supply chains resemble the stark world of Helpman and Krugman (1985) in which only large countries specialize in manufacturing.

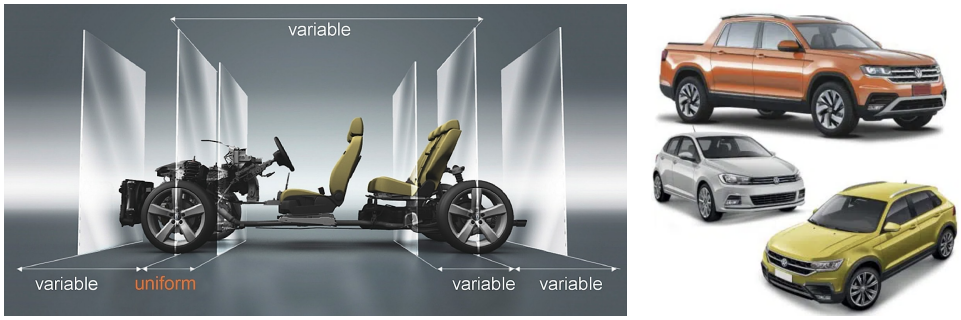
Finally, I relate to a long empirical literature on trade in the automotive sector (Goldberg, 1995; Goldberg and Verboven, 2001; Coşar et al., 2018; Castro-Vincenzi, 2022; Sabal, 2025; Castro-Vincenzi et al., 2024; Barwick et al., 2025; Head and Mayer, 2019; Head et al., 2024) and particularly to work on shared assembly (Praetorius, 2025) and battery plants (Head et al., 2025). I complement this work in two ways. First, while these papers focus on trade and industrial policy effects, I study the diffusion of a new organizational technology. Second, these papers model plant-level fixed costs and then use numerical methods to solve for counterfactual equilibria. In contrast, I study a different force (platform-level increasing returns), derive analytical predictions for the effects of modular design, and test the predictions using observed design and supply chain decisions. My approach thus resembles recent empirical work on automotive and garment supply chains (Bai et al., 2025; Helper and Munasib, 2022; Cajal-Grossi et al., 2023) showing that organizational technologies (e.g., joint ventures and "Japanese" sourcing) matter for supply chain decisions.

Structure This paper proceeds as follows. Section 2 defines modularity. Section 3 presents the data and key facts. Section 4 develops the model. Section 5 describes the empirical strategy and results. Section 6 presents the quantification procedure and counterfactuals. Section 7 concludes.

2 What is modular design? Definitions and context

Definition and example Modular design is the imposition of common *design rules*, also called *platforms*, on product teams within a firm. These rules standardize certain technical features of a product while allowing others to vary. Doing so allows products with different end requirements to nevertheless share core inputs, or *modules*, that embed key technologies.

Figure 1: Modular design rules (left) and compatible products (right) at *Volkswagen*



Notes: Left panel shows the physical design rules for the *Volkswagen MQB* platform, which is shared by cars of multiple lengths. Right panel shows three cars designed using the *MQB* platform. Source for both images is *Volkswagen AG*.

Design rules often concern physical dimensions. For example, *Volkswagen* uses its *Modularer Querbaukasten* (MQB) platform to design over 200 region-specific car models. As shown in the left panel of Figure 1, the MQB standardizes the length of one key section: the engine package compartment. However, the front, seating, and back sections can vary in length across models.

Volkswagen constrains designers in this way to enable input-sharing. Because the engine package area is standardized, all *MQB*-compatible vehicles can use the same engines, transmissions, and electronics. The result is shown in the right panel of Figure 1: *VW*'s cars around the world – ranging from European city cars to Mexican pickup trucks – now use identical components, even while remaining differentiated for consumers. In doing so, *VW* saves on both R&D and production costs: engineers no longer invent customized inputs for each individual car, and each component is produced at greater scale.

Early applications and diffusion While interchangeable parts were a central part of the American system of manufacturing developed starting in the 1800s (Hounshell, 1984), the use of prespecified design rules (platforms) to facilitate interchangeability is more recent. Modular design platforms were first widely used in the electronics industry, including for mainframe computers in 1965 via the *IBM System/360* platform (Baldwin and Clark, 2000), and for semiconductors in 1978 via the *Mead-Conway* design rules (Kuan and West, 2023). I provide a full history of these platforms and associated products in Appendix Subsection B.1, including images in Figure B.1.

Across industries, design platforms are now used for cars, commercial airplanes (Boeing), military airplanes (Lockheed Martin), wind turbines (Vestas), refrigerators (Whirlpool), elevators (Thyssen), satellites (Planet Labs), spacecraft (SpaceX), AI (all major labs, via common foundation models), mobile phones (Samsung and Nokia, see Thun et al., 2025), and by 2019 act of Congress, any weapons program of the US Department of Defense (OUSD, 2025).

Within the automotive industry, design platforms were widely shared by the 1990s by products of multiple within a firm, such as *Volkswagen* and *Audi* cars (Cusumano and Nobeoka, 1998; Sako and Murray, 1999); have gradually expanded in scope to cover multiple destinations (e.g., North

America and China) and product segments (e.g., subcompact cars and pickup trucks); and are especially important for electric vehicles: the world's largest EV producer (*BYD*) uses a single platform for all vehicles.

Reasons for diffusion: complementary technologies and knowledge-sharing The diffusion of modular design across firms and industries is largely driven by parallel technological changes related to electrification and computerization. A first force is the rise of electronic control techniques, which allow mechanically identical components to be programmed to operate differently across products.[4] A second force is the energy transition; while engines, power plants, and other fuel-powered technologies must be custom-designed for heat and size requirements, batteries (including for EVs) and solar cells can be easily stacked and rearranged for different final uses. A third reason is that technology firms now use knowledge-sharing technologies, such as computer-assisted design (CAD), digital collaboration tools, and parts databases, that enable designers to understand inputs invented elsewhere within the firm.[5]

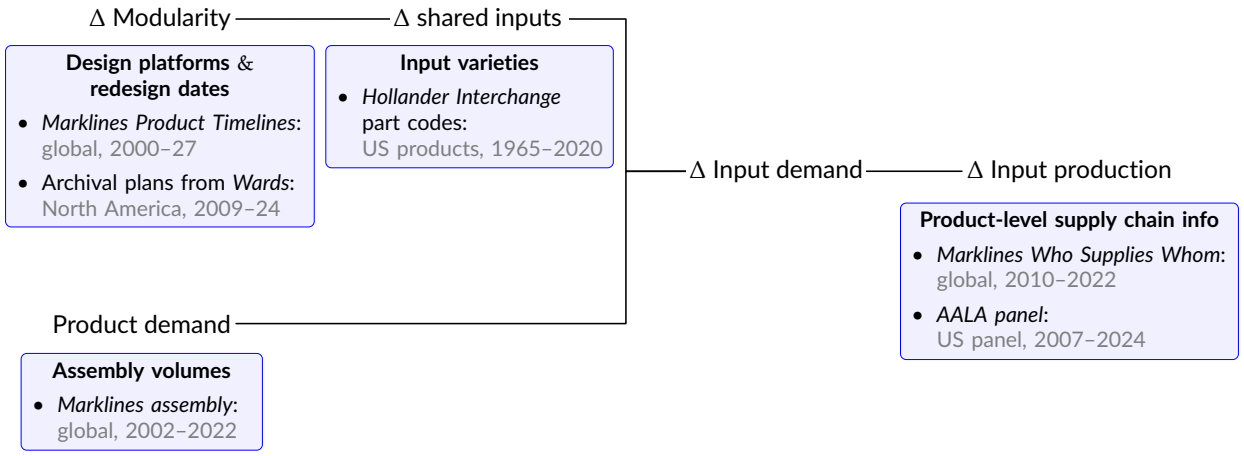
These technological forces have been augmented by knowledge diffusion efforts of academic institutions. The seminal books *Design Rules* (Baldwin and Clark, 2000) and *Product Design and Development* (Ulrich and Eppinger, 1995) codified the lessons of electronics-sector modularity for managers and engineers respectively; and the *MIT International Motor Vehicle Project* directly shared lessons from early modularization efforts at European firms (*Fiat* and *Volkswagen*) with American and Japanese automakers (Sako and Warburton, 1999; Helper et al., 2003).

Two key benefits: averted R&D costs and input scale economies The first benefit of modular design is that total R&D costs automatically fall, holding product variety fixed, because engineers no longer invent customized inputs for each product. For instance, as described by Ford's head of EV development, "*without platforms, software engineers face having to re-develop the same features for different customers and vehicles*" (Field, 2025).

A second force is more relevant for industrial location: if physical input production (e.g., of engines, batteries, and electronics) features economies of scale, firms can reduce average production costs by producing shared inputs on common production lines in fewer locations.[6] To save transport costs, this geographic concentration is especially likely to occur near large markets, as formalized by Krugman (1979, 1980). The central goal of this paper is to understand how modular design platforms (and thus, the ability to share inputs) interact with these underlying scale economies to determine global trade patterns.

[4]The ability to program common inputs to customize performance enabled the 2015 *Volkswagen* emissions scandal.
[5]For instance, until 1989 Volkswagen did not have a parts databases, so two teams using common design rules would be unaware of compatible components designed elsewhere within the firm (Lexcom 2008).
[6]In manufacturing, location-specific input scale economies may arise due to (i) *plant fixed costs*, as documented by Head et al. (2025) for EV batteries and Antràs et al. (2024) across sectors; (ii) *learning-by-doing*, as documented in automotive (Levitt et al., 2013; Adhvaryu et al., 2023), battery (Barwick et al., 2025), semiconductor (Irwin and Klenow, 1994; Goldberg et al., 2024) and other electronics production (Egelman et al., 2017); and within-plant *agglomeration effects* due to thicker internal labor markets (Doeringer and Piore, 1971; Adhvaryu et al., 2024) or larger knowledge spillovers. Case et al. (2023) argue that, in automotive input production, platform-level increasing returns are primarily driven by learning-by-doing (motive (ii)).

Figure 2: Collection and harmonization of six car model-by-year datasets



Notes: Figure shows the six car model-by-year datasets that I harmonize for the global automotive industry. Data from *Marklines* and *Wards* product timelines provide design platform names, *Hollander Interchange* database provides parts lists (i.e. specific input varieties), *Marklines Assembly* provides product-by-country assembly volumes (i.e. input demand), and *Marklines Who Supplies Whom (WSW)* and the American Automotive Labeling Act (AALA) manufacturer filings provide input origin locations (i.e. input supply). The supplier-reported *WSW* covers all cars globally but is updated once per generation; the *AALA* administrative data covers only cars sold in the United States but is updated annually.

3 Measurement and descriptive facts

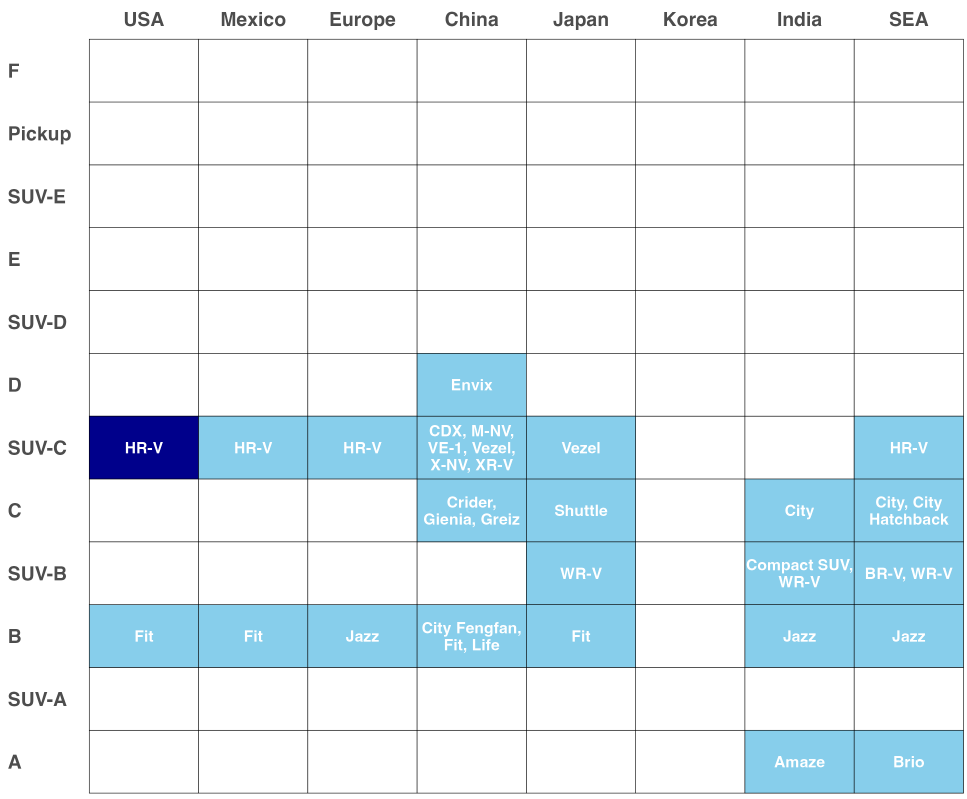
Economic research on modularity faces two key challenges. First, product-level design rules, input choices, and supply chain information are proprietary firm information, and thus rarely made public.[7] Second, even with such data, an attribution problem remains: products that share a design platform often have similar characteristics, and thus may share inputs and source locations even if design platforms are irrelevant.

To overcome these data and identification issues, I turn to data from the global automotive sector. A combination of history, cultural salience, and regulation imply that substantial public information exists on individual car models. I thus collect and harmonize twenty years of annual *car model*-level data, including design platforms from internal firm plans and announcements, assembly volumes from car registrations, supply chain data from regulatory filings and supplier self-reports, and input varieties from a parts database used by car junkyards. These data include cases where existing products are redesigned to use new design platforms, enabling my empirical strategy for separating design from product characteristics. Figure 2 summarizes the six datasets and the economic objects to which they correspond.

In the remainder of this section, I introduce these datasets and provide several descriptive facts on the diffusion, effects on car production functions (i.e. on input-sharing), and potential geographic implications of automotive modular design platforms.

[7] In many cases, firms hide input-sharing from consumers to maintain markups. For instance, Audi and VW cars often use identical components, but identical Audi-branded parts may be sold at higher prices with a different SKU number.

Figure 3: Example of platform-sharing with the US-market *Honda HR-V*



Notes: Figure shows all car models that share a design platform with the 2020 *Honda HR-V*, organized by location in the product space. Each row is a product segment, ordered from largest to smallest (from full-size luxury sedans in the F segment to micro cars in the A segment). SUVs are larger than sedans in the corresponding segment. Each column indicates the destination market (region) in which the car model is sold. SEA refers to Southeast Asia. The focal model (the American-market *HR-V* is indicated in dark blue). All other models are indicated in light blue. Data from *Marklines Automotive* and verified by manual search.

3.1 How has modularity diffused within automakers over time? Data and facts

Data: automotive design platforms from *Marklines*. I obtain design platform data for 2000-2027, which are in PDF images and require extraction via a neural network, from automotive firms' global product portfolio timelines as collected by *Marklines Automotive*.[8] See Appendix Figure B.2 for a sample timeline. For each year, the timeline indicates the firm, make (brand), model name, destination region, segment, and design platform for all car models offered. For instance, one observation in the dataset is *(Honda, Honda, HR-V, 2019, USA, SUV-C, Honda Global Small Car Platform)*. To my knowledge, this dataset is the first to track product-level design information over time for an entire industry.

Timelines embed three critical pieces of information. The first is the *position of each car model in*

[8]2023-2027 data reflect product plans at the time of data download in 2022. I extract these data using two neural networks: one for OCR (to recognize text) and one for object detection (to identify the years, indicated by lines and diamonds, in which each car is sold). I then manually verify and augment these data. This is important because similar names refer to entirely different car models, such as the *BMW 2 Series Coupe* and *BMW 2 Series Gran Coupe* in Appendix Figure B.2.

the firm's product space. As shown in Figure 3, cars are differentiated by product segment (row) and destination (column).[9] *Product segments* reflect physical characteristics, typically different dimensions of car size. For instance, poorer, denser, and high-fuel price countries prefer smaller cars, and Americans strongly prefer pickup trucks.[10] In contrast, *destinations* reflect divergent regional regulations.[11] For instance, US-market cars are illegal to drive in the EU, because European safety rules mandate smaller bumpers and different headlights.

For example, in Figure 3, the focus vehicle is the US-market *Honda HR-V*, located in the dark blue cell. To create this vehicle, designers had to meet both US regulations and the physical requirements of a C-segment compact SUV.

The second feature is the *adoption of shared design platforms* – in other words, the set of cars that use each design platform, and their locations in the product space. This information is shown for the *HR-V*'s platform (the *Honda Global Small Car Platform*) in light blue in Figure 3. 35 different models in six product segment and seven regions share a design platform with the Honda *HR-V*.[12]

The third feature is a *list of platform adoption years.* To maintain consumer recognition, popular models are redesigned every 5-10 years without changing the model name, segment, or end market. These redesigns – at discrete dates – allow a car model to adopt a new design platform without changing its location in the product space. I leverage these within-product design changes, shown as diamonds in Figure B.2, in my primary event-study empirical strategy.

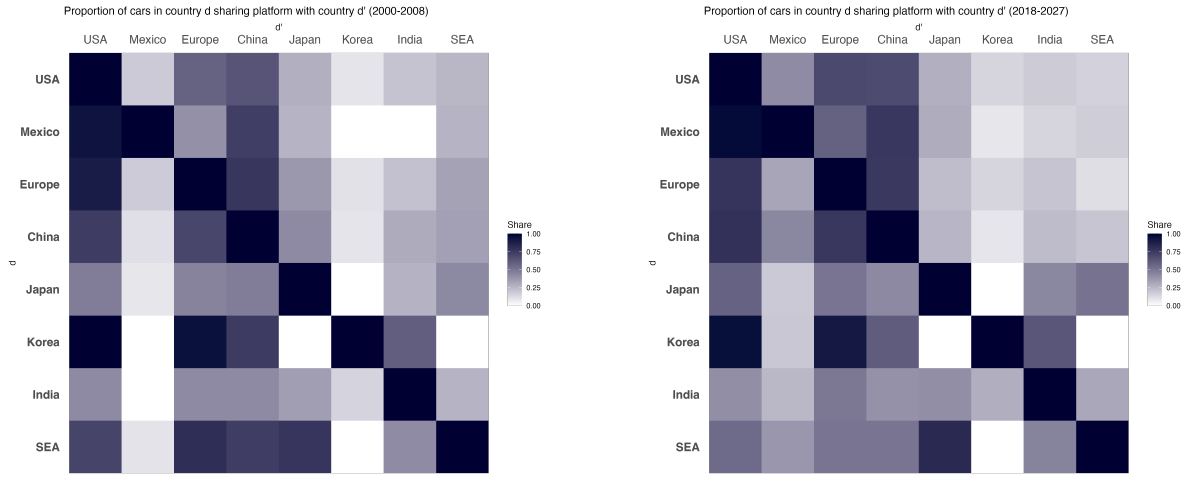
Fact: design platforms are shared in two phases, and will become universal due to EVs. Using the product timeline data, I next document how modular design has diffused across the automotive industry. To quantify platform-sharing across destinations (segments), I calculate the probability that a car model sold in destination (segment) x shares a platform with at least one car sold in destination (segment) x' by the same firm. I describe the full procedure in Appendix Subsection B.3. Repeating this procedure for two time periods (2000-2008 and 2018-2027) with no overlap in platforms, I show the within-firm diffusion of modular design follows certain systematic patterns.

First, *platform-sharing across destinations was common by 2000.* As indicated by the off-diagonal shading in Panel A of Figure 4, by the early 2000s, products sold in different regions already used shared design platforms, with only small increases thru the present.[13]

[9]Further brand differentiation (e.g., *Audi* vs. *VW* vs. *Skoda*) is largely cosmetic and often called *badge engineering*.
[10]All firms use the same size classifications due to size-dependent emissions and safety regulations. The length categories are A (smallest) to F (largest), and the body type (reflecting height) is typically a car (sedan/hatchback) or SUVs. For instance, C-segment SUVs are taller but not longer than C-segment sedans. There are also two niche segments: pickup trucks and MPVs (vans).
[11]See Freund and Oliver (2015); Lamy (2016); Grossman et al. (2021) and Maggi and Mrázová (2024) for discussions of the prevalence and trade policy implications of regional regulatory divergence.
[12]Despite being technologically compatible, these products appear very different to consumers. For instance, as depicted in Appendix Figure A.1, while the *Honda HR-V* is an SUV marketed for North American suburban families, its platform is also used for the *Honda City*, a popular sedan marketed towards urban use in India and Southeast Asia.
[13]While *Marklines* does not report platform data before 2000, historical cases suggests that – prior to the 1990s – design rules were not always shared across destinations. For instance, from the 1930s until the early 2000s, Ford and *GM* ran

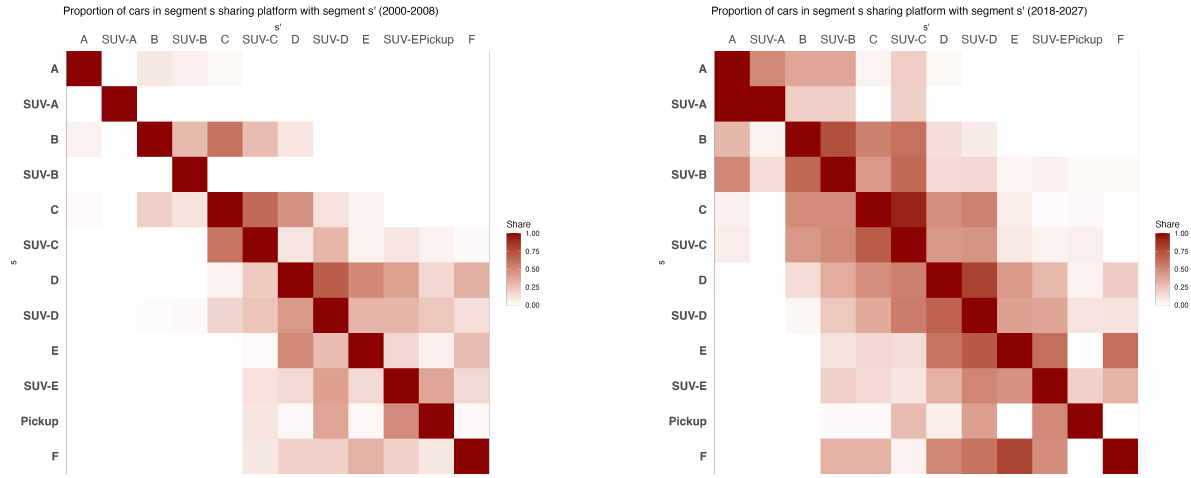
Figure 4: Platform-sharing across destinations and segments, 2000-2008 and 2018-2027

Panel A: across destinations



2000–2008 2018–2027

Panel B: across segments



2000–2008 2018–2027

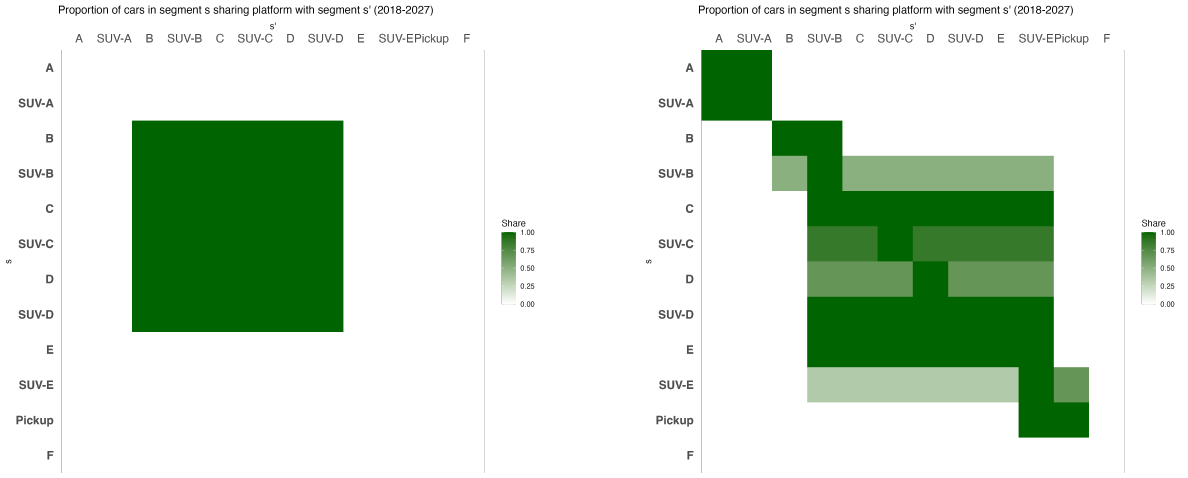
Notes: Panel A shows the probability of platform-sharing across destination regions for two time periods, 2000-2008 and 2018-2027, with no overlap in design platform names. Destinations are ordered by distance to Detroit. Panel A similarly shows the probability of platform-sharing across product segments. Segments are ordered from smallest (A) to largest (F). 2023-2027 data reflect plans for vehicles under development at the time of data download (late 2023). Darker colors indicate that a higher share of car models in each row share a platform with at least one car in the column. See Appendix Subsection B.3 for full procedure. Data from *Marklines Automotive*.

Second, *platform-sharing across product segments has increased significantly since 2000, but only between similar segments.* Panel B of Figure 4 shows that, in both time periods, design rules are largely shared with segments that are nearby (of similar size) in the product space.[14] Furthermore, a comparison of the two time periods in Figure 4 Panel B reveals that platform-sharing across segments has increased since 2000, as indicated by the darker off-diagonal shading in the right

their European units as separate firms with independent design teams (Wilkins and Hill, 1964).

[14] I order segments from smallest (A) to largest (F) by length, with SUVs considered larger than corresponding sedan.

Figure 5: Current plans suggest that electric vehicles will use universal design platforms



BYD EVs only | *General Motors* EVs only

Notes: Figures show patterns of platform-sharing for current and announced electric vehicles only. Left panel subsets to platforms used by *BYD*, and right panel subsets to platforms used by *GM*. Darker colors indicate that a higher share of models in segment (row) s share a platform with at least one car in segment (column) s'. Data from *Marklines Automotive*.

panel (2018-2027) vis-a-vis the left (2000-2008).

Third, looking forward, *future generations of electric vehicles (EVs) will use universal design platforms* shared by all destinations and segments. The left panel of Figure 5 shows that *BYD*, the world's largest EV producer, uses a single platform (the *e-Platform 3.0*) for all cars. The right panel shows that *GM*'s planned EV design platforms are shared across most products, with the exception of the smallest (A- and B-segment cars) and largest vehicles (pickups).[15]

This paper's focus is the geographic effects of each phase of platform-sharing – first across destinations, then across product segments – and is thus agnostic as to *why* these phases emerge. However, for the interested reader, in Appendix Subsection B.4 I argue these phases likely reflect higher costs to designers of accommodating product heterogeneity across segments relative to destinations, and for internal combustion engine (ICE) vehicles relative to EVs.

3.2 Do design platforms change production functions? Which inputs are shared?

Data: the *Hollander Interchange* automotive junkyard parts database. To observe how modular design changes car production functions (i.e. the extent of input-sharing), I merge in parts lists by car model from the *Hollander Interchange*, an independent parts database used by automotive recyclers since 1934. For all car models sold in the United States each year, the *Hollander* database includes a full parts list, with each part assigned a unique ID, allowing recyclers to verify if parts in scrapped cars can be sold as spares for other vehicles. For instance, the database shows that

[15]BYD and GM had announced future EV development plans at the time of data download (fall 2023). Recent announcement suggest other firms will follow: in August 2025, *Ford* CEO Jim Farley announced *Ford* will develop a universal EV platform explicitly to compete with *BYD* and other Chinese EV manufacturers.

the engine assembly for the 2005 *Ford Mustang*, with code *300-09140*, is found in three other *Ford* models. This level of detail – the input variety-level production functions for every product – is exceedingly rare in any industry.

Fact: cars on same platform share internal inputs but remain visibly differentiated. I test which kinds of parts are actually shared within the firm due to common design platforms. To do this, I identify all pairs of US-sold car models (m, m') that adopt a common platform in the *Marklines* data. There are 56 such pairs. I then calculate the part type-wise Jaccard similarity in each year t for the pair:

$$Similarity_{mm'tp} = \frac{\textbf{PartsList}_{mtp} \cap \textbf{PartsList}_{m'tp}}{\textbf{PartsList}_{mtp} \cup \textbf{PartsList}_{m'tp}}$$

where $\textbf{PartsList}_{mtp}$ is the list of type-c part codes for model m in year t. The ratio captures the share of model m's type-p part varieties also used in model m'. This takes value 1 if the two cars m and m' are entirely identical, and value 0 if no common parts are used.

Parts types p fall into two broad categories. The body and interior are *external* inputs directly visible to the consumer. Engines, drivetrains, chassis, and electronics are *internal* inputs: households do not see them, yet they embed 80 percent of value-added and enable the car's movement.[16]

Specification. I then estimate descriptive event studies in which the event is the year $\bar{t}(m, m')$ of adoption of a common platform by models m and m', and the outcome is the similarity of the parts lists in year t.[17] These pairwise regressions are estimated separately by car section s. The regression is:

$$Similarity_{mm'tp} = \alpha_{mm'p} + \gamma_{pt} + \sum_{\tau \neq -1} \beta^{\tau} \mathbb{1}[t - \bar{t}(m, m') = \tau] + \varepsilon_{mm'tp} \quad (1)$$

where α is a fixed effect for the pair and component category and γ is a time fixed effect. Equation (1) leverages *within-model pair* variation in platform-sharing over time, controlling (via $\alpha_{mm'p}$) for the propensity of some model pairs to share parts (of a specific type) in all years. The coefficient of interest β^{τ} for $\tau \geq 0$ therefore captures the increase in input-sharing relative to the year before platform-sharing begins, net of the average increase for all included model pairs over the same period.[18]

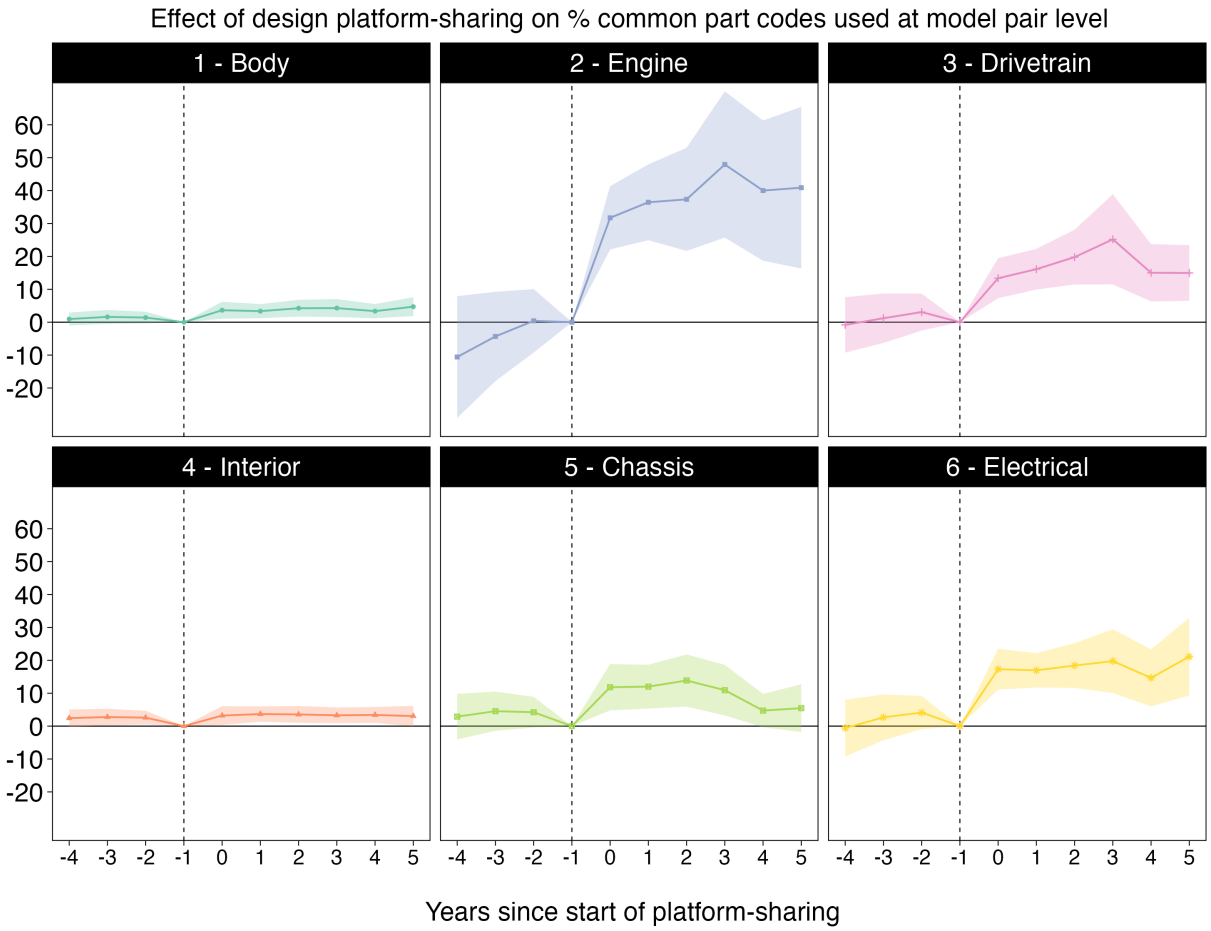
Result: shared design rules associated with higher internal (but not external) part commonality. Event study effects by car section are in Figure 6. The left panels show that model pairs sharing a platform exhibit a small but significant 3 percent increase in parts similarity for the body and interior. This small magnitude is consistent with the fact that car bodies and interiors are external

[16] Parts classified into one of 158 *part types*, which are in turn grouped into six broad *sections*: body, interior, engine, drivetrain, chassis, and electrical. Example part types for each section: bumpers, headlights, windshield, doors (body); steering wheel, seats, dashboard (interior); engine assembly, carburetor, fuel injection parts (engine); transmission, suspension (drivetrain); brakes, wheels, chassis frame (chassis); air conditioner, GPS, motors, battery (electrical).

[17] If the two cars adopt the platform in different years, I use the latter of the two years.

[18] For panel balance, I restrict to events in which the models m and m' are continuously produced 5 years before and 5 years after the adoption date.

Figure 6: Cars start sharing *internal* parts after adopting common platforms



Notes: events are pairwise adoption of common platforms, based on platform names in Marklines and Wards Automotive. Regression is at the model pair-part type-year level. For any model pair and part type, the outcome is the share of part codes used by the first model that the second model also uses. Regressions are estimated separately for six car sections: body, interior, engine, drivetrain, chassis, and electrical parts. Lists of unique part codes come from the Hollander Interchange.

parts, and therefore can be used to visibly differentiate car models for consumers.

Parts-sharing for internal parts increases by much larger magnitudes. When two models m and m' start sharing a platform, the share of m's parts also found in m' increases by 40 percent for engines, 15 percent for the drivetrain, 20 percent for electrical components (including batteries), and 10 percent for chassis parts.

Implication: value-added input demand becoming more homogeneous. The fact that internal (but not external) parts are shared suggests that *the adoption of modular design allows firms to maintain product differentiation while standardizing input demand across countries*. In other words, even if households demand different final products, the plants around the world that assemble those products may in practice demand the same engines, batteries, and other value-added inputs.

3.3 Does input-sharing affect production locations? Data and suggestive evidence

Having verified that shared design platforms leads firms to share internal inputs across products, I now provide suggestive evidence that platform-sharing influences where value-added inputs are *produced*. To do so, I collect three additional datasets that allow me to observe supply and demand for each input.

Input supply data [sourcing panel]: *American Automotive Labeling Act* manufacturer filings. I obtain car model-by-year input sourcing decisions from automakers' annual filings for model years 2005 to 2024 under the *American Automotive Labeling Act* (AALA).[19] Combined AALA data are only publicly available since 2011. I therefore digitize and manually harmonize over 500 original manufacturer filings for 2005 to 2010, which I obtain directly from the US *NHTSA*. A sample filing is shown in Appendix Figure B.3.

Because AALA reporting is mandatory, takes place annually, and covers all US-sold cars, I use these data to estimate the event studies in Section 5, which examine within-product changes in sourcing over time. Specifically, I construct two forms of sourcing outcomes: the share of all parts that are domestically-produced (i.e. from the country of assembly); and the country of origin of two specific parts – the engine and transmission – that together comprise 30 percent of value-added. Appendix Subsection B.5 describes how I construct these variables.

Input supply data [global cross-section]: *Who Supplies Whom* supplier relationships. To observe supply chain decisions for all cars globally, I scrape and clean the *Marklines Who Supplies Whom* (WSW) database, which contains 320,000 unique sourcing relationships, for 300 unique part types, for cars sold in all countries. These relationships are compiled from both supplier self-reports (in the hope of obtaining future business) and teardowns of individual vehicles. To my knowledge, no other industrywide dataset exists – for cars or any other sector – of product-level supplier information for covering all firms in all countries.[20]

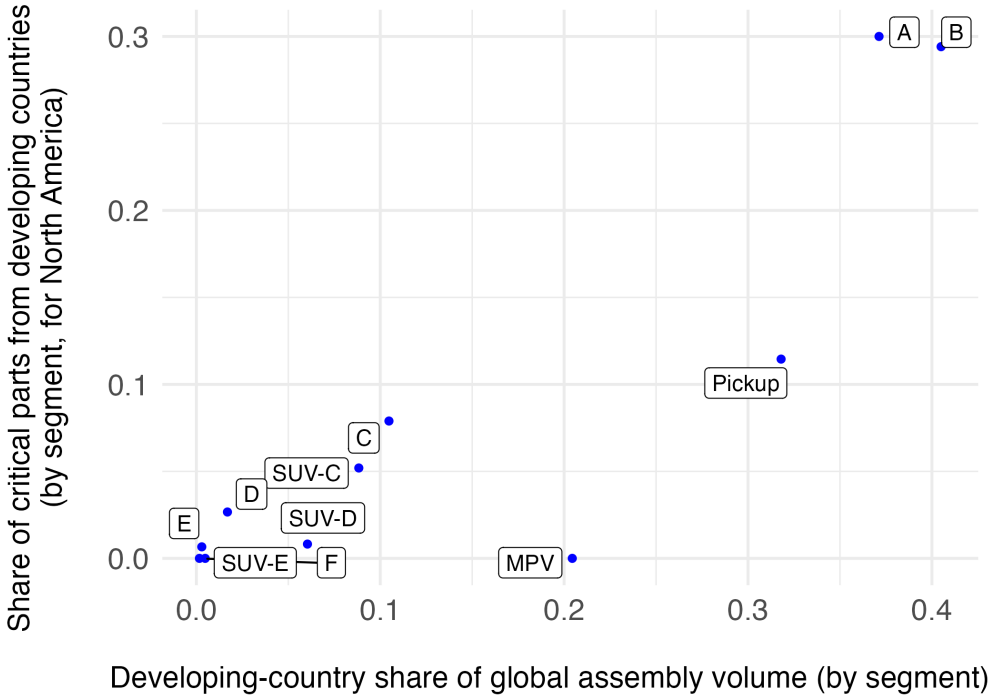
Because the WSW has global coverage, and contains the location of each supplier, I harmonize and aggregate the data to construct input sourcing flows (relationship counts) at the origin-destination-part type level, where the destination is a product (ex: the US-built *Ford Focus*) rather than a country. This is because platform-sharing is a product-level decision. I then use these input flows to study how patterns of cross-country specialization vary with modularity in Section 5, and to estimate the quantitative model in Section 6. Appendix Subsection B.6 describes how I clean and aggregate the WSW data.

Input demand data: *Marklines* product assembly volumes. The supply chain data above allow me to observe where inputs for each platform are produced. To observe where each input is

[19] Head et al. (2024) use a cross-section of the AALA data to study local content requirements. To my knowledge, I am the first to leverage the product-level panel dimension of these data.

[20] Neither the well-known *Factset Revere* database of suppliers to publicly traded US firms, nor the administrative VAT datasets found in many countries, report suppliers separately by product variety within a firm; and supplier information from trade data (for instance for pharmaceuticals in Galdin (2024)) only cover imports to one country.

Figure 7: Developing countries export inputs to North America for locally-assembled segments



Notes: Horizontal axis shows, for each product segment, the share of all cars from 2012-2022 assembled in developing countries. Vertical axis shows, for each product segment, the probability that critical inputs (engines and transmission) for car models sold in the United States and assembled in North America in that product segment were produced in developing countries. Developing countries are all countries except the USMCA countries (USA, Canada, Mexico), the EU, China, Japan, and Korea. Assembly data by segment and country from *Marklines Production* database. Sourcing information for North America-assembled cars from *American Automotive Labeling Act* reports. Segment for North America-assembled cars from *Marklines Product Portfolio* database.

demanded, and in what magnitudes, I combine data on each model's platform (as described earlier, and shown in Figure 3) with information on model-by-country assembly volumes from the *Marklines Production* database. The production database records annual volumes, by car model and country, for all car-assembling locations. These data are constructed from car registrations and direct communication with manufacturers.[21]

I observe assembly volume from 2002 to 2022, which I merge by car model, year, and assembly country to the *Marklines* WSW and timeline datasets. These merges typically require manual inspection because model names often differ across countries and there are no unique identifiers.

Fact: patterns of specialization within the product space mirror design choices. Using the assembly and supply chain data introduced above, I provide a first piece of evidence that input-sharing and production locations may be related: *design platforms and cross-country specialization patterns are split along the same dimension (by segment)*.

Recall from Figure 4 that shows that platform-sharing is common across countries but limited

[21] Unlike the *S&P Global* car registration database used in recent trade papers, including Head and Mayer (2019) and Sabal (2025), assembly volumes are not disaggregated by sale location in these data. This detail is not necessary for my research question since input demand depends on total assembly volume.

across segments. Inputs are thus largely segment-specific. If input production features economies of scale and high transport costs, then countries should specialize in production of inputs that they demand – in other words, in inputs for product segments that they assemble.

To test this using the assembly and supply chain data collected above, I calculate two statistics for each product segment: the share of final cars which are sold in developing countries (from the assembly data), and the share of all inputs used by North America-assembled cars that are produced in developing countries. These respectively capture developing countries' demand for and (revealed) cost advantage in inputs for each product segment.[22]

Figure 7 shows that the three segments in which developing countries have the highest share of total assembly volume – micro (A-segment) and subcompact (B-segment) cars, followed by pickup trucks – are exactly the three segments in which parts of North America-assembled cars are most likely to come from developing countries. In other words, poor countries specialize in locally-assembled (small) product segments. This is exactly the pattern that we would expect in a world with segment-specific inputs, as shown in Figure 4.

Implication: modularity may reshape export opportunities for many countries. Figure 7, combined with the fact that platforms are segment-specific (Figure 4) and enable shared inputs (Figure 6) suggests that current patterns of design platform-sharing may create export opportunities for many countries. For instance, shared platforms across destinations may enable developing countries to export engines for small A- and B-segment cars to other countries. My goal in the model, empirics, and counterfactuals below are to understand (i) if the adoption of modular design til date enables these patterns of specialization, and (ii) how specialization will change as platforms are also shared across product segments (as show in Figure 5) due to the EV transition.

4 Theoretical framework

To study the geographic effects of modularity, I next incorporate modular design into a model of intrafirm trade. To do so, I leverage an insight from design theory: *design platforms function as input markets* with separate supplier plants, economies of scale, and locations of demand (Baldwin and Clark, 2000). The production location consequences of firm design choices then follow from standard models of increasing returns (Krugman, 1980): to increase production scale while reducing trade costs, firms concentrate input production in each platform's largest market.

I choose to model platforms as input "sectors" because standard representations of technologies as cost-shifters poorly characterize product design. In particular, a firm's design technology is a partition: products are grouped into design platforms that use common value-added inputs. Modularity creates coarser partitions in which many distinct products, first within and then across

[22]Because final cars face high tariffs and are customized for each destination, the segments that each country assembles often reflect its domestic preferences (as well as the preferences of nearby export markets for final goods).

product segments, use a single platform. As a result, changes in platforms neither augment a country's factors as in neoclassical models nor change task-specific productivities. Put differently, increased modularity affects whether a car engine is used in the United States, India, or both; but does not directly improve American or Indian supplier productivity in producing that car engine.

4.1 Setup

There are two countries $\{H, F\}$ indexed o for origin and d for destination. The economy consists of many nontraded manufactured products $j = (f, s, d)$, which are differentiated by destination d, segment $s \in \{1...S\}$, and firm $f \in \{1...F\}$; and a freely traded outside good indexed 0.

Preferences. The representative household in d is of size L_d and has preferences:

$$U_d \;=\; Q_{0d} \;+\; \gamma_d \cdot \log\left(\prod_s \left(\sum_f Q_{fsd}^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}\cdot\beta_{sd}}\right) \tag{2}$$

Households maximize utility subject to spending total income $w_d L_d$ taking wages w_d as given. This yields residual demand curves $Q_j = \gamma_d \cdot \beta_{sd} \cdot P_{sd}^{\varepsilon-1} \cdot P_j^{-\varepsilon}$ with segment price index $P_{sd} = (\sum_f P_{fsd}^{1-\varepsilon})^{\frac{1}{1-\varepsilon}}$. In other words, households spend γ_d on manufactures, with conditional shares β_{sd} on segment s and $P_j^{1-\varepsilon}/P_{sd}^{1-\varepsilon}$ on product j.

Technology. We drop subscripts f for readability. Each integrated firm assembles and sells multiple nontraded final products ("cars") j in destinations $d(j)$ using inputs ("engines") that are produced in-house.[23]

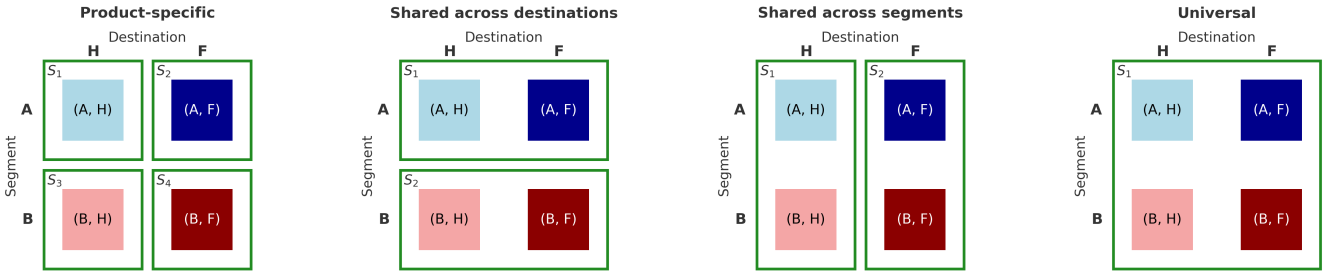
Design technology. To model shared platforms, I assume each product j is exogenously assigned to a platform. Formally, each firm has an exogenous *design technology*, which is a partition $\{S_g\}$ of products $j = (s, d)$ into design platforms (blocks) indexed by g. S_g is the set of products j that use platform g. The total number of platforms is $G = |\{S_g\}|$. For tractability, we restrict the possible $\{S_g\}$ to be a grid in which segments s and destinations d are partitioned independently.[24] For instance, as shown in Figure 8, with two segments $s = \{A, B\}$ there are four possible partitions: platforms can be shared across segments, destinations, both, or neither.

Phases of modularity. Because automotive platform-sharing occurs in two phases as shown in Figure 4, in comparative statics I consider two changes in design technology. The first is from

[23]We technically only require joint optimization: if input production and final assembly were done by separate firms, with rent-sharing via Nash bargaining, we would obtain the same solutions as for the integrated firm. In the automotive industry, value-added input production (especially of engines) is either vertically integrated or conducted within long-term relationships for which Nash bargaining is reasonable.

[24]In other words, certain segments use the same inputs, and certain destinations use the same inputs; and each platform is formed by the intersection of these decisions: $S_g = S_{(g_s,g_d)} = S_{g_s} \times S_{g_d}$. This assumption reflects the fact that differentiation along each dimension typically affects separate modules. Destination-specific differentiation mainly involves regulatory adaptations (e.g., in seatbelts, emissions controls, drive orientation) that leave core technologies unchanged, while segment differentiation requires designers of many core subsystems to coordinate.

Figure 8: Possible partitions (design technologies)



Notes: Each panel shows a possible partition for the case of two segments $\{A, B\}$ and two destinations $\{H, F\}$. Each block S_g is indicated by a green rectangle.

product-specific platforms to platform-sharing across destinations (the first to second panel of Figure 8), which I refer to as *modularity across destinations*. The second is from platform-sharing across destinations to universal platforms (the second to fourth panel), which I refer to as *modularity across segments*.

In Appendix Subsection C.1 I provide a microfoundation for the two phases based on several conversations with automotive engineers. In particular, suppose engineers face coordination costs in product development that increase in within-g product heterogeneity (i.e. in how dissimilar the cars using platform g are from each other). Then, as electrification and computerization enable engineers to coordinate more efficiently over time, automotive firms first share platforms where it is easiest and then expand out: first across destinations within a segment, then between similar segments, and then across all segments.

Final assembly. Conditional on design technology $\{S_g\}$, each platform g functions as a separate input market. Let the function $g = G(j)$ be the index of the block S_g that contains product j. To assemble q_j of final good j for local sale, the firm combines inputs from both countries with CES production function:

$$\forall j: \quad Q_j = \left((q_j^{Hg})^{\frac{\sigma-1}{\sigma}} + (q_j^{Fg})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{where} \quad g = G(j) \tag{3}$$

where q_j^{og} is use of the platform-g origin-o input for the segment-s destination-d output.[25] Importantly, each product j only uses inputs from the single platform g to which it is assigned. We impose $\sigma > \varepsilon > 1$ so that inputs from two origins remain net substitutes, and $\sigma < \frac{1}{\eta}$ to ensure an interior solution in which both countries produce on all platforms (Kucheryavyy et al., 2023).

Final products $j = (f, s, d)$ are nontraded.[26] Due to CES product demand, firms choose prices P_j

[25] For instance, $q_{Small,USA}^{China,1}$ is the quantity of the China-produced variety of the platform-1 input used to assemble small cars in the USA.

[26] I assume nontradability so that domestic preferences directly map to local product demand, which simplifies the exposition. Nontradability of final goods (especially across regions) is reasonable in the automotive sector for several reasons. First, final cars are costly (relative to value) to ship because they are heavy and are too large for standard

to be constant markups $\varepsilon/(\varepsilon-1)$ above (shadow) marginal cost.

Input production and trade. Core inputs for platform g are produced only with labor:

$$y^{og} = ((1-\eta)\cdot\alpha^{og}\cdot l^{og})^{\frac{1}{1-\eta}} \tag{4}$$

where y^{og} is output of design platform g-specific inputs in origin o. l^{og} are labor inputs hired competitively at wage w_o. α^o is origin o's productivity in input production (relative to the outside good) which in the baseline model is constant across platforms.

Because $\eta > 0$, input production features platform-level *increasing returns to scale*. This is because each input needs its own production line, and firms benefit from production line-level learning-by-doing (learning curves) over the lifetime of a platform (see Case et al. (2023) for a full discussion of learning in automotive input production). We use an isoelastic function for consistency with the learning-by-doing literature (Arrow, 1962; Thompson, 2012; Case et al., 2023).[27]

Equation (4) embeds two critical assumptions. First, scale economies are specific to a platform g but not necessarily to a product j. This means that, if platforms are shared, inputs for many products can be produced together with *shared economies of scale*, consistent with the stated motivation for many platform consolidations (Ulrich and Eppinger, 1995). Second, scale economies are origin o-specific, so cost savings are compounded if firms concentrate production in a single location.[28]

Finally, once produced, inputs are exported within-the-firm to assembly locations with symmetric[29] iceberg trade costs $\tau_d^o = \tau_o^d$. For each origin-o and platform-g input producer, the intrafirm resource constraint (analogous to an internal market-clearing condition) is:

$$y^{og} = \sum_{j\in S_g} \tau_{d(j)}^o q_j^{og} \tag{5}$$

i.e. that all inputs are used.

Firm profit-maximization The firm chooses input production $\{y^{og}\}$ and labor use $\{l^{og}\}$ by origin and platform, input use by product $\{q_j^{og}\}$, product volumes $\{Q_j\}$, and product prices $\{P_j\}$ to maximize revenues from final sales net of labor costs for inputs:

$$\max_{\{l^{og}\},\{y^{og}\},\{q_j^{og}\},\{Q_j\},\{P_j\}} \quad \underbrace{\sum_j P_j Q_j}_{\text{product revenues}} - \underbrace{\sum_{o,g} w_o l^{og}}_{\text{input labor costs}} \tag{6}$$

shipping containers. Second, tariffs are substantially higher on final cars than intermediates; for instance, India imposes a 125% tariff on final cars and a 15% tariff on most parts. Third, consumers have preferences for locally-assembled final cars (Coşar et al., 2018) even from foreign brands. Fourth, final assembly plant closures are highly controversial and often resisted by policymakers.

[27] I assume that both countries o (H and F) produce nonzero differentiated inputs (a la Arlington) for all platforms g, regardless of where platform g is used.

[28] If scale economies were entirely global, firms could instead maintain dispersed production facilities near each source of demand without losing scale economies.

[29] The symmetry assumption is used for the analytical results but not the quantitative model and counterfactuals.

subject to design technology $\{S_g\}$, demand from household maximization of (2) (taking competitor prices P_{sd} as given), the assembly production function (3), input production function (4), and internal input market-clearing (5).

Outside good and market clearing I assume that the outside good is produced one-for-one with labor in each country ($y_0^o = l_0^0$) and freely traded. Using the price of the outside good as the numeraire, it follows that domestic labor markets clear in both countries at wage equal to 1.

4.2 Solution

Solving for optimal production. I first derive the optimal input production choices (y^{Hg}, y^{Fg}) for each platform g that maximize (6) conditional on the set of products $\{j \in S_g\}$ using the platform. Appendix Subsection C.2 gives the equations that implicitly define the solution.

Platform-level home market effects. The key property of the solution is that input production for platform g disproportionately locates in countries d that demand products $\{j \in S_g\}$, either because of large total demand for manufactures in d, strong preferences in d for the segments s contained in g, or high local competitor prices for those segments. Formally, for any x, define $\widehat{x} = x^F / x^H$ as the ratio of values in countries F to H. Then:

Lemma 1 (Platform market size effects)**.** *Wlog, F's relative production share $\widehat{y}^g$ is increasing in the ratio of total demand-shifters $\widehat{D}^g$ for products contained in S_g, which equals the ratio of total expenditures when in a single-firm economy ($F = 1$):*

$$\widehat{D}^g \equiv \frac{\overbrace{\sum_{s:\,(f,s,F)\in \mathbf{S}_g}}^{\text{made in } F \text{ using } g} \gamma_F \beta_{sF} (P_{sF})^{\varepsilon-1}}{\underbrace{\sum_{s:\,(f,s,H)\in \mathbf{S}_g}}_{\text{made in } H \text{ using } g} \gamma_H \beta_{sH} (P_{sH})^{\varepsilon-1}} = \Big|_{F=1} \quad \widehat{\gamma} \cdot \widehat{\beta^g} \tag{7}$$

where $\widehat{\beta^g} \equiv \left(\frac{\sum_{s:\,(f,s,F)\in \mathbf{S}_g} \beta_{sF}}{\sum_{s:\,(f,s,H)\in \mathbf{S}_g} \beta_{sH}} \right)$ is F's relative preference (expenditure share) for segments on platform g.

Proof. See Appendix Subsection C.3. □

Lemma 1 shows that design affects production location by shaping relative demand ($\widehat{D}^g$) for platform-compatible inputs. This is because the set of products using platform g, and underlying demand for each product, together determine *which* country is the larger market for platform g's inputs. Input production for each platform then locates in the larger market, because concentrating production enables economies of scale, and producing near demand saves trade costs.

For instance, suppose that two products j, some assembled in India and some in the United States, use a common platform g. Lemma 1 states that, if total Indian assembly volume *on the platform* is

much larger (smaller) than the USA, car engines for g are more (less) likely to be produced in India.

This within-the-firm market size effect resembles the sector-level home market effects that many authors have documented in general-equilibrium settings with constant-elasticity demand and increasing returns (Krugman, 1980; Costinot et al., 2019; Dingel et al., 2023). This similarity is not a coincidence: the optimal solution to (6) is isomorphic to the equilibrium of an internal input market in which input prices $\{p^{og}\}$ (which in the firm's optimization problem are the Lagrangian multipliers on output y^{og}) equate input supply and use. The key difference from standard models is that input market size depends both on final product demand and on how product demand is split across platforms.

4.3 Comparative static effects on input sourcing and patterns of specialization

In this section I show that the production location effects of the two phases of modularity – first across destinations d with a product segment s, then across segments s – have distinct effects on the global geography of input production. In particular, the first phase increases trade and leads countries to specialize in inputs for preferred product segments as in Krugman (1980), creating a segment-level home-market effect. The second phase then creates winner-take-all supply chains as in Helpman and Krugman (1985): large final markets specialize in value-added input production, while small final markets specialize in the constant-returns outside good.

To do so, I derive testable predictions for the effects of the two phases on input production locations for products j. I consider input production locations (i.e. sourcing decisions) for individual products $j = (f, s, d)$ because final products (unlike inputs) are consistently defined before and after changes in platform-sharing, and because my event study empirical design leverages product-level platform adoption events and sourcing changes.

I first consider the effects of sharing a design platform across *destinations* within a product segment.

Proposition 1 (Phase 1: platform-sharing across destinations)**.** *Adoption of modularity across destinations d within each segment s has the following effects on the input production location o for product $j = (f, s, d)$:*

- ***Trade shares increase.*** *For products j made in $d = H$ (without loss of generality), adoption increases imported input shares:*

$$\Delta \frac{q_j^F}{Q_j} \equiv \left.\frac{q_j^F}{Q_j}\right|_{post} - \left.\frac{q_j^F}{Q_j}\right|_{pre} > 0$$

- ***Sourcing shifts to large foreign markets for segment.*** *For products j made in $d = H$ (wlog), the*

change in import shares $\Delta \frac{q_j^F}{Q_j}$ *is increasing in F's relative demand*[30] *for segment* $s(j)$*:*

$$\widehat{D^{g(j)}}\Big|_{post} = \widehat{\gamma} \cdot \widehat{\beta_{s(j)}} \cdot \widehat{P_{s(j)}}$$

i.e. in F's relative expenditure on manufactures $\widehat{\gamma}$*, preferences for segment* $\widehat{\beta_{s(j)}}$*, and competitor prices* $\widehat{P_{s(j)}}$ *at firms* $f' \neq f$*.*

Proof. See Appendix Subsection C.4. □

As shown in Proposition 1, platform-sharing across destinations reshapes production and trade in a manner typically associated with changes in trade costs. First, *trade increases*, in that imported content shares rise for all products j *holding trade costs fixed*. Second, *production is offshored to large markets*: specifically, j's value-added inputs are sourced from relatively large markets for j's platform $g(j)$, and thus for segment $s(j)$. Because countries demand different segments, and platforms are segment-specific, either overall expenditure on manufactures (γ) or preferences for j's segment $s(j)$ can generate a relative market size advantage. In other words, modularity across destinations may lead small American cars to use India-made components in place of American-made ones, either because India assembles many cars, or because Indian consumers prefer small cars.

These effects stem from the way that relative demand on j's platform (formally, the ratio of demand-shifters $\widehat{D^{g(j)}}$ in Lemma 1) is affected by product j's adoption of a platform that is shared with the other destination $d' \neq d$. In particular, when design platforms are not shared, products have unique inputs, and thus there is no demand in F (formally, relative demand in F ($\widehat{D^{g(j)}}$) is *zero*) for any inputs used by products in H. From Lemma 1, a world without design platforms thus features substantial home bias: if j is built in H, then to save trade costs, input production for platform $g(j)$ will concentrate in H because there is no other market for those inputs. For example, if an American *Ford* car is built in Detroit, and uses its own platform (and thus inputs), a larger share of that car's inputs will be produced in Detroit to save trade costs to the sole user.

Platform-sharing across destinations increases trade because each input is demanded in both locations. Formally, for a product j built in H, relative demand in F $\widehat{D^{g(j)}}$ must be positive. With nonzero use of type-g inputs in F, the firm has some incentives to produce in F to save trade costs to assembly locations for other products j', unambiguously increasing the imported input share for product j. The intensity of offshoring scales with relative foreign market size: incentives to save trade costs to F – by producing there – are especially strong when relative demand in F for type-g inputs ($\widehat{D^{g(j)}}$) is large. For instance, the import share for a Detroit-built Ford car j will increase more from platform-sharing with China than with Rwanda, because China's larger assembly volumes $D^g_{\text{China}}/D^g_{\text{USA}}$ imply larger transport cost savings from concentrating production there.

[30] I use the term relative demand to refer to all shifters of quantity demanded beyond price.

Importantly, because platforms remain segment s-specific and the change in $\widehat{D^{g(j)}}$ is heterogeneous by segment $s(j)$, countries can still produce and export inputs for different segments. As a result, the first phase of platform-sharing across destinations need not concentrate aggregate input production. For instance, because *Ford* builds few small cars and many large SUVs in the United States, *Ford* may concentrate production of inputs for large SUV segments in the USA (including for export), while concentrating production of small car inputs in (for instance) India.[31]

I next show that the second phase of modularity, in which design platforms are shared across product segments s and s', eliminates these segment specialization opportunities.

Proposition 2 (Phase 2: platform-sharing across product segments)**.** *Assume within-segment modularity is already adopted. Then the adoption of cross-segment modularity has the following effects on the input production location o for product $j = (f, s, d)$, and on country H's cost advantage in inputs for product j relative to j'* $\left(CA^{HF}_{jj'} = \frac{\widehat{p^{g(j)}}}{\widehat{p^{g(j')}}}\right)$*:*

- ***Production of inputs for a segment shifts away from countries that prefer that segment.*** *Define demand for segment s in destination d as $D_{sd} \equiv \gamma_d \cdot \beta_{sd} P_{sd}^{\varepsilon-1}$ so that the share of segment s in destination d total demand is $\omega_{sd} = \frac{D_{sd}}{\sum_{s'} D_{s'd}}$ and relative demand for segment s (in destination F vs. H) is $\widehat{D_s} = \frac{D_{sF}}{D_{sH}}$. Then, for product j in segment s built in country H (wlog), adoption of modular design across segments decreases the probability of sourcing from F if and only if F has stronger-than-average demand for segment s, i.e.:*

$$\Delta \frac{q_j^F}{Q_j} < 0 \quad \textit{iff} \quad \widehat{D_s} > \sum_{s'} \omega_{s'H} \widehat{D_{s'}}$$

  Furthermore, as $\varepsilon \to 1$, then sourcing from F only decreases in segments it relatively prefers, i.e.:

$$\Delta \frac{q_j^F}{Q_j} < 0 \quad \textit{iff} \quad \beta_{sF} > \beta_{sH}$$

- ***Cross-country specialization in inputs for preferred segments s ceases to exist.*** *Consider two products in different segments $j = (f, s, d)$ and $j' = (f, s', d)$. Before (after) adoption, H's cost advantage in j is increasing in (unrelated to) H's relative preference for s, i.e.:*

$$\left.\frac{\partial\left(CA^{HF}_{jj'}\right)}{\partial\left(\frac{\beta_{s(j)H}}{\beta_{s(j)F}}\right)}\right|_{pre} > 0 \quad , \quad \left.\frac{\partial\left(CA^{HF}_{jj'}\right)}{\partial\left(\frac{\beta_{s(j)H}}{\beta_{s(j)F}}\right)}\right|_{post} = 0$$

[31] These results still go through if countries also have pre-existing productivity differences (along with demand differences) inputs for specific segments s (ex: in engines for cheap cars). Proposition 1 still holds as long as these skill advantages do not depend on the set of countries that demand each input (e.g., India can be more productive in small car inputs relative to large car inputs, but not in small car inputs for the USA relative to small car inputs for Europe).

Proof. See Appendix Subsection C.5. □

Proposition 2 shows that modular design across segments eliminates production opportunities in inputs for preferred (and thus locally-assembled) product segments. This restriction on how countries specialize emerges because, with common inputs across segments, the two key determinant of relative input prices $\widehat{p^{g(j)}}$ – countries' relative demand $\{\widehat{D^g}\}$ on each design platform – become constant within a firm. Specifically, following Lemma 1, relative demand is fixed at $\widehat{D^1} = \sum_s \gamma_F \beta_{sF} P_{sF}^{\varepsilon-1} / \sum_s \gamma_H \beta_{sH} P_{sH}^{\varepsilon-1}$ for all cars $\{j\}$ of firm f.

For instance, suppose India strongly prefers small cars (i.e. $\beta_{\text{small,India}}$ is large). Without platform-sharing across segments, this preference would make India a large assembler of small cars, thus a large market for small-car inputs, and thus a large producer of small-car inputs (due to cost advantages through scale). However, with platform-sharing across segments, India only produces inputs for any car if it is the cheapest or largest *overall* market for that firm's inputs.

In other words, Proposition 2 shows that platform-sharing across segments induces a *de facto* reduction in product variety: countries no longer specialize in technology-intensive inputs for different product segments because all segments use common inputs. In Section 6, I show in counterfactuals that this reduction in differentiation leads aggregate input production to concentrate in fewer countries.[32]

Parametrized example. To build intuition for how the two phases of platform-sharing change relative market size, Figure C.1 in Appendix C shows a special case of the model with a single firm ($F = 1$), two segments (A and B), no skill differences ($\alpha^H = \alpha^F$), and selected values for product demand, accompanied with a full explanation in Appendix Subsection C.6.

Interactions with trade and industrial policies. While modular design is a technological rather than policy change, it nevertheless shapes the effects of trade and industrial policies in several ways. First, *modular design across destinations increases the gains from trade* because it increases baseline trade shares holding fixed trade costs (Arkolakis et al., 2012).

Second, *modular design across destinations enables 'import protection as export promotion'.* When domestic and foreign products use separate platforms (and thus separate inputs), import tariffs create cost advantages in inputs for which there is no demand abroad. In contrast, with shared platforms across destinations, the same inputs are both used domestically and exported, so import tariffs can create a scale advantage (and thus a cost advantage) in inputs that the country exports.

Third, *modular design across segments makes it harder for industrial policy to attract some input production, but easier to capture all production.* Specifically, platform-sharing across segments increases

[32]The change in Proposition 2 intuitively resembles a reduction the variance of idiosyncratic productivities in Eaton and Kortum (2002)-style mode or a shortening of an industry's quality ladder (Khandelwal, 2010). The key differences are that in this model, (i) input variety is a technology – that can change over time – with direct effects on trade, not an exogenous industry characteristic to be interacted with other shocks; and (ii) because of increasing returns, the effects of reduced input variety on production location depend on the details of where each input is demanded.

(decreases) the production subsidy required to create a cost advantage in inputs for at least one product (for all products) relative to the outside good.

4.4 Taking stock

The model predicts that the two phases of modular design adoption – first across *destinations* and then *product segments* – reshape global trade in distinct ways. In the first phase (Proposition 1), trade increases and production shifts to large foreign markets for each segment, allowing many countries to simultaneously produce and export in preferred segments. In the second phase (Proposition 2), segment specialization weakens and countries specialize in inputs for an entire firm or in the outside good. These effects are driven by changes in platform market size (Lemma 1) holding underlying productivities and trade costs fixed.

5 Empirical effects of modularity on production location

In this section, I construct empirical tests for the changes in production location predicted by Proposition 1 and Proposition 2: increased foreign sourcing, especially from a segment's largest foreign markets; followed by a weakening of segment-specific specialization patterns.

The ideal experiment to test these predictions is the random assignment of design platforms to products, within otherwise-identical firms. I approximate this ideal experiment using three complementary research designs: (i) event studies of staggered rollouts of new design platforms; (ii) changes in platform scale induced by differential exposure to mergers; and (iii) gravity regressions in which the strength of countries' segment-specific specialization are allowed to vary with modular design.

5.1 Does modularity increase foreign sourcing?

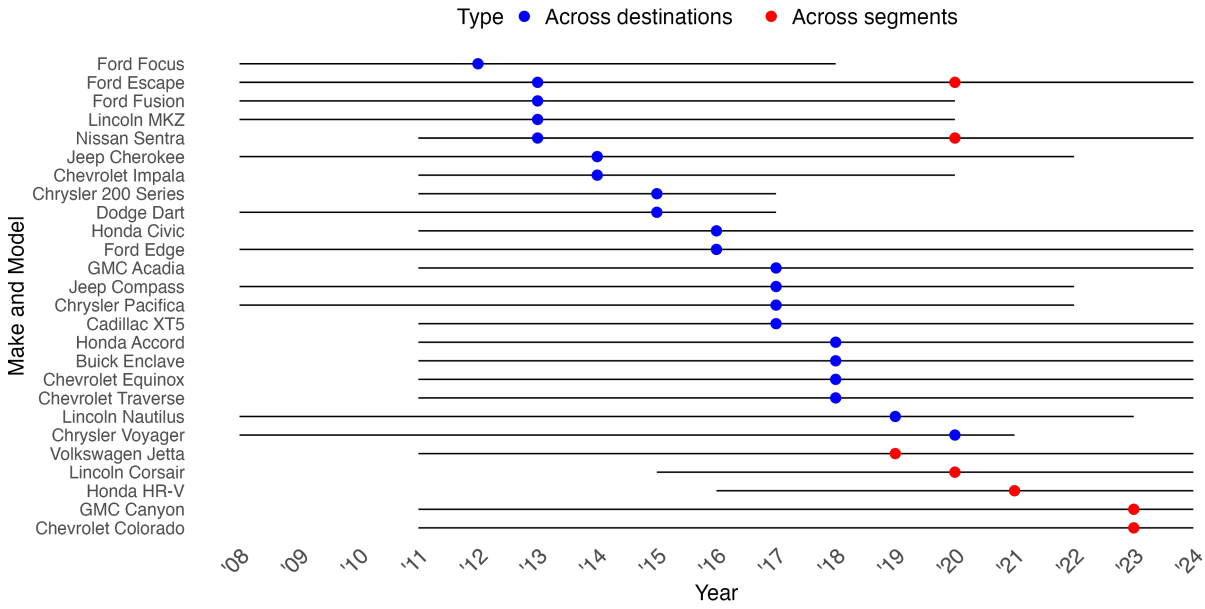
I first test if the adoption of modularity across destinations by an individual product j in year t changes the share of j's parts that are imported, as predicted by Proposition 1. To to do so, I construct a panel of design and sourcing choices for each car model j and year t, restricting to North America-assembled vehicles for which such information is consistently available.[33]

Events: product redesigns to platforms shared across destinations. I aim to estimate the input sourcing effects of shared design platforms holding product characteristics fixed.[34] To do so, I leverage that – because car models are culturally salient products with established individual

[33] I treat models assembled in different countries (in the same year or over time) as separate panel units because assembly location changes will affect sourcing. For instance, the Chevrolet Silverado is assembled in both the USA and Mexico, with different levels of US-originating content.

[34] In contrast, many researchers use similarity in characteristics as a direct measure of design similarity (e.g., Argente et al. (2025)).

Figure 9: Modular design platform adoption dates



Notes: Figure shows all events in which an existing North America-assembled model is redesigned, and the new design platform (but not the old one) is shared with models assembled in another region (in blue) or product segment (in red). Black lines indicate years in which each model is assembled. Event years and redesign type are from *Wards* and *Marklines Automotive*. Sourcing information is from *American Automotive Labeling Act* reports for 2008-2024.

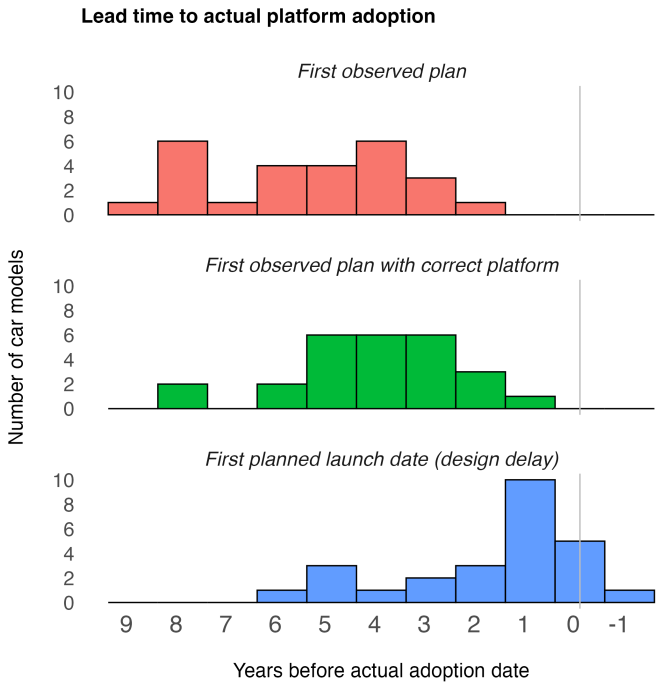
reputations – automotive firms redesign car models at regular intervals, changing the car's platform while keeping basic product characteristics fixed for consumers.[35]

Using the platform timeline data, I identify events in which a car model first adopts a platform that is shared with car models in other destinations or segments (see Appendix Subsection D.1 for the full procedure). The set of identified events, and panel data coverage around each event, are shown in Figure 9. There are 20 events (in blue) at which a car model j first shares platforms with (models j' in) other destinations, facilitating testing of Proposition 1. These events are from 11 unique makes (brands) of 5 unique firms. I also identify 7 additional events (in red) for which a model's platform is first shared across product segments, which I return to in Section 5.3 to test Proposition 2. Finally, I identify 43 *placebo redesign events*, which take place during the sample period and inside the same firms, but *do not change* the set of segments and destinations with which model j shares a platform. These events are shown in Appendix Figure A.2.

Empirical design: rollouts of pre-existing design plans. Event studies of within-product redesign events face two sources of endogeneity. First, redesign choices and timing may be endogenous to changes in countries' production costs: products may only be redesigned when an opportunity arises to source cheaper foreign inputs, in response to competition from a foreign firm, or in

[35]For instance, the *Toyota Corolla* has been redesigned eight times since its launch in the U.S. in 1983, each time remaining in the same segment (a C-segment sedan). Its most recent redesign, in 2018, led the American *Corolla*, Japanese *Corolla*, and American *Toyota Prius* to share a platform for the first time.

Figure 10: Redesigns to more modular platforms are planned several years in advance



Notes: Top panel shows histogram of years from the date in which redesign intent first appears for model j (in a *Wards* archival redesign plan) to the realized redesign date for model j. Middle panel shows histogram of years from the date that the name of the eventually-adopted platform first appears for model j to the realized redesign date for model j. Bottom panel shows histogram of years from the listed planned redesign date (in the plan in which the redesign is first observed) to the realized redesign date for model j. All events from Figure 9 are included. Data from January 2008, January 2011, January 2014, January 2022, and January 2023 versions of firm-specific product redesign plans purchased from *Wards Automotive*.

response to aggregate demand shocks.[36] Second, the decision to increase platform-sharing at redesign time (rather than leaving platform scope unchanged) is a strategic firm decision that may respond to sudden changes in trade barriers or production costs. For example, in response to a tariff reform, the firm may immediately increase platform-sharing across destination regions because concentrating input production in one region (and then exporting) is less costly.

To overcome these issues, I exploit new historical data on automotive firms' product redesign plans. Specifically, I obtain archives of firm-specific design plan as they existed in 2008, 2011, and 2014 from *Wards Automotive*. These design plans project both the *timing* and *intended platform* of all of a firm's redesigns in the following 5-7 years.[37]

As shown in Figure 10, these plans reveal that automotive design occurs in a manner similar to government policy rollouts commonly exploited in economics: decisions to adopt modular design are made several years in advance, and then implementation is staggered within each firm due to engineer capacity constraints. Specifically, the top two panels of Figure 10 show that – across all realized events – event dates are first planned a median of five years in advance, and the scope of

[36]For instance, many product redesigns were paused during the 2008 recession and 2020 COVID lockdowns.
[37]For example, the Ford Fusion, which was redesigned in 2012 as shown in Figure 9, appears in both the 2008 and 2011 plans with a planned redesign year of 2012 and a planned platform (the *Ford Global C1 platform*, which is a "global" platform shared by products in all major *Ford* assembly locations.)

Table 1: Summary statistics for AALA trade outcomes (out of 100)

	% Import	% USMCA	% from home		% from Europe		% from dev ctry	
			Engine	Trans	Engine	Trans	Engine	Trans
Treated	37.3	79.6	62.9	68.3	12.0	5.4	0.4	0.9
	(15.2)	(14.0)						
Untreated	45.4	68.8	58.2	44.5	13.7	15.4	1.6	2.9
	(14.8)	(18.6)						

Notes: data from *Marklines Automotive*. Treated units are events where car is redesigned to more modular platform as shown in Figure 9. All proportions out of 100. Standard errors shown in parentheses. Developing countries are all countries excluding the USA, Canada, Mexico, the EU, China, Korea, Japan, and Australia. Trans is short for transmission. Averages are unweighted across car models. Source is the 2007-2024 American Automotive Labeling Act Reports.

each new platform (e.g., whether the platform covers multiple destinations, segments, both or neither) is announced a median of four years in advance.[38] These long lead times suggest that neither the choice to redesign a a model j nor the platform that model j adopts due to the redesign are endogenous to instantaneous changes in production or trade costs, because automotive firms cannot instantly redesign products in response to shocks.[39] Finally, the bottom panel shows that archival plans include intended redesign dates (henceforth *planned adoption dates*), and that redesigns are difficult to accelerate relative to these initial plans. Specifically, while 11 of 27 redesign events occur three or more years after their first planned launch date, only one event occurs ahead of schedule.[40]

The information in Figure 10 suggest that modularity adoption events are not endogenous to short-run shocks that occur after redesign plans are made. The remaining confounders, addressed below, stem from the potential endogeneity of *planned* redesign dates and platform choices.

Summary statistics. Table 1 shows average import shares and critical input sourcing locations separately for treated (adopting modularity across destinations) and untreated units. Treated cars have lower import shares in levels. This fact suggests that adopters are not ex-ante more likely to source inputs from abroad. I control for these level differences via fixed effects α_j in all regressions.

Differences-in-differences event study specification. I now present the event-study specification. For car model j in year t produced by firm f in assembly country a,[41] I estimate:

$$ImportShare_{jt} = \alpha_j + \gamma_{fat} + \sum_{\tau \neq -1} \beta^{\tau} \mathbb{1}[t - \bar{t}_j^{new} = \tau] + \varepsilon_{jt} \tag{8}$$

[38] As discussed in (Cusumano and Nobeoka, 1998), platform development takes 2-3 years and model development takes a further 1-2 years. These lead times reflect the complexity of coordinating engineers that design different subsystems.
[39] For instance, the events in Figure 9 are unlikely to be contemporaneous responses a tariff reduction, because that tariff reduction would have to be anticipated four years beforehand.
[40] A substantial management literature argues that these delays are due to coordination failures that inhibit engineers from reaching launch deadlines, and are especially pronounced among American automakers that comprise most of the events in my sample (Cusumano and Nobeoka, 1998).
[41] Either USA/Canada or Mexico.

where $ImportShare_{jt}$ is the share of model j's parts that are imported. α_j captures time-invariant characteristics of a car associated with its sourcing decisions. γ_{fat} are firm-assembly location-year trends that control for any trends that affect all models within a firm, including changes in firm strategy (e.g., Ford engaging more with China), trade costs (e.g., US tariffs on China), and relative productivity (e.g., Chinese suppliers getting cheaper). The event time τ is the time elapsed since the year ($\bar{t}_j^{new}$) in which product j adopts a new platform.

The coefficients of interests $\{\beta^\tau\}$ capture the change in import shares for product j relative to the year before design platform adoption ($\bar{t}_j^{new} - 1$) net of the average change in sourcing over that time period for all other products j' within the same firm. For instance, effects on the United States-assembled *Ford Escape*, which adopts modularity across destinations in 2012, are obtained by comparing its sourcing behavior only to other Ford vehicles assembled in the USA and Canada (but not Mexico) during the same time period.[42]

I estimate (8) both via TWFE (which efficiently estimates group-specific trends γ_{fatp}) and several alternative difference-in-difference methods (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Gardner, 2022; Borusyak et al., 2024) that are robust to bias from staggered treatment. These biases are small in practice because, within each firm, each car model that adopts modular design in Figure 9 is compared to a much larger pool of never-treated car models. I cluster by car model j in all specifications since treatment assignment (adoption of a design platform shared across destinations) occurs at that level.

The key identification assumption is that, in absence of treatment (i.e. of platform-sharing across destinations), adopters and non-adopters should exhibit parallel trends in input sourcing. Since design technology adoption is in general endogenous, to rule out that treatment effects reflect immediate responses to product-specific shocks (e.g., tariffs, management changes, or competition), *I confirm that all events in my sample are pre-planned*, appearing in firm design plans with the correct platform name at least two years before treatment (as shown in Figure 10).

I take several additional steps to rule out specific classes of potential confounders. First, to control for any firmwide trends in sourcing behavior as well as destination-wide shocks such as input tariffs, *I include firm-assembly location-time fixed effects*, thus only comparing car models within the same firm and location. Second, to rule out firmwide sourcing changes that are realized at redesign time, *I estimate effects of placebo events* in which redesigns do occur, but the extent of modularity does not change (see Appendix Figure A.2).[43] Third, to rule out that planned dates

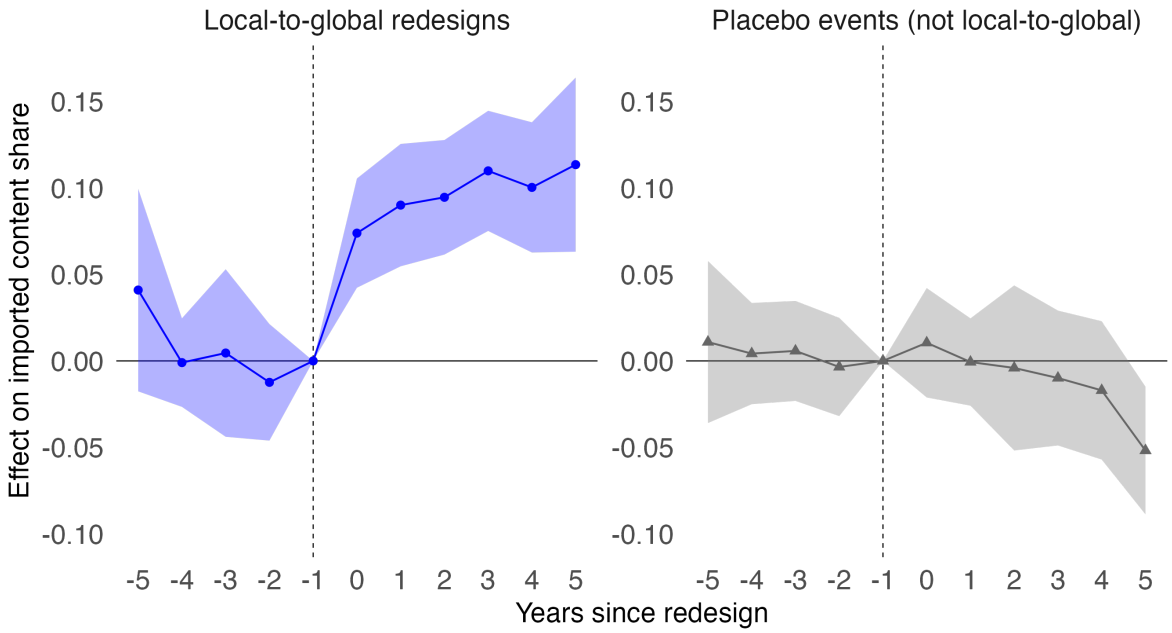
[42]To test for increases in foreign sourcing specifically for value-added inputs (engines and transmissions), I estimate a similar regression to (8):

$$Imported_{jtp} = \alpha_{jp} + \gamma_{fatp} + \sum_{\tau \neq -1} \beta^\tau \mathbb{1}[t - \bar{t}_j^{new} = \tau] + \varepsilon_{jtp} \tag{9}$$

i.e. in which unit and firm-assembly-time fixed effects are part-specific, and $Imported_{jtp}$ takes value 1 if product j sources part $p \in \{\text{engine, transmission}\}$ from abroad.

[43]For instance, Ford might implement plans to offshore input production to Mexico (for all products) at major redesign dates for reasons unrelated to platform-sharing

Figure 11: Modularity across destinations increases import shares by 11 p.p. (24%)



Notes: Left panel shows effect of adopting a platform that is modular across countries on the import share at the car model-year level. Right panel shows effect of adopting a new platform (i.e. a major redesign) that does *not* lead to platform-sharing with cars in another region or segment. All effects at the model-year level. Import share data from the American Automotive Labeling Act. Redesign year, platform information, and locations of platform use from *Marklines*.

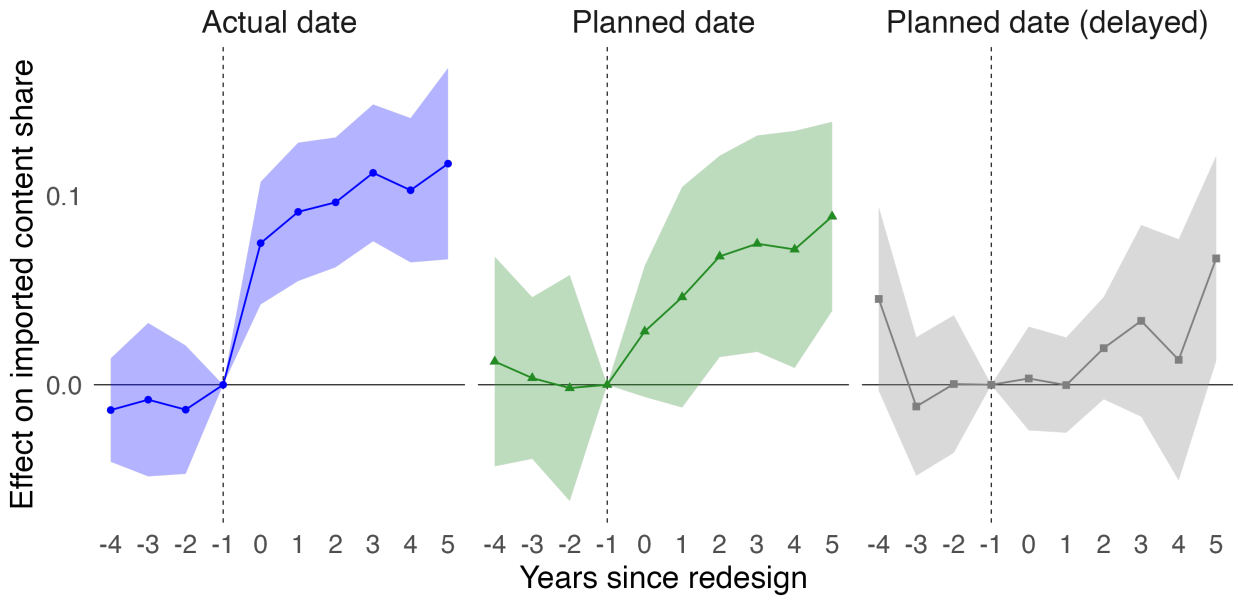
are set in anticipation of cost changes, *I verify that planned-but-delayed adoption events in the bottom panel of Figure 10 have no effect.* Fourth, to avoid any bias from delays relative to plan, which may reflect strategic firm decisions, *I re-estimate equation* (8) *using planned in place of actual dates*, and verify that effects persist.

Hypothesis. Proposition 1 predicts that *platform-sharing across destinations increases import shares* (i.e. that $\beta^\tau > 0$ for $\tau \geq 0$) because foreign demand for model j's inputs increases, so production of some inputs concentrates in foreign locations $d' \neq d(j)$ rather than near j's assembly plant.

Results. I first estimate equation (8), and plot event-study estimates in the left panel of Figure 11.

I find that product-level import shares increase substantially for car models that adopt modularity across destinations, and therefore share inputs with cars assembled in other destinations. The effects are economically significant: an 11 p.p. increase in the import share is 55% of the aggregate increase in U.S. imported auto parts content from 1997-2016 (Klier and Rubenstein, 2019). They also are consistent with Proposition 1, which suggests that import shares should increase because (as tested in the next subsection) foreign relative market size has increased from zero. Pre-trends are zero and statistically insignificant, consistent with the idea that the timing of new design platforms is staggered due to engineer time constraints rather than in response to pre-existing trend in foreign sourcing costs.

Figure 12: Effects around planned adoption dates



Notes: Panels show effect of adopting a platform that is modular across countries on the import share at the car model-year level. Left panel shows uses realized event dates. Middle panel uses planned event dates. Middle panel uses the first planned event date that appears in the parent firm's product timeline. Right panel uses planned dates, restricting to events in which the redesign occurred 3 or more years after of the first planned date. Import share data are from the American Automotive Labeling Act. Redesign year (planned and actual), platform information, and locations of platform use are from *Marklines*.

Importantly, in the right panel, I find no effect of placebo events in which products are redesigned *without* changes in the extent of platform-sharing (see Appendix Figure A.2 for a list of events). This suggests that I am estimating the effects of redesigns to platforms that are modular across destinations, not of redesigns more generally during this time period.

In Appendix Figure A.3 I show that the two-way fixed effects results in Figure 11 are robust to the use of several recent staggered difference-in-difference methods (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Gardner, 2022; Borusyak et al., 2024). For consistent comparison across these methods, in Figure A.3 I use a standard model-by-year panel, without firm-assembly-specific trends. Effects are broadly similar in magnitude and similarly display no pre-trends.

In Figure 12, I next rule out two adoption timing-related concerns. First, the effects in Figure 11 use actual redesign dates, and thus may be endogenous to shocks (e.g., tariffs, competitor product launches, etc.) correlated with trends in relative foreign sourcing costs for treated products. As a result, in the middle panel I estimate equation (8) using *planned dates* from archival redesign plans. While effects are smaller in magnitude (because some launches are delayed), increases in foreign sourcing remain positive and significant.

Second, *planned* adoption timings themselves may be endogenous to future cost shocks, for instance, if firms can forecast future foreign input cost reductions (e.g., due to patent expiry) and time adoption accordingly. The results in the right panel of Figure 12 are inconsistent with such a

threat: planned-but-delayed events have null effects around the planned event date, suggesting that the realized adoption of a modular design platform – not its plan – drive the results in Figure 11.

I conduct several additional robustness checks in Appendix Figure A.4. To ensure that my results speak to industry-wide trends, Panel A shows that production volume-weighted event studies have similar effects. To verify that modularity increases imports specifically of R&D-intensive parts, which are disproportionately shared on common platforms as shown in Figure 6, Panel B shows that adoption also increases the probability of importing engines and transmissions. Panel C shows that platform-sharing across countries increases the share of parts from outside North America, ruling out that foreign sourcing reflects US offshoring to Mexico (either secularly or in response to the USMCA). Finally, Panel D shows that effects are robust to restricting to firms with US headquarters (Ford and GM). This rules out an alternative mechanism: that global platforms simply reflect the centralization of research and production at the headquarters location, rather than (as I model in Section 4) changes in relative input demand or supplier productivities.

5.2 Does modularity shift production to large markets for shared inputs? How large are the scale economies?

I next examine if the increases in offshoring observed in Figure 11 specifically reflect production concentration due to platform-specific economies of scale. To do so, I test if products that adopt modularity across destinations shift sourcing to the shared platform's largest markets, as predicted by Proposition 1.[44]

To do so, I construct a model-year-country (j, t, o) panel of sourcing decisions for two critical parts: the engine and transmission. With this unit of analysis, input supply decisions can be represented by indicators $Source^{o}_{jtp} = \mathbb{1}[Origin_{jtp} = o]$ for sourcing from country o.[45]

To observe changes in input demand by location, I merge in yearly assembly volumes $Q_{j,t}$ from *Marklines* for all products globally, not just those in North America, and sum to obtain total assembly volumes by country o on product j's assigned platform $g(j,t)$:

$$PlatVol^{o}_{jt} = PlatVol^{o,g(j,t)} = \sum_{j' \in S_{g(j,t)}} \mathbb{1}[d(j') = o] \cdot Q_{j,t} \tag{10}$$

For instance, if a car j is built in the USA but on a platform also used to build cars in Germany, then $PlatVol^{\text{Germany}}_{jt} > 0$ even though $d(j) = \text{USA}$. In words, Germany can have positive input demand for car j's inputs without building car j. Importantly, I use product volumes $Q_{j,t}$ for the

[44]Other mechanisms would not necessarily induce substitution to a platform's largest markets. For instance, if modularization leads to better codification of production specifications as in Juhász and Steinwender (2018), and therefore easier offshoring, production may shift to low-cost outsourcing locations rather than large markets for inputs.
[45]There are 31 countries (any that appear for at least two different cars) and a rest-of-world category.

five years after platform launch[46] rather than over a single year, since platform-level increasing returns may accrue from cumulative platform production through learning-by-doing, and firms internalize these gains upfront.

Multi-country event study specification. Since I am testing for platform-level market size effect – in other words, a (weak) platform home market effect – I begin with a standard gravity regression in which sourcing choices are regressed on some shifter of platform-origin-level input demand (here equivalent to assembly volume), as in Costinot et al. (2019):[47]

$$Source^{o}_{jtp} = \alpha^{o}_{jp} + \gamma^{o}_{fatp} + \beta \cdot f\left(PlatVol^{o,g(j,t)}\right) + \varepsilon^{o}_{jtp} \tag{11}$$

where f is some function of platform market size. $Source^{o}_{jtp}$ is an indicator for product j sourcing part p from origin o in year t. In Appendix D I show that Equation (11) can be formally derived from the Section 4 theoretical framework. Since platform market size $PlatVol^{o,g(j,t)}$ only changes when cars are redesigned, for all cars j, we can define the change in (any function of) platform market size:

$$\Delta f\left(PlatVol^{o}_{j}\right) = \begin{cases} f\left(PlatVol^{o,g(j,t^{NEW}_{j})}\right) - f\left(PlatVol^{o,g(j,t^{NEW}_{j}-1)}\right) & \text{if } j \text{ is redesigned} \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

We can thus rewrite (11) as an event study in which the time since new platform adoption is interacted with the change in platform market size in *each country* due to the redesign:

$$Source^{o}_{jtp} = \alpha^{o}_{jp} + \gamma^{o}_{fatp} + \sum_{\tau \neq -1} \beta^{\tau} \left(\underbrace{\mathbb{1}[t - \bar{t}^{new}_{j} = \tau]}_{\text{time since redesign}} \cdot \underbrace{\Delta f\left(PlatVol^{o}_{j}\right)}_{\Delta \text{ mkt size in } o \text{ post-redesign}} \right) + \varepsilon^{o}_{jtp} \tag{13}$$

for car j, year t, origin o, part category p (engine or transmission), firm f, and assembly location a. In (13), α^{o}_{jp} are now product-part-origin fixed effects, capturing time-invariant factors that would, for instance, lead Ford to source engines for the *Fusion* from China. Meanwhile, γ^{o}_{fatp} are firm-assembly-year-part-origin fixed effects, capturing any trends in sourcing from China over time that all to all Ford US-built vehicle models.

The coefficients of interest $\{\beta^{\tau}\}$ now capture the effect of a unit increase in some function $f\left(\Delta PlatVol^{o}_{j}\right)$ of relative platform size. For example, if the Ford *Fusion* adopts a design platform on which one million cars are assembled in Germany,[48] and $f(.)$ is the identity function, then

[46] This is because platforms are redesigned every five to ten years.

[47] This specification diverges from recent tests of market size (e.g., Costinot et al. (2019); Dingel et al. (2023)) in several respects. First, because platform-sharing is a firm decision (see Section 2), design choices can change the products in each platform, and thus the set of countries that are large markets for each platform's inputs. In contrast, most papers consider drivers of market size – population size, density, income, and age – that shift final preferences. Second, because redesigns change platform market size for the same product, and are staggered within a firm, I use exclusively within-product j panel variation in input demand rather than cross-sectional variation.

[48] To avoid endogenous volume changes, I add total assembled quantities of each product $j = (f, s, d)$ in the five years

Table 2: Post-Great Recession changes in firm boundaries

	Realized Merger		Alternative Merger	
Acquirer	**Target**	**Primary Countries**	**Target**	**Primary Countries**
Chrysler	Fiat	Italy	Renault, Nissan	France, Japan
Ford	Ford of Europe, (-)Mazda	Germany, Spain, UK, (-)Japan	Volkswagen	Germany, Brazil
GM	Opel	Germany, Hungary	Peugeot S.A.	France

Notes: Table shows realized and alternative (proposed but not realized) changes in firm boundaries during the 2008 recession. I refer to the firm or affiliate with U.S. production as the "acquirer" for convenience. Ford of Europe and GM's European unit Opel were technically owned by Ford and GM respectively, but were managed separately (and used separate design platforms) until 2008. A (-) indicates a spinoff.

$f\left(\Delta PlatVol_j^o\right) = 1$ million.[49]

Equation (13) thus tests if cars that adopt new design platforms disproportionally increase input sourcing from foreign locations that – as a result of shared platform adoption – now demand their inputs in large volumes.

Firm boundaries as alternative market size shifters. The key endogeneity issue in (13) is that production costs may be falling for other reasons in origins o that demand product j's inputs in large volumes after shared platform adoption. For example, Ford might move the *Fusion* (a midsize car) to a platform shared with a Chinese midsize car exactly because Chinese suppliers are improving at producing midsize car parts over time.

To isolate plausibly exogenous variation over time in platform market size, I leverage the fact that the 2008 global recession induced several *de facto* changes in American auto firm boundaries, and therefore in the set of products (and thus countries) using each design platform. I focus on Great Recession-induced mergers because they were motivated by a short-term need to avert capital outlays (and therefore avoid bankruptcy) rather than a desire to achieve long-term production cost savings through production synergies.[50] Table 2 shows the mergers that occurred, as well as *alternative* mergers that were proposed but did not take place.[51] Appendix Subsection D.2 discusses the specific motives for each event.

Because mergers also directly affect relative production costs (e.g., because affiliates can access the parent's supplier base), I exploit that *the change in platform market size induced by a merger differs across products within each firm*. Specifically, because most automakers used platforms that were specific to a size category (i.e. to compact or midsize cars, see Figure 4), and cars of different size categories were assembled in different locations, car models in different segments face differential

before adoption occurs to calculate predicted post-treatment volumes.

[49] I consider both the level $f(x) = x$ and the inverse hyperbolic sine $f(x) = sinh^{-1}(x)$ of assembly volume.

[50] While automakers often acquire each other (or participate in joint ventures and alliances) explicitly to share design platforms (Praetorius, 2025), these decisions are endogenous to production costs and demand: for instance, in recent years, many global auto firms' joint ventures with Chinese automakers (Bai et al., 2025) were created explicitly so that firms could both assemble cars at scale in China and source from lower-cost Chinese suppliers.

[51] Because first- and second-choice mergers are with different firms – with assembly plants in different locations – realized mergers led certain American car models to share platforms and inputs with a *different* set of locations than in the second-choice merger scenarios.

changes in foreign input demand. For example, when Ford integrated its formerly-independent European affiliate, Ford's American pickup trucks faced no change in platform-sharing, Ford's American subcompact cars and SUVs adopted a platform used in Poland, and Ford's American compact cars and SUVs adopted a platform also used in Germany. Relative demand from Poland (Germany) thus increased for inputs that go into Ford's subcompact (compact) car built in the United States.

I represent these changes in input demand as follows. Suppose firm f acquires firm f'. For a product j of firm f, redesigned in year $\bar{t}_j^{new}$ after a merger, the predicted change in input demand in o on j's platform is:

$$\underbrace{\Delta PlatVolMerger_j^o}_{\text{Added demand in } o \text{ for } j\text{'s inputs due to merger}} = \sum_{j'} \underbrace{Q_{j'}}_{\text{Cars assembled}} \cdot \underbrace{\mathbb{1}[f(j') = f']}_{\text{at target firm } f'} \cdot \underbrace{\mathbb{1}[s(j') = s(j)]}_{\text{in the same segment as } j} \cdot \underbrace{\mathbb{1}[d(j') = o]}_{\text{in origin } o} \tag{14}$$

i.e. total pre-merger assembly volumes in country o, in the same segment $s(j)$ of product j, at the *other* firm f' that is being acquired.

I then re-estimate the event study in equation (13) using the predicted change in platform assembly volume due to the merger $\Delta f\left(PlatVolMerger_{s(j)}^o\right)$:

$$Source_{jtp}^o = \alpha_{jp}^o + \gamma_{fatp}^o + \sum_{\tau \neq -1} \beta^\tau \left(\underbrace{\mathbb{1}[t - \bar{t}_j^{new} = \tau]}_{\text{time since redesign}} \cdot \underbrace{\Delta f\left(PlatVolMerger_{s(j)}^o\right)}_{\Delta\text{new assembly on segment } s(j) \text{ in } o \text{ in due to merger}} \right) + \varepsilon_{jtp}^o \tag{15}$$

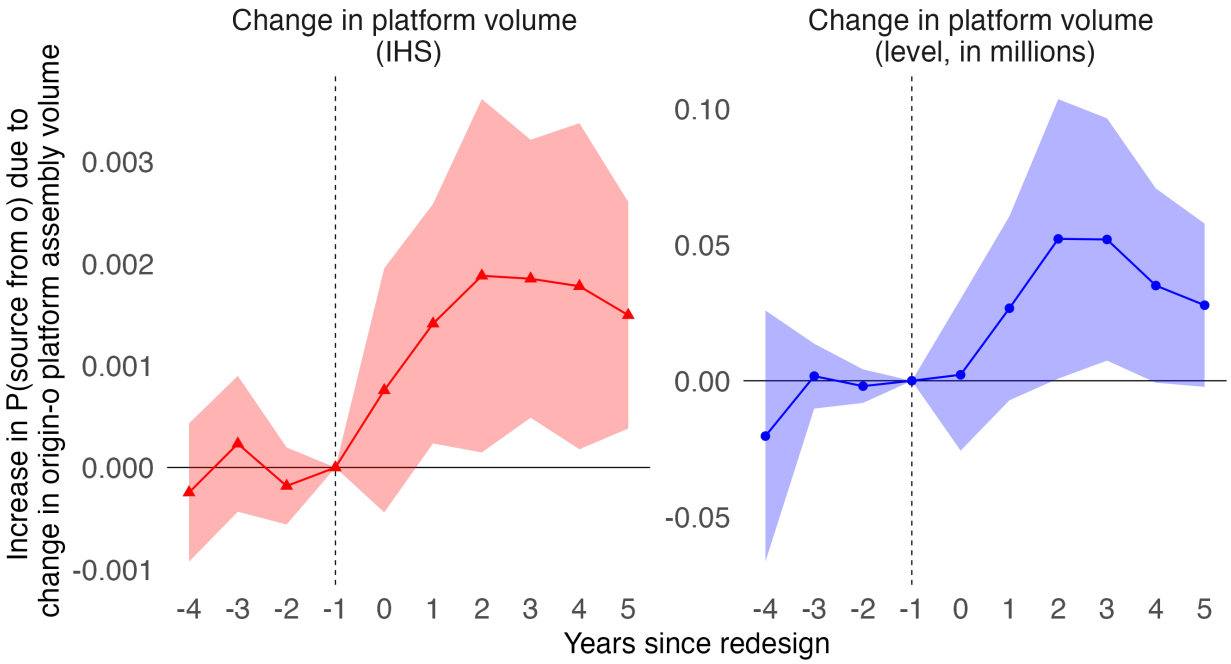
where γ_{fatp}^o now controls for the "direct" firmwide sourcing effects of the merger, for instance from access to the acquired firm's foreign engine suppliers.

Two key assumptions underly the use of $\Delta PlatVolMerger_j^o$ as a causal shifter of platform market size across products within a firm. The first is that after the merger, design platforms are shared within but not across size categories.[52] This assumption is consistent with the patterns shown in Figure 4 in which before the Great Recession (in 2000-2007), the average firm used shared platforms across destinations and body types but within a size category. The second is that product demand conditions facing the target firm before and after the merger are roughly unchanged, so that we can sum *pre-merger* product assembly volumes to obtain predicted changes in platform market size. This is important because *post-merger* assembly volumes are endogenous to car prices, and thus to any changes in input costs.

Because $\Delta f\left(PlatVolMerger_{s(j)}^o\right)$ includes predicted platform market size changes only due to firm boundary changes, it is exogenous to sourcing trends as long as the target firm's largest assembly locations *for each size category* are not also experiencing downward trends in size category-specific

[52]For example, when Chrysler and Fiat merger, I assume that Fiat and Chrysler's compact (C-segment) vehicles will use a common platform, but compact and midsize (C-segment and D-segment) vehicles will not. Size categories are: A to F, pickups, or MPVs. SUVs and sedans of the same size (e.g., C and SUV-C products) are in the same category.

Figure 13: Modularity across destinations increases sourcing from large markets for platform



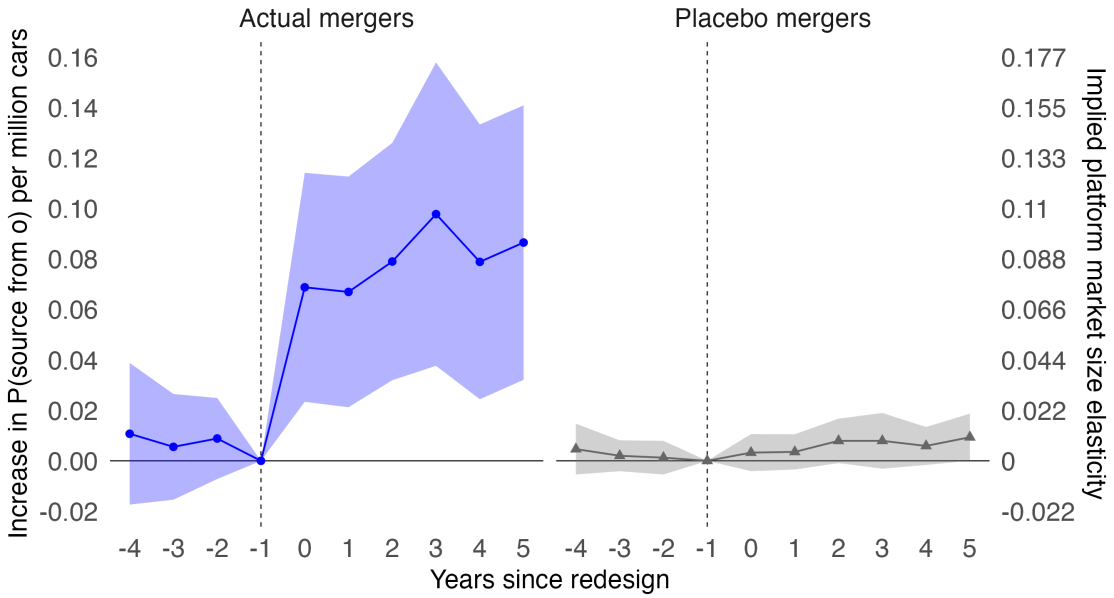
Notes: Panels show effect of adopting a platform that is used at larger scale in other countries on the probability of importing a core input (engine and/or transmission) from those countries. All effects at the model-year-origin country level. Coefficients interact time-since-redesign with the difference in platform assembly volume in each country between the old and the new platforms. The left panel uses the inverse hyperbolic sine of total platform volume, and the right panel uses the level (in millions of cars). Import share data from the American Automotive Labeling Act. Redesign year, platform information, and locations of platform use are from *Marklines*.

input production costs over time. Two facts support this assumption. First, acquirers had little choice of "targets": Ford and General Motors (GM) integrated their European units into their North American operations; and Chrysler was sold to the first buyer (Fiat) who had sufficient financing. Second, European input production costs were (if anything) increasing relative to Chinese competitors during this period. As further evidence that relative production costs are not falling in merger-exposed locations, as a robustness check I re-estimate (15) using data from second-choice *placebo mergers* – with different European firms – seriously considered by American automakers at the time of the financial crisis.[53]

Hypotheses Because I assume that firms adopt modular design across all *destinations* (but not all product segments) to construct the merger-derived platform market size changes, the changes in platform-sharing due to merging with foreign firms mirrors the first phase of modular design described in Proposition 1. I therefore expect that *modular design across destinations shifts production to large foreign (input) markets*: cars that adopt new design platforms, which are modular across

[53]For instance, Fiat purchased Chrysler a few weeks after Peugeot (which had expressed interest) backed out due to an unexpected change in its own financial position. Such a deal which would have increased platform scale in France (rather than Italy) for many US-built cars. If alternative mergers are selected similarly to realized ones, and not correlated with trends in post-recession production costs, then these "placebo" exposure measures should have no effect.

Figure 14: Firms source inputs from countries with increased platform market size due to merger



Notes: Panels show effect of changes in predicted platform market size (for each car j at firm f, proxied by the pre-merger assembly volume in segment $s(j)$ in origin o within the target firm f') on imports a core input (engine and/or transmission). All effects at the model-year-origin country level. In the left panel, coefficients interact time-since-redesign with the difference in predicted platform scale due to the mergers, assuming that platforms remain largely within-segment. In the right panel, coefficients interact time-since-redesign with differences due to placebo mergers that were proposed but did not take place. Import share data from the American Automotive Labeling Act. Redesign year, platform information, and locations of platform use are from *Marklines*.

countries, should disproportionately source from foreign countries with large assembly volumes on the shared platform.

Results I first estimate Equation (13), which uses redesign-induced (rather than merger-induced) changes in platform market size. The coefficients from this regression are shown in Figure 13.

Figure 13 shows that sourcing does increase from the shared platform's largest foreign market (specifically, from newly added countries with the largest assembly volumes on j's platform), consistent with Proposition 1. This effect is economically significant: in the left panel (which uses an IHS scale) a log point increase in platform scale increases the probability of sourcing from there by 0.21 p.p (10% from a mean of 3.2 p.p.). The right panel shows that effects are qualitatively similar and remain significant when exposure is measured as the *level* of input demand (in millions of cars assembled).

To isolate the causal effect of platform market size, I next estimate Equation 15.

The results are shown in Figure 14. In the left panel, using an entirely different source of variation from Figure 13 (platform market size changes from mergers, assuming that platforms are size-category-specific), I find effects of a similar time-path and magnitude: there is an immediate 8 p.p. increases in the probability of sourcing from country o for each 1 million cars assembled in country o in the target firm. The right panel shows that the change in country-segment-level input

demand that would have resulted from alternative "placebo" mergers have no effect, consistent with the idea that 2008 recession-era mergers were not selected due to their potential cost-savings.

We can interpret the coefficient in the left panel of Figure 14 as a platform-level market size (or weak home market) effect. As shown using the right axis of Figure 14, rescaling the estimated effects yields a platform market size elasticity of approximately 0.11.[54] In other words, for a given car j, a log point increase in platform market size in o – in other words, in the number of cars assembled in origin o on model j's platform – due to the merger leads to a 11% increase in the probability that j sources inputs from o.

In Appendix Figure A.5 I interact the event with the inverse hyperbolic sine (IHS) of the merger-predicted quantity rather than the level. The time-path of effects are qualitatively similar to Figure 14, and the implied platform market size elasticity is also approximately 0.10.

While significant, this platform home market elasticity of 0.11 is smaller than sector-wide estimates obtained for other sectors. For instance, in Costinot et al. (2019) the equivalent estimate for the pharmaceutical sector is 0.9. Furthermore, this estimate is likely subject to measurement error since pre-merger assembly volumes only provide an imperfect prediction of post-merger volume. The positive estimate nevertheless implies that countries that assemble cars in large volumes on a platform have a cost advantage in producing its inputs. This implies, for instance, that poor countries will specialize in inputs for platforms that are small car-specific.

Implications Together these results suggest that, consistent with Proposition 1, modularity across destinations increases import shares and leads input production to locate in large markets for each product segment.

5.3 Does modular design weaken specialization in product segments?

I now test if the second phase of adoption – modularity across product segments – makes firm sourcing behavior more homogeneous, as in Proposition 2. Specifically, I test if countries specialize more weakly in product segments that they demand when inputs are no longer segment-specific.

5.3.1 Event study approach

In a first approach, I estimate event studies in which a product adopts modularity across segments. Motivated by the evidence in Figure 7 that developing countries both demand and specialize in technology-intensive inputs (engines and transmissions) for small vehicles[55], I focus on redesigns of small cars in which *the post-redesign platform is shared with at least one additional segment, while the pre-redesign platform is not*.[56]

[54] I multiply each coefficient by the mean change in platform market size (the treatment) and divide by the mean sourcing probability (the outcome).
[55] Length classes A, B, and C (including SUVs) as well as pickups.
[56] Formally, the set of segments $\{s : (f, s, d) \in S_{g(j)^{POST}}\}$ is strictly larger than the set $\{s : (f, s, d) \in S_{g(j)^{PRE}}\}$.

The seven events are shown in Figure 9. For example, in 2020, the American *Ford Escape*, which previously shared a platform with other *C*-segment cars, started sharing inputs with *D*- and *E*-segment cars (such as the Chinese-made *Ford Edge* and *Lincoln Z*) as well as pickup trucks (the *Ford Maverick*).

Proposition 2 states that, if inputs for small cars are also used by larger vehicles, then small cars should be *less likely* to source inputs from countries that prefer them. To test this, in the AALA sourcing panel, I define the outcome:

$$OriginPrefersSmallCars_{jtp} = \mathbb{1}\left[\frac{\sum_{j'} Q_{j'} \cdot \mathbb{1}[s(j') \in \{\text{mini, subcompact}\}] \cdot \mathbb{1}[d(j') = Origin_{jtp}]}{\sum_{j'} Q_{j'} \cdot \mathbb{1}[d(j') = Origin_{jtp}]} > \bar{k}\right]$$

which takes value 1 if part p for product j is sourced from an origin $Origin_{jtp}$ for which the small car share of total local assembly exceeds some cutoff k. For consistency with Proposition 2, I use $k = 0.21$ (the global mean). I then estimate the effect of platform-sharing across segments on the probability of sourcing from a small car-preferring country:

$$OriginPrefersSmallCars_{jtp} = \alpha_{jp} + \gamma_{fatp} + \sum_{\tau \neq -1} \beta^{\tau} \mathbb{1}[t - \bar{t}_j^{new} = \tau] + \varepsilon_{jtp} \tag{16}$$

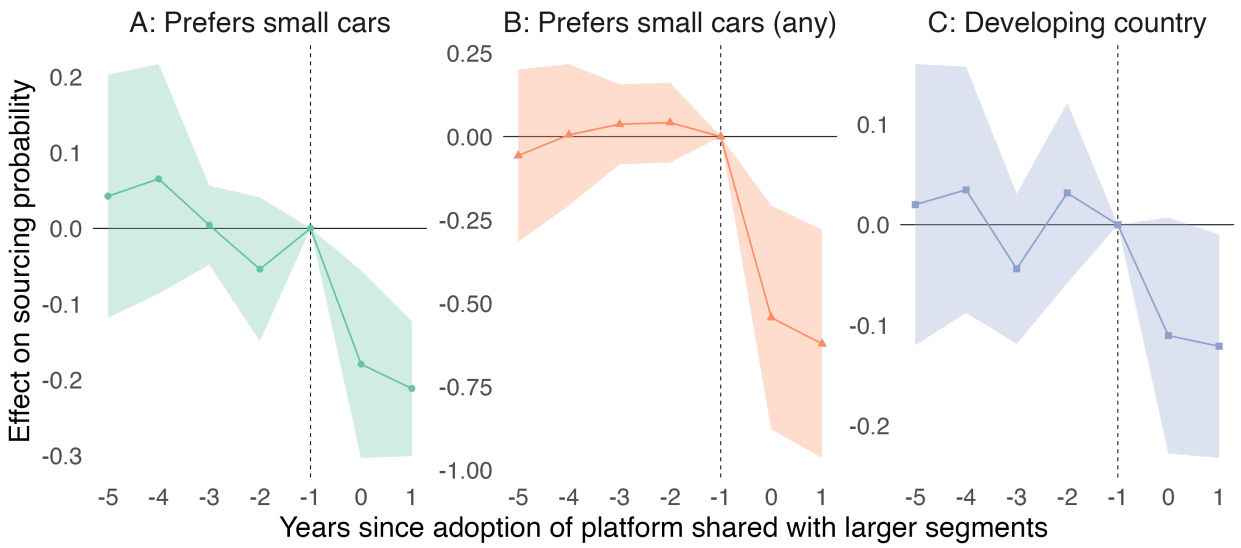
controlling for product-part (ex: Ford *Focus* transmission) fixed effects α_{jp} and firm-assembly-time-part fixed effects γ_{fatp}. Because Figure 7 suggests that developing countries prefer small cars, I also estimate the effect on $DevCountry_{jtp}$, which takes value 1 if product j sources part p from a developing country in year t.

Results. The results in Figure 15 are consistent with Proposition 2: *platform-sharing (and thus input-sharing) with larger segments leads small cars to source less from countries that prefer them*. As shown in Panel A, adopters of cross-segment modularity source 20 percent less on average from countries that prefer them. This effect is driven by the extensive margin: there is a significant 65 percent decline in the probability of a model having *any* core inputs from countries that assemble small cars (Panel B).[57] Finally, Panel C shows that platform-sharing across segments leads to a significant 11 percent decline in the probability that small cars assembled in North America source from poor countries.

The results in Figure 15 suggest that the pattern of cross-country specialization observed in Figure 7, in which poor countries export value-added inputs to North America in segments that they prefer, may be weakened by modular design. In other words, poor countries may no longer be able to specialize in (inputs for) for affordable products, because those inputs are no longer specific to affordable final goods.

[57] We can estimate this regression because each model j has multiple transmission and engine locations reported.

Figure 15: When segments share platforms, small car assemblers export fewer small car inputs



Notes: Panels show effect of adopting a platform that is shared across segments (with larger cars) on sourcing decisions for small cars, relative to all cars within the firm. All effects at the model-year level. There are seven treated car models, all of which use small-car platforms before the event: the VW Jetta, Ford Escape, Lincoln Corsair, Nissan Sentra, Honda HR-V, GMC Canyon, and Chevrolet Colorado. Controls consist of all other cars within the same firm. Outcomes are: the share of core inputs that are sourced from a country that prefers small cars (Panel A), an indicator for this share being greater than 0 (Panel B), and the share sourced from developing countries (Panel C). Small cars are segments A, B, SUV-A, and SUV-B. Countries that prefer small cars are locations in which least 20.4 percent of total assembly volume (the 33rd percentile) consists of small cars. Developing countries are all countries excluding the USA, Canada, Mexico, the EU, Japan, Korea, and China. Import share data from the American Automotive Labeling Act. Redesign year, platform information, and segments of platform use are from *Marklines*.

5.3.2 Gravity regression approach

In a final reduced-form analysis, I complement the event studies with cross-sectional gravity regressions that use data on supplier relationships for *all car models* in *all countries*.

This specification has two advantages relative to the event studies, which focus on a small set of adopters (the events in red in Figure 9) and two specific components (the engine and transmission). The first is generalizability: I aim to test if modularity influences input trade in the full sample of countries and components. The second is greater variation in modularity across *segments* (ex: between large and small cars), which is relatively limited at present, but expected to increase substantially as a result of the transition to electric vehicles.

I use the *Marklines Who Supplies Whom* supplier relationship database, which covers input sourcing for cars assembled in all countries rather than solely in North America. To understand how specialization varies with modularity, I merge in three additional pieces of information. First, to proxy for a country's demand for inputs for cars in a particular product segment, I use *Marklines*

assembly volumes by country, and calculate the share of all local assembly in each segment:

$$AssemblyShare^o_s = \frac{\sum_{j:s(j)=s,d(j)=o} Q_j}{\sum_{j:d(j)=o}} \tag{17}$$

Second, to understand the extent to which each car uses the same inputs as *other segments* – which should weaken segment specialization, as shown in Proposition 2 – I merge in platform information for each product from the Marklines design plans, and calculate the share of all products designed on the platform that are in *other* segments:

$$SharePlatOtherSegment_j = \frac{\sum_{j'} \mathbb{1}[g(j') = g(j), s(j') \neq s(j)]}{\sum_{j'} \mathbb{1}[g(j') = g(j)]} \tag{18}$$

From the perspective of a car j in segment $s(j)$, $SharePlatOtherSegment_j$ is a summary measure (from 0 to 1) of the extent to which the platform is non-specific, i.e. used for products in other segments $s' \neq s(j)$. If a platform is only used for cars in the same segment $s(j)$, then $SharePlatOtherSegment_j 0$ is 0. If a platform is used for many cars in other segments s', then $SharePlatOtherSegment_j$ approaches 1.

$SharePlatOtherSegment_j$ provides continuous variation in the extent of platform-sharing across segments. For interpretation, I also use the binary measure $\mathbb{1}[SharePlatOtherSegment_j > 0]$ as a measure of any platform-sharing across product segments.

Specification. For origin o, car model j of firm f assembled in country d, and part category p[58], I estimate:

$$\begin{aligned} \log NumRels^o_{jp} = \gamma_{jp} + \delta^o_{dfp} + \beta^1 \left(AssemblyShare^o_{s(j)} \right) + \\ \beta^2 \left(AssemblyShare^o_{s(j)} \cdot SharePlatOtherSegment_j \right) + \varepsilon^o_{jp} \end{aligned} \tag{19}$$

where $NumRels^o_{jp}$ is the number of supply agreements with suppliers in country o, within part category p, for final product j. The product-part category fixed effect γ_{jp} controls for the fact that some cars have more relationships in general.[59] The origin-destination country-firm-part category fixed effect δ^o_{dfp} accounts for firmwide sourcing strategies that apply to all of an automaker's products. By including this term, we ensure that we study patterns of specialization conditional on supplying products to the same firm and assembly location – in other words, whether *within the set of Ford's US-built cars*, certain countries specialize in supplying parts for smaller or larger car models. The same fixed effect also controls for any (proportional) trade costs that are constant within a firm and product, such as tariffs or physical shipping costs.

There are two coefficients of interest. β^1 is an estimate of the association between the share of country-o assembly that is in segment s and the number of supply relationships in which

[58] Examples include "engine core parts" and "air conditioner parts".
[59] This is analogous to controlling for multilateral resistance in a standard gravity model.

country-o suppliers produce inputs (including for export) for segment-s cars. For instance, if India both disproportionately assembles small cars and produces small-car inputs, then β^1 will be positive. Larger (positive) values of β^1 imply that *countries specialize more strongly in product segments that they demand*. Importantly, this is a statement about the correlation of input demand composition with input production – in other words, about *observed segment-level specialization* – and is not a causal statement about whether demand is the structural driver of specialization.

In contrast, β^2 captures heterogeneity in segment-level specializations due to differences in modularity. Specifically, β^2 captures how much more (or less) strongly a country specializes in inputs for segment-s cars when the car uses a platform that is largely used for other segments s', relative to a platform that is only used for segment s. If estimated β^2 is zero, then segment-level specialization is unrelated to the extent of design commonality. If $\beta^2 = -\beta^1$ (and $\beta^1 > 0$), then segment-level specialization exists when platforms are segment-specific, but nonexistent when a car platform is largely used for other segments. In other words, a negative β^2 would imply that platform-sharing (and thus input-sharing) with other segments weakens the strength of a country's specialization in inputs for preferred segments as predicted in Proposition 2.

One concern with equation (19) is that null effects (β^1 and $\beta^2 = 0$) may occur because – even if countries specialize in inputs for different segments – $AssemblyShare^o_s$ may mismeasure input demand, either due to measurement error or because it does not incorporate the segment composition of input demand in nearby countries.[60] To test if modularity weakens *any* segment-level specialization patterns, whether or not due to local demand, I re-estimate equation (19) using an alternative measure of countries' revealed cost advantages. Specifically, I estimate a first-stage regression:

$$\log NumRels^o_{jp} \;=\; \alpha^o_{s(j)p} \;+\; \gamma_{jp} \;+\; \delta^o_{dfp} \;+\; \varepsilon^o_{jp} \tag{20}$$

in which $\{\alpha^o_{s(j)p}\}$ are measures of origin-segment-part relative productivity. The estimated fixed effects $\{\widehat{\alpha^o_{s(j)p}}\}$ capture, for instance, a country's relative cost advantage in making brakes for small cars, both because of local demand for small cars and other reasons. I then re-estimate (19) with estimates $\{\widehat{\alpha^o_{s(j)p}}\}$ in place of $AssemblyShare^o_{s(j)}$:

$$\log NumRels^o_{jp} \;=\; \gamma_{jp} \;+\; \delta^o_{dfp} \;+\; \beta^1 \left(\widehat{\alpha^o_{s(j)p}}\right) \;+\; \beta^2 \left(\widehat{\alpha^o_{s(j)p}} \cdot SharePlatOtherSegment_j\right) \;+\; \varepsilon^o_{jp} \tag{21}$$

In (21), if $\beta^2 < 0$, then a country's general segment-level cost advantage weakens as platforms become less specific to a segment.

Hypotheses Equation (19) and (21) are tests of Proposition 2, in which the adoption of cross-segment modularity eliminates segment-level specialization due to preferences.. We therefore

[60] For instance, Canada might enjoy a market size (and therefore scale) advantage in inputs for pickup trucks because its closest export (the USA) assembles many pickups.

expect that *modularity across segments weakens specialization in inputs for locally-assembled segments.* Specifically, in (19), we expect that $\beta^1 > 0$ (countries specialize in preferred segments) and $\beta^2 < 0$ (weakening due to modularity). We further expect that modularity across segments weakens pre-existing revealed cost advantages, so that $\beta^2 < 0$ in (21) as well.

These predictions are especially of interest because segment specialization allow many countries to export value-added automotive inputs. If within-sector segment specialization modularity (i.e. if $\beta^2 < 0$) then many countries lose these opportunities, and may specialize in other sectors that lack the increasing returns, potential for knowledge transfer, and other benefits of value-added automotive manufacturing. In contrast, most standard models of cross-country specialization that do not account for the extent of input-sharing would predict $\beta^2 = 0$ in both regressions.

Table 3: Modularity across segments weakens specialization in inputs for locally-assembled segments

Dependent Variable:	# of supply relationships with origin o firms for model m (PPML)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Segment s share of origin o assembly	0.409***	0.799***	0.907***			
	(0.098)	(0.253)	(0.197)			
Segment s share of origin o assembly × 𝟙[Platform has ≥ 1 other s′ ≠ s]		-0.448*				
		(0.269)				
Segment s share of origin o assembly × Share of platform models in s′ ≠ s			-0.966***			
			(0.332)			
Origin o RCA in segment s inputs				1.009***	1.177***	1.163***
				(0.153)	(0.147)	(0.153)
Origin o RCA in segment s inputs × Share of platform models in s′ ≠ s					-0.354***	
					(0.078)	
Origin o RCA in segment s inputs × 𝟙[Platform has ≥ 1 other s′ ≠ s]						-0.196***
						(0.062)
Fixed-effects						
Model-destination-part group	Yes	Yes	Yes	Yes	Yes	Yes
Origin-destination-firm-part group	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Dependent variable mean	2.0616	2.0616	2.0616	1.8460	1.8460	1.8460
Observations	29,205	29,205	29,205	33,056	33,056	33,056

Clustered (Origin-segment-platform) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows estimated effects of gravity specifications (19) and (21) in which proxies for origin-segment cost advantages are interacted with two proxies for cross-segment modularity. The proxies are the share of origin-o assembly in segment s (columns 1-3), which capture relative input demand; and an origin-segment fixed effect (columns 4-6), estimated from a first-stage regression without covariates, capturing origin o's revealed cost advantage (RCA) in inputs for segment s. All data from *Marklines Automotive*.

Results. PPML estimates of equations (19) and (21) are given in Table 3. The results imply that *the effects of two different proxies for country-segment-level cost advantages weaken on platforms that cover more segments.*

In Columns 1-3, I first use a proxy for demand-based cost advantages: the share of all cars assembled in origin o (by any firm) that are in segment s.[61] Column 1 shows that countries do indeed specialize in inputs for product segments that they demand: a 10 p.p. increase in assembly share on a segment predicts a 4.1% increase in supply relationships (i.e. input production opportunities) for o. Column 2 shows that this specialization is almost twice as large (at 8.0%) on platforms that are only used by one segment; and weakens significantly (by 4.5%, a proportional decline of 56 percent) when other the platform is shared with at least one other segment $s' \neq s$. Column 3 shows that a similar weakening of specialization occurs when this indicator for platform-sharing is replaced with a continuous measure of the *share* of cars on model m's platform that are in other segments. In other words, patterns of demand-based specialization across product segments of varying affordabilities and sizes within an industry – a central driver of manufacturing production and trade (Linder, 1961; Schott, 2004; Khandelwal, 2010) – are heterogeneous, weakening on design platforms that cover multiple segments.

Proposition 2 also has the implication that the overall strength of segment-level specialization patterns should weaken with modular design. To test this, in Columns 4-6 of Table 3 I replace the segment s share of cars assembled with the revealed cost advantage of country o in input production for segment s, which I obtain by estimating an origin-segment fixed effect in a first stage on the full sample. By construction, higher values of the fixed effect predict higher sourcing, as shown in Column 4. Columns 5 and 6 show that – as with the demand-based proxy – this fixed effect-based revealed cost advantage measure has weaker effects on observed sourcing when platforms cover multiple segments. In other words, Columns 5 and 6 show that cross-country specialization in product segments is driven largely by platforms that cover a single segment; as shared platforms are adopted across segments, these specialized export opportunities weaken.

In other words, controlling for any shifters that apply to all segments, both the share of local assembly in each segment (a proxy for country-segment preferences) and the model-implied revealed cost advantage are less strongly associated with production (including exports) in that segment. These patterns are consistent with the predictions in Proposition 2 that countries can only specialize and export in inputs for different product segments (such as in parts of big or small cars) when those inputs come from *different* platforms.

I next explore heterogeneity by component, leveraging that the *Who Supplies Whom* data contains sourcing information for 300 separate part types (not just engines and transmissions). In Appendix Table A.1, I show that the weakening of segment specialization in Table 3 holds for both specialization proxies when restricting to core components. This is consistent with the stylized fact in Figure 6, which showed that input-sharing disproportionately increases sharing of core

[61] Because of trade costs for final cars, these composition effects often reflect consumer preferences.

components (engines, drivetrains, and electronics). Results are more mixed for differentiators, as shown in Appendix Table A.2: cross-segment modularity does weaken demand-based specialization in production of interior and exterior parts, but does not have any effect on the overall strength of segment specialization.

Implications A broader point is that Table 3 suggests that the relative strength of within-sector specialization can change within an industry over time due to changes in firm design technology. In contrast, most drivers of within-sector specialization that economists typically model – consumer love-of-variety or preferences (Krugman, 1979, 1980), idiosyncratic productivity differences (Eaton and Kortum, 2002), input customization as in Rauch (1999), or quality ladders as in Khandelwal (2010) – are typically assumed to be fixed characteristics of a product, with implications for the long-run concentration of production, but not for changes over time.

5.4 Taking stock

I have shown that the adoption of modular design within a product segment (i) increases trade, without changing trade costs, and (ii) shifts production towards large foreign countries, without increasing the elasticity of returns of scale; and that modularity across product segments (iii) weakens segment-level specialization, such as for poor countries in value-added inputs for small cars. All of the changes I document are design changes alone, which I isolate by leveraging staggered rollouts of pre-planned redesigns, changes in platform market size from firm mergers, and cross-sectional variation in design in gravity regressions.

These three predictions are consistent with a model of design platform-level economies of scale. An aggregate implication of such a model is that increases in modularity (in platform-sharing) will progressively concentrate input production in large assembly locations for each design platform, and eventually for each firm. I next quantify the model to explore these potential aggregate effects.

6 Quantification

I have shown that the adoption of modular design leads firms to increase foreign input sourcing, especially from large markets for common inputs; and then source from a common set of origins rather than separate origins for each product segment. To explore the aggregate implications of these sourcing shifts, I now quantify the model from Section 4. In particular, I examine if modularity leads to a more unequal distribution of industrial activity, because countries can no longer exploit opportunities to specialize in inputs for certain destinations or product segments.

6.1 Approach

Modifications I modify the model in two ways to allow for quantification. First, there are 28 countries rather than two. This allows a country's input market size advantage to depend on its proximity to large foreign assembly locations, along with its own domestic assembly. Second, I assume that consumers have logit product demand (rather than CES) with elasticity ε over products j within each segment; and that car j assemblers have logit input demand (rather than CES) with elasticity σ over inputs for platform $g(j)$ from different origins o. These assumptions allow me to calibrate the model directly using final car assembly quantities and input sourcing choice shares. In contrast, under CES demand I would require origin-platform-level input prices, which are unobservable because automotive inputs are either traded intrafirm or in long-term relationships; as well as product-level final car prices, which unlike quantities are not systematically collected via car registrations.

Method To conduct counterfactuals, I first show that the exact-hat method of Dekle et al. (2007) also admits changes in design platforms alongside changes in trade costs. To see this, recall that equilibrium production locations depend only on the shadow price p^{og}, which are in turn defined by the internal market-clearing condition:

$$y^{og} \;=\; \sum_{j \in \mathbf{S_g}} q_j^o \cdot \tau_j^o \;=\; \sum_{j \in \mathbf{S_g}} \pi_j^o \cdot Q_j \cdot \tau_j^o \tag{22}$$

i.e. that total input production y^{og} equals total input use q_j^o for all models j on platform g, which in turn is the product of model j assembly volume and the share of model j parts sourced from o. Following Dekle et al. (2007), we can denote counterfactual objects as x' and proportional changes as $\widehat{x} = \frac{x'}{x}$. We also index counterfactual platform assignments by g' instead of g. In Appendix Subsection D.3. I show that proportional changes $\widehat{y^{og}}$, $\widehat{Q_j}$ and $\widehat{\pi_j^o}$ can be written solely in terms of vectors of their baseline quantities and three elasticities: the scale elasticity η, sourcing elasticity σ, and final demand elasticity ε. As a result, the *counterfactual* market-clearing condition for each o and g' can be written as:

$$\forall (o,g): \quad y^{og} \cdot \widehat{y}^{og}(\widehat{p}^{og};\eta) \;=\; \sum_{j \in \mathbf{S}'_g} \pi_j^o \cdot Q_j \cdot \widehat{\pi}_j^o\big(\widehat{p}^{og};\{q_j^{og}\},\{\pi_j^o\},\varepsilon,\sigma\big) \cdot \widehat{Q}_j\big(\widehat{p}^{og};\{q_j^{og}\},\{\pi_j^o\},\varepsilon,\sigma\big) \tag{23}$$

which differs from the baseline condition (22) in that input demand is summed over all products j using platform g in the counterfactual scenario rather than at baseline (all $j \in \mathbf{S}'_\mathbf{g}$ rather than $j \in \mathbf{S_g}$). By solving (23), I therefore obtain the proportional changes in equilibrium input shadow prices $\widehat{p^{og}}$ that allow counterfactual input markets to clear.

Equation (23) allows for any counterfactual in which products are removed or added to a platform. For example, suppose platform g covers both small and large cars at baseline, but is split in the counterfactual so that (wlog) small cars remain on platform g while large cars are moved to a

different platform h. Given baseline quantities and elasticities ($\{q_j^{og}\}$, $\{\pi_j^o\}$, ε, σ and η), equation (23) implicitly defines the proportional change in input prices $\{\widehat{p^{og}}\}$ for small cars that remain on the platform.[62]

My key outcome of interest is the proportional change in country o's share of global input production:

$$\widehat{ProdShare}_o = \frac{\widehat{\sum_g y^{og}}}{\sum_g \sum_o y^{og}} = \frac{\frac{\sum_{g'} (y^{og'})'}{\sum_{g'} \sum_o (y^{og'})'}}{\frac{\sum_g y^{og}}{\sum_g \sum_o y^{og}}} \tag{24}$$

so that (for example) if the United States produces 20 percent of all inputs at baseline, and 30 percent in the counterfactual, then $\widehat{ProdShare}_o = 1.5$. In Appendix Subsection D.3, I show that $\widehat{ProdShare}_o$ can be calculated directly from optimal solutions $\{\widehat{p^{og}}\}$, baseline quantities, and the three elasticities.

The key simplification in my model relative to reality is that I *do not allow for trade in final goods.* In a world of free trade, we would expect modularity to relocate final assembly locations for each product, as studied by Praetorius (2025), Head et al. (2024), and Head et al. (2025). However, empirically, adjustments in assembly locations tend to be slow for political economy reasons. As a result, an alternative interpretation of my counterfactuals is that they are medium-run effects in which input production adjusts taking product assembly locations and (residual) demand functions as given. The assumption of local final assembly is especially reasonable, even in the long-run, for the three largest car assemblers markets – the US, China, and India – that are most important in my upcoming empirical results.[63]

Estimation of scale elasticity. To estimate the strength of scale economies η that lead to platform market size effects, I instrument for the quantity of platform g-specific input *production* in country o with the merger-predicted change in platform g-specific input *demand* in that country as in Equation 13. Since I do not observe input production directly, I assume that input demand is exactly equal[64] to final assembly volume, which I observe; and assign input demand to origin locations using 2010-2022 product-platform-origin-specific input sourcing shares (π_{jt}^o) from the *Marklines Who Supplies Whom* database.[65] To increase power and avoid biases from many instruments, I estimate a pooled DID regression rather than an event study. As shown in Appendix D, interpreted through the model, this procedure yields the product of two elasticities: the *scale elasticity* η, or the elasticity of input production costs to input production scale; and the *trade elasticity* σ, or the elasticity of substitution across origins.

Coefficient estimates and corresponding elasticities are shown in Table 4. In the preferred

[62] An identical procedure applies to merges of platforms: volume from additional products is summed.
[63] In particular, Chinese imports of final cars are negligible, and as of August 2025 the US and India impose tariffs of 25% and 110% respectively on cars assembled abroad.
[64] In other words, each car requires exactly one core input.
[65] For instance, if a product j's first post-merger redesign was in 2014, I use 2010-2013 and 2014-2018 reported relationships to calculate separate sourcing shares $\pi_j^{og(j,t-1)}$ $\pi_j^{og(j,t)}$ for the pre- and post-redesign platforms.

Table 4: Firms source inputs from countries with larger input production scale

Dependent Variable:	=1 if model m sources part p from country o				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Δ Production scale in origin o (IHS)	0.009**	0.066*	0.044**	0.063	0.065*
	(0.004)	(0.037)	(0.018)	(0.043)	(0.037)
Specification	OLS	IV	IV	IV	IV
Functional form for merger-predicted Δ market size		Linear	Quadratic	Linear	Linear
Origins	All	All	All	All	Foreign only
Fixed-effects					
Part-make-model-origin	Yes	Yes	Yes	Yes	Yes
Part-firm-destination-origin-year	Yes	Yes	Yes		Yes
Part-segment-origin-year	Yes	Yes	Yes	Yes	Yes
Part-destination-origin-year				Yes	
Fit statistics					
Observations	272,961	272,961	272,961	273,357	264,637
R^2	0.751	0.740	0.747	0.698	0.655
F-test (1st stage), Δ Production scale in origin o (IHS)		2,888.6	2,163.7	2,480.8	3,068.9
Dep. var. mean for treated units	0.183	0.183	0.183	0.183	0.183
Implied scale x trade elasticity ($\sigma \cdot \eta$)	0.047	0.361	0.240	0.346	0.357
Implied scale elasticity (η) at $\sigma = 6$	0.008	0.060	0.040	0.058	0.059

Clustered (model-origin) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows the estimated change in probability of sourcing part p (engine or transmission) for model j from origin o due to a log point increase in total input production (Δ Production Scale) in origin o on platform $g(j)$. Input production volume is calculated using input sourcing shares, assembly volumes, platform assignments from *Marklines* (the regressor). Changes in platform market size (the instrument) are constructed as in Equation (12). Engine and transmission sourcing origin data are from the *American Automotive Labeling Act*.

specification (column 2), a log point increase in platform-specific input production scale in a country leads to a 6.6 percent increase in the probability of sourcing from that country. This estimate is net of any firmwide effects of the merger on sourcing for all products, and is also net of segment-specific trends, such as changes over time in the US' comparative advantage in small cars. Estimates remain significant (though slightly smaller in magnitude) when using a second-order polynomial in the change in market size as instruments (Column 3); when removing firm-specific time trends, thereby allowing for cross-firm rather than only within-firm comparisons (Column 4); and when estimating effects on sourcing decisions *conditional* on importing, i.e. removing local sourcing (Column 5). All effects remain significant other than in Column 4, which nevertheless is essentially identical in magnitude. Importantly, these IV estimates are significantly larger than the uninstrumented OLS effects (Column 1), implying that production scale and other determinants of a country's input production costs (such as wages and producer productivities) are *negatively correlated*.

Dividing the Column 2 estimates the mean sourcing probability for treated units, I estimate that

Table 5: Parameters for quantification

Symbol	Name	Value	Source
σ	Input demand elasticity over origins o	6	Fontagné et al. (2022) [HS code 8708], Head et al. (2024)
η	Scale elasticity per origin o and platform g	0.06	Table 4 Column 2 ($\sigma\eta = 0.36$)
ε	Demand elasticity over cars j	4	Conlon and Gortmaker (2020), Head and Mayer (2019)
$\{S_g\}$	Baseline platform assignments	–	*Marklines* Product Timelines 2010–2022
$\{Q_j\}$	Baseline assembly volumes	–	*Marklines* Production 2010–2022
$\{\pi_j^g\}$	Baseline input sourcing shares	–	*Marklines* Who Supplies Whom 2010-2022
$\tau^o_{d(j)f}$	Baseline input trade costs	–	$\{\tau^{-\sigma}\}$ estimated from gravity equation

the product of scale and trade elasticities is $\sigma\eta = 0.36$. In absence of input price data or tariff variation during my sample period, I follow Head et al. (2024) and use a trade elasticity of $\sigma = 6$ for automotive parts.[66] This yields a scale elasticity of $\sigma = 0.06$ (i.e. of 6%) in automotive input (engine and transmission) production, as shown in the bottom row of Column 2.

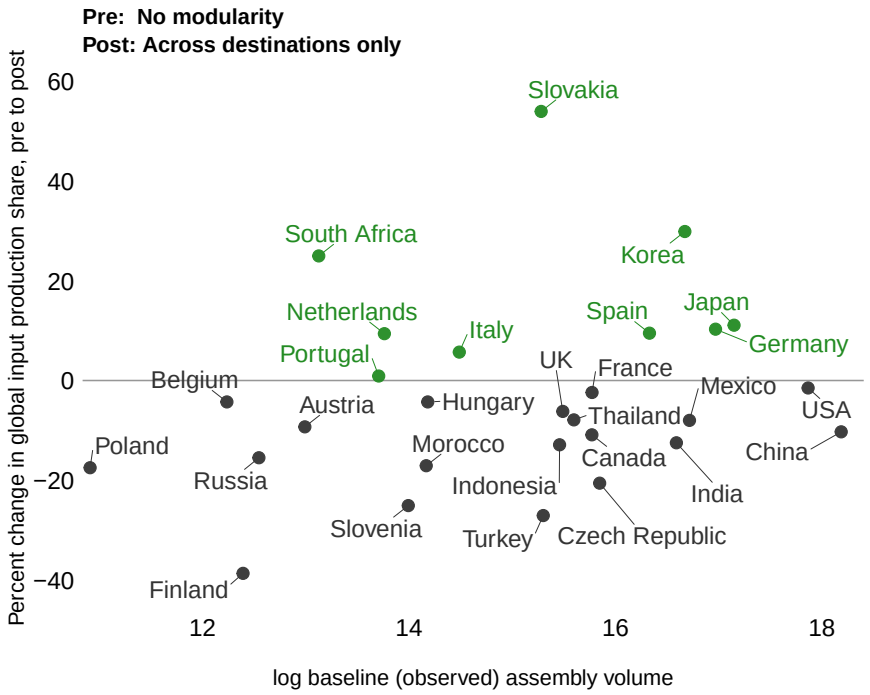
My estimate of $\eta = 0.06$ is larger than the estimates from Head and Mayer (2019) and Praetorius (2025), who estimate external economies of scale (with respect to total industry size) of 0.035 and 0.022 respectively for final assembly using a similar approach.[67] The smaller magnitudes in these papers, using a similar methodology, are consistent with the idea that technology-intensive input production features stronger returns to scale than final assembly.

Calibration of remaining parameters The parameters and data used to calibrate the model are shown in Table 5. I obtain the key parameter that governs scale-induced production relocation – the product of scale and trade elasticities $\eta \cdot \sigma$ – from the IV estimates in Table 4 as described above. Without data on prices of inputs and final cars, I use demand elasticities for auto parts sourcing and final automobile sales from the literature (assuming no trade in final assembled cars). With these estimates, the model can be exactly calibrated using data on car assembly volumes $\{Q_j\}$ and sourcing probabilities π_j^o, where j is the triplet of segment s, destination d, and firm f.

Importantly, we estimated $\sigma\eta = 0.36$ in our preferred specification (Column 2 of Table 4). This is significantly less than 1, so following Kucheryavyy et al. (2023) the model yields a unique interior solution for shadow input prices (and thus production) for each platform and location. In other words, even though input prices fall in a location due to larger production scale (supply curves are downward-sloping), because input demand curves is sufficiently inelastic with respect to prices (demand curves are steeper than supply curves), large markets but does not *fully* concentrate in one location.

[66] Head et al. (2024), who also use the AALA data, show estimates for both $\sigma = 4$ and $\sigma = 6$, where the latter is estimated on global trade data by Fontagné et al. (2022). Meanwhile, Head and Mayer (2019) and Praetorius (2025) estimate input demand elasticities (assuming that inputs only come from headquarters) of $\sigma = 2.87$ and $\sigma = 3.31$. Because I estimate $\sigma \cdot \eta = 0.36$, the later $\sigma = 6$ yields a more conservative (i.e. smaller) estimate of the scale elasticity.
[67] Head and Mayer (2019) and (Praetorius, 2025) estimate $\sigma\eta = 0.21$ and $\sigma\eta = 0.27$ respectively, as well as final assembly tariff elasticities of 7.70 and 9.36.

Figure 16: Change in input production shares due to modularity across destinations



Notes: Figure shows the percent change in each country's share of total global input production due to the adoption of platform-sharing across destinations, plotted against log total cars assembled in the observed baseline equilibrium. Production share changes are obtained by comparing no-modularity and across-destinations-only counterfactual scenarios. Countries for which production shares rise (fall) are colored in green (brown). All data from *Marklines Automotive.*

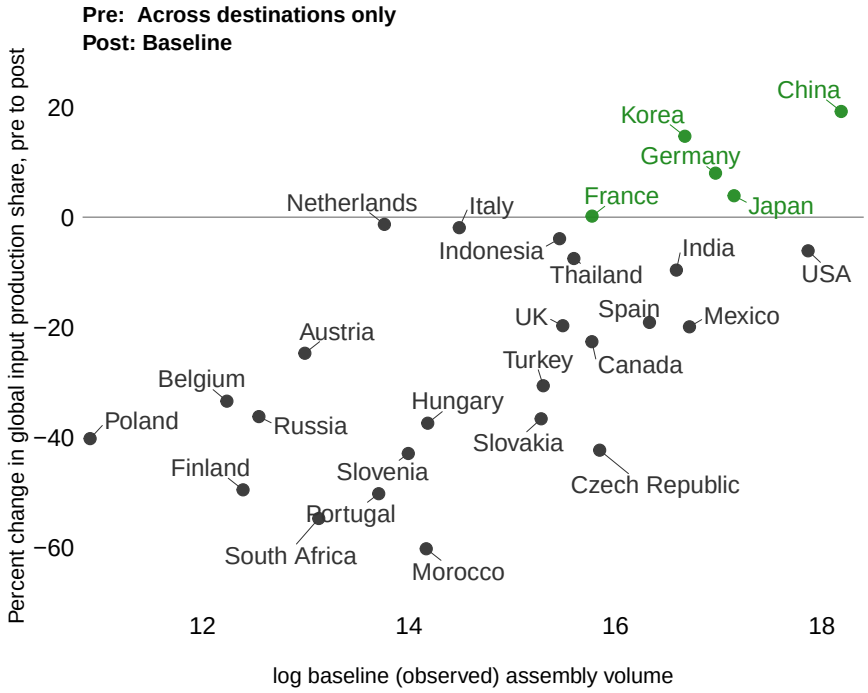
6.2 Quantification exercise: effects of observed modularity

To understand how observed levels of modularity change each country's share of global automotive input production, I compare three scenarios: (i) baseline adoption, (ii) a *destinations-only* counterfactual in which platforms that cover multiple segments are split into one platform per product segment, and (iii) a *no-modularity* counterfactual in which these segment-specific platforms are further split by destination. Comparing (ii) to (iii) gives the effects of the first phase (modularity across *destinations* within each segment), and comparing (i) to (ii) gives the effects of the second phase (modularity across *product segments*).

To understand if modularity concentrates aggregate input production *fewer* countries, consistent with increasing returns to scale, in each phase I calculate the share of countries for which input production shares rise (rather than fall). To understand if this production takes place in *large economies*, I correlate the change in production share with each country's input market size. To measure input market size, I use total baseline *assembly* volume summed across all firms, which is large both for large final car markets (such as United States and China) and trading partners that export many final cars to those economies (such as Mexico and Germany).

Figure 16 shows that *the first phase (modularity across destinations) creates large winners and losers,*

Figure 17: Change in input production shares due to modularity across product segments



Notes: Figure shows the percent change in each country's share of total global input production due to the adoption of platform-sharing across product segments, plotted against log total cars assembled in the observed baseline equilibrium. Production share changes are obtained by comparing baseline shares to an across-destinations-only counterfactual scenario. Countries for which production shares rise (fall) are colored in green (brown). All data from *Marklines Automotive.*

moderately increasing geographic concentration. 9 of 28 countries have larger input production shares (in green). At the extremes, three countries (South Africa, Slovakia, and Korea) experience proportional increases in market shares of above 20%, while three others (Finland, Slovenia, and Turkey) experience equally large declines. In the model, these shifts occur because countries no longer have a relative market size advantage in inputs for local assembly. Instead, production concentrates in large and productive assemblers in each firm-segment pair.

Because cars in each product segment are built in different countries (often due to local preferences), *production does not systematically shift to large markets.* For instance, because South Africa is a large assembler of Mercedes compact sedans and Toyota's pickup trucks, it produces and exports inputs for these firm-segment pairs despite having a small total market size. In contrast, two largest input markets (China and the USA) both lose production shares,

In contrast, Figure 17 shows that *the ongoing second phase (modularity across product segments) substantially increases geographic concentration in large markets.* 5 of 28 countries have larger input production shares – France, Germany, Korea, Japan, and China – and across all countries, the change between the destination-only scenario and baseline input production share is strongly correlated with market size.[68] Consistent with Proposition 2 in Section 4, which predicts that

[68]Results are qualitatively similar when plotting against counterfactual destinations-only assembly volume.

concentration occurs in each *firm's* largest markets, the five "winners" are all large assembly location for at least one firm. In contrast, the biggest losers are smaller countries (such as South Africa, Portugal, and Morocco) that are large in a firm-segment but not large markets (or productive input producers) overall. For instance, Morocco assembles a high share of affordable *Peugeot*-brand cars, but not larger vehicles (either for *Peugeot* or for other brands under its parent firm *Stellantis*).

Despite being a large market, the USA does not see higher shares in Figure 17 because, modularity across segments allows the small and midsize vehicles assembled in most countries to use common inputs, the car segments primarily built in the United States full-size SUVs and pickup trucks) use *body-on-frame* platforms due to their weight, rather than the *unibody* platforms of other vehicles. As a result, the US' large assembly volumes of pickup trucks do not give it a scale advantage in inputs used by other countries.

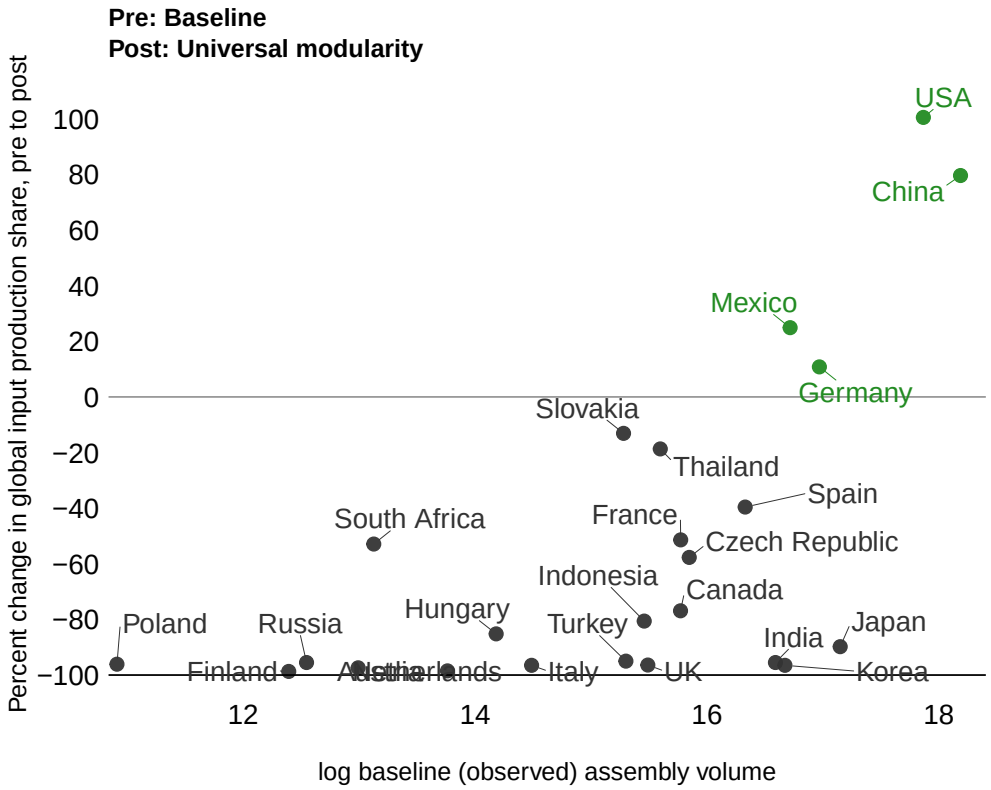
The effects in Figure 17 are intuitively similar to a de facto reduction in product variety. Even as final goods remain differentiated, inputs are not, so many countries source from the same large-scale (and therefore low-cost) suppliers.

6.3 Technological change counterfactual: effects of universal EV design platforms

I now show that future electrification-induced increases in modularity significantly magnify the concentration of global automotive supply chains in the world's largest economies. Specifically, relative to internal combustion engine (ICE) vehicles, electric vehicles have fewer distinct subsystems, more easily customized via software rather than hardware, and easier to balance because batteries lie under the car rather than in the front. Product differentiation – in size, shape, and performance – is consequently easier for EV platform designers to accommodate. Reflecting these features, the world's largest EV producer (*BYD*) uses a firmwide *universal* design platform called the *BYD e-Platform 3.0*. Similar universal platforms, such as the *Ford Universal EV Platform* and Volkswagen *Scalable Systems Platform*, will be introduced in 2027 and 2029 respectively.

Figure 18 shows that, in my baseline (largely pre-EV) sample, *universal design platforms significantly concentrate production in the United States and China* while reducing production shares in most other nations by over 50%. Only four of 28 countries see larger input production shares: the USA (101%), China (86%), Mexico (28%) which primarily assembles cars for Americans, and Germany (5%). These larger shares occur for two reasons. The first is total market size: for instance, 44 percent of all cars worldwide are built in China or the United States, and Germany is the largest assembler for many automakers (BMW, Mercedes, Volkswagen, and Stellantis). The second is composition: reflecting consumer tastes and firm comparative advantages, these countries strongly specialize in assembly of certain physically large segments (pickup trucks in the USA and Mexico, full-size SUVs in China, and executive cars in Germany) that use different platforms in the baseline data. With universal design platforms, assembly volume in these segments results in

Figure 18: Effects of universal EV design platforms on input production by country



Notes: Figure shows the percent change in each country's share of total global input production due to the adoption of firmwide universal design platforms, plotted against log total cars assembled in 2012-2022 in the observed baseline equilibrium. Production share changes are obtained by comparing a counterfactual with one platform per firm to baseline shares. Countries for which production shares rise (fall) are colored in green (brown). All data from *Marklines Automotive.*

greater scale advantages in producing inputs for all of a firm's segments. The flipside is that many smaller countries in which automotive manufacturing comprise a large share of manufacturing output – including many developing countries, such as Hungary, Thailand, Turkey, and India – see production shares fall by over 80%.

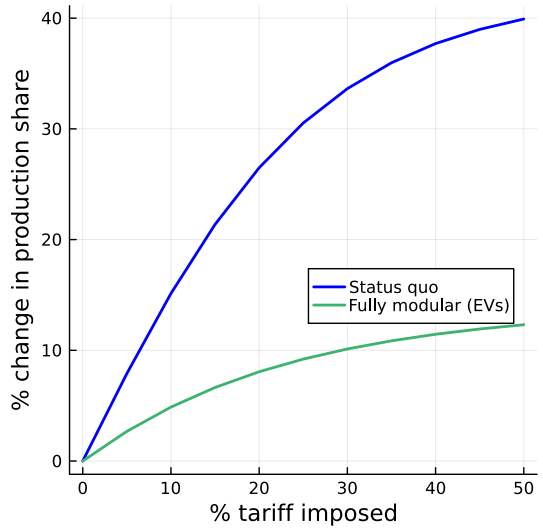
Figure 18 therefore suggests that *new ways of organizing EV production have winner-take-all implications*, rapidly induce industrialization in China and the United States while deindustrializing smaller markets.

6.4 Reshoring industrial policies: effects at baseline vs. universal design platforms

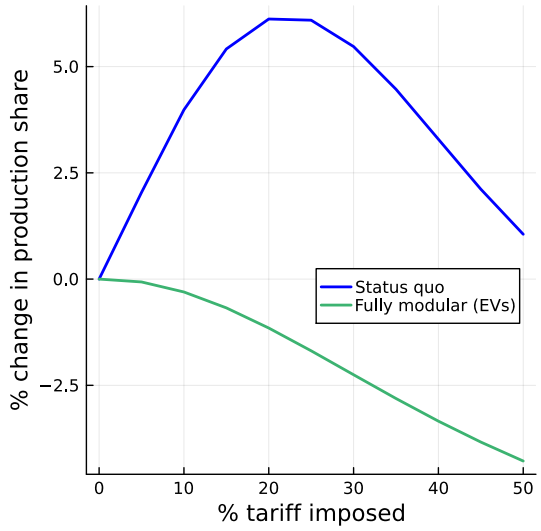
Does modularity also change the effectiveness of industrial policies intended to "reshore" automotive supply chains? In a final analysis I show that the answer is yes. Specifically, *universal platforms increase the relative effectiveness of supply chain "reshoring" policies that target China rather than all trading partners.*

Figure 19: Effects of tariffs on U.S. share of global input production (baseline vs. EV counterfactual)

Panel A: tariff on all countries

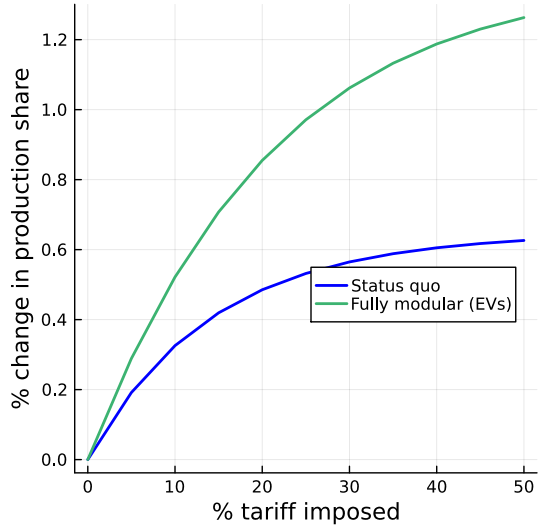


Imposed by USA only

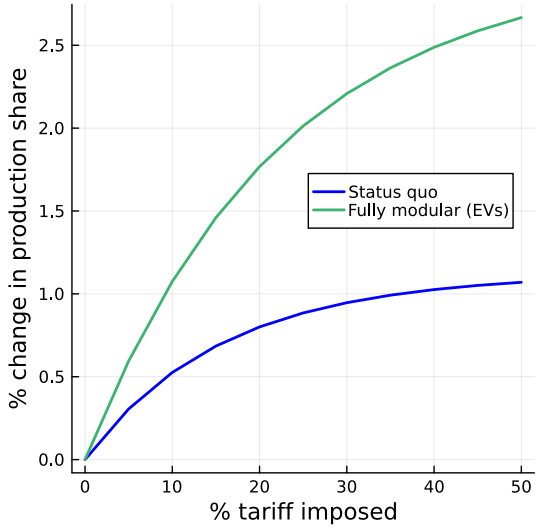


Imposed by all countries (trade war)

Panel B: tariff on China only



Imposed by USA only



Imposed by all other countries (coordinated)

Notes: Each plot shows the counterfactual percent change in the U.S. share of global input production (y-axis) due to tariff (x-axis), with no change in modularity relative to 2012-2022 (blue line) and in a counterfactual scenario in which each firm uses a single universal design platform (red line). Panel A considers a tariff on all countries, imposed by the U.S. only (left) or by all countries (right). Panel B considers a tariff on China only, imposed by the U.S. only (left) or by all other countries (right).

To show this, in Panel A of Figure 19 I study how US input production shares change due to counterfactual ad-valorem tariffs imposed by the US only (left) or all countries (right). The right panel shows that relative to baseline modularity (in blue), universal modularity (in green) weakens

the production relocation effects of US input tariffs on all countries. The right panel shows that the US can lose production share if all countries match the US tariff (a trade war scenario). Intuitively, tariffs weaken platform market size effects, which favor the US under universal modularity.

In contrast, universal modularity magnifies the effects on US input production of tariffs specifically on China. This is true for unilateral tariffs imposed only by the US (left panel); as well as for coordinated tariffs imposed by all countries (right panel), simulating the cascading protectionism that is occurring in many countries to avoid diversion of Chinese imports goods from the US to third countries. This effect occurs because China and the US are the two large-scale (and thus low-cost) producers, so input production that relocates from China due to tariffs (for final products both in the US and third countries) is especially likely to concentrate in the US, which further reinforces its relative scale advantage.

7 Conclusion

Using the global automotive sector as a laboratory, this paper shows that modular product design magnifies the scale advantages of large economies in high-tech input production. Production shifts in two phases: modularity first *increases trade*, concentrating input production in large and productive markets for each product segment; and then *homogenizes* sourcing, weakens specialization in locally-assembled segments. These shifts occur because inputs are platform- rather than product-specific. As a result, capitalizing on economies of scale, countries specialize in platforms for which they are large assemblers (and thus large input markets).

In aggregate, these forces imply that modularity concentrates production in large assemblers, and that impending universal EV design platforms will concentrate automotive supply chains in China and the United States while deindustrializing most other countries. As a result, modularity alters the returns to industrial policy: universal modularity weakens the production relocation effects of broad U.S. tariffs on all countries, but magnifies the effects of tariffs targeting China alone.

My results suggest that the ongoing transformation of complex product design may substantially reshape opportunities for industrial development. For instance, two traditional pathways to industrialization – backward linkages from local assembly to input production, and export-led growth via specialization in affordable product segments – become more difficult as inputs are no longer specific to countries or segments. Understanding how new design technologies interact with countries' existing demand structures and industrial policies will therefore be critical for policymakers navigating the transition to electric vehicles, as well as to other new technologies – from artificial intelligence to mRNA vaccines to solar panels – that follow principles of modular design.

References

Adhvaryu, A., V. Bassi, A. Nyshadham, J. A. Tamayo, and N. Torres (2023): "Organizational responses to product cycles," Tech. rep., National Bureau of Economic Research.

Adhvaryu, A., J.-F. Gauthier, A. Nyshadham, and J. Tamayo (2024): "Absenteeism, Productivity, and Relational Contracts Inside the Firm," *Journal of the European Economic Association*, 22, 1628–1677.

Agrawal, A., J. S. Gans, and A. Goldfarb (2024): "Artificial intelligence adoption and system-wide change," *Journal of Economics & Management Strategy*, 33, 327–337.

Alesina, A. and E. Spolaore (1997): "On the Number and Size of Nations," *The Quarterly Journal of Economics*, 112, 1027–1056.

Alesina, A., E. Spolaore, and R. Wacziarg (2005): *Trade, Growth and the Size of Countries*, Elsevier, 1499–1542.

Alfaro, L., P. Conconi, F. Kamal, and Z. Kroff (2025): *Trade within Multinational Boundaries*.

Alfaro-Ureña, A., I. Manelici, and J. P. Vasquez (2022): "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages," *The Quarterly Journal of Economics*, 137, 1495–1552.

Antràs, P., E. Fadeev, T. C. Fort, and F. Tintelnot (2024): "Exporting, Global Sourcing, and Multinational Activity: Theory and Evidence from the United States," *The Review of Economics and Statistics*, 1–48.

Argente, D., S. Moreira, E. Oberfield, and V. Venkateswaran (2025): *Scalable Expertise: How Standardization Drives Scale and Scope*.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.

Arrow, K. J. (1962): "The Economic Implications of Learning by Doing," *The Review of Economic Studies*, 29, 155–173.

Atalay, E., A. Hortaçsu, and C. Syverson (2014): "Vertical Integration and Input Flows," *American Economic Review*, 104, 1120–48.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The Fall of the Labor Share and the Rise of Superstar Firms*," *The Quarterly Journal of Economics*, 135, 645–709.

Bai, J., P. J. Barwick, S. Cao, and S. Li (2025): "Quid Pro Quo, Knowledge Spillovers, and Industrial Quality Upgrading: Evidence from the Chinese Auto Industry," *American Economic Review*, 115, 3825–3852.

Balassa, B. (1978): "Exports and economic growth: further evidence," *Journal of Development Economics*, 5, 181–189.

Baldwin, C. Y. and K. B. Clark (2000): *Design Rules: The Power of Modularity*, The MIT Press.

Barwick, P. J., H.-S. Kwon, S. Li, and N. Zahur (2025): *Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry*.

Bassi, V., J. H. Lee, A. Peter, T. Porzio, R. Sen, and E. Tugume (2023): *Self-Employment Within the Firm*.

Boehm, J., S. Dhingra, and J. Morrow (2022): "The Comparative Advantage of Firms," *Journal of Political Economy*, 130, 3025–3100.

Borusyak, K., X. Jaravel, and J. Spiess (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 91, 3253–3285.

Cajal-Grossi, J., R. Macchiavello, and G. Noguera (2023): "Buyers' Sourcing Strategies and Suppliers' Markups in Bangladeshi Garments," *The Quarterly Journal of Economics*, 138,

2391â€“2450.

Callaway, B. and P. H. C. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225, 200–230.

Case, S., J. J. Michalek, and K. S. Whitefoot (2023): "Global Product Design Platforming: A Comparison of Two Equilibrium Solution Methods," *Journal of Mechanical Design*, 145, 061702.

Castro-Vincenzi, J. (2022): "Climate hazards and resilience in the global car industry," *Princeton University manuscript*.

Castro-Vincenzi, J. M., E. Menaguale, E. Morales, and A. Sabal (2024): "Market Entry and Plant Location in Multi-Product Firms," Working paper / technical report, University of Chicago and Princeton University, available at: https://static1.squarespace.com/static/5fbd5c064c271a353f8a9840/t/67dc5f3d98a97e61e20e8b3f/1742495550442/cmms.pdf.

Conlon, C. and J. Gortmaker (2020): "Best practices for differentiated products demand estimation with PyBLP," *RAND Journal of Economics*, 51, 1108–1161.

Costinot, A., D. Donaldson, M. Kyle, and H. Williams (2019): "The More We Die, The More We Sell? A Simple Test of the Home-Market Effect*," *The Quarterly Journal of Economics*, 134, 843–894.

Coşar, A. K., P. L. Grieco, S. Li, and F. Tintelnot (2018): "What drives home market advantage?" *Journal of International Economics*, 110, 135–150.

Cusumano, M. and K. Nobeoka (1998): *Thinking Beyond Lean: How Multi-project Management is Transforming Product Development at Toyota and Other Companies*, MIT international motor vehicle program, Free Press.

de Weck, O. L., E. S. Suh, and D. Chang (2003): "Product Family and Platform Portfolio Optimization," in *Proceedings of the ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Chicago, IL, USA: ASME, DETC2003/DAC-48721, 175–185.

Dekle, R., J. Eaton, and S. Kortum (2007): "Unbalanced Trade," *American Economic Review*, 97, 351–355.

Diao, X., M. Ellis, M. McMillan, and D. Rodrik (2024): "Africa's Manufacturing Puzzle: Evidence from Tanzanian and Ethiopian Firms," *The World Bank Economic Review*, 39, 308–340.

Ding, X. (2023): *Industry linkages from joint production*, US Census Bureau, Center for Economic Studies.

Dingel, J., J. Gottlieb, M. Lozinski, and P. Mourot (2023): *Market Size and Trade in Medical Services*.

Dingel, J. I. (2016): "The Determinants of Quality Specialization," *The Review of Economic Studies*, rdw054.

Doeringer, P. B. and M. J. Piore (1971): *Internal Labor Markets and Manpower Analysis*, D.C. Heath and Co.

Eaton, J. and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

Egelman, C. D., D. Epple, L. Argote, and E. R. H. Fuchs (2017): "Learning by Doing in Multiproduct Manufacturing: Variety, Customizations, and Overlapping Product Generations," *Management Science*, 63, 405–423.

Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011): "Income Distribution, Product Quality, and International Trade," *Journal of Political Economy*, 119, 721–765.

Field, D. (2025): "How Ford's Software Vision is Redefining the Road Ahead for Customers," https://www.fromtheroad.ford.com/eur/en/articles/2025/ford-vehicles-software-vision, accessed: YYYY-MM-DD.

Fontagné, L., H. Guimbard, and G. Orefice (2022): "Tariff-based product-level trade elasticities," *Journal of International Economics*, 137, 103593.

Freund, C. and S. Oliver (2015): "Gains from harmonizing us and eu auto regulations under the transatlantic trade and investment partnership," *Policy Brief*, 15–10.

Galdin, A. (2024): "Resilience of global supply chains and generic drug shortages," *Princeton University manuscript*.

Gardner, J. (2022): "Two-Stage Differences in Differences," *The R Journal*, 14, 162–173.

Goldberg, P. K. (1995): "Product differentiation and oligopoly in international markets: The case of the US automobile industry," *Econometrica: Journal of the Econometric Society*, 891–951.

Goldberg, P. K., R. Juhász, N. J. Lane, G. L. Forte, and J. Thurk (2024): "Industrial policy in the global semiconductor sector," Tech. rep., National Bureau of Economic Research.

Goldberg, P. K. and T. Reed (2023): "Presidential Address: Demand-Side Constraints in Development. The Role of Market Size, Trade, and (In)Equality," *Econometrica*, 91, 1915–1950.

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

Grossman, G. M., P. McCalman, and R. W. Staiger (2021): "The "New" Economics of Trade Agreements: From Trade Liberalization to Regulatory Convergence?" *Econometrica*, 89, 215–249.

Grossman, G. M. and E. Rossi-Hansberg (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978â€"1997.

Hanson, G. H. and C. Xiang (2004): "The Home-Market Effect and Bilateral Trade Patterns," *American Economic Review*, 94, 1108–1129.

Head, K. and T. Mayer (2019): "Brands in Motion: How Frictions Shape Multinational Production," *American Economic Review*, 109, 3073–3124.

Head, K., T. Mayer, and M. Melitz (2024): "The Laffer curve for rules of origin," *Journal of International Economics*, 150, 103911.

Head, K., T. Mayer, M. Melitz, and C. Yang (2025): "Industrial policies for multi-stage production: The battle for battery-powered vehicles," Tech. rep., Working Paper.

Helper, S., J. P. MacDuffie, F. Pil, M. Sako, A. Takeishi, and M. Warburton (2003): "Modularization and Outsourcing: Implications for the Future of Automotive Assembly," Tech. rep., Massachusetts Institute of Technology, Project Report to IMVP.

Helper, S. and A. Munasib (2022): "Economies of scope and relational contracts: Exploring global value chains in the automotive industry," BEA Working Papers 0195, Bureau of Economic Analysis.

Helpman, E. and P. R. Krugman (1985): *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*, Cambridge, MA: MIT Press.

Henderson, R. M. and K. B. Clark (1990): "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms," *Administrative Science Quarterly*, 35, 9.

Hirschman, A. O. (1958): "The strategy of economic development," .

Hounshell, D. A. (1984): *From the American System to Mass Production, 1800–1932: The Development of Manufacturing Technology in the United States*, Studies in Industry and Society, Baltimore: Johns Hopkins University Press.

Irwin, D. A. and P. J. Klenow (1994): "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, 102, 1200–1227.

Jacobides, M. G., C. Cennamo, and A. Gawer (2018): "Towards a theory of ecosystems," *Strategic Management Journal*, 39, 2255–2276.

Juhász, R. and C. Steinwender (2018): "Spinning the web: The impact of ICT on trade in intermediates and technology diffusion," Tech. rep., National Bureau of Economic Research.

Khandelwal, A. K. (2010): "The Long and Short (of) Quality Ladders," *The Review of Economic Studies*, 77, 1450–1476.

Kikuchi, S. (2025): "Does Skill Abundance Still Matter? The Evolution of Comparative Advantage in the 21st Century," SSRN working paper.

Klier, T. H. and J. M. Rubenstein (2019): "What Do We Know About the Origin of Parts in Vehicles Sold in the U.S. Market?" Blog post.

Krugman, P. (1980): "Scale economies, product differentiation, and the pattern of trade," *American economic review*, 70, 950–959.

Krugman, P. R. (1979): "Increasing returns, monopolistic competition, and international trade," *Journal of International Economics*, 9, 469–479.

Kuan, J. and J. West (2023): "Interfaces, Modularity and Ecosystem Emergence: How DARPA Modularized the Semiconductor Ecosystem," *Research Policy*, 52, 104789.

Kucheryavyy, K., G. Lyn, and A. Rodríguez-Clare (2023): "Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale," *American Economic Journal: Macroeconomics*, 15, 372–412.

Lamy, P. (2016): "The Changing Landscape of International Trade: The Frank D. Graham Lecture, Princeton University – 7 April 2016," https://pascallamyeu.files.wordpress.com/2017/02/2016-04-07-lamy-princeton-graham-lecture-final.pdf.

Levitt, S. D., J. A. List, and C. Syverson (2013): "Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant," *Journal of Political Economy*, 121, 643–681.

Linder, S. (1961): *An Essay on Trade and Transformation*, Almqvist & Wiksell Boktr.

MacDuffie, J. P. (2013): "Modularity-as-Property, Modularization-as-Process, and 'Modularity'-as-Frame: Lessons from Product Architecture Initiatives in the Global Automotive Industry," *Global Strategy Journal*, 3, 8–40.

Maggi, G. and M. Mrázová (2024): "Harmonization... What Else? The Role for International Regulatory Agreements," Tech. rep., National Bureau of Economic Research.

Matouschek, N., M. Powell, and B. Reich (2025): "Organizing Modular Production," *Journal of Political Economy*, 133, 986–1046.

Murphy, K. M., A. Shleifer, and R. W. Vishny (1989): "Industrialization and the Big Push," *Journal of Political Economy*, 97, 1003–1026.

OUSD (2025): "Technical Highlight: Modular Open Systems Approach (MOSA)," Se&a info sheet, Office of the Under Secretary of Defense for Research & Engineering, Washington, DC, distribution Statement A. DOPSR Case # 25-T-1439.

Praetorius, S. (2025): "Collaboration in Technology and Multinational Production," .

Rauch, J. E. (1999): "Networks versus markets in international trade," *Journal of International Economics*, 48, 7–35.

Rodrik, D. (2015): "Premature deindustrialization," *Journal of Economic Growth*, 21, 1–33.

Sabal, A. (2025): "Product Entry in the Global Automobile Industry," Tech. rep., January 2025. Working paper.

Sako, M. (2005): *Modularity and Outsourcing: The Nature of Co-evolution of Product Architecture and Organization Architecture in the Global Automotive Industry*, Oxford University PressOxford, 229–253.

Sako, M. and F. Murray (1999): "Modular strategies in cars and computers," .

Sako, M. and M. Warburton (1999): "Modularization and Outsourcing Project: Preliminary

Report of European Research Team," Tech. rep., Massachusetts Institute of Technology, International Motor Vehicle Programme (IMVP).

Schott, P. K. (2004): "Across-Product Versus Within-Product Specialization in International Trade," *The Quarterly Journal of Economics*, 119, 647–678.

Sturgeon, T. J. (2002): "Modular production networks: a new American model of industrial organization," *Industrial and Corporate Change*, 11, 451–496.

Sun, L. and S. Abraham (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225, 175–199.

Thompson, P. (2012): "The Relationship between Unit Cost and Cumulative Quantity and the Evidence for Organizational Learning-by-Doing," *Journal of Economic Perspectives*, 26, 203–224.

Thun, E., D. Taglioni, T. J. Sturgeon, and M. P. Dallas (2025): "Massive Modular Ecosystems: A Framework for Understanding Complex Industries in the Digital Age," Policy Research Working Paper Series 11106, The World Bank.

Ulrich, K. and S. D. Eppinger (1995): "Product design and development," .

Wilkins, M. and F. E. Hill (1964): *American business abroad: Ford on six continents*, Cambridge University Press.

A Appendix Tables and Figures

Figure A.1: Example of divergent marketing materials for two cars that share a design platform



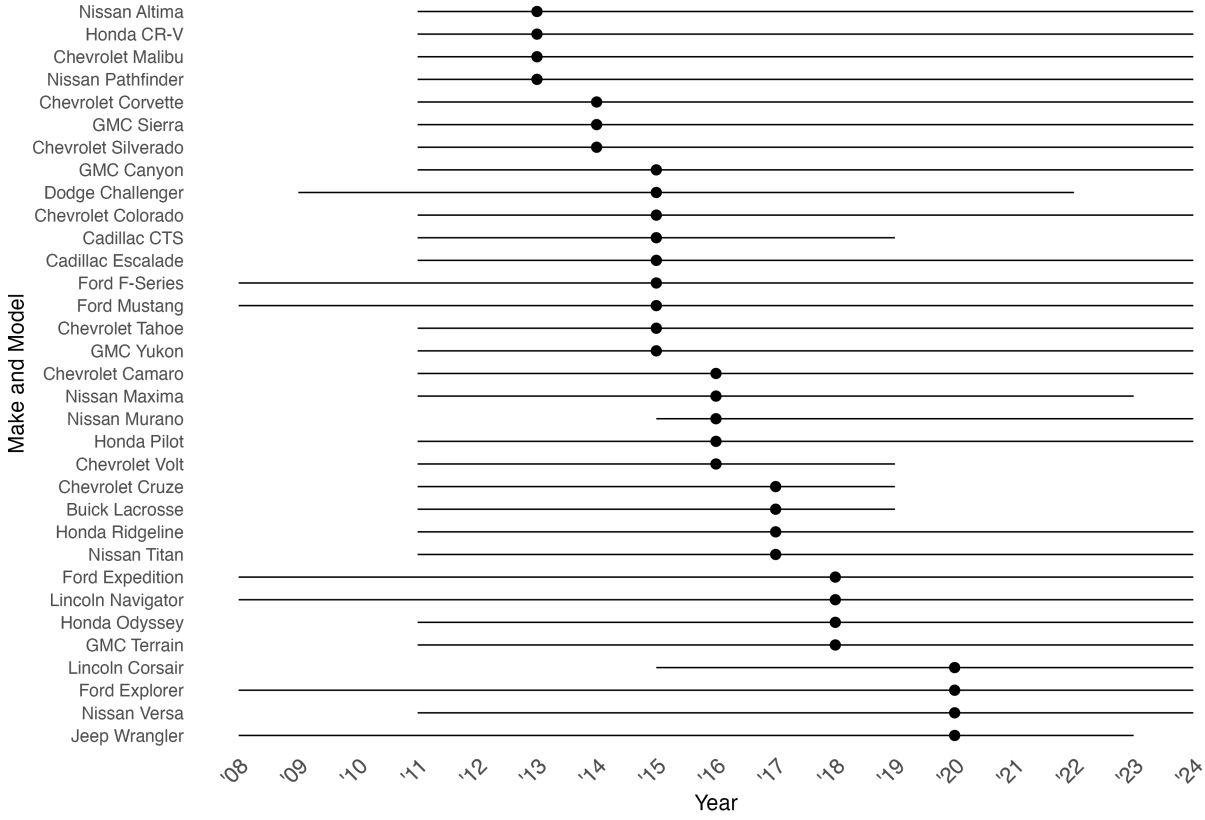
HR-V (USA)



City (India and SEA)

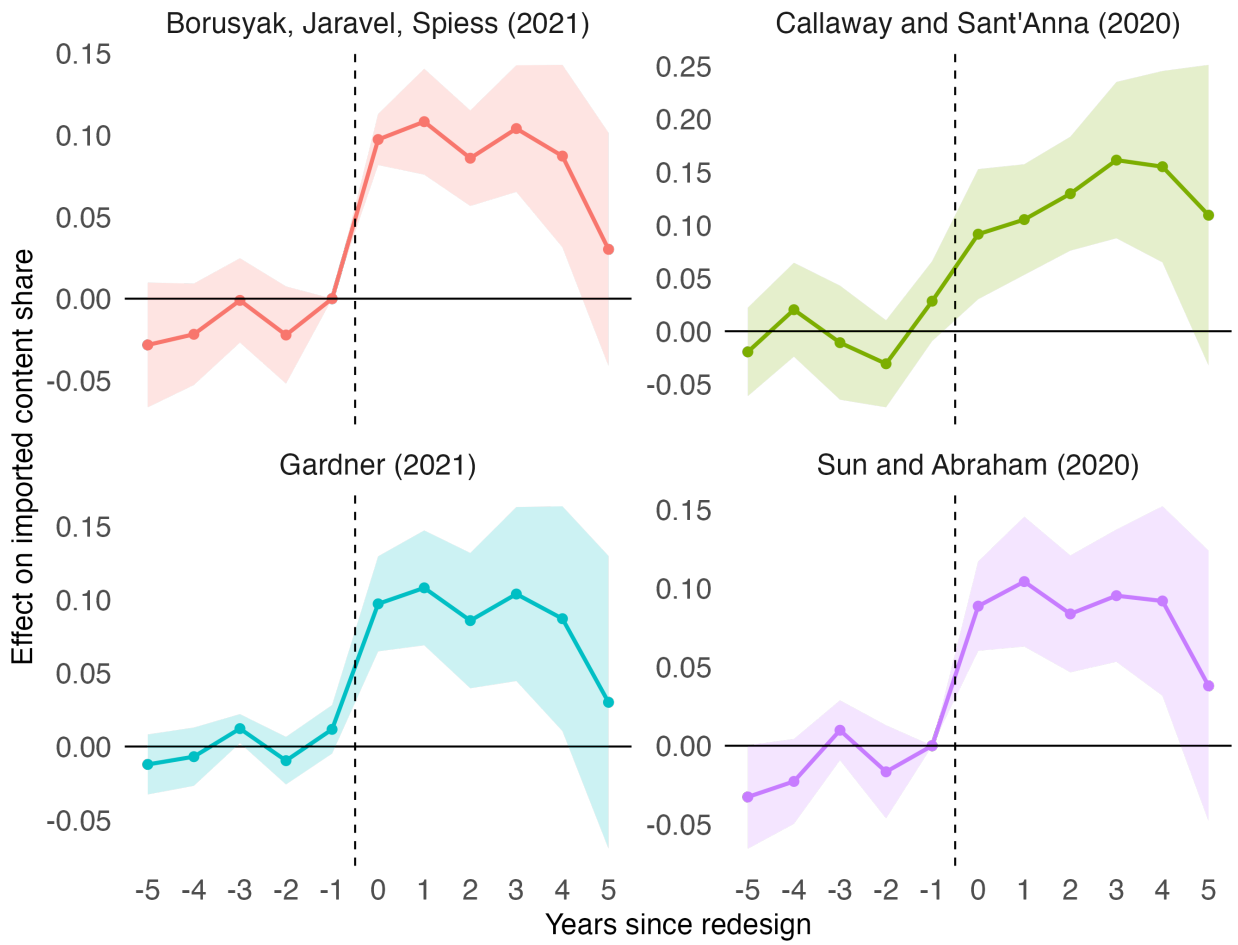
Notes: Figure shows images from official *Honda* marketing materials for the US-market *Honda HR-V* SUV and the Thailand-market *Honda City* sedan. These two cars share the *Honda Global Small Car Platform* despite being differentiated by product segment and destination.

Figure A.2: Placebo modular design adoption dates



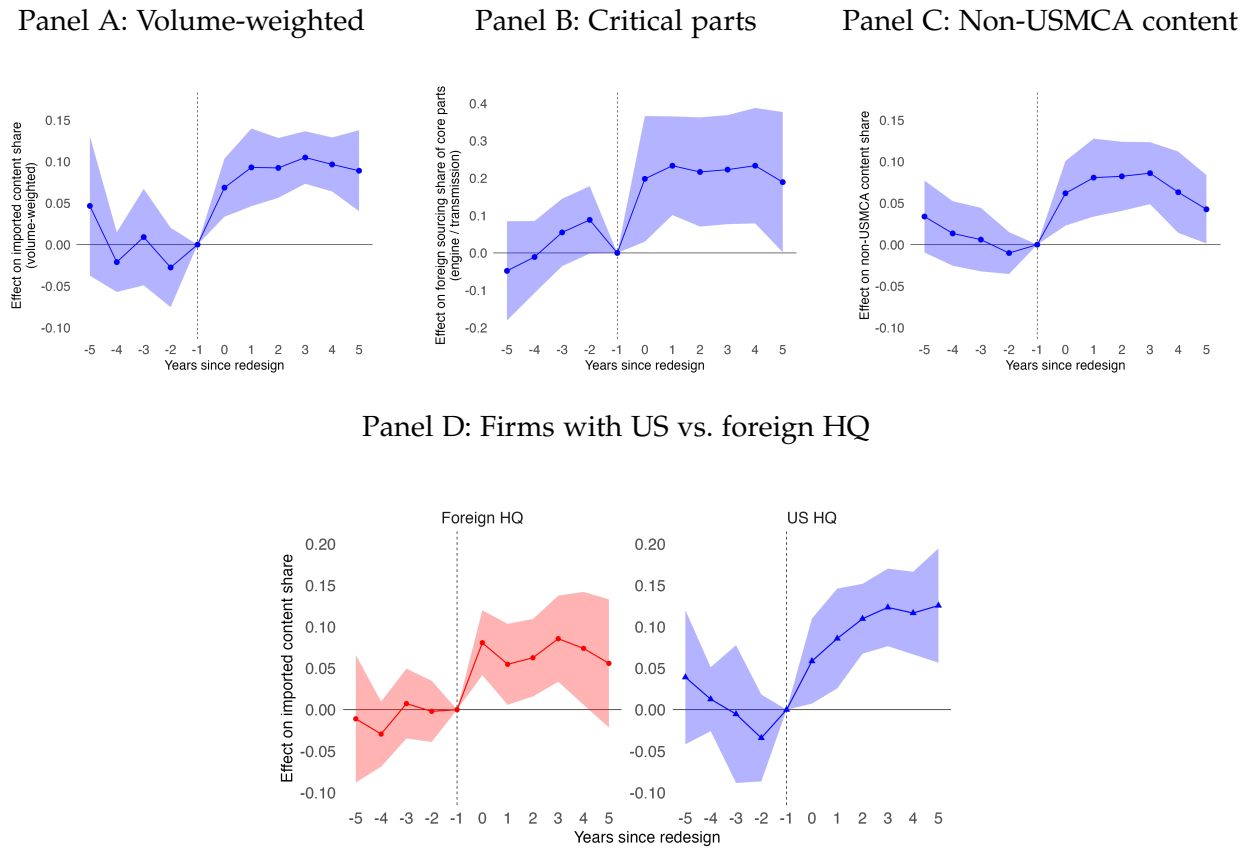
Notes: Figure shows all placebo adoption events in which an existing North America-assembled model is redesigned, and the new design platform is *not* shared with models assembled in an additional region (in blue) or product segment (in red). Black lines indicate years in which each model is assembled. Sample is restricted to models with at least one non-placebo adoption event as in Figure 9. Event years and redesign type are from *Wards* and *Marklines Automotive*. Sourcing information is from *American Automotive Labeling Act* reports for 2008-2024.

Figure A.3: Effects of modularity across destinations using staggered treatment-robust DID methods



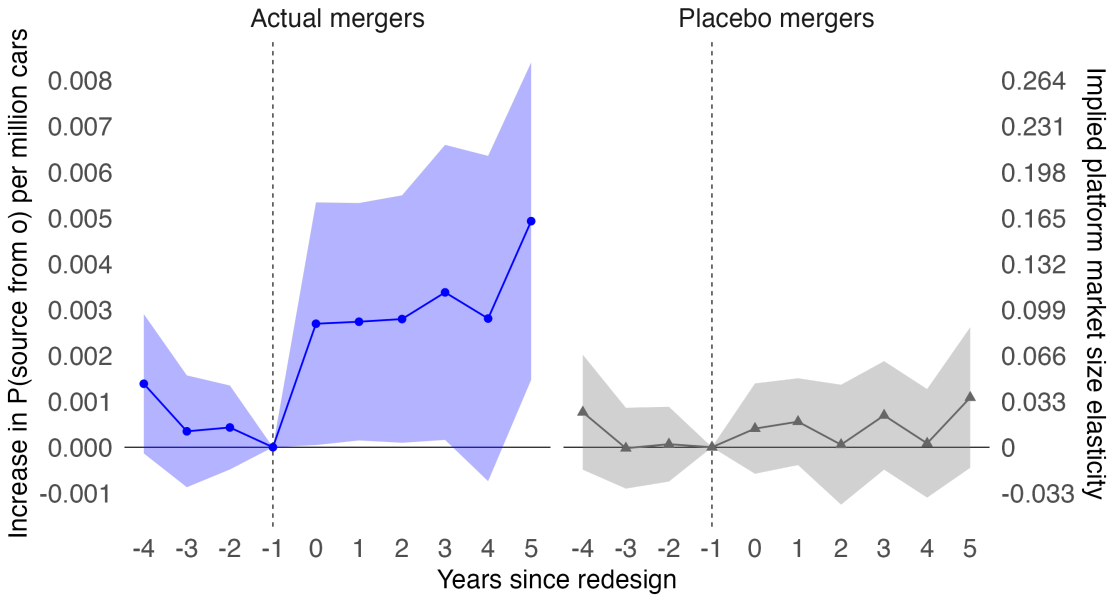
Notes: Panels show effect of adopting a platform that is shared by car models assembled in other destinations on the import share at the car model-year level using several different difference-in-difference estimators that allow for treatment effect heterogeneity and staggered treatment. All effects at the model-year level. Import share data from the American Automotive Labeling Act. Redesign year, platform information, and locations of platform use from *Marklines*.

Figure A.4: Effects of modularity across destinations on import share: robustness



Notes: Panels show effect of adopting a platform that is modular across destinations on the import share. In Panel A, each panel unit (car model) is weighted by total assembly volume over the analysis period. In Panel B, the regression is at the model-year-critical part (engine or transmission) level, and the outcome is the probability of sourcing the critical part from a foreign origin (rather than the foreign share of all parts). In Panel C, the outcome is the share of content from outside the USMCA region (i.e. from neither USA and Canada nor Mexico). In Panel D, the sample is split into firms with US headquarters (Ford and GM) and all other firms. Model-level import share and engine and transmission origin data are from the American Automotive Labeling Act. Redesign years and platform information are from *Marklines Automotive* and *Wards Automotive*.

Figure A.5: Sourcing increases from origins with higher platform market size due to merger (asinh)



Notes: Panels show effect of changes in the inverse hyperbolic sine (IHS) of predicted platform market size (for each car j at firm f, proxied by the pre-merger assembly volume in segment $s(j)$ in origin o within the other firm f') on imports a core input (engine and/or transmission). All effects at the model-year-origin country level. In the left panel, coefficients interact time-since-redesign with the difference in predicted platform scale due to the mergers, assuming that platforms remain largely within-segment. In the right panel, coefficients interact time-since-redesign with differences due to placebo mergers that were proposed but did not take place. Import share data from the American Automotive Labeling Act. Redesign year, platform information, and locations of platform use are from *Marklines*.

Table A.1: Modularity weakens segment-level specialization: core components only

Dependent Variable:	# of supply relationships with origin o firms for model m (PPML)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Segment s share of origin o assembly	0.424***	0.697**	0.876***			
	(0.111)	(0.288)	(0.220)			
Segment s share of origin o assembly × 𝟙[Platform has ≥ 1 other s′ ≠ s]		-0.313				
		(0.307)				
Segment s share of origin o assembly × Share of platform models in s′ ≠ s			-0.876**			
			(0.373)			
Origin o RCA in segment s inputs				0.869***	1.048***	1.033***
				(0.114)	(0.112)	(0.120)
Origin o RCA in segment s inputs × Share of platform models in s′ ≠ s					-0.365***	
					(0.081)	
Origin o RCA in segment s inputs × 𝟙[Platform has ≥ 1 other s′ ≠ s]						-0.205***
						(0.066)
Fixed-effects						
Model-destination-part group	Yes	Yes	Yes	Yes	Yes	Yes
Origin-destination-firm-part group	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Dependent variable mean	1.7946	1.7946	1.7946	1.6242	1.6242	1.6242
Observations	26,362	26,362	26,362	29,481	29,481	29,481

Clustered (Origin-segment-platform) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows results from gravity regressions (Equations (19) and (21)) restricted to engine, drivetrain, and chassis components. Origin *o* RCA is the revealed cost advantage of country *o* in segment *s* as estimated in Equation (20). All data from *Marklines Automotive*.

Table A.2: Modularity weakens segment-level specialization: differentiator components only

Dependent Variable:	# of supply relationships with origin o firms for model m (PPML)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Segment s share of origin o assembly	0.356**	0.740*	0.836***			
	(0.152)	(0.392)	(0.315)			
Segment s share of origin o assembly × 𝟙[Platform has ≥ 1 other s′ ≠ s]		-0.444				
		(0.417)				
Segment s share of origin o assembly × Share of platform models in s′ ≠ s			-0.928*			
			(0.521)			
Origin o RCA in segment s inputs				0.754***	0.767***	0.738***
				(0.137)	(0.154)	(0.166)
Origin o RCA in segment s inputs × Share of platform models in s′ ≠ s					-0.027	
					(0.139)	
Origin o RCA in segment s inputs × 𝟙[Platform has ≥ 1 other s′ ≠ s]						0.019
						(0.112)
Fixed-effects						
Model-destination-part group	Yes	Yes	Yes	Yes	Yes	Yes
Origin-destination-firm-part group	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Dependent variable mean	1.1508	1.1508	1.1508	1.0712	1.0712	1.0712
Observations	15,128	15,128	15,128	16,527	16,527	16,527

Clustered (Origin-segment-platform) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

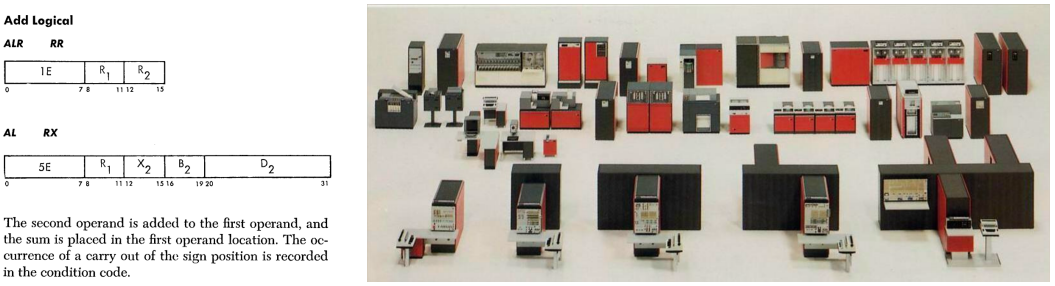
Notes: Table shows results form gravity regressions (Equations (19) and (21)) restricted to interior, exterior, and generic small parts. Origin *o* RCA is the revealed cost advantage of country *o* in segment *s* as estimated in Equation (20). All data from *Marklines Automotive*.

B Additional context on design platforms and data collection

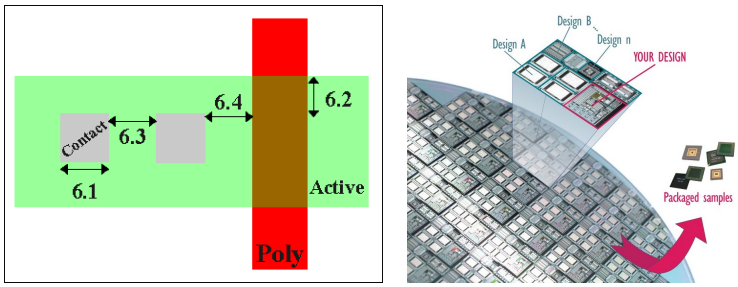
B.1 History of modular design in mainframe computers (*IBM System/360*) and semiconductors (*Mead-Conway Rules*)

Figure B.1: Modular design rules (left) and products (right) in two industries

Panel A: Mainframe computers (IBM System/360)



Panel B: Semiconductors (Mead-Conway rules)



Notes: Panel A shows a sample instruction set for the IBM System/360 architecture (left) and a marketing image of five System/360 mainframes of different sizes (right, foreground) as well as associated interchangeable components (right, background). Panel B shows an example of the Mead-Conway λ integrated circuit design rules (left) and a marketing image of a multi-project wafer service in which multiple Mead-Conway rule-compliant chips are fabricated on the same wafer.

The key benefit of modular design for multiproduct firms – shared economies of scale across products – is best illustrated by its early application at *IBM* in the 1960s (Baldwin and Clark, 2000). As the leading producer of mainframe computers, *IBM* sold separate mainframes for accounting, scientific, factory, and educational use. Each product was designed from scratch by a separate engineering team.[69] As a result, *IBM* could not leverage shared increasing returns, or synergies, across products because the core components of each were technically incompatible.

This changed in 1964, when *IBM* redesigned all products to use common data types, instruction sets, and other rules collectively called the *System/360* architecture, and a compatible set of chips, storage units, user manuals, and other components. These are shown in Panel A of Figure B.1. By sharing these *System/360*-compatible components across its products, *IBM* both avoided fixed costs of inventing technology-intensive inputs (such as software), and could produce these inputs with greater economies of scale.[70]

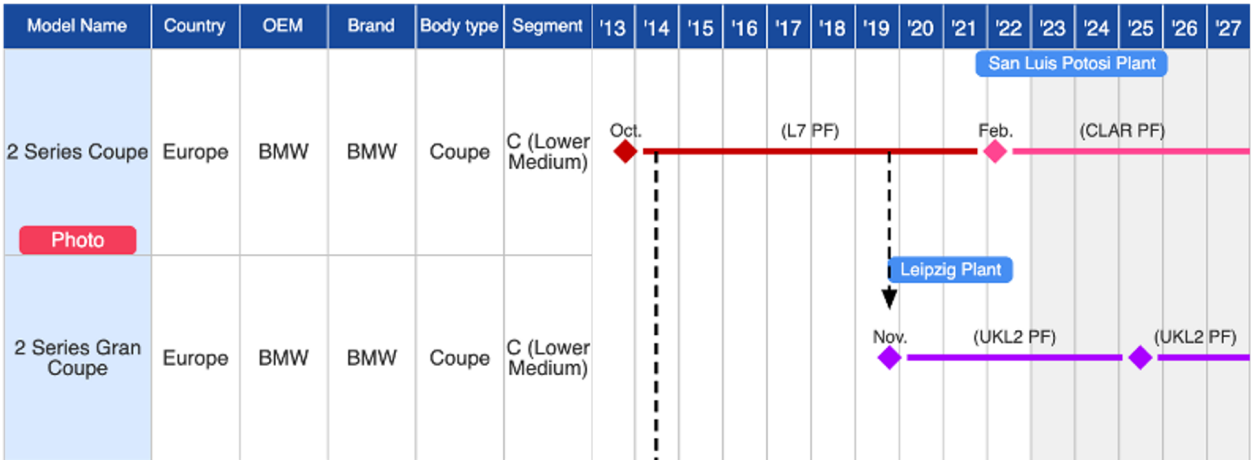
[69] This bespoke non-modular design process is known as *integral design*.

[70] Modular design also confers several other benefits beyond the sharing of input scale economies, including easier outsourcing through codification (Baldwin 2008), increases in product variety, and reduced buyer bargaining power

At another early adopter – the semiconductor industry – modularity also transformed where production took place. Design rules and associated products for the semiconductor industry are shown in Panel B of Figure B.1. Semiconductor design rules were developed in the 1970s, not by a firm, but by professors Carver Mead and Lynn Conway to teach engineering students at Caltech and MIT (Kuan and West, 2023).[71] These design rules transformed the geography of the chip industry: many distinct chip designs could now be fabricated using the same silicon wafers and during the same production run. These commonalities allowed chip designers to share learning-by-doing (Irwin and Klenow, 1994) and other production scale benefits as long as production took place in a single location (often TSMC's plant in Taiwan).

B.2 Product timeline example

Figure B.2: *Marklines* product timeline example



Notes: Figure shows original product timeline images from *Marklines* for two entirely different products, the *BMW 2 Series Coupe* and *BMW 2 Series Gran Coupe*. Diamonds indicate years of major product redesigns, and text in parentheses indicate platform names. For instance, the *2 Series Coupe* adopts the *CLAR Platform* in 2022.

B.3 Procedure for calculation of platform-sharing probabilities

This section describes how I obtain the platform-sharing probabilities plotted in Figure 4 and Figure 5. Formally, for any firm f, characteristic k (either segment or destination), and two values x and x' of that characteristic, I calculate the share of products with value x that share a design platform with at least one product with value x':

$$Share_k(x, x') = \frac{1}{|\{m : X_{mk} = x\}|} \sum_{m: X_{mk}=x} \mathbb{1}\Big[\big|\{m' : X_{m'k} = x', G(m) = G(m')\}\big| > 0\Big] \quad (25)$$

because inputs can be used by other buyers. These forces are less relevant in my context: single-model design rules, though not shared platforms, were widely used by automakers by the late 1990s (Sako 2003, Sako and Murray 1999), the most popular car models are kept consistent over time to maintain consumer recognition, and platform-sharing across firms is rare and requires a formal licensing agreement (Praetorius 2025) for intellectual property reasons.

[71] The Mead-Conway rules were scalable: all geometric distances were expressed in terms of a factor λ. This meant that designers working in isolation could ensure allowing chips of any size simply by adjusting λ.

where X_{jk} is the value of characteristic k for car model m, and $g = G(m)$ is the platform (or "group") used by model m.

For instance, $Share_{\text{destination}}(\text{USA}, \text{China})$ denotes the probability that a US-market car shares a platform with at least one car sold in China. Similarly, $Share_{\text{segment}}(\text{C}, \text{SUV-D})$ is the probability that a compact sedan shares a platform with at least one midsize SUV.

B.4 Discussion of the costs to engineers of multi-destination and multi-segment design platforms, and of EV vs. ICE platforms

The observed patterns of adoption in Figure 4 and Figure 5 – first across destinations within a segment, then with nearby segments, and finally across all segments due to EVs – are broadly consistent with details of automotive engineering. Specifically, *segment (destination) dissimilarity is more (less) difficult for engineers to accommodate.* Multi-destination design platforms are relatively simple to develop because destination-specific regulations on bumpers, seatbelts, airbags, and lights do not interact with the complex subsystems that enable a car's movement (Freund and Oliver, 2015). In contrast, multi-segment design platforms are more challenging, because segments differ in many fundamental ways – e.g., size, weight, and engine power – and thus require (i) physically distinct drivetrains (engines) and (ii) significant coordination across subsystem teams.[72]

These difficulties are significantly lower for EVs, which have (i) drivetrains (i.e. batteries and motors) that can be programmed for product-specific requirements, and (ii) have 40 percent fewer unique parts than ICE vehicles, which reduces coordination costs.[73] A final difference is that EVs use "skateboard" designs with batteries under the vehicle, so heavier vehicles simply stack more batteries without affecting balance, something that is impossible for ICE vehicles with engines in the front of the car.

As a result, as the costs to engineers of codifying design rules for heterogeneous products has gradually fallen, automakers first "merged" platforms for products in different destinations (but the same product segment) because doing so was not too onerous for engineers, then merging platforms for similar (but not all) product segments in later model generations, and finally are merging all platforms into one (per firm) for electric vehicles.

B.5 *American Automotive Labeling Act* panel construction

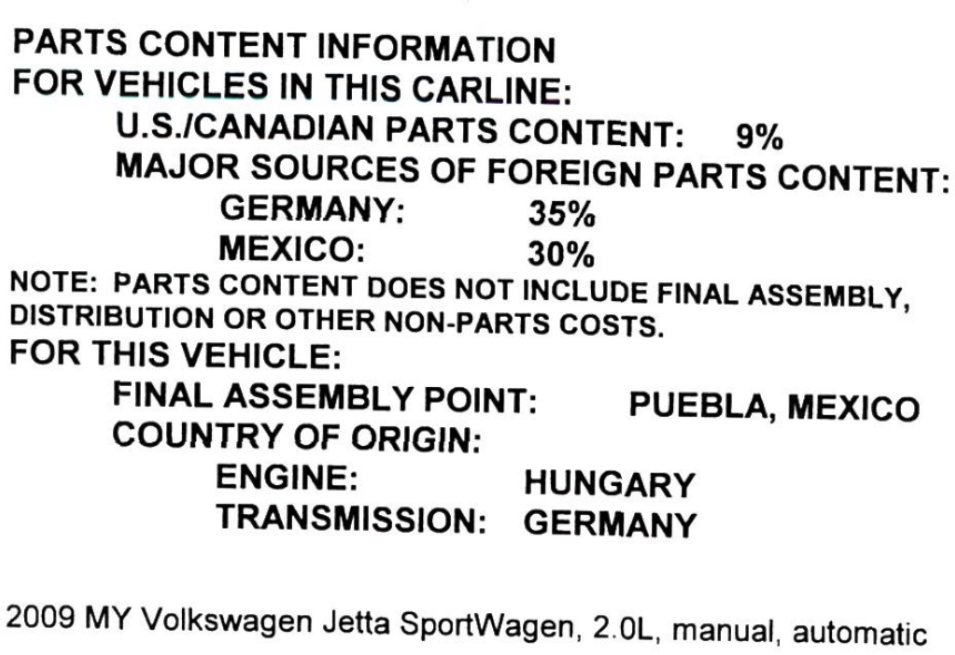
In this section I detail how I construct car model-by-year sourcing outcomes from the raw AALA data.

Since 1994, car manufacturers have been required to report certain country-of-origin information (both for final assembly and certain parts) to the US National Highway Traffic Safety Administration. This information is posted on the *Monroney* sticker that accompanies cars for sale at all US dealerships. Since 2011, the individual filings have been aggregated into annual spreadsheets as PDFs and posted on the NHTSA website (these are known as *Part 583* reports). I digitize and harmonize these reports across years.

[72] In economics research, a similar experience occurs with regression specifications: a single code file that runs a regression analysis is more difficult to design if it must accommodate a wide variety of robustness checks (different samples, specifications, bins, standard errors, etc.).

[73] I thank Jeremy Michalek at Carnegie Mellon University for clarifying the differences between EV and ICE platforms.

Figure B.3: *AALA* manufacturer filing example



Notes: Figure shows example of raw manufacturer filing for one model and year (the 2009 *Volkswagen Jetta* built in Mexico) under the *American Automotive Labeling Act*. I manually digitize filings as each firm uses a separate format.

I obtain the 2005-2010 individual manufacturer filings directly via request to the *NHTSA*. Automakers are free to choose the format in which data are reported (see Figure B.3 for an example), so each firm and year files different reports. As a result, I manually digitized each individual filing.

The main statistic reported is the domestic content share $DomesticShare_{jt}$, which is the share of inputs from the assembly location, and ranges from 0 to 100 percent. The AALA only reports content shares from individual countries that exceed 15 percent of total content; in practice the assembly country always satisfies this restriction. For USA and Canada-assembled vehicles, I use the share of inputs from USA and Canada combined (since the AALA requires automakers to report pooled USA/Canada content); for Mexico-assembled vehicles, I use the share of inputs from Mexico.

The AALA filings also reports the origin country $Origin_{jtp}$ for two critical parts p, engines and transmissions, which together typically comprise 20 to 30 percent of total vehicle content. This origin country information allows me to construct indicators for sourcing from particular locations. One such indicator is

$$Imported_{jtp} = \mathbb{1}[Origin_{jtp} \neq d(j)]$$

which takes value 1 if critical part p is imported, and serves as a different measure of importing to the continuous $ImportShare_{jt}$.

To avoid selection into the panel, I restrict to cars also assembled in North America (USA, Canada, and Mexico). 95% of car models built in North America have nonzero sales in the USA (and therefore appear in the AALA data). In contrast, only 4% of car models assembled outside North

America are exported to the USA, and the set of car models that are exported (and thus appear in the panel) are both highly selected and change over time conditional on the car being assembled elsewhere. This margin of adjustment is explored extensively in Sabal (2025).

B.6 *Marklines Who Supplies Whom* input sourcing flows construction

Each row in the WSW is a new sourcing agreement for a particular part type p made by a supplier s, for a particular model m, generation z (i.e. platform), and assembly location d, starting in a particular year t. For instance, one row might report that *Bosch (Thailand)* supplies the left headlight for the *Honda Civic (Japan)* starting in 2022. Because suppliers in *Marklines* do not update records each year (unlike under the AALA), I restrict to relationships starting between 2010-2022 and treat the sample as a single cross-section. Under this sampling restriction, sourcing decisions for each model generation should appear at least once. I do not use the exact year recorded because it corresponds to the year in which *Marklines* first discovered the relationship, not to when the relationship started.

The first step is to determine the make and model for each reported agreement. This is challenging because strings are not standardized by *Marklines*; for instance four suppliers might report "MINI Cooper", "Cooper", "Cooper Hardtop", and "Cooper Hardtop 2-Door" to refer to the same product.

To aggregate from individual relationships to bilateral input flows, I count the number of sourcing relationships by make, model, assembly location, and origin country. In other words, the origin is a *country*, as in a typical gravity regression; but the destination is a *product* (i.e. a make, model, and assembly location), and the sector is a *part type*. I then construct a balanced panel at the origin o-destination product j-part p level, therefore interpreting each relationship as a separate sourcing choice (over potential origins) and assigning each choice equal weight.[74]

An additional concern is that suppliers could report their headquarters locations rather than where production actually takes place for an input. Fortunately, *Marklines* requires different regional units of a supplier to have separate subscriptions, and thus to report independently, so that (for instance) Bosch's plants in Thailand and Germany are reported separately.

[74]Without price data, this procedure will underweight flows of higher-value complex modules such as engines and batteries. In the data, this bias is counteracted by the fact that complex modules have many components; for instance, suppliers of three engine module components (pistons, spark plugs, and sensors) are typically reported separately.

C Theory

C.1 Microfoundation for two phases of design platform-sharing

I assume that, to adopt modular design, firms incur *design costs* in engineer time that scale with within-platform product dissimilarity. This is because heterogeneity makes design rules harder to specify and compatible inputs more difficult to engineer.[75] To model design costs, I assume that each product $j = (s, d)$ is located at a point (X_s, X_d) in characteristic space. For example X_s could be product size. For any design technology $\{S_g\}$,

$$A\left(\{S_g\}\right) \quad \equiv \quad \sum_{g=1}^{G} \sum_{(s,d)\in S_g} \left(\theta \left\|X_d - \mu_d^g\right\|^2 \; + \; \theta\,\phi \left\|X_s - \mu_s^g\right\|^2\right) \tag{26}$$

where $\mu^g = (\mu_s^g, \mu_d^g)$ are platform ideal points in characteristic space that engineers target when designing inputs. Because distances are quadratic, the firm will endogenously choose ideal points to be platform *centroids*, i.e. $\mu_d^g = \mathbb{E}[X_d|(s,d) \in S_g]$ and similar for μ_s^g. This means that design costs are a weighted sum of within-platform variance in characteristics.[76] [77]

Design costs depend on two parameters. The first is the *marginal cost of dissimilarity* (θ) in terms of the outside good. Falls in θ over time represent secular technological changes discussed in Section 2 that enable R&D departments to more easily share design rules and compatible inputs. The second parameter is the *relative cost of cross-segment differentiation* (ϕ). This is the ex-ante difficulty of imposing common design rules across segments with different physical requirements, relative to within a segment. In many complex good settings, ϕ is high because physical size or power are linked to several interdependent engineering requirements, and therefore require specification of additional rules.[78]

Firm profit-maximization problem with endogenous design The firm chooses a design technology $\{S_g\}$, platform ideal points (μ_s^g, μ_d^g), input production $\{y^{og}\}$ and labor use $\{l^{og}\}$, input use by product $\{q_j^{og}\}$, product volumes $\{Q_j\}$, and product prices $\{P_j\}$ to maximize profits net of labor costs and

[75] For instance, two cars with common lengths but different weights might share a platform. This requires that engineers specify rules for the shared characteristic (length) separately from other dimensions (width and height), design an engine that somewhat satisfies both cars' weight requirements, and build in multiple settings (either physically or via software) that allow the engine to pass the distinct emissions checks of multiple regulators without sacrificing performance. With fully customized inputs, these *design costs* would not exist.

[76] Because we have two countries, if we consider each segment a separate 'variety', then the design cost function $A(.)$ is similar to that of Grossman et al. (2021), who allow for countries to have different ideal points for a characteristic of each variety. If (X_s, X_d) were uniformly distributed and the loss function was linear (rather than quadratic), $A(.)$ would be isomorphic to the cost function of Alesina and Spolaore (1997) faced by governments in serving heterogeneous citizens.

[77] Formally, in terms of variances, design costs are:

$$A\left(\{S_g\}\right) \; = \; A\left(\{S_g\}; \mu_s^g, \mu_d^g\right) \; = \; \sum_{g=1}^{G} \left|S_g\right| \; \left(\theta \cdot \mathrm{Var}(X_d \mid (s,d) \in S_g) \; + \; \phi \cdot \theta \cdot \mathrm{Var}(X_s \mid (s,d) \in S_g)\right) \tag{27}$$

[78] Intuitively, ϕ is often high because cars in different segments have different core technical requirements (computing power, engine size, etc.), while cars in different countries must only satisfy certain idiosyncratic regulator needs.

design costs:

$$\max_{\{S_g\},\{l^{og}\},\{y^{og}\},\{q_j^{og}\},\{Q_j\},\{P_j\},\mu_s^g,\mu_d^g} \quad \underbrace{\sum_j P_j Q_j \; - \; \sum_{o,g} w_o l^{og}}_{\text{variable profits}} \; - \; \underbrace{A(\{S_g\}; \mu_s^g, \mu_d^g)}_{\text{design costs}} \tag{28}$$

subject to demand from household maximization of (2) (taking competitor prices P_{sd} as given), design costs (27), the assembly production function (3), input production function (4), and internal input market-clearing (5).

Equation (28) highlights the basic tradeoff facing each firm. On one hand, modular design *always increases variable profits*. This is clear from equation (4): due to economies of scale, firms can (at worst) save labor costs, while keeping input production fixed, by combining plants for two platforms g and g' in any country o. On the other hand, it follows from equation (26) that *modular platforms always increase design costs*, because merging two clusters (i.e. reducing the number of platforms) always increases within-cluster variance. This tradeoff between scale economies and design costs determines how (and to what extent) modular design is adopted.

Because value-added production and trade within-the-firm take place conditional on design, we can write firm profits as:

$$\max_{\{S_g\}} \quad \underbrace{V(\{S_g\})}_{\text{variable profits}} \; - \; \underbrace{\sum_{g=1}^{G} \sum_{(s,d)\in S_g} \left(\theta \, \|X_d - \mu_d^{g*}\|^2 \; + \; \theta\,\phi \, \|X_s - \mu_s^{g*}\|^2\right)}_{\text{design costs}} \tag{29}$$

where $j = (f, s, d)$ indexes each of firm f's products, $V(\{S_g\})$ is the firm's maximized variable profits conditional on design $\{S_g\}$, and (μ_d^{g*}, μ_d^{s*}) are centroids in characteristic space for platform g.

Because each partition is a (two-way) clustering of the data, I cannot solve analytically[79] for the optimal partition $\{S_g\}$ without strong distributional assumptions on characteristics.[80] However, given a firm's existing design technology $\{S_g\}$ and some knowledge of design costs (in particular segment differentiation ϕ), I can still identify the dimensions (destinations d or segments s) across which modularity will first be adopted as marginal design costs θ over time.

Comparative statics In this section I study the effects on design technology $\{S_g\}$ of the gradual technological improvements – in electronic control, electric power, and within-firm communication – that enable engineers to more efficiently accommodate within-platform product dissimilarity. These improvements are represented by falls in θ. I show that – under realistic assumptions on the relative costs of each dimension of product differentiation – modular design is adopted in two phases: first within product segments s with similar technical characteristics, and then across segments s.

[79] In general, optimal clusterings are data-dependent, and known solution methods such as the k-means algorithm are not globally optimal.

[80] The exception is in extreme cases in which some model forces are inactive. In particular, (i) without scale economies ($\eta = 0$), input-sharing does not reduce costs, so optimal design $\{S_g\}^*$ is *non-modular* in that no two products j and j' share a design platform; (ii) without input trade ($\{\tau_d^o\} \to \infty$), production cannot concentrate in one country and then be exported elsewhere, so optimal design $\{S_g\}^*$ is non-modular across countries in that no two products in different countries (with $d \neq d'$) share a design platform; and (iii) without design costs ($\theta = 0$), within-platform product heterogeneity is not costly for engineers, so optimal design $\{S_g\}^*$ is universally modular in that all products $\{j\}$ share a single design platform.

Proposition 3 (Two phases). *There exists a cutoff value $\bar{\theta}$ of marginal dissimilarity costs above which design is non-modular. When designing over segment characteristics is sufficiently difficult ($\phi > \bar{\phi}$ for some cutoff $\bar{\phi}$), then as θ falls from $\bar{\theta}$, design technology $\{S_g\}$ evolves in two sequential phases:*

1. ***Modularity within segment** s: Firms adopt modular design across destinations d but remain non-modular across segments s. In other words, firms move from $G = 2S$ platforms (one per product j) to $G = S$ platforms (one per segment s). Products in H and F that share a segment s use common inputs. Within a firm, countries specialize in segments s rather than products $j = (s, d)$.*
2. ***Modularity across segments** s: Firms adopt modular design across segments s, moving from $G = S$ platforms (one per segment s) to $G = 1$ platforms (one firm-wide). All products in H and F in all segments $\{s\}$ use common inputs. Within a firm, countries do not specialize at all, rather than specializing in segments s.*

Proof. We use a monotone comparative statics argument. Variable profits are independent of design cost parameters (θ, ϕ), and the loss function $A(\{S_g\})$ has strictly increasing differences in both $(\theta, \sum ||X_d - \mu_d^g||^2)$ and $(\theta\phi, \sum ||X_s - \mu_s^g||^2)$. Thus the within-platform heterogeneity in characteristics ($\sum ||X_d - \mu_d^g||^2$ and $\sum ||X_s - \mu_s^g||^2$) that maximize the objective (29) are strictly decreasing in θ and $\theta\phi$, respectively. There thus exist cutoffs θ_d^{max} ($\phi \cdot \theta_s^{max}$) above which each destination (segment) uses its own platform, and cutoffs θ_d^{min} ($\phi \cdot \theta_s^{min}$) below which all destinations (segments) use one platform. For sufficiently large but finite ϕ, which we denote $\phi > \bar{\phi}$, it holds that

$$\theta_d^{max} > \theta_d^{min} > \theta_s^{max} > \theta_s^{min}$$

so that as design costs θ fall, we obtain the phases described. □

Proposition 3 shows that patterns of adoption of modular design within each firm mirror the technical challenges faced by product engineers. In particular, platform-sharing is especially costly (in that rules and inputs are more difficult to develop) across product segments with different physical and technical characteristics. As a result, as marginal design costs (θ) fall, modular design platforms are first common within product segment but across destinations; and then across product segments within a firm. For instance, inputs might shared between Ford's American and non-American small car models; and then between smaller and larger car segments.

C.2 Formal solution to firm optimization problem

This section gives the equations that implicitly define the solution to Equation 6. We first rearrange the input production function Equation 4 to write the wage bill in terms of input production y^{og} (for any platform g, and either origin $o \in H, F$):

$$w^o l^{og} = \frac{w^o}{\alpha^o(1-\eta)} (y^{og})^{1-\eta} \tag{30}$$

and substitute Equation 30 into the objective Equation 6. Writing the full Lagrangian and differentiating with respect to y^{og} gives the inverse supply curve:

$$p^{og} = \frac{w^o}{\alpha^o} \cdot (y^{og})^{-\eta} \tag{31}$$

where p^{og} be the Lagrangian multiplier on input production y^{og} (i.e. the shadow price of an additional unit of the (o,g)-specific input). Next, we write down the platform-g resource constraints for both countries H and F, then express them in "expenditure" form by multiplying both sides by p^{og}:

$$\begin{aligned} p^{Hg} y^{Hg} &= \sum_{j \in \mathbf{S_g}} q_j^{Hg} \cdot \tau_d^H \cdot p^{Hg} \\ \underbrace{p^{Fg}}_{\text{Shadow price}} \cdot \underbrace{y^{Fg}}_{\text{Supply}} &= \sum_{j \in \mathbf{S_g}} \underbrace{q_j^{Fg} \cdot \tau_d^F}_{\text{Demand}} \cdot \underbrace{p^{Fg}}_{\text{Shadow price}} \end{aligned} \tag{32}$$

and use Equation 31 to substitute in for y^{og} in terms of p^{og} on the LHS. To write the RHS in terms of p^{og}, from differentiating Equation 6 (after substituting Equation 30) with respect to Q_j, P_j, and q_j^{og} and rearranging, and using that trade costs are symmetric ($\tau_H^F = \tau_F^H = \tau$), we obtain that for any product j sold in H:

$$P_j = \frac{\varepsilon}{\varepsilon - 1} \left((p^{Hg(j)})^{1-\sigma} + (\tau p^{Fg(j)})^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{33}$$

and symmetrically for j in F. In other words, final goods prices are a constant markup over the "shadow" marginal cost of an additional unit of final good j, which takes the form of a CES price index for the platform $g(j)$ to which product j belongs. Combining this with standard CES expenditure shares over products j (with prices P_j) and over inputs (with shadow prices p^{og}), we can rewrite Equation 32 as:

$$\begin{aligned} (w^H/\alpha^H)^{\frac{1}{\eta}} (p^{Hg})^{1-\frac{1}{\eta}} &= D_H^g \cdot \kappa \cdot \frac{(p^{Hg})^{1-\sigma}}{\left((p_H^g)^{1-\sigma} + (\tau \cdot p_F^g)^{1-\sigma}\right)^{\varepsilon-\sigma}} \\ &+ D_F^g \cdot \kappa \cdot \frac{(\tau \cdot p^{Hg})^{1-\sigma}}{\left((\tau \cdot p_H^g)^{1-\sigma} + (p_F^g)^{1-\sigma}\right)^{\varepsilon-\sigma}} \\ (w^F/\alpha^F)^{\frac{1}{\eta}} (p^{Fg})^{1-\frac{1}{\eta}} &= D_H^g \cdot \kappa \cdot \frac{(\tau \cdot p^{Fg})^{1-\sigma}}{\left((p_H^g)^{1-\sigma} + (\tau \cdot p_F^g)^{1-\sigma}\right)^{\varepsilon-\sigma}} \\ &+ D_F^g \cdot \kappa \cdot \frac{(p^{Fg})^{1-\sigma}}{\left((\tau \cdot p_H^g)^{1-\sigma} + (p_F^g)^{1-\sigma}\right)^{\varepsilon-\sigma}} \end{aligned} \tag{34}$$

where $\kappa = (\frac{\varepsilon}{\varepsilon=1})^{-\varepsilon}$ is a constant, and

$$D_H^g \equiv \sum_{j \in \mathbf{S_g}: d(j)=H} \gamma_H \, \beta_{s(j)H} \, P_{s(j)H}^{\varepsilon-1} \quad , \quad D_F^g \equiv \sum_{j \in \mathbf{S_g}: d(j)=F} \gamma_F \, \beta_{s(j)F} \, P_{s(j)F}^{\varepsilon-1} \tag{35}$$

are the sum of all demand-shifters for products assembled using platform g in destinations H and F respectively. This system of two equations in two unknowns (p^{Hg}, p^{Fg}) implicitly defines the firm's optimal sourcing decision for platform g. After solving for shadow prices, production in both countries immediately follows from Equation 31.

C.3 Proof of Lemma 1

Proof. First observe that:

$$\begin{aligned}\left(p^{Hg}\right)^{1-\sigma}+\left(\tau p^{Fg}\right)^{1-\sigma} &= \left(p^{Hg}\right)^{1-\sigma}\cdot\left(1+\left(\tau\widehat{p^g}\right)^{1-\sigma}\right) = \left(p^{Fg}\right)^{1-\sigma}\cdot\left(\left(\widehat{p^g}\right)^{\sigma-1}+\tau^{1-\sigma}\right)\\ \left(\tau p^{Hg}\right)^{1-\sigma}+\left(p^{Fg}\right)^{1-\sigma} &= \left(p^{Hg}\right)^{1-\sigma}\cdot\left(\left(\widehat{p^g}\right)^{1-\sigma}+\tau^{1-\sigma}\right) = \left(p^{Fg}\right)^{1-\sigma}\cdot\left(1+\left(\tau\widehat{p^g}\right)^{\sigma-1}\right)\end{aligned} \tag{36}$$

We substitute these definitions into the first equation in Equation 34, using the definitions with $(p^{Hg})^{1-\sigma}$ factored out, and then divide both sides by $(p^{Fg})^{1-\sigma}$.[81] We similarly substitute into the second equation in Equation 34 using definitions with $(p^{Fg})^{1-\sigma}$. Next, divide the second equation in Equation 34 by the first. Then factor out $(\widehat{p^g})^{\sigma-1}$ from the numerator and divide both sides through by $(\widehat{p^g})^{\sigma-1}$. This gives:

$$\frac{\widehat{w}}{\widehat{\alpha}}\left(\widehat{p^g}\right)^{\sigma-\frac{1}{\eta}} = \frac{D_H^g\cdot\tau^{1-\sigma}\cdot R+D_F^g}{D_H^g\cdot R+D_F^g\cdot\tau^{1-\sigma}} = \frac{R\cdot\tau^{1-\sigma}+\widehat{D^g}}{R+\widehat{D^g}\cdot\tau^{1-\sigma}} \tag{37}$$

where D_H^g and D_F^g are as in Equation 35, and where we have defined:

$$R \equiv \left(\frac{1+\left(\widehat{p^g}\tau\right)^{1-\sigma}}{\tau^{1-\sigma}+\left(\widehat{p^g}\right)^{1-\sigma}}\right)^{-\alpha} \quad \text{where} \quad \alpha\equiv\frac{\sigma-\varepsilon}{\sigma-1}$$

We aim to show that $\widehat{p^g}$ is decreasing in $\widehat{D^g}$. Observe first that we can rewrite the RHS of Equation 37 as:

$$\tau^{1-\sigma}\cdot\frac{R+\widehat{D^g}\cdot\tau^{\sigma-1}}{R+\widehat{D^g}\cdot\tau^{1-\sigma}}$$

Clearly, because $\tau>1$, this term is increasing in $\widehat{D^g}$. Starting from values where Equation 37 holds, increasing $\widehat{D^g}$ will increase the RHS and keep the LHS unchanged. From the implicit function theorem, to prove that we must *reduce* shadow prices $\widehat{p^g}$ to restore equilibrium, it suffices to show:

(i) The LHS is decreasing in $\widehat{p^g}$; and
(ii) The RHS is increasing in $\widehat{p^g}$.

Statement (i) follows from the fact that we assumed $\sigma\eta<1$, so the LHS is decreasing in $\widehat{p^g}$. To establish (ii), we first rewrite term R as:

$$R = \tau^{\sigma-1}\cdot\left(\frac{1+\tau^{1-\sigma}\left(\widehat{p^g}\right)^{1-\sigma}}{1+\tau^{\sigma-1}\left(\widehat{p^g}\right)^{1-\sigma}}\right)^{-\alpha}$$

and observe that:

(1) The term in parentheses is decreasing in $\widehat{p^g}^{1-\sigma}$ because $\tau>1$

[81] The ratio captures marginal 'excess demand', which should decrease as price falls.

(2) Because we assumed $\sigma > \varepsilon$, we have $\alpha > 0$, so the exponent on the parentheses term is negative. Thus R is increasing in $\left(\widehat{p^g}\right)^{1-\sigma}$.

(3) Because $\sigma > 1$, $\widehat{p^g}^{1-\sigma}$ is decreasing in $\widehat{p^g}$. Thus R **is decreasing in** $\widehat{p^g}$.

To complete the proof it suffices to show that the RHS is decreasing in R, which combined with the fact that R is decreasing in $\widehat{p^g}$, establishes statement (ii). Since we can rewrite the RHS of Equation 37 as

$$\tau^{\sigma-1} \cdot \frac{R \cdot \tau^{1-\sigma} + \widehat{D^g}}{R \cdot \tau^{\sigma-1} + \widehat{D^g}}$$

and because $\tau > 1$, it follows that the RHS is decreasing in R. We have shown that the RHS of Equation 37 is increasing in $\widehat{D^g}$ and in $\widehat{p^g}$, and that the LHS is decreasing in $\widehat{p^g}$. By the implicit function theorem, $\widehat{p^g}$ is decreasing in $\widehat{D^g}$.

Finally, using Equation 31 and rearranging, we have that:

$$\widehat{y^g} = \left(\frac{\widehat{w}}{\widehat{\alpha}}\right)^{\frac{1}{\eta}} \left(\widehat{p^g}\right)^{-\frac{1}{\eta}} \tag{38}$$

Since $\eta > 0$, relative production $\widehat{y^g}$ decreases in $\widehat{p^g}$, thus increases in relative demand $\widehat{D^g}$. Finally, from the CES demand function, for any product j using platform $g(j)$ and assembled in H (without loss of generality), the share of inputs sourced from F is:

$$\pi_j^F = \frac{(\tau \widehat{p}^{g(j)})^{1-\sigma}}{1 + (\tau \widehat{p}^{g(j)})^{1-\sigma}} \tag{39}$$

which is decreasing in $\widehat{p^{g(j)}}$, thus increasing in $\widehat{D^{g(j)}}$. □

C.4 Proof of Proposition 1

Proof. Consider without loss of generality the sourcing decision for a product j assembled in $d(j) = H$. With no platform-sharing, product j's platform $g(j)$ is not used for any products in F, so:

$$\widehat{D}^{g(j)} = \frac{0}{\gamma_H \beta_{s(j)H} P_{s(j)H}^{\sigma-1}} = 0$$

i.e. relative foreign market size for j's inputs is 0. With platform-sharing across destinations (phase 1), other products j' in the same segment (with $s(j') = s(j)$) use the same platform, so:

$$\widehat{D}^{g(j)} = \frac{\gamma_F \beta_{s(j)F} P_{s(j)F}^{\sigma-1}}{\gamma_H \beta_{s(j)H} P_{s(j)H}^{\sigma-1}} = \widehat{\gamma}\widehat{\beta}_{s(j)}\widehat{P}_{s(j)} > 0 \tag{40}$$

i.e. relative foreign market size for j's inputs is positive. Platform-sharing across destinations thus increases $\widehat{D}^{g(j)}$ for any product assembled in H. Because the foreign sourcing share π_j^F) as defined in Equation 39) is increasing in $\widehat{D}^{g(j)}$ per Subsection C.3, it follows that *platform-sharing across destinations increases foreign sourcing shares.*

Furthermore, per Equation 40, the change in foreign input sourcing shares is increasing in the foreign countries' relative total expenditure $\widehat{\gamma}$, relative preference for j's segment ($\widehat{\beta}_{s(j)}$), and relative price of foreign competitors $\widehat{P}_{s(j)}$. If there is a single firm (i.e. no competitors), this simplifies to:

$$\widehat{D}^{g(j)} = \widehat{\gamma}\widehat{\beta}_{s(j)}$$

. □

C.5 Proof of Proposition 2

Proof. We first appeal to Lemma 1. Without platform-sharing across segments (i.e. in phase 1) platforms are segment-specific, so relative demand for inputs for product j are:

$$\widehat{D}_j^{PRE} \equiv \frac{\gamma_F \beta_{s(j)F} \bar{P}_{s(j)F}^{\varepsilon-1}}{\gamma_H \beta_{s(j)H} \bar{P}_{s(j)H}^{\varepsilon-1}} \tag{41}$$

With a single platform, demand-shifters are *summed* across platforms, so relative demand is the same across all j:.

$$\widehat{D}^{POST} = \frac{\sum_s \gamma_F \cdot \sum_s \beta_{sF} \bar{P}_{sF}^{\varepsilon-1}}{\sum_s \gamma_H \cdot \sum_s \beta_{sH} \bar{P}_{sH}^{\varepsilon-1}} = \sum_s \widehat{D}_s^{PRE} \omega_s \tag{42}$$

where the weights ω_s are the importance of segment s in country H's total demand:

$$\omega_s = \frac{\gamma_H \beta_{sH} \bar{P}_{sH}^{\varepsilon-1}}{\sum_{s'} \gamma_H \beta_{s'H} \bar{P}_{s'H}^{\varepsilon-1}}$$

Thus for product j, adoption of platform-sharing across segments (phase 2) increases relative shadow prices in F ($\widehat{p^{g(j)}}$), and therefore *reduces* the share of inputs sourced from F, iff:

$$\widehat{D}_{s(j)}^{PRE} > \sum_{s'} \widehat{D}_{s'}^{PRE} \omega_{s'} \tag{43}$$

i.e. country F's relative demand for j's segment exceeds the weighted average of country F's relative demand advantages across *all segments*. Thus, with a single integrated platform, F loses its cost advantages in inputs for segments where it had a particularly large relative market size.

As $\varepsilon \to 1$ (or equivalently, if the model features a single integrated firm) the expressions simplify, and relative demand $\widehat{D^g}$ changes from $\frac{\gamma_F}{\gamma_H}\frac{\beta_{sF}}{\beta_{sH}}$ to $\frac{\gamma_F}{\gamma_H}$. Thus, if and only if $\beta_{sF} > \beta_{sH}$ (F more strongly prefers segment $s(j)$), platform-sharing across segments increases relative shadow costs in F ($\widehat{p^g}$ increases), and thus sourcing for good j substitutes away from F.

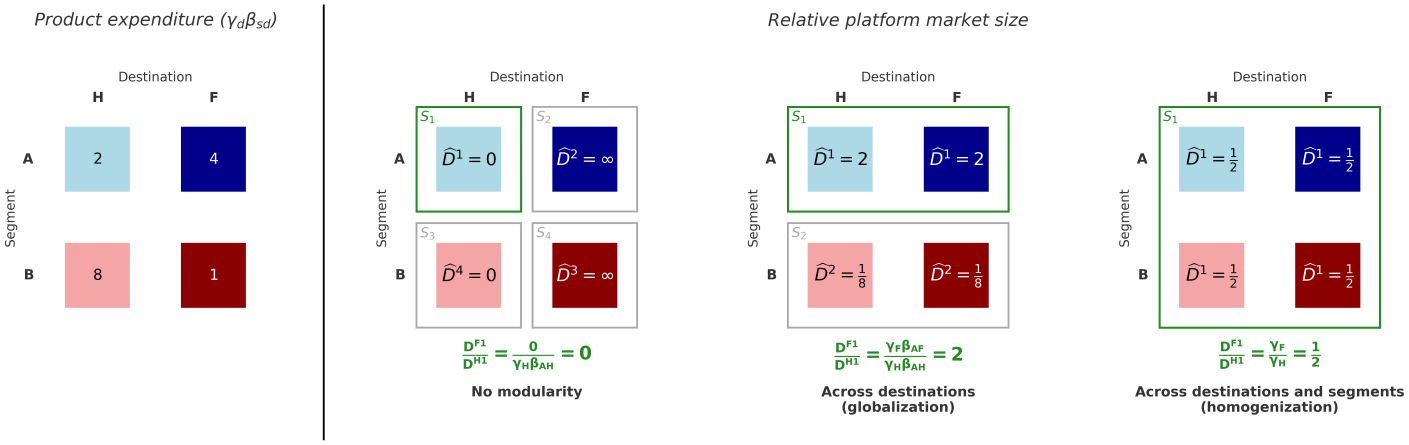
Finally, cost advantages (and therefore cross-country specialization) in inputs for j relative to j' cease to exist with platform-sharing across segments, because $g(j) = 1$ for all products j (all products use a single platform) with relative demand $\widehat{D}^{POST}$.

□

C.6 Parametrized example of theory results

This section provides special case of the Section 4 model with a single firm ($F = 1$), two segments (A and B), and no skill differences ($\alpha^H = \alpha^F$). H is the larger market, with total expenditures of $\gamma^H = 10$ vs. $\gamma^F = 5$. The countries have symmetric preferences ($\beta_{AH} = 0.2, \beta_{BH} = 0.8$) and ($\beta_{AF} = 0.2, \beta_{BF} = 0.8$) so that H prefers segment B. The *total expenditure* on each product is shown in the left panel.

Figure C.1: Platform market size changes due to the two phases of modularity



Notes: The leftmost panel shows total expenditure on each product $j = (s, d)$ in a special case of the model with a single firm ($F = 1$), two segments (A and B), no skill differences ($\alpha^H = \alpha^F$), total expenditure $\gamma_H = 10$ and $\gamma_F = 5$, and preferences $\beta_{AH} = 0.2$, $\beta_{BH} = 0.8$, $\beta_{AF} = 0.8$, and $\beta_{BF} = 0.2$. The right 3 panels show relative *platform* market size $\widehat{D}^{G(j)}$, as defined in Lemma 1, for product j's platform $G(j)$ under three different design technologies: no modularity, modularity across destinations (within a segment), and modularity across segments. Platform assignments $\{S_g\}$ under each design technology are indicated by rectangles. The calculation of market size $\widehat{D}^{G(j)}$ for product $j = (s, d)$, which is always assigned platform $g = 1$, is shown in green below each panel.

The remaining panels show how, given the same product demand, the two phases of modularity each result in very different platform market size advantages, and therefore production locations. With no modularity, H has a market size advantage of $\widehat{D}^g = 0$ (disadvantage of $\widehat{D}^g = \infty$) in inputs for local products (foreign products), so input production locates near local assembly. As a result, without modularity, supply chains exhibit substantial home bias. With modularity across destinations but within a segment (Proposition 1), country H (F) has a market size advantage in its preferred segment B (segment A) products, including for export, which increases trade. Finally, with modularity across destinations and segments (Proposition 2), H has an identical market size advantage in inputs for all products, which generates geographic concentration in H.

Figure C.1 also illustrates how I test the model: changes in design technology in turn reshape sourcing decisions for individual products j. For instance, as shown in the green boxes, in the globalization phase relative market size for the inputs used in product $j = (A, H)$ increases from 0 to 2. As a result, *import shares increase* because there is now twice as much demand in F than H for j's input, relative to zero demand in the non-modular world, which shifts production of j's inputs to F. This increase is especially pronounced because j is in segment A, which F prefers (with expenditure share $\beta_{AF} = 0.8$). In contrast, in the subsequent homogenization phase, relative market size for j's input falls from 2 to 0.5 because the platform now encompasses both segments, so sourcing for product $j = (A, H)$ substitutes away from the country that prefers segment A, which in this case is F.

D Empirical analysis details

D.1 Procedure for identifying platform adoption events

To identify events in which a product first adopts a platform that is shared with products in other destinations, I use *Wards* and *Marklines* automaker design timelines to identify the platform name $g(j,t)$ for each model and year in the AALA data, and the years of redesigns to new platforms $\bar{t}_j^{new} = \{t : g(j,t) \neq g(j,t-1)\}$. To identify the subset of product redesigns that are modular design adoption events (i.e. of increased platform-sharing), I also merge in the set of products $j' \in S_g$ that use each platform, and the assembly locations $d(j')$ for each product.

Using this data, I identify events in which the post-redesign platform $g(j,t)$ is shared with all regions within a segment, while the pre-redesign platform $g(j,t-1)$ is not. Formally, these are events for which (i) in the years when the pre-event platform $g(j)^{PRE}$ is in use, some products j' in the same segment s but different region d' are assigned to a different platform $g' \neq g(j)^{PRE}$; and (ii) in the years when the post-event platform $g(j)^{POST}$ is in use, all products j' in the same segment s, in all regions $\{d\}$, are assigned to platform $g(j)^{POST}$. Following *Marklines* and industry convention, regions are: North America, Europe, China, Japan, Korea, India, Southeast Asia, and South America.

To ensure sufficient panel balance, I restrict to events between 2012 and 2023 in which there are at least four years of pre-treatment data and two years of post-treatment data.

D.2 Motives for 2008 recession-induced firm boundary changes

The reason for all changes were similar: an urgent need to avoid bankruptcy by averting short-term capital outlays.

Chrysler, the smallest of the three, was first rescued by the U.S. federal government, and immediately merged with Fiat after the initial mooted buyer (Renault) withdrew mid-talks to conserve cash as financial conditions worsened in 2008.[82] Chrysler then coordinated with Fiat to develop a shared set of platforms and inputs for the next generation of vehicles.

While General Motors and Ford were not sold to other buyers, both firms moved to integrate their European units – which historically had operated as separate firms as a legacy of their presence in Europe since the 1920s – with their American operations. General Motors was bailed out by the United States federal government, which directly instructed it to save capital costs. Meanwhile, Ford remained private but faced investor pressure to consolidate operations. Both firms therefore chose to combine platforms to immediately avert platform development costs. In both cases, the adoption of cross-country modularity between the U.S. and Europe was behind rather than ahead of industry norms; Volkswagen and Toyota, for instance, had used global design platforms since the mid-1990s (Cusumano and Nobeoka, 1998).

As a result, cars assembled in North America by all of the Big 3 automakers were gradually redesigned to share platforms with similar-segment cars assembled in Europe, something that was largely untrue before 2008.

A final change was that Ford divested its longtime share in the Japanese producer Mazda, again to immediately raise capital, leading Ford's American cars to *stop* platform-sharing with Japan.

[82] See Reuters (2008).

D.3 Counterfactuals exact-hat system of equations

The procedure described below applies to both:

- "Platform-splitting" counterfactuals in which observed platforms (blocks) S_g of each firm's baseline design technology $\{S_g\}$ are either unchanged or *split* into two or more platforms h
- "Platform-merging" counterfactuals in which products using a platform g' are added to a platform g (so that $h = g$, but g' no longer exists) rather than using a separate platform.

The present-day quantification exercise is a platform-splitting counterfactual and the universal platform scenario is a platform-merging counterfactual.

(1) Product-level logit price index For each product j, with destination $d(j)$, segment $s(j)$, and assigned platform $h(j)$, the proportional change in output price is:

$$\boxed{\widehat{P}_j^{-\sigma} \;=\; \sum_{o=1}^{N} \pi_j^o \left(\widehat{\tau}^o_{d(j)k(j)}\, \widehat{\tau}^o_{d(j)k(j)}\, \widehat{p}^{\,oh(j)}\right)^{-\sigma}}$$

(2) Segment–destination logit price index For each segment–destination pair (s, d), the proportional change in price index is then:

$$\boxed{\widehat{P}_{sd}^{-\varepsilon} \;=\; \sum_{f=1}^{F} \chi_{j(sdf)}\, \widehat{P}^{-\varepsilon}_{j(sdf)}}$$

where $j(sdf)$ is the unique product with indices (s, d, f) and $\chi_j = Q_{j(sdf)} / \sum_f Q_{j(sdf)}$ is product j's market share within segment s in destination d.

(3) Input-market clearing / input-price equations Let

$$y^{oh}_{\text{PRE}} \;=\; \sum_{j:\, h(j)=h} Q_j\, \pi_j^o\, \tau^o_{d(j)f(j)}$$

be *baseline* input usage from origin o for all products j that use platform h in the *counterfactual* scenario. Then for each (o, h), for any proportional tariff change $\{\widehat{\tau}\}$ and we can write the counterfactual resource constraint as:

$$\boxed{y^{oh}_{\text{PRE}} \left(\widehat{\alpha}^{\,oh}\, \widehat{p}^{\,oh}\right)^{-1/\eta} \;=\; \sum_{j:\, h(j)=h} \underbrace{z_j^h}_{\text{new platform}} \underbrace{Q_j\, \pi_j^o\, \tau^o_{d(j)k(j)}}_{\text{baseline flows}} \underbrace{\chi_j\, \widehat{P}^{\varepsilon-1}_{s(j)d(j)}\, \widehat{P}_j^{\sigma-\varepsilon} \left(\widehat{p}^{\,oh}\right)^{-\sigma} \underbrace{\left(\widehat{\tau}^o_{d(j)k(j)}\right)^{1-\sigma}}_{\text{tariff change}}}_{\text{proportional change in flows}}}$$

where z_j^h takes value 1 if product j uses platform h in the counterfactual scenario and 0 otherwise. This system provides $N \times H$ equations in the unknowns $\{\widehat{p}^{\,oh}\}$, which are the proportional changes in shadow input prices.

Aggregate counterfactual outcomes

The counterfactual level of aggregate input production in country o is:

$$Y'_o \;=\; \sum_h y^{oh}_{\text{POST}}, \qquad y^{oh}_{\text{POST}} \;\equiv\; \sum_{j:\, h(j)=h} Q_j \, \widehat{Q}_j \, \pi^o_j \, \widehat{\pi}^o_j \, \tau^o_{d(j)}.$$

and in counterfactuals, I plot the percent change in input production shares from the baseline to counterfactual scenario by country o:

$$\widehat{\text{ProdShare}}_o \;-\; 1 \equiv \frac{\dfrac{Y'_o}{\sum_o Y'_o} \;-\; \dfrac{Y_o}{\sum_o Y_o}}{\dfrac{Y_o}{\sum_o Y_o}}.$$