Conduct and Scale Economies: Evaluating Tariffs in the US Automobile Market

Marco Duarte[*] Lorenzo Magnolfi[†] Daniel Quint[‡] Mikkel Sølvsten[§]

Christopher Sullivan[¶]

July 2025

Abstract

We evaluate tariffs' effects in the US automobile market, accounting for global value chains. The medium-run effects depend on manufacturers' strategic responses and scale economies. We develop a data-driven procedure which selects Cournot quantity-setting with substantial scale economies as the best-fitting supply model. Counterfactuals reveal that 25% tariffs on imported cars induce pass-through slightly above one for foreign manufacturers, while domestic firms decrease prices; adding parts tariffs increases domestic prices by 6.5%. Consumer welfare losses double when tariffs extend to parts, with total surplus losses exceeding $30 billion in 2018. We evaluate these losses against employment gains in domestic manufacturing.

[*]Department of Economics, University of North Carolina at Chapel Hill. Email: duartema@unc.edu
[†]Department of Economics, University of Wisconsin-Madison, Madison. Email: magnolfi@wisc.edu
[‡]Department of Economics, University of Wisconsin-Madison, Madison. Email: dquint@ssc.wisc.edu
[§]Department of Economics and Business Economics, Aarhus University. Email: miso@econ.au.dk.
[¶]Department of Economics, University of Calgary. Email: christopher.sullivan1@ucalgary.ca

We thank Panle Barwick, Agustín Gutiérrez, and seminar participants at Berry Fest, IIOC 2025, and the UBC Summer IO Conference for insightful comments. Mikkel Sølvsten acknowledges financial support from the Danish National and Aarhus University Research Foundations (DNRF Chair #DNRF154 and AUFF Grant #AUFF-E-2022-7-3). All errors are our own.

1 Introduction

In the last decade, trade policy has reemerged as a dominant form of economic strategy, with tariffs now serving as its primary instrument. In the US, both Trump administrations have proposed and enacted tariffs on imported goods, prompting retaliatory measures that further complicate international commerce. Analyzing these interventions requires understanding how tariffs propagate through complex global value chains (Antràs and Chor, 2022), particularly in industries where components may cross borders before reaching consumers.

The automobile industry exemplifies this complexity. Targeted for tariff intervention by both Trump administrations, this sector features intricate production networks crossing national borders. Foreign components constitute over 40% of US-assembled vehicles, while US-made parts feature prominently in foreign-made cars. This web of interdependencies, fostered by trade agreements like the North American Free Trade Agreement (NAFTA) and its successor, the United States-Mexico-Canada Agreement (USMCA), creates layered effects when tariffs are applied. In particular, both Trump administrations have threatened stacked tariffs (tariffs applied to both final goods and intermediate inputs) on foreign-assembled cars and foreign car parts used in domestic assembly, magnifying potential impacts throughout the supply chain. The stated goals of such policies often involve reshaping the geography of production to boost domestic employment. However, the high degree of associated policy uncertainty tends to slow down investment (see, e.g., Caldara, Iacoviello, Molligo, Prestipino, and Raffo, 2020; Handley and Limão, 2022), reducing firms' ability to relocate production. In turn, this makes it useful to assess the medium-run (i.e, a one year time frame) effects of tariffs on consumers and firms while taking as given the locations of production and market structure.

Such evaluation requires a methodological framework that accounts for two critical but often neglected features of industry structure. First, the literature has long recognized that firm conduct matters significantly for how trade policy affects market outcomes (Dixit, 1984; Brander and Spencer, 1985; Dixit and Grossman, 1986),[1] as the degree of strategic substitutability or complementarity guides domestic firms' responses to foreign competitors' tariff-induced price changes. Second, cost structures—particularly returns to scale—shape these responses (see, e.g., Antràs, Fort, Gutiérrez, and Tintelnot, 2024), potentially allowing domestic producers to reduce their costs as they gain market share following tariff implementation. Crucially, neither conduct nor cost structure is known a priori by researchers or policymakers, calling for a data-driven approach to accurately predict the effects of policy.

We address these challenges by developing a new empirical framework and applying it

[1] See Head and Spencer (2017) for a discussion of how oligopoly models have been used in the international trade literature over the last few decades. Bian, Head, and Orr (2025) show in simulations that oligopoly conduct affects optimal trade policy.

to the US automobile market using a uniquely comprehensive dataset. Our approach, while building on a long line of related studies that use industrial organization (IO) tools to study industry-specific trade policy (e.g., Goldberg, 1995; Berry, Levinsohn, and Pakes, 1999), makes two key contributions. First, we assemble novel data that not only records equilibrium outcomes but also tracks both country-model level production and the assembly location and parts origin of every vehicle model sold in the US market. This granular information on global value chains allows us to precisely model how stacked tariffs cascade through production networks, affecting prices and welfare. Second, we develop an econometric procedure for testing firm conduct that accommodates non-constant marginal costs—thus going beyond existing methods (Backus, Conlon, and Sinkinson, 2021; Duarte, Magnolfi, Sølvsten, and Sullivan, 2024) that typically assume away economies of scale despite their importance in many manufacturing industries.

Leveraging state-of-the-art demand estimates from Grieco, Murry, and Yurukoglu (2024) that credibly capture consumers' substitution patterns, we find that Nash-Cournot quantity-setting competition (hereafter Cournot) with substantial economies of scale best characterizes the US automobile market. This finding aligns with institutional features of the industry, where production targets are set well in advance and used in negotiations with suppliers. In terms of scale economies, we estimate that a 10% increase in production reduces marginal costs by a bit more than 1%. This estimate is broadly in line with previous literature, and has a meaningful effect on our counterfactual predictions.

When we simulate tariff scenarios under the conduct model and scale economies selected by our framework, we find striking results that underscore the importance of both accounting for global value chains and correctly modeling conduct and costs. Specifically, we perform three counterfactuals in the 2018 US car market where we levy a 25% tariff on (i) imported cars alone, (ii) imported cars and parts, and (iii) cars and parts imported to the US as well as parts exported from the US for foreign assembly. In our first counterfactual, levying a 25% tariff on imported cars causes the price of foreign-assembled vehicles to increase by 24.2% while the price of US-assembled cars decreases by 0.9%. Adding tariffs on imported parts fundamentally reverses these patterns. US-assembled cars, which initially benefited from car-only tariffs, face substantial cost increases due to their reliance on imported components, causing prices to rise by 6.5% rather than fall. This reversal highlights the crucial importance of accounting for global value chains in policy evaluation. Foreign car prices remain high but slightly lower than in the previous counterfactual (22.8%), as some substitution back to foreign car models occurs. Reciprocal tariffs lead to further price increases, driven by higher marginal costs for foreign manufacturers.

The predicted price effects illustrate the importance of learning the model of conduct and cost when seeking to learn the pass-through of tariffs. In all three counterfactuals,

we find that the pass-through of tariffs to the price of foreign-assembled cars under our preferred model (Cournot competition with economies of scale) are close to but exceed one (a 25% tariff increases the average port cost - analogous to a wholesale price - by 21.3% under our model, but leads to a 24.2% increase in retail price). This finding is broadly consistent with the empirical literature on the 2018 Trump administration tariffs: Amiti, Redding, and Weinstein (2019), Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020), Flaaen, Hortaçsu, and Tintelnot (2020), and Cavallo, Gopinath, Neiman, and Tang (2021) document complete or near-complete pass-through of tariffs to import prices across a wide range of products, with tariff-inclusive prices rising roughly one-for-one with tariff rates. Furthermore, we find in our first counterfactual that pass-through of tariffs to domestically produced cars is negative under our preferred model. Instead, imposing the standard assumption in the IO literature – Bertrand-Nash price competition (Bertrand, hereafter) with constant marginal costs – implies a positive pass-through of tariffs to the prices of domestically produced cars. These differences in pass-throughs have material implications for the predicted effects of tariffs on profits and welfare.

Tariffs in this industry generate heterogeneous profit impacts across manufacturers based on their integration in global value chains. With car-only tariffs, large importers like the Volkswagen Group suffer disproportionate losses, while firms with mainly domestic production using US parts, such as Tesla and Honda, see substantial gains. When parts tariffs are added, firms with large US assembly operations but low US parts content face significant losses instead. Reciprocal tariffs most severely impact manufacturers with production strategies relying on foreign assembly using US parts.

Consumer welfare effects show similar heterogeneity and sensitivity to tariff structure. With car-only tariffs, high-income consumers and foreign car buyers experience the largest welfare losses, while buyers of US cars see small welfare gains on average due to price decreases. Adding parts tariffs nearly doubles the overall welfare loss, with additional losses now disproportionately falling on purchasers of American cars as their prices rise. Further welfare losses from reciprocal tariffs are relatively small, highlighting the limited ability of US trade partners to hurt the average American consumer in this market.

Aggregating these effects, we find substantial welfare losses that far outweigh government revenue gains. Car-only tariffs result in a loss of combined consumer surplus and profits of $31.4 billion, with only $13.5 billion offset by tariff revenue. Stacked tariffs on both cars and parts nearly double the loss in total surplus from $18 billion to $31.6 billion in 2018. Notably, revenue from parts tariffs exceeds revenue from car tariffs, highlighting the outsized importance of intermediate goods in global value chains. Reciprocal tariffs lead to an additional net surplus loss of $3.6 billion.

Beyond these immediate welfare effects, we can extend our analysis to reveal additional

consequences of tariffs, e.g., on manufacturing employment. While tariffs create domestic jobs, generating approximately 107,000 jobs under car-only tariffs and 44,000 jobs under the full reciprocal scenario, these gains come at a substantial cost, ranging from $168,000 to $876,000 per year in welfare loss per job created.

Our study is related to an important paper by Head and Mayer (2019), which also analyzes counterfactual tariffs in the US auto market. Our approach differs in two key dimensions, offering complementary evaluations of this policy. First, while they study long-run effects allowing for global production reallocation and product entry or exit, we examine medium-term impacts where market outcomes adjust but production locations and product portfolios remain fixed. Second, while they model the global industry structure, we focus exclusively on the US market with granular data on parts content, enabling a detailed account of how stacked tariffs propagate through value chains.

Our medium-term approach follows other studies that evaluate trade policy using IO methods in the automobile industry,[2] including analyses of voluntary export restraints (Goldberg, 1995; Berry et al., 1999) and tariff reductions (Fershtman, Gandal, and Markovich, 1999; Tovar, 2012).[3] We contribute to this literature by pursuing greater flexibility on the supply side—specifically addressing questions of firm conduct and cost structure raised by earlier pioneering studies such as Feenstra and Levinsohn (1995) and Verboven (1996). Rather than assuming a particular form of competition, we develop new methods to test alternative models of conduct under non-constant marginal costs.

Methodologically, our paper expands the current conduct testing toolkit developed by Backus et al. (2021) and Duarte et al. (2024) by accommodating non-constant marginal costs in the Rivers and Vuong (2002) framework.[4] This provides new tools for policy evaluation in industries characterized by economies of scale, and is portable to a wide range of applications. We also develop novel insights regarding instrument relevance, showing that the identification challenges highlighted by Bresnahan (1982) and Lau (1982) are specific to homogeneous goods settings. Specifically, we build on the results in Dearing, Magnolfi, Quint, Sullivan, and Waldfogel (2024) to show that the data on differentiated products can be used to construct *economically distinct* instruments. By affecting market outcomes through different economic channels, these instruments can both distinguish conduct and identify the cost structure. This insight substantially broadens the set of relevant instruments, making falsification of supply models feasible with standard industry data.

Recent studies have examined the impact of realized Trump-era tariffs (e.g., Amiti et al.,

[2] See Biesebroeck and Verboven (2025) for a comprehensive survey of empirical work on the automobile industry.

[3] Verboven (1996), Goldberg and Verboven (2001), and Coşar, Grieco, Li, and Tintelnot (2018) study related questions on international price discrimination and home-market advantage in the auto industry.

[4] Bian et al. (2025) propose an alternative testing approach, also under constant marginal cost.

2019; Fajgelbaum et al., 2020; Flaaen et al., 2020; Cavallo et al., 2021). We focus specifically on the automobile industry—a strategically important sector with a particularly complex value chain—which was threatened with but ultimately not subjected to broad tariffs during the first Trump administration. Within this industry, we are able to dissect the effect of different counterfactual policies that target either assembled cars only, or also parts.

The paper proceeds as follows. In the next section, we present facts illustrating the dependence of the US auto market on global value chains and our data sources. In Section 3, we introduce our empirical model of demand and supply. In Sections 4 and 5, we develop a general method for distinguishing models of conduct allowing for non-constant marginal costs and implement our testing procedure in the auto industry. Section 6 presents our counterfactual analysis of tariff policies. Section 7 concludes.

2 Background and Data

2.1 The Threatened 2018 Auto Trade War

In May 2018, the Trump administration launched a Section 232 investigation into whether automotive imports posed a threat to national security. The investigation potentially jeopardized nearly $300 billion in imports, as the administration threatened to impose tariffs of up to 25% on imported cars and auto parts. Despite significant concerns raised by industry stakeholders and trading partners, the investigation continued through 2019. Ultimately, broad tariffs on imported vehicles were never implemented during the first Trump administration, though the threat remained a source of uncertainty for the industry. The second Trump administration has also threatened, imposed, and paused stacked tariffs on cars and car parts.

Our research examines a counterfactual question: what effect would the earlier threatened tariffs have had on welfare in the US automobile market in 2018? To understand the potential impact, we must first recognize that the automotive industry is characterized by a highly integrated global value chain, developed in part through free trade agreements such as NAFTA and its successor, the USMCA. This integration creates complex effects when tariffs are stacked on assembled vehicles and car parts: imported vehicles face direct tariff levies, domestically assembled vehicles experience higher costs due to tariffs on imported parts, and reciprocal measures by trading partners can further raise costs for manufacturers using US-made components abroad. Understanding these interconnections requires granular data on both assembly locations and parts sourcing—information that becomes available through mandatory automotive labeling requirements. We next examine the data that allows us to quantify these forces, highlighting three key stylized facts that motivate our modeling approach.

2.2 Stylized Facts from Automotive Labeling Requirements

To understand how stacked tariffs would propagate through automotive value chains, we use data collected by the National Highway Traffic Safety Administration (NHTSA) under the American Automobile Labeling Act (AALA).[5] The AALA requires vehicle window stickers to display assembly location and parts sourcing information, including the percentage of parts value produced domestically versus abroad.[6] Combining AALA data with US car sales data from Grieco et al. (2024), three main stylized facts emerge:

Fact 1: A Substantial Share of Cars Sold In the US Are Foreign-Assembled In 2018, foreign-assembled vehicles accounted for 39.5% of new car sales in the US. This means that tariffs on assembled vehicles would affect a large portion of the market. These imports originate from diverse locations: while traditional suppliers like Japan and Germany remain important, Mexico has emerged as a major source, accounting for almost 10% of US car sales.

Fact 2: US-Assembled Cars Rely Heavily on Imported Parts For vehicles assembled in the United States, foreign parts represented, on average, 50% of total vehicle value in 2018. This import dependence for intermediate inputs means that tariffs on auto parts would significantly increase costs for domestic manufacturers. The reliance on imported parts varies considerably across manufacturers and models, with many US-assembled vehicles containing over 50% foreign content: parts tariffs would thus have differential effects across domestic producers, depending on their supply chain strategies.

Fact 3: Foreign-Assembled Cars Rely Substantially on US Parts Cars assembled abroad for the US market contain on average 9% US-made components. However, there is substantial variation across countries of assembly: US parts constitute approximately 38% of the value of Mexican car imports to the US (De Gortari, 2019). This reverse integration means that reciprocal tariffs imposed by trading partners would harm US parts exporters and could increase costs for foreign manufacturers selling in the US.

These three facts highlight the interdependencies in modern automotive production. Building on these three facts, our counterfactual analysis will leverage our granular data to

[5] Other studies using this source include Klier and Rubenstein (2007) and Head, Mayer, and Melitz (2024).

[6] Due to the high degree of economic integration between Canada and the US, the NHTSA considers parts sourced from both these countries as domestic. We manually disentangle, for each car model produced in 2018, the value of parts sourced from the US and Canada as follows. As transmissions and engines are the two highest-value automotive parts, we find all the plants manufacturing these parts in Canada. We then search for the car models that use the parts produced at these plants in their assembly.

assess the counterfactual effects of car tariffs. To inform the model that will predict counterfactual effects, we use a broader database on the US car industry, which we describe next.

2.3 Data

For equilibrium outcomes in the US automobile market and car model characteristics, we rely on the dataset produced in Grieco et al. (2024). This dataset contains 5,046 car model-year observations from 2002 to 2018. For each observation, we observe the manufacturer's suggested retail price (MSRP which we refer to as price throughout), sales, and product characteristics, including size dimensions (height, width, length), curb weight, horsepower, and fuel efficiency measured as miles per dollar (MPD) and miles per gallon (MPG). We also observe the number of available trims for each model[7] and the years since the current design was introduced. We also observe the real exchange rate (RXR) for the country of production. Vehicle segments in our data include sedans (42% of observations), Sport Utility Vehicles (SUVs) (23%), trucks (7%), and vans (7%).

To the Grieco et al. (2024) dataset, we merge both the AALA data described above and annual production data at the car model-country level from Marklines, an automotive industry data provider. The Marklines data, available from 2002 onward, enables us to capture economies of scale more comprehensively than previous studies of this market. Our final dataset contains 3,929 observations at the car model-year level.

Table 1 presents summary statistics, which reveal two key features of the industry that inform our modeling approach. First, the data exhibit substantial variation in production scale, with annual production ranging from 8,000 units at the 25th percentile to 74,000 units at the 75th percentile. Sales also show considerable dispersion, with a mean of approximately 60,000 units but a maximum of 891,000 units annually. As economies of scale could be an essential determinant of production costs, we must account for scale effects when modeling firm behavior and predicting the effects of trade policy.

Second, the market displays important product differentiation. Prices average $39,000 with substantial variation, possibly reflecting not only cost differences but also significant product differentiation across multiple dimensions. Vehicle characteristics vary considerably - for example, horsepower ranges from 173 to 296 from the 25th to 75th percentile, while fuel economy ranges from 17 to 23 miles per gallon; the average model offers 7 different trim levels, with some popular models available in over 20 configurations. This multidimensional differentiation may contribute to creating market power for manufacturers within their product niches.

[7] Car models may have multiple "trim levels" - for example, the 2018 Toyota Camry was available as a L, LE, SE, XLE, and XSE trim level, each with increasing numbers of options and higher price. The Grieco et al. (2024) dataset, aggregates observations across trims.

TABLE 1: Summary Statistics

	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
Prices (thousands $)	39	16	13	27	46	100
Sales (thousands)	60	91	0.01	8.3	74	891
RXR	0.98	0.17	0.39	0.9	1.1	1.6
Height (inches)	64	8.1	45	57	69	107
Width (inches)	74	4	61	71	76	89
Length (inches)	190	18	106	179	200	274
Curb weight (lbs)	3,966	900	1,808	3,362	4,456	7,230
Horsepower	240	82	66	173	296	645
MPD	7.6	3.4	2.8	5.3	8.9	28
MPG	21	7.8	10	17	23	50
Number of trims	6.7	14	1	2	7	210
Years since design	4.2	1.9	3	3	5	12
Production by car model (thousands) - US made	166	167	2.7	47	249	1,023
Production by car model (thousands) - Foreign made	122	138	2.6	30	169	761
Number of car models per category			2002	2007	2013	2018
Sedan			51	89	146	150
SUV			39	82	94	106
Truck			17	18	12	10
Van			12	17	12	15

The table reports summary statistics for our new car sales sample with 3,929 car model-year observations from 2002-2018. RXR is the real exchange rate for the country of production. MPD is miles per dollar, calculated using contemporaneous local gas prices. MPG is miles per gallon.

Given this variation in scale and complex product differentiation, different assumptions on cost functions and conduct (e.g., about whether firms compete on prices or quantities) can lead to very different predictions about pass-through rates, production reallocations, and consumer welfare effects after tariffs are levied. Thus, our empirical approach develops a testing procedure that evaluates alternative models of supply rather than imposing a pre-selected one. As we seek to learn features of supply from the data, there are important sources of variation that we can leverage. For instance, we observe significant variation over time in the number of models per category: the SUV category grows dramatically in terms of models offered (and sales), while the van category declines. This variation will be useful for testing conduct.

3 Model

To capture the medium-run effects of tariffs on the US automobile market, we develop a static equilibrium model of the industry. We discuss demand and supply in turn.

3.1 Demand

We use the demand system estimated in Grieco et al. (2024). This state-of-the-art demand system builds on seminal work on the auto industry (e.g. Berry, Levinsohn, and Pakes, 1995; Goldberg, 1995). The model allows for flexible substitution patterns, informed by detailed microdata on consumer choices and second-choice survey responses.[8] We provide a brief summary here, which allows us to develop the notation used throughout the paper. In what follows, for a generic variable a_{jt}, a_t denotes a'_{jt} stacked across products j within market t and a denotes a_t stacked across markets t.

Firms offer a set of products $\mathcal{J}_t$ in each market t. Consumers in market t, indexed by i receive indirect utility from each new car model j according to:

$$u_{ijt} = \alpha_{it} p_{jt} + x'_{jt}\beta_{it} + \xi_{jt} + \epsilon_{ijt}.$$

Here, consumer utility depends on the product's price p_{jt}, a vector of observed car characteristics x_{jt}, an unobserved car attribute ξ_{jt} and an idiosyncratic shock ϵ_{ijt} which is assumed to follow a Type 1 extreme value distribution. The model permits rich and time-varying consumer heterogeneity by allowing the preference parameters α_{it} and β_{it} to depend on observed demographics and consumer-level shocks.[9] Consumers maximize their utility by choosing either a single car model or the outside option of no new car purchase, with the utility of the latter varying over time (η_t will capture the average utility of the outside good). The market share of each product in each market s_{jt} takes on the familiar form:

$$s_{jt} = \int \frac{\exp(\alpha_{it} p_{jt} + x'_{jt}\beta_{it} + \xi_{jt})}{\exp \eta_t + \sum_{k \in \mathcal{J}_t} \exp(\alpha_{it} p_{kt} + x'_{kt}\beta_{it} + \xi_{kt})} \mathrm{d}F_t(i),$$

where F_t is the distribution of the random coefficients. Thus market shares stacked across products in market t can be expressed as the function $s_t = \mathcal{s}\left(p_t, x_t, \xi_t, \theta_0^D\right)$ where θ_0^D is the vector of demand parameters. For a market t of size M_t, the equilibrium quantity of model j is given as $q_{jt} = M_t s_{jt}$.

This demand specification produces realistic substitution patterns, with similar vehicle styles serving as close substitutes and strong correlations between purchased and second-choice vehicles in characteristics such as size, horsepower, and fuel economy. Given the rigorous estimation strategy developed in Grieco et al. (2024) and their extensive microdata, we adopt their parameter estimates as θ_0^D without modification.

[8] Head and Mayer (2023) find that CES and monopolistic competition can approximate a DGP with BLP demand, Bertrand conduct and constant cost when the implied pass-through is close to one. In this paper, we want to allow for flexible substitution patterns under alternative supply models.

[9] Importantly, α_i contains both consumer-level income and income squared. This flexibility in the mixing distribution helps alleviate the concerns about curvature raised by Birchall, Mohapatra, and Verboven (2024) and Miravete, Seim, and Thurk (2025).

3.2 Supply

The trade literature has shown that assumptions on the supply side—the imposed model of conduct (see e.g., Dixit, 1984; Brander and Spencer, 1985; Dixit and Grossman, 1986; Eaton and Grossman, 1986) and functional form of marginal costs (see e.g., Antràs et al., 2024)—have ramifications for the effects of tariffs. Simulations in Bian et al. (2025) also show that the oligopoly model crucially affects conclusions about optimal trade policy. In empirical IO, researchers typically assume that firms face constant marginal costs and compete according to Bertrand pricing. While these assumptions may hold in many settings, there is reason to question whether they are appropriate for studying the US automobile market. In particular, Berry et al. (1995) provide suggestive evidence that economies of scale may exist in the assembly of cars. Furthermore, the long production lead times and capacity constraints in automobile manufacturing suggest that output decisions may be more rigid than prices; in the US market, Feenstra and Levinsohn (1995) and Berry et al. (1999) consider alternatives to Bertrand pricing, which include Cournot quantity setting and mixed models where some firms set prices and others set quantities.

Thus, instead of assuming a particular parametric model and using it to measure the effect of tariffs, we seek to use the data to guide our assumptions on conduct and cost. To do so, we begin with a general framework where the data in each market t are generated by equilibrium play in some static model of supply where prices and quantities are endogenous. A system of first-order conditions characterizes the true supply model,

$$p_t = \Delta_{0t} + c_{0t},$$

where $\Delta_{0t} = \Delta_0\left(p_t, s_t, \theta_0^D\right)$ is the true vector of markups in market t and c_{0t} is the true vector of marginal costs. Following Berry et al. (1995), Verboven (1996), Goldberg and Verboven (2001), Coşar et al. (2018), and Grieco et al. (2024), we specify our model for the log of marginal cost. Further, we assume a Cobb-Douglas production function,[10] resulting in the functional form in Berry et al. (1995), Verboven (1996), and Goldberg and Verboven (2001):[11]

$$\log(c_{0jt}) = \gamma_0 \log(q_{jt}^p) + \mathrm{w}_{jt}'\tau_0 + \omega_{0jt}, \tag{1}$$

where w_{jt} is a vector of observed cost shifters that affects the product's marginal cost, q_{jt}^p reflects the production quantity corresponding to the level at which economies of scale accrue, and ω_{0jt} is an unobserved shock.

The choice of q_{jt}^p deserves some discussion. In our main specification, we assume q_{jt}^p

[10] See Khmelnitskaya, Marshall, and Orr (2025) for a recent empirical approach that also allows for economies of scope.

[11] While other papers have assumed constant marginal costs (e.g. Goldberg, 1995; Berry et al., 1999; Van Biesebroeck, Gao, and Verboven, 2012), our functional form nests constant marginal cost.

corresponds to the total production of model j in the country that supplies the US market in year t. In this industry, this is similar to the assumption that economies of scale accrue at the plant-model level: in the vast majority of cases, car models sold in the US are sourced from one country (Head and Mayer, 2019), and produced in a single plant in that country. Our assumption stands between the approaches that scale economies accrue from global production across countries for a model (Verboven, 1996; Goldberg and Verboven, 2001), and those that assume "external" scale economies across all models produced by a firm within a country (Head and Mayer, 2019). We explore the robustness of our results to alternative assumptions in Appendix D.3, including platform-level economies of scale (which captures economies of scope among car models using the same engineering platform) and log-linear and quadratic functional forms.

Because the true markups (or, equivalently, true costs) are unobserved in the market for cars, we want to determine the supply model that best fits the data from a menu of candidates motivated by the literature. Distinguishing supply models requires instruments (see e.g., Berry and Haile, 2014; Dearing et al., 2024). We therefore maintain that instrumental variables z_{jt} exist such that, for the true model, the exclusion restriction $E[\omega_{0jt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$ holds. This assumption requires that the instruments are exogenous for supply, and therefore uncorrelated with the unobserved cost shifters for the true model. Berry and Haile (2014) provides standard sources of exogenous variation, such as variation in the set of rival firms and rival products, own and rival product characteristics, rival cost, and market demographics. These variables are available in our dataset and we now explore how they can be used to distinguish models of conduct under economies of scale.

4 Testing Conduct With Non-Constant Marginal Cost

Our application requires distinguishing between models of firm conduct while accounting for economies of scale in production. To do so, we must address two limitations in the IO literature. First, current testing methods (Backus et al., 2021; Duarte et al., 2024) assume constant marginal costs. Second, credible inference on conduct depends on instrument selection; while this problem is well understood under constant marginal costs (Dearing et al., 2024), concerns remain about finding relevant instruments when relaxing this assumption. Bresnahan (1982) and Lau (1982) demonstrate that homogeneous product markets require specific demand-rotating instruments to distinguish monopoly from perfect competition. Therefore, even if we extend the testing procedure to accommodate non-constant marginal cost, it is not immediate from the literature that the exogenous variation in standard differentiated products datasets like ours will distinguish models of conduct in practice.

We address these two concerns in this section. First, we propose a general procedure

to test models of conduct while also estimating economies of scale, which we apply to the car market in Section 5. Our procedure extends the methodology of Duarte et al. (2024), which adapts the RV non-nested model selection test (Rivers and Vuong, 2002) to implement the falsifiable restrictions of Berry and Haile (2014). This approach is appropriate because our candidate supply models are non-nested[12] and the RV test offers advantages over model assessment alternatives. Second, we examine instrument relevance for the RV test in our differentiated products setting with economies of scale. We argue that in such markets there is a broader set of potentially relevant instruments than in homogeneous products contexts, as researchers can leverage cross-product variation in cost and product characteristics unavailable in Bresnahan (1982). This insight makes testing conduct models under non-constant marginal cost feasible with standard datasets like ours. While we focus on the specific cost function in Equation (1), adopted for our application, Appendices A and B broaden our procedure and guidance for instrument relevance beyond this specific case.

4.1 Our General Procedure

Step 0. Construct Menu of Models: The researcher specifies plausible candidate models of conduct based on institutional knowledge. Our procedure applies to static conduct models where, for each model m, prices and quantities in market t satisfy model-specific first-order conditions:

$$p_t = \Delta_{mt} + c_{mt} \tag{2}$$

where Δ_{mt} is the vector of *implied markups* under model m in market t and c_{mt} is the vector of *implied marginal costs*.[13]

Step 1. Obtain Implied Markups and Marginal Costs Under Each Model: For each model we consider, Δ_{mt} can be expressed as a function of data and demand primitives. Thus, given data on equilibrium outcomes and the demand estimates (such as those in Grieco et al. (2024) for our application), Δ_{mt} can be computed in each market and the implied costs c_{mt} can be recovered as $c_{mt} = p_t - \Delta_{mt}$.

As an example, consider the canonical Bertrand pricing model, or $m = B$. The implied markups stacked across products j in market t can be expressed as:

$$\Delta_{Bt} = -(\Omega_t \odot D_t')^{-1} s_t,$$

[12] Even with nested models, testing can be preferred to estimation; see Magnolfi and Sullivan (2022).

[13] For the first-order conditions of any model m to characterize a well-defined empirical model, we require there exists a unique equilibrium, or the equilibrium selection rule is such that the same p_t arises whenever the vector (c_{mt}, x_t, ξ_t) is the same, analogous to Assumption 13 in Berry and Haile (2014).

where Ω_t is the ownership matrix in market t whose (j,k)-th element is an indicator that products j and k are produced by the same firm, D_t is the matrix of demand derivatives whose (j,k)-th element is $\frac{\partial s_{jt}}{\partial p_{kt}}$, and s_t is the vector of market shares. The symbol $\odot$ denotes Hadamard (element-wise) multiplication.

Step 2: Estimate Model-Implied Scale Economies and Cost Shocks Given the vector of implied markups Δ_{mt} and implied marginal costs $c_{mt} = p_t - \Delta_{mt}$, we can estimate the marginal cost function implied by model m:

$$\log(c_{mt}) = \gamma_m \log(q^p_{jt}) + \mathrm{w}'_{jt}\tau_m + \omega_{mjt}. \tag{3}$$

Remark 1. Compared to the case of constant marginal costs, the right-hand side of (3) contains endogenous log quantities. Thus, estimating γ_m and τ_m requires the use of instruments and estimation via two-stage least squares (2SLS).

We pin down the parameters in marginal cost as solutions to the standard two-stage least squares (2SLS) moment equations $E[\mathrm{w}_{jt}\omega_{mjt}] = 0$ and $E[\tilde{q}_{jt}\omega_{mjt}] = 0$ where $\tilde{q}_{jt}$ is the best linear predictor of log quantities using cost shifters and instruments,

$$\tilde{q}_{jt} = (z'_{jt}, \mathrm{w}'_{jt})E\left[(z,\mathrm{w})'(z,\mathrm{w})\right]^{-1} E\left[(z,\mathrm{w})'\log(q^p)\right].$$

This leads to standard 2SLS definitions of the parameters γ_m and τ_m,

$$\begin{pmatrix}\gamma_m \\ \tau_m\end{pmatrix} = E\big[(\tilde{q},\mathrm{w})'(\log(q^p),\mathrm{w})\big]^{-1} E\big[(\tilde{q},\mathrm{w})'\log(p-\Delta_m)\big].$$

The implied cost shocks ω_{mjt}, recovered as the 2SLS projection errors, rationalize the observed equilibria (p_t, s_t) under model m in market t.

Step 3. Construct Measure of Lack-of-Fit for Each Model: For the true model, we know that $E[\omega_{0jt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$. Therefore, given that we have recovered the model-implied cost shocks ω_{mjt}, we can define a measure of lack-of-fit based on the moment condition $E[z_{jt}\omega_{mjt}] = 0$, which will hold if model m is the correct model.

Remark 2. Pinning down the economies of scale parameter uses part of the variation in z_{jt}. Thus, the key deviation from the constant marginal cost case is that the effective instrument only contains $d_z - 1$ (instead of d_z) sources of linearly independent variation.

For example, if z_{jt} is a scalar, then the cost function would be exactly identified and, by definition, the variation in z_{jt} would be orthogonal to ω_{mjt} for any model. It is therefore useful to form our measure of lack-of-fit using only the variation that remains in the instruments after pinning down the cost function. This remaining variation in the instruments,

denoted z^e_{jt}, is the error in z_{jt} from a population projection on w_{jt} and $\tilde{q}_{jt}$,

$$z^e_{jt} = z_{jt} - E\left[z'\left(\tilde{q}, \mathrm{w}\right)\right] E\left[\left(\tilde{q}, \mathrm{w}\right)'\left(\tilde{q}, \mathrm{w}\right)\right]^{-1} \begin{pmatrix} \tilde{q}_{jt} \\ \mathrm{w}_{jt} \end{pmatrix}.$$

We consider the following generalized method of moments (GMM) objective function as our measure of lack-of-fit for a model m,

$$Q_m = g_m' W g_m,$$

where $g_m = E[z^e_{jt}\omega_{mjt}]$. In analogy with Duarte et al. (2024), the weight matrix is based on the $d_z \times d_z$ covariance matrix $E[z^e_{jt}z^{e\prime}_{jt}]$. However, z^e is obtained by residualizing z against a linear combination of its elements, i.e., $\tilde{q}$. As the rank of $E[z^e_{jt}z^{e\prime}_{jt}]$ is $d_z - 1$, we use as the weight matrix the Moore-Penrose inverse,[14] $W = E[z^e_{jt}z^{e\prime}_{jt}]^+$. Notice that, if $E[z^e_{jt}\omega_{mjt}] = 0$, model m is indistinguishable from the true model and $Q_m = 0$. Otherwise, if $E[z^e_{jt}\omega_{mjt}] \neq 0$, the instruments distinguish model m from the truth and $Q_m > 0$.

Step 4. Run RV Test, Obtain Model Confidence Set: For each model m, we have a measure of fit $Q_m = E[z^e_{jt}\omega_{mjt}]' W E[z^e_{jt}\omega_{mjt}]$. To make inferences about conduct, we adopt the non-nested model selection test in Rivers and Vuong (2002). We consider all pairs of models from our menu and, for each, we run the pairwise RV test of models $m = 1, 2$. The null hypothesis for the test is that the two competing models of conduct have the same fit,

$$H_0^{\mathrm{RV}} : \; Q_1 = Q_2.$$

Relative to this null, we define two alternative hypotheses corresponding to cases of better fit of one of the two models:

$$H_1^{\mathrm{RV}} : Q_1 < Q_2 \qquad \text{and} \qquad H_2^{\mathrm{RV}} : Q_2 < Q_1.$$

With this formulation of the null and alternative hypotheses, the statistical problem is to determine which of the two models has the best fit, or equivalently, the smallest lack of fit.

For the GMM measure of fit, the RV test statistic is then

$$T^{\mathrm{RV}} = \frac{\sqrt{n}(\hat{Q}_1 - \hat{Q}_2)}{\hat{\sigma}_{\mathrm{RV}}},$$

where $\hat{Q}_m$ is a sample analog of Q_m and $\hat{\sigma}^2_{\mathrm{RV}}$ is an estimator for the asymptotic variance of the scaled difference in the measures of fit appearing in the numerator of the test statistic. We denote this asymptotic variance by σ^2_{RV}.

Remark 3. Since the estimation of γ_m is accounted for in constructing Q_m, no adjustments

[14] Equivalently, one may redefine z^e to only include $d_z - 1$ of the residualized instruments, where the dropped instrument can be any instrument with a non-zero coefficient in $\tilde{q}$. In doing so, one obtains full rank of $E[z^e_{jt}z^{e\prime}_{jt}]$ and can let $W = E[z^e_{jt}z^{e\prime}_{jt}]^{-1}$.

are needed for $\hat{Q}_m$, defined in Appendix B. In constructing $\hat{\sigma}^2_{\text{RV}}$, it is important to account for estimation of the linear predictor $\tilde{q}$, which we do by modifying the variance estimator of Duarte et al. (2024) in Appendix B.

For a menu of two models, there is one RV test statistic and a researcher can conclude for a model if the test rejects in favor of that model. For more than two models, the researcher obtains multiple RV test statistics. For instance, in our application, the menu contains five candidate models, meaning that we have ten unique pairs, leading to ten RV test statistics. To adjust for multiple testing, Duarte et al. (2024) shows how all the pairwise T^{RV} statistics can be used to construct a model confidence set.

Step 5. Construct F-statistic to diagnose weak instruments for testing: For each pair of models, the RV test statistic, T^{RV}, is asymptotically standard normal under the null as long as the estimator of the asymptotic variance converges in probability to $\sigma^2_{\text{RV}} > 0$. If instead $\sigma^2_{\text{RV}} = 0$, the RV test is said to be degenerate, which Duarte et al. (2024) characterizes as a problem of irrelevant instruments for testing. Near the space of degeneracy, instruments may be weak, resulting in size distortions or low power. Duarte et al. (2024) provides a diagnostic for both size distortions and low power. Their diagnostic relies on a joint F-statistic of two "first stage" regressions, $\omega_{mjt} = z'_{jt}\pi_m + e_{mjt}$ for model $m = 1, 2$ in the pair being tested. They then provide critical values for worst-case size and maximal power to which the F-statistic can be compared in order to diagnose instrument strength.

Here, given that our measure of fit is formed with the residual variation in the instruments z^e, we modify the weak instruments diagnostic as a scaled F-statistic for a joint test of the hypotheses H_{0m} : $\pi_m = 0$ for $m = 1, 2$ in the regressions

$$\omega_{mjt} = z^{e\prime}_{jt}\pi_m + e_{mjt} \quad \text{with } E[z^e_{jt}e_{mjt}] = 0 \quad \text{for model } m = 1, 2. \tag{4}$$

Remark 4. To appropriately diagnose instrument strength, one can compare our F-statistic to the same critical values for size and power as reported in Duarte, Magnolfi, Sølvsten, Sullivan, and Tarascina (2022) with the value for the number of instruments being $d_z - 1$.

This remark holds because our modified RV test statistic T^{RV} and F-statistic converge to the same joint distribution as in Proposition 4 of Duarte et al. (2024) with d_z being adjusted to $d_z - 1$. Appendix B proves this and provides the definition of the F-statistic for our setting.

4.2 Instrument Relevance for Testing Conduct with Economies of Scale

The procedure outlined in Section 4.1 relies on the presence of relevant instruments to obtain sharp inference on conduct. In light of the discussion in Bresnahan (1982) and Lau (1982) for the homogeneous product case, a researcher may fear that such instruments are

difficult to form in practice, and special variation (such as rotation of demand) is needed to distinguish models of conduct in the presence of economies of scale. Thus, we investigate in this subsection what requirements instruments need to satisfy to be relevant for testing conduct in differentiated products settings such as the one in our application.

To develop an intuitive, but still practical, discussion, we consider a simplified environment where the true model of conduct is Bertrand, and the researcher wants to test Bertrand and Cournot. We further suppress the observed cost shifters, so that for either model $\log(c_{mj}) = \gamma_m \log(q_{jt}^p) + \omega_{mjt}$. Appendix A develops instead formal results for relevance under general economies of scale, and for a broader range of models.

The literature has considered relevance when testing models under constant marginal cost. Dearing et al. (2024) establish that instruments are relevant in this case if the features of pass-through matrices the instruments target differ across models. Thus, for the logit class of demand systems, one single instrument formed with standard sources of variation (e.g., product characteristics and rival cost shifters) satisfies the requirement for the Bertrand and Cournot models. Here we show that there are two additional requirements for instrument relevance.

Remark 5. A single instrument $z_{jt}^{(1)}$ is irrelevant for testing any pair of models of conduct with economies of scale.

If the researcher has a single available instrument, then in Step 2 of our procedure the 2SLS estimator for economies of scale would be given by the sample analogue of

$$\gamma_m^{(1)} = E[z_{jt}^{(1)} \log(q_{jt}^p)]^{-1} E[z_{jt}^{(1)} \log(p_{jt} - \Delta_{mjt})].$$

Because γ_m is just-identified, it follows that the implied 2SLS errors satisfy $E[\omega_{mjt} z_{jt}^{(1)}] = 0$. When forming a measure of lack of fit in Step 4, the researcher only has one instrument and thus has to form Q_m with the same $z^{(1)}$. Mechanically, then, without an overidentifying restriction, $Q_m = 0$ for any model m, and the RV test is degenerate with $F = 0$.

Remark 6. Two instruments $z_{jt}^{(1)}$ and $z_{jt}^{(2)}$ must be *economically distinct* to be relevant for testing model m against the true model, which we define as

$$\frac{E[z_{jt}^{(1)} \log(p_{jt} - \Delta_{mjt})]}{E[z_{jt}^{(2)} \log(p_{jt} - \Delta_{mjt})]} \neq \frac{E[z_{jt}^{(1)} \log(q_{jt}^p)]}{E[z_{jt}^{(2)} \log(q_{jt}^p)]}. \tag{5}$$

This requirement, which generalizes the results in Lau (1982), rules out that the two instruments have the same correlation with log of production quantities and with log of the implied costs under model m. As long as the two instruments correspond to sources of variation that have different economic effects, this statistical condition will generally hold.

To see where this condition comes from, suppose a researcher estimates economies of scale with one of the instruments $z^{(k)}$ for $k = 1, 2$, recovering $\omega_{mjt}^{(k)}$ as the error. While

mechanically $E[z_{jt}^{(k)}\omega_{mjt}^{(k)}] = 0$, the researcher can use $E[z_{jt}^{(-k)}\omega_{mjt}^{(k)}]$ to form a measure of fit $Q_m^{(-k)}$. If each instrument, however, implies the same economies of scale, then $\omega_{jt}^{(1)} = \omega_{jt}^{(2)}$. It follows then that $E[z_{jt}^{(-k)}\omega_{mjt}^{(k)}] = E[z_{jt}^{(-k)}\omega_{mjt}^{(-k)}] = 0$, and the instruments are irrelevant for testing any model m. The condition in (5) ensures that the economies of scale separately measured with the two instruments are different, thus preventing this mechanical channel for irrelevance.

Remark 7. Standard datasets from differentiated product markets typically contain economically distinct instruments.

In the representative-firm homogeneous goods setting in Bresnahan (1982), there are no rival firms and products, so that variation that both shifts and rotates the inverse demand curve—which can be achieved by a demand rotator—is the only relevant source of variation. With product differentiation, observing distinct product characteristics and cost shifters across firms provides sources of economically distinct instruments. For instance, in a market with two single-product firms and logit demand, economically distinct pairs leverage cross-firm/product variation. Examples include pairing an own product characteristic with a rival cost shifter, or pairing an own and a rival product characteristic. Not all instrument pairs work, however. Economically indistinct pairs use variation within the same firm/product. These include two cost shifters of the same rival, two characteristics of the same product, or a characteristic and cost shifter from the same rival.

Taken together, we have developed a general and portable procedure to test conduct under non-constant marginal cost. Appendix A includes a more formal discussion of the economic determinants of instrument relevance for more general models of conduct and cost functions. Appendix C develops a set of Monte Carlo simulations to illustrate the procedure in practice and the role of instrument relevance. We now turn to implement the procedure in the market for US automobiles

5 Implementing Our Procedure in the Market for Cars

Step 0. Construct Menu of Models: We consider a menu of five different models of firm conduct, motivated by Feenstra and Levinsohn (1995) and Berry et al. (1999).

1. *Bertrand pricing:* All car manufacturers choose prices to maximize profits taking competitors' prices as given
2. *Cournot quantity setting:* All car manufacturers choose quantities of each car model to sell, prices adjust to clear market
3. *Mixed Model:* Subset of manufacturers $\mathcal{B}$ set prices, remaining subset $\mathcal{C}$ set quantities.

   3.1 $\mathcal{B}$: Asian firms, $\mathcal{C}$: US and European firms

3.2 $\mathcal{B}$: US firms, $\mathcal{C}$: Asian and European firms

3.2 $\mathcal{B}$: European firms, $\mathcal{C}$: Asian and US firms

While Bertrand pricing is a standard assumption in the empirical IO literature, in the market for cars Feenstra and Levinsohn (1995) and Berry et al. (1999) have also considered Cournot conduct and *"mixed" models* where some manufacturers compete in prices while others compete in quantities. This specification reflects potential asymmetries between manufacturers - for instance, firms with more flexible production systems may be better positioned to compete on price, while those with rigid capacity may effectively compete in quantities. We build from Feenstra and Levinsohn (1995)'s results and consider three mixed models based on manufacturers' nationality (American, Asian, and European); (i) only Asian Manufacturers play Bertrand, (ii) only American manufacturers play Bertrand, and (iii) only European manufacturers play Cournot.[15]

Step 1. Obtain Implied Markups and Marginal Costs Under Each Model: For each model m we consider, Δ_{mt} can be computed in each market from the data on equilibrium outcomes and the demand estimates in Grieco et al. (2024). The implied costs can be recovered as $c_{mt} = p_t - \Delta_{mt}$.

Specifically, for Bertrand pricing, the stacked markups for each product j in market t can be expressed as:

$$\Delta_{Bt} = -(\Omega_t \odot D_t')^{-1} s_t,$$

where Ω_t is the ownership matrix and D_t is the matrix of demand derivatives as defined in section 4.1. Under Cournot competition, the implied markups take a similar form given as:

$$\Delta_{Ct} = -(\Omega_t \odot (D_t^{-1})') s_t. \tag{6}$$

To define the implied markups for the mixed models, we first establish some notation. For any matrix A_{Mt}, we order Bertrand and Cournot players in market t, and partition

$$A_{Mt} = \begin{bmatrix} A_{BBt} & A_{BCt} \\ A_{CBt} & A_{CCt} \end{bmatrix}.$$

Markups for the mixed model are thus:

$$\Delta_{Mt} = \begin{bmatrix} -(\Omega_{CCt} \odot (D_{CCt}^{-1})') s_{Ct} \\ -(\Omega_{BBt} \odot (D_{BCt}' (D_{CCt}^{-1})' D_{CBt}' + D_{BBt}'))^{-1} s_{Bt} \end{bmatrix}$$

where subscript C denotes firms competing in quantities and B denotes firms competing in prices.

[15] Using data for 1987, Feenstra and Levinsohn (1995) find evidence for a Nash equilibrium where American firms choose prices, European firms choose quantities, and Japanese firms choose either prices or quantities.

In Figure 1, Panel A we illustrate the distribution of markups (expressed as Lerner index $\frac{\Delta_{mjt}}{p_{jt}}$) in our data for the five candidate models. In line with theory (Magnolfi, Quint, Sullivan, and Waldfogel, 2022), Cournot conduct implies higher markups than Bertrand, and mixed models imply markups between these two extremes. While the different conduct assumptions yield different levels of markups, they all imply similar declining trends over time, in line with the findings of Grieco et al. (2024). Specifically, across all specifications we find declining markups from 1980 to 2018, consistent with the Bertrand case—this is shown in Figure 1, Panel B.

FIGURE 1: Markup Implications of Models of Conduct

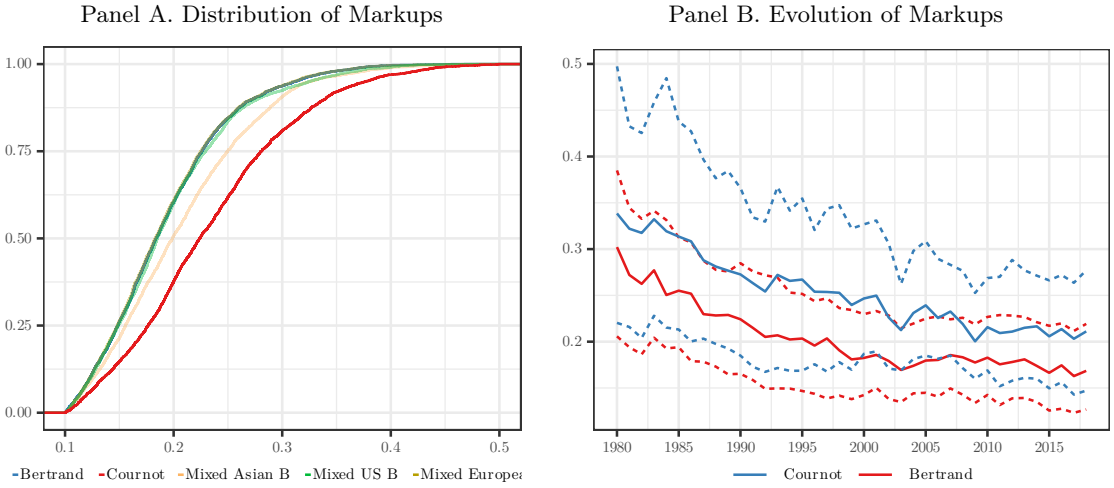


Figure illustrates the distribution of implied markups corresponding to different models (panel A), and the evolution of average markups over time (panel B).

We further explore the implications of conduct by examining in Table 2 the Bertrand and Cournot markups for the four best-selling models in the sedan, SUV, truck, and van categories. Across these car models, Bertrand markups are 4-13 percentage points below Cournot markups, which is substantial in this industry.

Step 2: Estimate Model-Implied Scale Economies and Cost Shocks To estimate the model-implied economies of scale, we need to specify a marginal cost function and construct instruments.

Instrument Choice: We construct instruments that satisfy our requirements for economic distinctness developed in Section 4.2. Following the insights from that section and from Appendix A, we select four sets of instruments that affect market outcomes through different economic channels, ensuring they generate distinct patterns of response in quantities and markup differences across models of conduct.

TABLE 2: Model-Implied Markups for Top Car Models by Segment

Sedans			SUVs		
Car Model	Bertrand	Cournot	Car Model	Bertrand	Cournot
Toyota Camry	23.1	28.7	Toyota RAV4	21.6	28.5
Honda Civic	28.7	33.7	Nissan Rogue	21.9	26.8
Honda Accord	20.7	24.3	Honda CR-V	20.1	24.8
Toyota Corolla	30.2	37.4	Chevrolet Equinox	23.8	32.8
Trucks			**Vans**		
Car Model	Bertrand	Cournot	Car Model	Bertrand	Cournot
Ford F Series	18.8	31.6	Dodge Caravan	26.4	37.7
Chevrolet Silverado	17.8	31.5	Ford Transit	21.2	32.8
Ram Pickup	15.1	24.6	Chrysler Pacifica	25.8	36.4
Toyota Tacoma	23.4	29.7	Honda Odyssey	17.1	21.6

The table displays estimated markups (as Lerner indices, Δ_{mjt}/p_{jt}) in percent units under Bertrand and Cournot conduct assumptions for the top-selling car models in each vehicle segment during 2018.

Our first instrument leverages market structure variation through the number of rival models within the segment (sedan, SUV, truck, van). This captures the level of competitive intensity, which impacts markups via demand and strategic interactions. The second set of instruments exploits variation in consumer demographics by using the average income and age, which affect consumers' price sensitivity, and thus demand. The third instrument is an SUV indicator interacted with a time trend. This instrument is particularly valuable for identifying economies of scale since it predicts demand-driven variation in production volumes over time within this growing segment. The fourth instrument is the average real exchange rate (RXR) of rivals. As demonstrated in Grieco et al. (2024), exchange rates for the country of assembly serve as effective cost shifters. Our segment-level aggregation generates opponent cost-shifter variation that differs fundamentally from own-product demand shifters.

These instruments satisfy our *economic distinctness* requirement because they operate through separate channels: market structure (intensity of competition), consumer preferences (demographic-driven taste), product-specific demand trends (SUV evolution), and rivals' production costs (exchange rate effects). Therefore, this diverse set of instruments generates the variation needed to separately identify economies of scale and distinguish between alternative models of conduct, as required by our theoretical framework.

Economies of Scale Estimates: Following Grieco et al. (2024), we specify the vector of cost shifters w_{jt} to include the exchange rate, a quadratic time trend, the log of continuous characteristics (height, footprint, horsepower, miles per gallon, curbweight, number of trims), and indicators for release year, segment, electric, sport, luxury, years since design,

and manufacturer fixed effects. Table 3 presents 2SLS estimates of the economies of scale parameter γ_m under different models of conduct while Table 11 in Appendix D.1 reports full regression results. Appendix Table 15 explores the robustness of these estimates to alternative cost specifications, including platform-level economies of scale and log-linear and quadratic functional forms.

TABLE 3: Economies of Scale Estimates

Model m	γ_m Estimate	Standard Error
Bertrand	−0.119	0.032
Cournot	−0.114	0.033
Mixed (Asian firms Bertrand)	−0.115	0.032
Mixed (US firms Bertrand)	−0.116	0.032
Mixed (European firms Cournot)	−0.118	0.032

The table reports estimates of parameter γ_m obtained by estimating Equation (3) via 2SLS under different models of conduct m (reported in different rows). All specifications include cost shifters and fixed effects (see Table 11 for full regression results). Standard errors are clustered by car model.

The estimates suggest meaningful economies of scale across all specifications, with scale elasticities ranging from −0.114 to −0.119. The magnitude of the economies of scale estimates is stable across different conduct assumptions, suggesting this finding is robust to how we model competitive interaction. These results are broadly in line with other findings in the literature: Fuss and Waverman (1990) find a coefficient of −0.07, Verboven (1996) finds −0.11 (essentially identical to our finding), while Goldberg and Verboven (2001) find −0.006 to −0.03, although the authors point out that these estimates may be attenuated due to the presence of a quota.[16] The finding of economies of scale is also consistent with industry evidence of economies of scale in car assembly and parts manufacturing.[17]

Steps 3-5. Conduct Testing: We implement the testing procedure to evaluate our different models of conduct. Table 4 presents the results.

The test results favor the model where all firms engage in Cournot competition. This model is included in the model confidence set with a p-value of 1.00, while all other specifications are rejected. Although the instruments lack power for some model pairs (as indicated by the F-statistics), they are sufficiently strong overall to deliver reliable inference, yielding a model confidence set containing only one model.

The finding that the Cournot model best fits the data in this market is consistent with several features of this industry. First, it can approximate yearly production targets, with

[16] In the trade literature, Head and Mayer (2019) find a coefficient for external scale economies of −0.035, while Bartelme, Costinot, Donaldson, and Rodriguez-Clare (2025) estimate an external scale elasticity for motor vehicle production of −0.16.

[17] For example, Doner, Noble, and Ravenhill (2021) state: "By 2023, profitable economies of scale [require] annual sales of 1 million vehicles using a given platform and 300,000 units for engines."

TABLE 4: Conduct Test Results

Models	1	2	3	4	5	MCS p-values
Panel A: Test Results (T^{RV})						
1. Bertrand		3.05	1.75	1.19	−0.84	0.005
2. Cournot			−3.16	−2.89	−3.18	1.000
3. B: Asia, C: US, EU				−1.27	−2.06	0.007
4. B: US, C: Asia, EU					−1.90	0.004
5. B: Asia, US, C: EU						0.009
Panel B: F-Statistics						
1. Bertrand		14.8	13.6	6.4†	7.2†	
2. Cournot			16.9	13.4	14.2	
3. B: Asia, C: US, EU				10.5	11.9	
4. B: US, C: Asia, EU					6.1†	
5. B: Asia, US, C: EU						

Panels A-B report the RV test statistics T^{RV} and the effective F-statistic for all pairs of models, and the MCS p-values. A negative RV test statistic suggests a better fit of the row model. F-statistics indicated with † are below the appropriate critical value for best-case power above 0.95. With MCS p-values below 0.05 a row model is rejected from the model confidence set. Both T^{RV} and the F-statistics account for clustering at the market level.

prices that adjust via dealers incentives (e.g., Tremblay, Tremblay, and Isariyawongse, 2013). Moreover, choosing quantities facilitates parts sourcing. Finally, capacity constraints are salient for automobile manufacturing and Cournot can approximate a two-stage capacity and price competition game (Kreps and Scheinkman, 1983; Hendel, 1994).

6 Evaluating Tariffs

We now want to evaluate the effects of imposing stacked and reciprocal tariffs in the 2018 US automobile market. To do so, we use the model chosen by our model selection procedure—Cournot quantity setting with economies of scale. To disentangle the effect of tariffs on finished cars and car parts, we perform three counterfactuals. In the first counterfactual, denoted C1, we impose a 25% tariff on all foreign-manufactured cars (including Mexico and Canada) imported to the US. Second, in C2, we impose a 25% tariff on both foreign-manufactured cars and foreign-manufactured car parts imported to US. Third, in C3, we impose 25% tariffs on all foreign imports and 25% reciprocal tariffs on US manufactured parts exported for foreign assembly. The chosen 25% tariff level is in line with previous proposals by President Trump in 2018.[18]

[18] E.g., https://www.piie.com/blogs/trade-and-investment-policy-watch/2018/trumps-proposed-auto-tariffs-would-throw-us-automakers.

6.1 Implementation of Counterfactuals

Under the Cournot model with economies of scale, the first order conditions in each market are given by Equation (2), where $\Delta_{mt} = \Delta_{Ct}$ takes the form specified in Equation (6) and the implied costs $c_{mt} = c_{Ct}$ are obtained by taking the exponential of Equation (3). To simulate counterfactual outcomes, we hold the demand and marginal cost parameters and estimated cost shock ω_{Ct} fixed at their estimated levels. Across the counterfactuals, we need to turn on and off specific tariffs depending on the country of assembly of the car model and the sourcing of model's parts. To simplify the exposition, we define an indicator variable for foreign assembly of model j, ϕ_j.

To perform C1, we need to modify the first order conditions to account for a 25% tariff on foreign assembled cars. Tariffs apply to the port cost, essentially an unobserved wholesale price between manufacturer and dealer. Since port costs are unobserved, we must model how they are set. We consider two approaches: (i) modeling the port price as a fraction λ of the observed retail price p_{jt} which the firm internalizes when making profit maximizing choices, or (ii) modeling the port cost as a fraction ν of the implied marginal cost c_{mjt} which the firm takes as given with respect to its action. In the market for cars, Goldberg (1995) and Goldberg and Verboven (2001) model unobserved wholesale prices by adopting the first approach. Meanwhile, Coşar et al. (2018) adopt the second approach.

To account for the tariff, we need to appropriately modify the first-order conditions for foreign assembled car models. Following the first approach, we could impose the 25% tariff as an $(25 \times \lambda)$% ad-valorem tax levied on the producers of foreign assembled cars; in the second approach, we could scale νc_{mjt} by 25%. Notice, that for $\lambda = \nu(1+\phi_j 0.25\nu)^{-1}$, these two approaches are equivalent and the modified first order conditions with tariffs become[19]

$$p_{jt} = \Delta_{mjt} + (1 + \nu 0.25)^{\phi_j} c_{mjt}(\cdot). \tag{7}$$

As our main specification for C1, we set $\lambda = 0.8$ (or equivalently $\nu = 1$), which is broadly in line with Goldberg (1995). As a robustness check we set $\lambda = 0.58$ (equivalent to $\nu = 0.68$, corresponding to the estimate from Coşar et al. (2018))—results are in Appendix D.3.

To perform C2 and C3, we need to further modify the first order conditions to account for tariffs on car parts. To perform these counterfactuals, we maintain several assumptions. First, we assume for both foreign and domestically assembled cars that car assembly is 29% of marginal cost while parts represent the remaining 71% (Menk, Chen, and Cregger, 2012). We also assume a fixed value for the pass-through of the tariff to price of parts, denoted Λ for all cars. From the AALA data, we know for each model of car j the fraction of the

[19] Dividing both sides of Equation (7) by $(1 + 0.25\phi_j\nu)$ and plugging in $\nu = \phi(1 - \phi_j 0.25\nu)^{-1}$ yields $(1 - \phi_j\lambda \times 0.25)p_{jt} = (1 - \phi_j\lambda \times 0.25)\Delta_{mjt} + c_{mjt}(\cdot)$ which can be interpreted as levying an ad valorem tax on the firms.

total value of parts that were produced outside the US, which we denote μ_j and the fraction produced in the US, denoted $(1-\mu_j)$. Thus, across both C2 and C3 we can augment the first order conditions to account for parts tariffs by scaling the appropriate part of marginal cost. Specifically, in C2, we add tariffs on the foreign parts used in US assembled vehicles so that the the first order conditions for each product j become

$$p_{jt} = \Delta_{mjt} + \left(1 + \nu 0.25\right)^{\phi_j} \left(1 + \mu_j \Lambda 0.71 \cdot 0.25\right)^{(1-\phi_j)} c_{mjt}(\cdot).$$

In C3, we now include reciprocal tariffs on US parts used in foreign production so that first order conditions for each product j become

$$p_{jt} = \Delta_{mjt} + \left(\left(1 + \nu 0.25\right)\left(1 + (1-\mu_j)\Lambda 0.71 \cdot 0.25\right)\right)^{\phi_j} \left(1 + \mu_j \Lambda 0.71 \cdot 0.25\right)^{(1-\phi_j)} c_{mjt}(\cdot).$$

In our main specifications, we set $\Lambda = 1$, which is akin to assuming in-house production of all parts. This is also in line with the finding in Ganapati and Hottman (2025) of near-unit pass-through of tariffs within supplier relationships. In Appendix D.3, we consider alternative values of Λ.

With each counterfactual's modified first-order conditions, we solve for equilibrium prices using the generalization of the Morrow and Skerlos (2011) procedure developed in Duarte (2025). There is a caveat to our analysis. We model scale economies arising from total production, so that changes in foreign sales in reaction to US tariffs could, in principle, impact firms' marginal costs and US market outcomes. However, we only model car sales in the US market. Thus, we keep foreign sales constant in our counterfactuals, and recompute global production figures based on changes in US sales predicted by our model.

Finally, we note that our counterfactual analysis extrapolates the estimated cost function to production levels that differ substantially from the baseline, with decreases in import volumes but also meaningful increases in domestic production for certain car models (with the median model-level increase reaching 16.2% in C3). A natural concern is whether such production increases are feasible while achieving the same returns to scale we estimate in-sample. We maintain that this assumption is plausible given industry capacity constraints in 2018. According to Federal Reserve data, capacity utilization in the US automotive sector was approximately 75% in 2018.[20] This means that virtually all domestic factories could plausibly add production shifts to accommodate increased demand in the short term while preserving similar cost structures.

[20] See https://fred.stlouisfed.org/series/CAPUTLG33611SQ.

6.2 Counterfactual Results

The counterfactual analysis reveals substantial effects from stacked tariffs. Table 5 shows average price increases of 9.05% across all models in C1. This average masks heterogeneity by assembly location: foreign-assembled cars experience 24.21% price increases (a 25% tariff increases the average port cost by 21.3% under our model, but leads to a 24.2% increase in retail price), while the prices of US-assembled cars decrease by 0.85%.

Our finding of pass-through exceeding one for foreign-assembled vehicles is broadly consistent with the empirical literature on the 2018 Trump administration tariffs. Amiti et al. (2019), Fajgelbaum et al. (2020), and Cavallo et al. (2021) documented complete or near-complete pass-through of tariffs to import prices across a wide range of products, with tariff-inclusive prices rising roughly one-for-one with tariff rates. While our analysis focuses on a counterfactual scenario rather than realized tariffs, and the automobile industry has specific characteristics that may generate different pass-through rates, our estimated pass-through of 1.14 falls within the range documented in that literature.

However, our findings reveal important differences from predictions that would emerge under standard assumptions in the industrial organization literature. As we show in Appendix D.3, a model assuming Bertrand pricing with constant marginal costs predicts that foreign manufacturers would exhibit sub-unit pass-through, while domestic manufacturers would increase prices in response to the tariff. In contrast, our preferred model of Cournot competition with economies of scale generates the opposite pattern: slightly above-unit pass-through for foreign firms and price decreases for domestic producers.[21] These differences have substantial welfare implications.[22]

The price reduction by domestic manufacturers in our model reflects strategic substitutability between foreign and domestic producers under Cournot competition with economies of scale. As tariffs raise foreign car prices, domestic firms gain market share and benefit from lower marginal costs due to increased production volumes. This strategic substitutability is reinforced by the quantity-setting nature of Cournot competition, where firms respond to competitors' higher costs by expanding output rather than by raising prices. Our finding echoes results in Berry et al. (1999), who documented that European car prices fell in response to voluntary export restraints on Japanese vehicles, even under their assumption of Bertrand competition with constant costs. The Cournot assumption combined with economies of scale further amplifies this strategic substitutability, leading domestic firms to

[21] The fact that different models of supply imply different pass-through rates suggests that, when variation in tariffs or tax rates is available in the data, pass-through may be used to distinguish models of supply. This is an old idea in IO (e.g., Sumner, 1981); Dearing et al. (2024) show that instrument-based methods, such as the ones we use, essentially generalize this intuition and make it more broadly applicable.

[22] We find that the Bertrand-constant cost model understates the total welfare loss by over $6 billion in our most comprehensive tariff scenario.

reduce prices as they capture market share from constrained foreign competitors.

TABLE 5: Counterfactual Results – Prices

		Fraction US		% Change Prices		
	# Models	Assembly	Parts	C1	C2	C3
All	314	0.35	0.23	+9.05	+12.96	+14.56
				(0.79)	(0.86)	(0.99)
US	111	1.00	0.50	−0.85	+6.52	+6.39
				(0.41)	(0.1)	(0.1)
Non-US	203	0.00	0.09	+24.21	+22.83	+27.07
				(2.52)	(2.11)	(2.51)

The table reports the sales-weighted average of the percentage change in prices corresponding to tariff counterfactuals C1-C3.

Table 5 also shows the effect of stacked and reciprocal tariffs on car prices in the US market. When tariffs are also levied on car parts, the marginal cost of assembling cars in the US increases. For cars assembled in the US, stacking tariffs on parts causes prices to rise on average by 6.52%, as opposed to falling by 0.85%. As this increase generates some substitution back to foreign firms, their overall price change under stacked tariffs is slightly lower (22.83% vs 24.21%), due to economies of scale and strategic substitutability via Cournot competition. Finally, comparing C2 and C3 shows the limited effect that reciprocal tariffs on car parts have on car prices to US consumers. Overall, the response of foreign manufacturers is to increase prices by an additional 4.24 percentage points, enabling US assembled cars to lower prices by a modest 0.13 percentage points. Accounting for all three types of tariffs, the average price of cars is increased by 14.56% in the US market.

Levying stacked and reciprocal tariffs on cars and car parts has differential effects on profits across firms, depending on their choices of assembly locations and parts sourcing, as we show in Table 6. When the tariffs are levied only on imported cars (C1) large importers like VW see disproportionate losses, while firms with mostly domestic production (e.g., Tesla and Honda) see large gains. The Big 3 automakers (GM, Ford, Fiat Chrysler), with intermediate levels of US production, see their profits almost unaffected. Stacking tariffs on parts imported to the US (C2) changes the relative effect on profits across firms. Now, relative to C1, losses accrue for firms with large US assembly and low US parts, including GM, Ford, and Tesla. While reciprocal tariffs on US parts (C3) have a limited impact for most manufacturers compared to C2, Fiat Chrysler – a firm with disproportionate foreign assembly and US parts – sees its losses double. In Appendix D.2, we further explore changes in profits on a model-specific basis, investigating whether manufacturers have incentive to discontinue certain car models.

TABLE 6: Counterfactual Results – Profits

		Fraction US		% Change Variable Profits		
	# Models	Assembly	Parts	C1	C2	C3
All	314	0.35	0.23	−3.68	−10.88	−11.43
				(0.77)	(0.41)	(0.54)
GM	42	0.76	0.41	−0.97	−13.66	−14.61
				(1.16)	(1.00)	(1.43)
Toyota	39	0.58	0.44	−10.66	−11.72	−13.58
				(0.37)	(1.15)	(1.36)
Ford	26	0.74	0.47	−0.90	−6.31	−6.43
				(2.45)	(0.49)	(0.56)
Fiat-Chrysler	27	0.65	0.57	−2.64	−0.79	−4.01
				(1.73)	(0.69)	(0.70)
Honda	17	0.91	0.49	+31.14	+19.05	+21.64
				(1.70)	(4.72)	(4.55)
Volkswagen	29	0.16	0.11	−51.05	−50.91	−52.13
				(2.10)	(2.96)	(2.98)
Tesla	3	1.00	0.51	+11.04	−3.11	−2.60
				(0.32)	(0.39)	(0.43)

The table reports, for a subset of manufacturers, the percent change in total variable profits corresponding to tariff counterfactuals C1-C3.

Tariffs on cars and car parts have sizable effects on consumer welfare. In Table 7, we explore changes in household-level welfare, measured in US dollars, both by demographic types and by the purchase decision made by the household. In C1, we see that tariffs on imported cars (C1) hurt consumers overall. However, the welfare losses are larger for high-income consumers (as they are more likely to stay in the market when prices increase) and buyers of foreign cars. In fact, there are small welfare gains for buyers of US cars, as these cars are sold at lower prices when the prices of foreign cars increase due to economies of scale and Cournot strategic substitution.

Comparing C1 to C2 shows the important effect that stacked tariffs have in this market. The additional losses to households are large – overall, the welfare loss nearly doubles. Now, since the cost and prices of US assembled cars increase, the losses compared to C1 disproportionally fall on purchasers of American cars. Finally, as the price effects of reciprocal tariffs are minimal, these responses by foreign governments in C3 have limited effects on US consumers of cars.

TABLE 7: Counterfactual Results – Household-Level Consumer Surplus

	Change in $				Change in $		
	C1	C2	C3		C1	C2	C3
All	−508	−1,115	−1,161	Purchase Type			
	(68.4)	(9.9)	(11.2)	Buy any car	−1422	−2604	−2694
Demog. Type					(95.7)	(13.5)	(18.1)
Income							
1st Q	−112	−198	−204	Low income	−113	−199	−204
	(9.5)	(1.4)	(2.2)		(10.0)	(1.6)	(2.6)
2nd Q	−210	−527	−543	Buy American	309	−1,633	−1,598
	(46.7)	(12.7)	(17.0)		(216.2)	(96.8)	(107.7)
3rd Q	−568	−1,225	−1,276	Buy Import	−3,514	−3,880	−4,146
	(86.9)	(12.0)	(15.5)		(54.8)	(121.7)	(91.4)
4th Q	−1,179	−2,593	−2,712	Inside Option	−1,422	−2,604	−2,694
	(134.8)	(30.3)	(22.7)		(95.7)	(13.5)	(18.1)
HH Size							
1	−250	−599	−615	Buy Amer. Low Inc.	248	−1,365	−1,338
	(48.2)	(8.5)	(12.0)		(182.2)	(87.3)	(97.1)
3-4	−684	−1,373	−1,445	Buy Amer. Low Inc. x Rural	256	−1,120	−1,101
	(74.8)	(11.3)	(13.8)		(175.9)	(75.0)	(84.0)
Rural	−436	−995	−1,030				
	(78.1)	(13.6)	(17.7)				
Urban	−518	−1,132	−1,181				
	(66.9)	(10.3)	(10.6)				

The table reports the average compensating variation across households within a group for counterfactuals C1-C3. The last three columns report changes based on pre-tariff purchase decisions.

In Table 8, we measure total welfare changes (in billions $) from the three counterfactuals. In each scenario, the welfare losses to consumers and firms easily outweigh government revenues. Levying taxes only on imported cars in C1 results in a total welfare loss to consumers and firms of $31.42 billion, of which only $13.47 billion is offset by tax revenue. Comparing C2 to C1, stacked tariffs magnify the welfare losses. The losses to consumers and firms more than double, while the total net surplus lost in the US market for cars increases to more than $31 billion. Interestingly, the revenue from tariffs on parts exceeds the revenue on car imports themselves. Finally, comparing C2 to C3, we find additional losses from reciprocal tariffs of $3.57 billion.

TABLE 8: Counterfactual Results – Total Surplus

	Change in Billions $		
	C1	C2	C3
Consumer Surplus	-25.91	-56.89	-59.27
	(3.49)	(0.51)	(0.57)
Profit	-5.51	-16.31	-17.13
	(1.15)	(0.62)	(0.80)
Tax Revenue:			
Car imports	13.47	17.93	15.57
	(1.63)	(1.48)	(1.39)
Parts imports		23.65	25.64
		(0.09)	(0.18)
Total Net Surplus	-17.95	-31.62	-35.19
	(3.12)	(1.26)	(1.10)

The table reports overall welfare changes and tax revenue corresponding to tariff counterfactuals C1-C3.

6.3 Beyond Welfare: Effects on Employment

While our counterfactual analysis reveals substantial welfare losses from tariffs, these policies may be motivated by objectives beyond consumer welfare—particularly employment creation in domestic manufacturing and the reshoring of production. Our model does not directly capture non- price and quantity decisions of auto manufacturers. However, we can use our counterfactual predictions to perform a back-of-the-envelope quantification of the employment effects of each tariff scenario to provide a more comprehensive assessment of trade-offs. We develop this analysis in this subsection, and provide in Appendix D.2 some results on how tariffs provide incentives for the reshoring of parts production.

To estimate employment effects in both vehicle assembly and parts manufacturing for the US in 2018, we assume that the annual ratio of labor expenditure to total revenue remains constant after tariff imposition.[23] For each of the two manufacturing sectors, indexed by k, this ratio is defined as:

$$\kappa_k = \frac{\bar{w}_k L_k}{R_k},$$

[23] This assumption is in line with the Cobb-Douglas production function that underlies our cost function specification.

where L_k represents total US employment in sector k, $\bar{w}_k$ indicates the sector's average wage, and R_k denotes total sector revenue. Further assuming that wages are unaffected in the medium-run, employment in counterfactuals C1-C3 (L'_k) can then be computed from counterfactual revenue R'_k as:

$$L'_k = \frac{\kappa_k}{\bar{w}_k} R'_k.$$

To implement this approach, we calculate κ_k for both assembly and parts sectors using data from the Bureau of Labor Statistics' Current Employment Statistics (for employment and wages) and the Census Bureau's Annual Survey of Manufactures (for total revenue).[24] In 2018, κ_k is 0.042 for assembly, and 0.096 for parts.[25]

In each counterfactual C1-C3, our model predicts counterfactual revenues for cars assembled in the US as $R'_{\text{assembly}} = \sum_{j \text{ assembled in US}} q'_j p'_j$, where q'_j is the counterfactual quantity of model j in 2018 and p'_j is the counterfactual price. To compute counterfactual revenues in parts manufacturing, we scale the appropriate fraction of the counterfactual marginal cost (c'_j) corresponding to US produced parts by the counterfactual quantity. R'_{parts} sums this model-level revenue over all models sold in the US, regardless of country of assembly. We calculate it as $R'_{\text{parts}} = \sum_j q'_j \times c'_j \times 0.71 \times (1 - \mu_j)$.

Table 9 presents our employment impact estimates across the three tariff scenarios. The results reveal that tariffs would indeed create domestic jobs in both assembly and parts manufacturing, though with varying effectiveness across scenarios.

TABLE 9: Employment Effects by Tariff Scenario

	C1	C2	C3
Assembly Employment:			
Added jobs (thousands)	68.7	32.65	40.28
	(10.32)	(4.21)	(3.93)
Percentage change (%)	29.41	13.98	17.25
	(4.42)	(1.80)	(1.68)
Parts Employment:			
Added jobs (thousands)	38.31	3.44	3.33
	(11.02)	(4.98)	(5.26)
Percentage change (%)	6.39	0.57	0.56
	(1.84)	(0.83)	(0.88)
Total Jobs Created (thousands)	107.01	36.09	43.61
Annual Welfare Loss per Job (thousands $)	167.7	875.9	807.1

The table reports estimated employment changes in the US automotive assembly and parts manufacturing sectors under tariff scenarios C1-C3. The annual welfare loss per job is calculated by dividing the total welfare loss from Table 8 by the total number of jobs created.

[24] The vehicle assembly industry corresponds to NAICS 3361, while parts manufacturing is represented by NAICS 3363.

[25] Both κ_k have remained stable over time in the data.

Under the car-only tariff scenario (C1), employment increases significantly in both sectors, with approximately 68,700 new assembly jobs (a 29.4% increase) and 38,300 new parts manufacturing jobs (a 6.4% increase). When tariffs are extended to imported parts (C2), the employment gains in assembly and parts are substantially reduced to 32,700 jobs (14% increase) and 3,400 jobs (0.6%), respectively. This stark reduction demonstrates how tariffs on intermediate inputs can undermine the employment benefits of final goods protection. The reciprocal tariff scenario (C3) shows slight changes to the employment effects: 40,300 new jobs (17.3% increase) are created in assembly while 3,300 jobs (0.6%) are created in parts manufacturing.

To assess the efficiency of tariffs as an employment creation tool, we calculate the welfare cost per job created, shown in the bottom row of Table 9. Each job created under the car-only tariff scenario (C1) costs $167,700 in lost welfare per year. When tariffs extend to parts (C2), this cost rises to $875,900 per job. Under the reciprocal tariff scenario (C3), the cost remains high at $807,100 per job.

These figures significantly exceed the average annual compensation in the automotive sector (approximately $75,000 in 2018), suggesting that tariffs represent a highly inefficient mechanism for job creation. The costs are particularly high under the stacked tariffs scenarios, where increased input costs severely undermine the employment benefits while magnifying welfare losses.

7 Conclusion

This paper makes two key contributions. First, we develop a new method for testing firm conduct under economies of scale, revealing that Cournot competition with substantial scale effects best characterizes the US automobile market. Our approach provides a framework applicable to other manufacturing sectors where economies of scale matter, and distinguishing the correct model of supply (conduct and cost) from data may help shape credible policy analysis. Second, we demonstrate how stacked tariffs propagate through global value chains, producing effects that would be hard to predict without a granular evaluation of the foreign parts content of domestic vehicles. Our counterfactual analysis shows how tariffs on imported vehicles alone induce price decreases for domestic cars and welfare gains for their buyers, while additional tariffs on parts reverse these effects, causing domestic prices to rise by 6.5% and nearly doubling total welfare losses to $31 billion. The employment gains from such policies come at a substantial welfare cost-rising from $107,000 per job under car-only tariffs to $875,900 per job when parts tariffs are included. These findings highlight the importance of accounting for both conduct and global value chains when evaluating trade policy, especially in industries with complex international production networks.

Our analysis is subject to several important caveats that echo concerns raised in Berry et al. (1999). First, we employ a static model on both demand and supply, abstracting from potentially important dynamic considerations. On the demand side, consumers may anticipate or delay purchases in anticipation of future price changes. On the supply side, important dynamic aspects of manufacturers' decision-making include capacity adjustments and product development. Second, we do not model the used car market, which may serve as an important substitute for new vehicles. For both limitations, our focus on medium-run effects provides some justification: production relocation and product development face substantial frictions in this time frame, while the used vehicle stock remains relatively fixed. Third, our counterfactuals involve substantial price changes, requiring extrapolation beyond observed variation. While such extrapolation always requires caution, we take some comfort in the fact that the demand system from Grieco et al. (2024) was estimated using long panel variation capturing substantial price movements, and our assumption of stable preferences is standard. These limitations suggest our results should be interpreted as informative but incomplete estimates of automotive trade policy effects.

References

Amiti, M., S. Redding, and D. Weinstein (2019): "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 33, 187–210.

Antrás, P. and D. Chor (2022): "Global Value Chains," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Amsterdam: Elsevier, vol. 5, chap. 5, 297–376.

Antràs, P., T. Fort, A. Gutiérrez, and F. Tintelnot (2024): "Trade Policy and Global Sourcing: An Efficiency Rationale for Tariff Escalation," *Journal of Political Economy Macroeconomics*, 2, 1–44.

Backus, M., C. Conlon, and M. Sinkinson (2021): "Common Ownership and Competition in the Ready-To-Eat Cereal Industry," NBER working paper #28350.

Bartelme, D., A. Costinot, D. Donaldson, and A. Rodriguez-Clare (2025): "The Textbook Case for Industrial Policy: Theory Meets Data," *Journal of Political Economy*, 133, 1401–1704.

Berry, S. and P. Haile (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

——— (1999): "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, 89, 400–431.

Bian, C., K. Head, and S. Orr (2025): "Inferring Conduct to Guide Strategic Trade Policy," *Canadian Journal of Economics*, forthcoming.

Biesebroeck, J. V. and F. Verboven (2025): "Demand, Competition and Public Policy in the Automobile Industry," Working paper.

Birchall, C., D. Mohapatra, and F. Verboven (2024): "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation," *Review of Economics and Statistics*, 1–40.

Brander, J. and B. Spencer (1985): "Export Subsidies and International Market Share Rivalry," *Journal of International Economics*, 18, 83–100.

Bresnahan, T. (1982): "The Oligopoly Solution Concept is Identified," *Economics Letters*, 10, 87–92.

Caldara, D., M. Iacoviello, P. Molligo, A. Prestipino, and A. Raffo (2020): "The Economic Effects of Trade Policy Uncertainty," *Journal of Monetary Economics*, 109, 38–59.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021): "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy," *American Economic Review: Insights*, 3, 19–34.

Conlon, C. and J. Gortmaker (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," *RAND Journal of Economics*, 51, 1108–1161.

Coşar, A., P. Grieco, S. Li, and F. Tintelnot (2018): "What Drives Home Market Advantage?" *Journal of International Economics*, 110, 135–150.

De Gortari, A. (2019): "Disentangling Global Value Chains," NBER working paper #25868.

Dearing, A., L. Magnolfi, D. Quint, C. Sullivan, and S. Waldfogel (2024): "Learning Firm Conduct: Pass-through as a Foundation for Instrument Relevance," Working paper.

Dixit, A. (1984): "International Trade Policy for Oligopolistic Industries," *Economic Journal*, 94, 1–16.

Dixit, A. and G. Grossman (1986): "Targeted Export Promotion with Several Oligopolistic Industries," *Journal of International Economics*, 21, 233–249.

Doner, R., G. Noble, and J. Ravenhill (2021): *The Political Economy of Automotive Industrialization in East Asia*, New York: Oxford University Press.

Duarte, M. (2025): "Extending Fixed-Point Methods for Equilibrium Computation in Markets with Differentiated Products," *Economics Letters*, 112275.

Duarte, M., L. Magnolfi, M. Sølvsten, and C. Sullivan (2024): "Testing Firm Conduct," *Quantitative Economics*, 15, 571–606.

Duarte, M., L. Magnolfi, M. Sølvsten, C. Sullivan, and A. Tarascina (2022): "pyRVtest: A Python package for testing firm conduct," https://github.com/anyatarascina/pyRVtest.

Eaton, J. and G. Grossman (1986): "Optimal Trade and Industrial Policy under Oligopoly," *Quarterly Journal of Economics*, 101, 383–406.

Fajgelbaum, P., P. Goldberg, P. Kennedy, and A. Khandelwal (2020): "The Return to Protectionism," *Quarterly Journal of Economics*, 135, 1–55.

Feenstra, R. and J. Levinsohn (1995): "Estimating Markups and Market Conduct with Multidimensional Product Attributes," *Review of Economic Studies*, 62, 19–52.

Fershtman, C., N. Gandal, and S. Markovich (1999): "Estimating the Effect of Tax Reform in Differentiated Product Oligopolistic Markets," *Journal of Public Economics*, 74, 151–170.

Flaaen, A., A. Hortaçsu, and F. Tintelnot (2020): "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines," *American Economic Review*, 110, 2103–2127.

Fuss, M. and L. Waverman (1990): "The Extent and Sources of Cost and Efficiency Differences between US and Japanese Motor Vehicle Producers," *Journal of the Japanese and International Economies*, 4, 219–256.

Ganapati, S. and C. Hottman (2025): "Quantity Discounts and Distortions: Scale Economies and Incomplete Tariff Pass-Through," Working paper.

Goldberg, P. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the US Automobile Industry," *Econometrica*, 891–951.

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

Grieco, P., C. Murry, and A. Yurukoglu (2024): "The Evolution of Market Power in the US Automobile Industry," *Quarterly Journal of Economics*, 139, 1201–1253.

Hall, A. and D. Pelletier (2011): "Nonnested Testing in Models Estimated via Generalized Method of Moments," *Econometric Theory*, 27, 443–456.

Handley, K. and N. Limão (2022): "Trade Policy Uncertainty," *Annual Review of Economics*, 14, 363–395.

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

——— (2023): "Poor Substitutes? Counterfactual Methods in IO and Trade Compared," *Review of Economics and Statistics*, 1–51.

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

Head, K. and B. Spencer (2017): "Oligopoly in International Trade: Rise, Fall and Resurgence," *Canadian Journal of Economics*, 50, 1414–1444.

Hendel, I. (1994): *Essays on Industrial Organization: Theory and Econometrics*, Harvard University.

Khmelnitskaya, E., G. Marshall, and S. Orr (2025): "Identifying Scale and Scope Economies using Demand-Side Data," *Working Paper*.

Klier, T. and J. Rubenstein (2007): "Whose Part Is It? Measuring Domestic Content of Vehicles," *Chicago Fed Letter*, 243.

Kreps, D. and J. Scheinkman (1983): "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14, 326–337.

Lau, L. (1982): "On Identifying the Degree of Competitiveness from Industry Price and Output Data," *Economics Letters*, 10, 93–99.

Magnolfi, L., D. Quint, C. Sullivan, and S. Waldfogel (2022): "Differentiated-Products Cournot Attributes Higher Markups than Bertrand-Nash," *Economics Letters*, 219, 110804.

Magnolfi, L. and C. Sullivan (2022): "A Comparison of Testing and Estimation of Firm Conduct," *Economics Letters*, 212, 110316.

Menk, D., Y. Chen, and J. Cregger (2012): "Methodology for Creating a Matrix to Assess the Domestic Content of a Vehicle by Make and Model ," Tech. rep., Center for Automotive Research.

Miravete, E., K. Seim, and J. Thurk (2025): "Elasticity and Curvature of Discrete Choice Demand Models," Working Paper.

Morrow, W. and S. Skerlos (2011): "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand," *Operations Research*, 59, 328–345.

Rivers, D. and Q. Vuong (2002): "Model Selection Tests for Nonlinear Dynamic Models," *Econometrics Journal*, 5, 1–39.

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

Sumner, D. (1981): "Measurement of Monopoly Behavior: An Application to the Cigarette Industry," *Journal of Political Economy*, 89, 1010–1019.

Tovar, J. (2012): "Consumers' Welfare and Trade Liberalization: Evidence from the Car Industry in Colombia," *World Development*, 40, 808–820.

Tremblay, V., C. Tremblay, and K. Isariyawongse (2013): "Cournot and Bertrand Competition when Advertising Rotates Demand: The Case of Honda and Scion," *International Journal of the Economics of Business*, 20, 125–141.

Van Biesebroeck, J. (2003): "Productivity dynamics with technology choice: An application to automobile assembly," *The Review of Economic Studies*, 70, 167–198.

Van Biesebroeck, J., H. Gao, and F. Verboven (2012): "Impact of FTAs on Canadian Auto Industry," *DFAIT Canada*, 31.

Verboven, F. (1996): "International Price Discrimination in the European Car Market," *The RAND Journal of Economics*, 240–268.

Online Appendix

Appendix A The General Falsification Problem: Non-constant Marginal Cost with Differentiated Products

In Section 4.2 of the paper, we focus on the problem of instrument relevance for testing models of conduct in the context of our application – distinguishing between particular conduct models (Bertrand, Cournot, or a hybrid) under a functional form assumption for marginal costs (linearity in quantity produced).

In this appendix, we develop a more formal and general analysis of what instruments are needed to distinguish models of conduct under scale economies. We adopt here a falsification lens: this is useful because Duarte et al. (2024) establishes that instruments are irrelevant for testing a pair of models with the RV test when neither model is falsified by the instruments, so that $E[\omega_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$ for $m = 1, 2$. Thus, to shed light on instrument relevance for a broader range of models under general economies of scale, we develop the condition for falsification of a conduct model, and the economic intuition surrounding it.

A.1 Preliminaries

As in the text, we assume data in each market are generated by equilibrium play in a true model of firm behavior $p_t = \Delta_{0t} + c_{0t}$. We normalize market size to 1, so that we may use s_{jt} and q_{jt} interchangeably. Each candidate model m yields its own set of implied markups Δ_{mt} as a function of observables and demand primitives; the marginal costs implied by that model, c_{mt}, are then calculated as $p_t - \Delta_{mt}$. We require the following for any model m to have well-defined first-order conditions:

Assumption 1. *(Equilibrium Uniqueness)* For any model m, including the true model, either: (i) A unique equilibrium exists, or (ii) The equilibrium selection rule consistently yields the same p_t for any given (c_{mt}, x_t, ξ_t).

This assumption is analogous to Assumption 13 in Berry and Haile (2014).

Demand and cost may include random shocks, varying across markets, that are unobserved by the researcher. We now assume that marginal costs follow a general functional form separable in the unobserved shock and a function of the observable cost shifters w_{jt} and own quantities, or $c_{0jt} = \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) + \omega_{0jt}$.[26] While this cost function departs from the functional form used in the body of the paper, which was chosen for its applicability

[26] In principle, the function $\bar{c}_{0j}$ can depend on the full vector q_t, or lags. Restricting $\bar{c}_{0j}$ to depend on own quantity simplifies the exposition, and is in line with our application. Extensions, including economies of scope, are straightforward, and generate no further qualitative insight.

to the car market, it allows us to maintain the falsifiable restriction in Berry and Haile (2014) which assumes that marginal cost is additively separable in ω_{0jt}. The researcher can observe w_{jt}, but does not know the function $\bar{c}_{0j}$ or observe ω_{0jt}. To falsify incorrect models, instruments z_{jt} are constructed, which are assumed to be mean independent of unobserved cost shocks under the true model:

Assumption 2. *(Instrument Exogeneity)* For each j, $c_{0jt} = \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) + \omega_{0jt}$, and z_{jt} is a vector of d_z excluded instruments such that $E[\omega_{0jt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$.

For the true model, $E[\omega_{0jt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$ by Assumption 2, or equivalently, $E[p_{jt} - \Delta_{0jt} - \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] = 0$ for the true cost function $\bar{c}_0(\cdot, \cdot)$. For a candidate model m and candidate cost function $\bar{c}_m$, then, defining $\omega_{mjt} = p_{jt} - \Delta_{mjt} - \bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt})$, the condition $E[\omega_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$[27] serves as a falsifiable restriction. If any cost function $\bar{c}_{mj}(\cdot, \cdot)$ allows this restriction to be satisfied, we say that model m is *not falsified by the instruments* z_{jt}. If no cost functions $\{\bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt})\}_{j=1}^{J}$ satisfy this restriction almost surely over w_{jt} and z_{jt}, then model m is falsified.

Since prices in the data are generated by the true model, we can rewrite the implied cost shock as $\omega_{mjt} = \Delta_{0jt} - \Delta_{mjt} + \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) - \bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt}) + \omega_{0jt}$. Thus, the condition for falsificaiton is as follows:[28]

Lemma 1. *Under Assumptions 1 and 2, model m is falsified by instruments z_{jt} if and only if for some j there exists no function $\bar{c}_{mj}$ such that*

$$E[\Delta_{0jt} - \Delta_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] = E[\bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt}) - \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] \qquad a.s. \tag{8}$$

Here, as in the rest of the section, the expectation is over unobservables, and *a.s. (almost surely)* is over realizations of the observables $(\mathrm{w}_{jt}, z_{jt})$.

A.2 Intuition: Constant vs Non-Constant Marginal Costs

The expectations in Equation (8) are taken conditional on the realization of w_{jt}, so a product's own cost shifters cannot offer any additional variation with which to form instruments. Since $\bar{c}_{mj}$ and $\bar{c}_{0j}$ are functions only of w_{jt} and q_{jt}, any instrument can only move $\bar{c}_{mj}$ and $\bar{c}_{0j}$ through its effect on own quantity q_{jt}.

To build intuition, we will first consider the case, explored in Dearing et al. (2024), where marginal costs do not depend on quantity produced. In that case, $\bar{c}_{mj}$ and $\bar{c}_{0j}$ can't depend on the instruments at all, so the right-hand side of Equation (8) is constant in z_{jt};

[27] This condition is analogous to the one in Theorem 9 in Berry and Haile (2014).

[28] Proofs of all lemmas, propositions, and corollaries are in Appendix A.6.

falsification therefore depends on whether an instrument moves Δ_{0jt} and Δ_{mjt} differentially. The following example illustrates how this would work.

Example 1: In this and the subsequent examples, we consider a simple environment with two single-product firms and logit demand, so that market shares $j \in \{1, 2\}$ are given by

$$s_{jt} = \frac{\exp(x_{jt}\beta - \alpha p_{jt})}{1 + \exp(x_{1t}\beta - \alpha p_{1t}) + \exp(x_{2t}\beta - \alpha p_{2t})},$$

where x_{jt} are characteristics of product j in market t and α and β are coefficients. We suppose that the true conduct model is Cournot, and we wish to falsify Bertrand competition.

For this example, consider the case where true marginal cost is constant in quantity and linear in a scalar w_{jt}, or $\bar{c}_{0j} = \mathrm{w}_{jt}\tau$, and hold unobservables fixed for ease of exposition. The researcher forms an instrument using variation in the cost shifter of firm 2, $z_{1t} = \mathrm{w}_{2t}$. Because cost shifters do not enter in the markup function of either model, z_{jt} effects Δ_{0jt} and Δ_{mjt} only through its effect on observed prices p_t.

Under the true (Cournot) model, the cost passthrough matrix can be calculated to be

$$P_{Ct} = \begin{bmatrix} \frac{s_{0t}}{1-s_{2t}} & 0 \\ 0 & \frac{s_{0t}}{1-s_{1t}} \end{bmatrix}$$

and therefore a change in $z_{1t} = \mathrm{w}_{2t}$ only changes p_{2t}, not p_{1t}. Firm 1's markup under the two models can be calculated to be $\Delta_{01t} = \Delta_{C1t} = \frac{1}{\alpha}\left(1 + \frac{s_{1t}}{s_{0t}}\right)$ and $\Delta_{m1t} = \Delta_{B1t} = \frac{1}{\alpha}\left(1 + \frac{s_{1t}}{1-s_{1t}}\right)$, where $s_{0t} = 1 - s_{1t} - s_{2t}$. It's useful to rewrite these as

$$\Delta_{01t} = \frac{1}{\alpha}\left(1 + \exp(x_{1t}\beta - \alpha p_{1t})\right) \qquad \text{and} \qquad \Delta_{m1t} = \frac{1}{\alpha}\left(1 + \frac{\exp(x_{1t}\beta - \alpha p_{1t})}{1 + \exp(x_{2t}\beta - \alpha p_{2t})}\right)$$

which makes it clear that an increase in p_{2t} (following an increase in c_{2t}) will increase Δ_{m1t} but not Δ_{01t}. Thus, the left-hand side of Equation (8) changes with z_{1t} while the right-hand side does not, so equality cannot be maintained for different values of z_{1t}, falsifying the wrong model. •

Falsification becomes more difficult when $\bar{c}_{0j}$ and $\bar{c}_{mj}$ are allowed to depend on the quantity produced: since q_{1t} changes in response to a change in the instrument, the change in the left-hand side of Equation (8) can potentially be matched by a change in the right-hand side due to the differential dependence of $\bar{c}_{mj}$ and $\bar{c}_{0j}$ on q_{jt}. In fact, with a single instrument, falsification will typically be impossible, as a cost function $\bar{c}_{mj}$ for a misspecified model $\Delta_m \neq \Delta_0$ can typically be constructed to satisfy Equation (9), as stated in Remark 5 for the log marginal cost case. This is further illustrated in the next example:

Example 2: Suppose that true marginal cost depends linearly on the quantity produced, $\bar{c}_{0j} = \mathrm{w}_{jt}\tau + q_{jt}\gamma$,[29] and consider again using a scalar rival cost shifter instrument $z_{1t} = \mathrm{w}_{2t}$ to falsify Bertrand competition when the true model is Cournot. The lack of falsification is easiest to see in the case where w_{2t} is the only variation across markets, i.e., unobservables and w_{1t} are held fixed. With only w_{2t} varying, there is a one-to-one mapping from the realization of w_{2t} to the realizations of both prices p_t and market shares q_t,[30] so whatever variation in c_{1t} is needed to rationalize the incorrect model m can be attributed to the dependence of $\bar{c}_{m1}$ on q_{1t}; letting $w_2(s)$ denote the value of w_{2t} which would lead to market share $q_{1t} = s$, and suppressing the dependence on w_{1t} since it is fixed, the cost function

$$\bar{c}_{m1}(q_{1t}) \quad = \quad \bar{c}_{01}(q_{1t}) + E(\Delta_{01t} - \Delta_{m1t}|\mathrm{w}_{2t} = w_2(q_{1t}))$$

along with

$$\bar{c}_{m2}(\mathrm{w}_{2t}, q_{2t}) \quad = \quad \bar{c}_{02}(\mathrm{w}_{2t}, q_{2t}) + E(\Delta_{02t} - \Delta_{m2t}|\mathrm{w}_{2t})$$

mechanically satisfies Equation (8), and therefore falsification fails (Lemma 1). (In the absence of variation in unobservables or w_{1t}, the difference in markups $\Delta_{0jt} - \Delta_{mjt}$ is a deterministic function of w_{2t}, and the expectations on the right-hand side are degenerate, but are written this way for consistency.)

With variation in both firms' cost shifters and unobservables, things get a little less transparent, as $\bar{c}_{mj}$ can only depend on w_{jt} and q_{jt}, and must match its "target" in expectation over everything else. However, the cost functions

$$\bar{c}_{mj}(\mathrm{w}_{jt}, q_{jt}) \quad = \quad \bar{c}_{0j}(\mathrm{w}_{jt}, q_{jt}) + E(\Delta_{0jt} - \Delta_{mjt}|\mathrm{w}_{jt}, q_{jt})$$

where the expectation on the right is now being taken over the distribution (conditional on the value of w_{jt}) of combinations of w_{-jt} and unobservables which would lead to the observed market share q_{jt}, will again satisfy the condition in Lemma 1 for falsification to fail. •

Now, the cost function $\bar{c}_{mj}$ that rationalizes the incorrect model might be ruled out in some other way: for example, it might exhibit diseconomies of scale, when economies of scale were expected. So falsification with a single instrument may still be possible in some instances; but it would be falsification of the combination of a conduct model and an additional assumption (beyond Assumption 2) about cost structure, not just the conduct model.

[29] While marginal cost is linear in own quantity, the researcher does not know this ex ante, and failure of linearity of a candidate cost function $\bar{c}_{mj}$ would not falsify the corresponding model m.

[30] This is particularly straightforward since passthrough under the true model (Cournot) is diagonal, and we can calculate $\frac{\mathrm{d}p_t}{\mathrm{d}z_{1t}} = \begin{bmatrix} 0 \\ \tau \frac{s_{0t}}{1-s_{1t}} \end{bmatrix}$, and $\frac{\mathrm{d}q_t}{\mathrm{d}z_{1t}} = \tau \frac{s_{0t}}{1-s_{1t}} \begin{bmatrix} \alpha s_{1t} s_{2t} \\ -\alpha s_{2t}(1-s_{2t}) \end{bmatrix}$ since $\frac{\partial \bar{c}_{2t}}{\partial \mathrm{w}_{2t}} = \tau$.

A.3 Recasting in terms of Marginal Effects

As in Dearing et al. (2024), to better understand the economic content of the falsifiable restriction and the determinants of falsification, it's useful to restate Lemma 1 in terms of the marginal impacts of the instruments on markups. This depends on an assumption that markups vary continuously in the instruments:[31]

Assumption 3. *(Continuous Markups)* For any model m under consideration, including the true one $m = 0$, $E[\Delta_{mjt} \mid \mathrm{w}_{jt}, z_{jt}]$ is absolutely continuous in z_{jt} for every j.

For Equation (8) to hold at different values of z_{jt}, a change in z_{jt} must have the same marginal impact on the two sides of the equation, leading to the following:

Proposition 1. *Suppose that Assumptions 1, 2, and 3 hold. Then, model m is falsified by the instruments z_{jt} if and only if for some j there exists no function $\bar{c}_{mj}$ such that for all k*

$$E\left[\frac{\mathrm{d}\Delta_{0jt}}{\mathrm{d}z_{jt}^{(k)}} - \frac{\mathrm{d}\Delta_{mjt}}{\mathrm{d}z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right] = E\left[\left(\frac{\partial \bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt})}{\partial q_{jt}} - \frac{\partial \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt})}{\partial q_{jt}}\right)\frac{\mathrm{d}q_{jt}}{\mathrm{d}z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right] \quad a.s. \tag{9}$$

As noted above, since the expectations in Equation (8) are conditional on a realization of w_{jt}, the instruments can only change the right-hand side through q_{jt}, hence the right-hand side of Equation (9). In the case of constant marginal costs, the right-hand side of Equation (9) is zero, and falsification depends on whether the left hand side is zero. As for the left-hand side, instruments can affect Δ_{0jt} and Δ_{mjt} in two ways. Markups are typically expressed as functions of prices, market shares, and price elasticities; depending on the instrument, it may have a direct effect on markups (a product characteristic affecting market share even at fixed prices), and an additional effect through its effect on prices. Defining P_{0t} and P_{mt} as the cost pass-through matrices of the two models, we can explicitly decompose the left-hand side of Equation (9) into these two effects, as

$$E\left[\frac{\partial \Delta_{0jt}}{\partial z_{jt}^{(k)}} - \frac{\partial \Delta_{mjt}}{\partial z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right] + E\left[\left(P_{mt}^{-1} - P_{0t}^{-1}\right)_j \frac{\mathrm{d}p_0}{\mathrm{d}z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right],$$

where the first term denotes the difference in the direct effects of instruments on markups, while the second denotes the difference in indirect effects through prices, which depends on the inverse pass-through matrices of the two models. Except in knife-edge cases where the two effects cancel out, then, falsification with constant marginal costs can be achieved if either the indirect or direct effect is non-zero. (In Example 1 above, there was no direct

[31] Assumption 3 holds for all models in our examples and application.

effect since costs do not enter directly into markups, but the indirect effect was nonzero, allowing falsification.) On the other hand, falsification with non-constant marginal costs will typically fail with one instrument, as in Example 2 above, since $\frac{\partial c_{mj}}{\partial q_{jt}}$ can be chosen judiciously to satisfy Equation (9).

A.4 Non-Constant Costs and Multiple Instruments

A possible solution when marginal costs depend on quantity is to use additional instruments. The next example shows why this can work:

Example 3: Suppose again that demand is logit, $\bar{c}_{0j} = \mathrm{w}_{jt}\tau + q_{jt}\gamma$, and the researcher wants to falsify the Bertrand model when the truth is Cournot. This time, the researcher chooses two instruments: a rival cost shifter, $z_{1t}^{(1)} = \mathrm{w}_{2t}^{(1)}$, and an own product characteristic which is excluded from cost, $z_{1t}^{(2)} = x_{1t}^{(1)}$. Let $\tau^{(1)}$ be the coefficient with which $\mathrm{w}_{2t}^{(1)}$ enters $\bar{c}_{02}$, and $\beta^{(1)}$ the coefficient with which $x_{1t}^{(1)}$ enters into logit demand, and assume for exposition that $\tau^{(1)}, \beta^{(1)} > 0$. From Proposition 1, we can falsify model m unless a cost function $\bar{c}_{m1}$ exists satisfying

$$E\left[\frac{\mathrm{d}\Delta_{01t}}{\mathrm{d}z_{1t}^{(k)}} - \frac{\mathrm{d}\Delta_{m1t}}{\mathrm{d}z_{1t}^{(k)}} \mid \mathrm{w}_{1t}, z_{1t}\right] = E\left[\left(\frac{\partial \bar{c}_{m1}}{\partial q_{1t}} - \frac{\partial \bar{c}_{01}}{\partial q_{1t}}\right)\frac{\mathrm{d}q_{1t}}{\mathrm{d}z_{1t}^{(k)}} \mid \mathrm{w}_{1t}, z_{1t}\right] \tag{10}$$

almost surely for both instruments. We'll show why no such cost function can exist, and therefore why the model is falsified.

First, note that $\mathrm{d}q_{1t}/dz_{1t}^{(k)} > 0$ for both instruments – under the true model (Cournot), an increase in a rival's cost, and an increase in one's own product quality, both result in a higher own market share in equilibrium. The first occurs because, under Cournot, passthrough is diagonal, so an increase in $\mathrm{w}_{2t}^{(1)}$ results in an increase in p_{2t} and no change in p_{1t}, increasing q_{1t}. The latter is more subtle, because an increase in own quality results in an increase in own price; but the equilibrium price adjustment is small enough that $x_{1t}\beta - \alpha p_{1t}$ increases with an increase in $x_{1t}^{(1)}$. With no change in p_{2t}, this then leads to an increase in q_{1t} as well.[32]

Again rewriting firm 1's markup under the two models as

$$\Delta_{01t} = \frac{1}{\alpha}\left(1 + \exp(x_{1t}\beta - \alpha p_{1t})\right) \qquad \text{and} \qquad \Delta_{m1t} = \frac{1}{\alpha}\left(1 + \frac{\exp(x_{1t}\beta - \alpha p_{1t})}{1 + \exp(x_{2t}\beta - \alpha p_{2t})}\right)$$

allows us to see the effect of each instrument on $\Delta_{01t} - \Delta_{m1t}$:

[32] To be more thorough, under Cournot, $\mathrm{d}p_t/\mathrm{dw}_{2t}^{(1)} = \tau^{(1)}\left[\begin{smallmatrix} 0 \\ s_{0t}/(1-s_{1t}) \end{smallmatrix}\right]$ and $\mathrm{d}p_t/\mathrm{dx}_{1t}^{(1)} = \frac{\beta^{(1)}}{\alpha}\left[\begin{smallmatrix} s_{1t}/(1-s_{2t}) \\ 0 \end{smallmatrix}\right]$. From the latter, $\mathrm{d}(x_{1t}\beta - \alpha p_{1t})/\mathrm{d}x_{1t}^{(1)} = \beta^{(1)} - \alpha\frac{\beta^{(1)}}{\alpha}s_{1t}/(s_{0t} + s_{1t}) > 0$.

- An increase in $w_{2t}^{(1)}$, by increasing p_{2t} without changing p_{1t}, does not change Δ_{01t}, but increases Δ_{m1t}. Thus, for the first instrument, the left-hand side of Equation (10) is negative.
- An increase in $x_{1t}^{(1)}$, by increasing $x_{1t}\beta - \alpha p_{1t}$ without changing $x_{2t}\beta - \alpha p_{2t}$, increases both Δ_{01t} and Δ_{m1t}, but it increases Δ_{01t} *by more*; so for the second instrument, the left-hand side of equation (10) is *positive*.

This makes it clear why the two instruments together suffice to falsify the wrong model: since $\frac{\mathrm{d}q_{1t}}{\mathrm{d}z_{1t}^{(k)}}$ is positive for both instruments and (by assumption) $\frac{\partial \bar{c}_{01}}{\partial q_{1t}} = \gamma$, for each realization of observables, the expected value of $\frac{\partial \bar{c}_{m1}}{\partial q_{1t}}$ would need to be less than γ to satisfy the exclusion restriction for the first instrument, but more than γ to satisfy the exclusion restriction for the second instrument.

It's useful to return to the economics of the situation for further intuition. An increase in firm 2's cost under the true model leads to an increase in p_{2t} but no change in p_{1t}; because rival's cost passthrough is positive under Bertrand, however, the Bertrand model "sees" the lack of a change in p_{1t} as evidence of a partly offsetting decrease in c_{1t}, expressed as an increase in firm 1's markup. Since firm 1's cost shifter did not change, this decrease in marginal costs must be attributed to the increased level of output, requiring $\frac{\partial \bar{c}_{m1}}{\partial q_{1t}}$ to be smaller than its true value. On the other hand, when firm 1's quality increases and price adjusts, this leads to a smaller change in the measured markup under Bertrand than under Cournot, or an *increase* in imputed marginal cost under Bertrand relative to Cournot; since (again) w_{1t} did not change, this greater marginal cost increase must now be attributed to the increase in output, requiring $\frac{\partial \bar{c}_{m1}}{\partial q_{1t}}$ to this time be *larger* than its true value. Since $\frac{\partial \bar{c}_{m1}}{\partial q_{1t}}$ can't simultaneously be less than and greater than γ, the model is falsified. •

This example makes falsification particularly transparent, because the left-hand side of Equation (10) is positive for one instrument and negative for the other, while $\frac{\mathrm{d}q_{1t}}{dz_{1t}^{(k)}}$ has the same sign for both instruments. Things won't always be this clean. Still, this illustrates the more general point that when two instruments effect the environment through economically distinct channels, there is no reason they should change the difference in measured markups between two models identically; and if they don't (relative to how they shift own market share), then the requirements for a cost function to fit the observed changes will be different for the two instruments, and no single cost function will be able to "satisfy" both, allowing the model to be falsified.

A.5 Redundant vs "Economically Distinct" Instruments

A key to this working, however, is that the two chosen instruments have economically distinct impacts on market outcomes. Multiple instruments won't allow for falsification

when they provide essentially the same information.

For example, suppose that w_{jt} includes two separate cost shifters, so that $\bar{c}_{0jt} = \mathrm{w}_{jt}^{(1)}\tau^{(1)} + \mathrm{w}_{jt}^{(2)}\tau^{(2)} + q_{jt}\gamma_0$. As before, the researcher wants to falsify the Bertrand model when the truth is Cournot. The researcher constructs two instruments: $z_{1t}^{(1)} = \mathrm{w}_{2t}^{(1)}$ and $z_{1t}^{(2)} = \mathrm{w}_{2t}^{(2)}$. However, the two instruments affect Δ_{01t}, Δ_{m1t}, and q_{1t} identically, just scaled in proportion to the corresponding coefficients $\tau^{(1)}$ and $\tau^{(2)}$. This implies that the same misspecified cost function $\bar{c}_{m1}$ that satisfies Equation (9) for $z_{jt}^{(1)}$, will also satisfy it for $z_{jt}^{(2)}$, making falsification impossible.[33]

We can formalize this insight somewhat by noting that in any setting where falsification is impossible with a single instrument (the typical case with economoies of scale), falsification remains impossible with $K > 1$ instruments if there exist a set of constants $\{\zeta_{jk}\}_{k>1}$ such that $\frac{\mathrm{d}q_{jt}}{\mathrm{d}z_{jt}^{(k)}} = \zeta_{jk}\frac{\mathrm{d}q_{jt}}{\mathrm{d}z_{jt}^{(1)}}$ and $\left(\frac{\mathrm{d}\Delta_{0jt}}{\mathrm{d}z_{jt}^{(k)}} - \frac{\mathrm{d}\Delta_{mjt}}{\mathrm{d}z_{jt}^{(k)}}\right) = \zeta_{jk}\left(\frac{\mathrm{d}\Delta_{0jt}}{\mathrm{d}z_{jt}^{(1)}} - \frac{\mathrm{d}\Delta_{mjt}}{\mathrm{d}z_{jt}^{(1)}}\right)$ for each $k > 1$. That is, additional instruments bring no additional power to falsify a model if their economic effects are simply rescalings of those of existing instruments. Notice that this is analogous to the condition (5) in Remark 6 in the main body of the paper.

This may sound so obvious that it doesn't bear discussion, but redundancy of instruments in this way can occasionally crop up in unexpected places. In our well-trodden example environment of simple logit demand, for example, instruments formed from a rival's cost shifter and a product characteristic of the same rival turn out to be redundant! That is, when we allow for economies of scale, falsifying the Bertrand model when true firm conduct is Cournot is impossible using the two instruments $z_{1t}^{(1)} = \mathrm{w}_{2t}^{(1)}$ and $z_{2t}^{(2)} = x_{2t}^{(1)}$. The failure of falsification, however, is due to a rather knife-edge fact about the particular demand system. In simple logit, demand depends on product characteristics and prices only through a single index $\delta_{jt} = x_{jt}\beta - \alpha p_{jt} + \xi_{jt}$ for each product. Given that fact, for most standard conduct models, equilibrium markups depend on product characteristics and marginal costs only through the terms $\left\{x_{jt}\beta - \alpha c_{jt}\right\}_{j=1,\ldots,J}$.[34] This means a change in a product characteristic has the same equilibrium impacts on market shares and implied markups as a change in its marginal cost, scaled by the appropriate coefficients. However, without the single index restriction, this would not be the case; in a richer demand system with random coefficients, for example, the effects of a product characteristic and a cost shifter of the same product would not simply be rescalings of each other, and we would

[33] That is, since both instruments affect outcomes only through c_{2t}, we will have $\frac{\mathrm{d}(\Delta_{01t}-\Delta_{m1t})}{\mathrm{d}z_{1t}^{(k)}} = \tau^{(k)}\frac{\mathrm{d}(\Delta_{01t}-\Delta_{m1t})}{\mathrm{d}c_{2t}}$ and $\frac{\mathrm{d}q_{1t}}{\mathrm{d}z_{1t}^{(k)}} = \tau^{(k)}\frac{\mathrm{d}q_{1t}}{\mathrm{d}c_{2t}}$ for $k = 1, 2$, so if (9) holds for $k = 1$ it will hold for $k = 2$.

[34] Appendix C of Dearing et al. (2024) explores the set of conduct models where this occurs; in brief, if each firm is maximizing profits under some assumption about how rivals will respond to their actions, their problem is $\max \sum_{j\in F}(p_{jt} - c_{jt})s_{jt}(p_t)$, we can think instead of the firms choosing markups to solve $\max \sum_{j\in F}\Delta_{jt}s_{jt}(\delta_t)$, and write $\delta_{jt} = x_{jt}\beta - \alpha p_{jt} + \xi_{jt}$ instead as $\delta_{jt} = \left(x_{jt}\beta - \alpha c_{jt} + \xi_{jt}\right) - \alpha\Delta_{jt}$, to make clear that the Δ_{jt} that form equilibrium can depend on x_t and c_t only through $x_t\beta - \alpha c_t + \xi_t$.

expect to be able to falsify the incorrect model of conduct.

This discussion is simply meant to highlight the fact that if multiple instruments are to allow falsification in the presence of economies of scale, it's essential that they be *economically distinct*, which we define informally as having equilibrium effects on market shares and implied markups which do not both vary (across instruments) by the same scalar multipliers. In particular, two economically distinct instruments either differentially move quantities under the true model or differentially move the difference in markups. As stated in Remark 7, it should be easy to select economically distinct instruments, meaning falsification under non-constant marginal costs is typically possible in a standard differentiated product environment. Example 3 illustrated this, pairing a characteristic of one product with a cost shifter of a rival product. We just noted that under certain demand systems such as simple logit, it is not sufficient to pair a cost shifter and product characteristic of the same rival product, but outside of the "single index" demand systems discussed above, we would expect cost side instruments and demand side instruments to typically be economically distinct. In addition, with more than two firms, instruments that shift the marginal costs of different rival products may in fact be economically distinct, particularly in demand models sufficiently rich for some products to be "closer substitutes" than others.

A.6 Proofs

Proof of Lemma 1. As we note in the text, in our parametric framework, the falsifiable restriction in Equation (28) of Berry and Haile (2014) is[35]

$$E[\omega_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] \quad = \quad E[p_{jt} - \Delta_{mjt} - \bar{c}_{mj}(\mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] \quad = \quad 0 \qquad a.s.$$

Since observed prices are generated under the true model as

$$p_{jt} \quad = \quad \Delta_{0jt} + c_{0jt} \quad = \quad \Delta_{0jt} + \bar{c}_{0j}(\mathrm{w}_{jt}) + \omega_{0jt}$$

and $E[\omega_{0jt} \mid \mathrm{w}_{jt}, z_{jt}] = 0$ under Assumption 2, the falsifiable restriction is equivalent to

$$E[\Delta_{0jt} + \bar{c}_{0j}(\mathrm{w}_{jt}) + \omega_{0jt} - \Delta_{mjt} - \bar{c}_{mj}(\mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] \quad = \quad 0 \qquad a.s.$$

or equivalently

$$E[\Delta_{0jt} - \Delta_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] \quad = \quad \bar{c}_{mj}(\mathrm{w}_{jt}) - \bar{c}_{0j}(\mathrm{w}_{jt}) \qquad a.s.$$

giving the result. □

[35] See Section 6, Case 2 in Berry and Haile (2014) for a discussion of their non-parametric environment.

Proof of Proposition 1. In our parametric framework, the falsifiable restriction in Equation (28) of Berry and Haile (2014) is that for all j there exists a cost function $\bar{c}_{mj}$ such that:

$$E[p_{jt} - \Delta_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] = E[\bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] \qquad a.s.$$

By plugging in for p_t as in the proof of Lemma 1, a model m is not falsified if for all j there exists a cost function $\bar{c}_{mj}$ such that

$$E[\Delta_{0jt} - \Delta_{mjt} \mid \mathrm{w}_{jt}, z_{jt}] = E[\bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt}) - \bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt}) \mid \mathrm{w}_{jt}, z_{jt}] \qquad a.s.$$

The result thus follows by extending Lemma 2 in Dearing et al. (2024) to this restriction, so that a model is falsified if for some j there exists no cost function $\bar{c}_{mj}$ such that for all k

$$E\left[\frac{\mathrm{d}\Delta_{0jt}}{\mathrm{d}z_{jt}^{(k)}} - \frac{\mathrm{d}\Delta_{mjt}}{\mathrm{d}z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right] = E\left[\frac{\mathrm{d}\bar{c}_{mj}(q_{jt}, \mathrm{w}_{jt})}{\mathrm{d}z_{jt}^{(k)}} - \frac{\mathrm{d}\bar{c}_{0j}(q_{jt}, \mathrm{w}_{jt})}{\mathrm{d}z_{jt}^{(k)}} \mid \mathrm{w}_{jt}, z_{jt}\right] \qquad a.s.$$

□

Appendix B Convergence of the Modified RV- and F-statistics

The derivations herein are not tied to the log-log specification of Equation (3) and apply for arbitrary, known transformations of the model-implied costs and relevant level of quantities. Hence, we adopt the notation $\mathfrak{c}_{mjt} = f(p_{jt} - \Delta_{mjt})$ and $\mathfrak{q}_{jt} = g(q_{jt}^p)$ for known functions f and g and extend the definitions of Section 4 from the log-transformed model considered there to these arbitrary, known transforms.

Test statistics We first give explicit definitions of the test statistics T^{RV} and F utilized in the paper. With data on n observations, the in-sample analogue of our measure of lack-of-fit is given as $\hat{Q}_m = \hat{g}_m' \hat{W} \hat{g}_m$ where $\hat{W} = (\frac{1}{n}\sum_{j,t} \hat{z}_{jt}^e \hat{z}_{jt}^{e\prime})^+$, $\hat{g}_m = \frac{1}{n}\sum_{j,t} \hat{z}_{jt}^e \hat{\omega}_{mjt}$, $\hat{z}_{jt}^e$ is the residual in a projection on $\hat{\tilde{q}}$ and w,

$$\hat{z}_{jt}^e = z_{jt} - \hat{\lambda}_q \hat{\tilde{q}} - \hat{\Lambda}_{\mathrm{w}} \mathrm{w}_{jt} \qquad \text{for} \quad \begin{bmatrix} \hat{\lambda}_q' \\ \hat{\Lambda}_{\mathrm{w}}' \end{bmatrix} = \left[\left(\hat{\tilde{q}}, \mathrm{w}\right)' \left(\hat{\tilde{q}}, \mathrm{w}\right)\right]^{-1} \left(\hat{\tilde{q}}, \mathrm{w}\right)' z.$$

the estimated cost error $\hat{\omega}_{mjt}$ is the residual in a sample 2SLS projection,

$$\hat{\omega}_{mjt} = \log\left(p_{jt} - \Delta_{mjt}\right) - \hat{\gamma}_m \log(q_{jt}^p) - \mathrm{w}_{jt}' \hat{\tau}_m$$

with

$$\begin{pmatrix}\hat{\gamma}_m\\ \hat{\tau}_m\end{pmatrix} = \left[\left(\hat{\tilde{q}}, \mathrm{w}\right)' \left(\log(q^p), \mathrm{w}\right)\right]^{-1} \left(\hat{\tilde{q}}, \mathrm{w}\right)' \log\left(p - \Delta_m\right),$$

where $\hat{\tilde{q}}_{jt}$ is an in-sample prediction, $\hat{\tilde{q}}_{jt} = z_{jt}'\hat{\zeta}_z + \mathrm{w}_{jt}'\hat{\zeta}_\mathrm{w}$, for $(\hat{\zeta}_z', \hat{\zeta}_\mathrm{w}')' = \left[(z, \mathrm{w})'(z, \mathrm{w})\right]^{-1}(z, \mathrm{w})'\mathfrak{q}$.

The standard error used in the RV test statistic is

$$\hat{\sigma}^2_{\mathrm{RV}} = 4\Big[\hat{g}_1'\hat{W}^{1/2}\hat{V}^{\mathrm{RV}}_{11}\hat{W}^{1/2}\hat{g}_1 + \hat{g}_2'\hat{W}^{1/2}\hat{V}^{\mathrm{RV}}_{22}\hat{W}^{1/2}\hat{g}_2 - 2\hat{g}_1'\hat{W}^{1/2}\hat{V}^{\mathrm{RV}}_{12}\hat{W}^{1/2}\hat{g}_2\Big]$$

where $\hat{V}^{\mathrm{RV}}_{\ell k}$ is an estimator of the covariance between $\sqrt{n}\hat{W}^{1/2}\hat{g}_\ell$ and $\sqrt{n}\hat{W}^{1/2}\hat{g}_k$. Our proposed $\hat{V}^{\mathrm{RV}}_{\ell k}$ is given by $\hat{V}^{\mathrm{RV}}_{\ell k} = n^{-1}\sum_{j,t}\hat{\psi}_{\ell jt}\hat{\psi}_{kjt}'$ where

$$\begin{aligned}\hat{\psi}_{mjt} = {}& \hat{W}^{1/2}\big(\hat{z}^e_{jt}\hat{\omega}_{mjt} - \hat{g}_m\big) - \tfrac{1}{2}\hat{W}^{3/4}\Big(\hat{z}^e_{jt}\hat{z}^{e\prime}_{jt} - \hat{W}^{+}\Big)\hat{W}^{3/4}\hat{g}_m\\ & + \tfrac{1}{2}\hat{W}^{3/4}\Big(\hat{W}^{+}\hat{Z}\hat{z}^r_{jt}\hat{q}^e_{jt}\hat{\lambda}_q' + \hat{\lambda}_q\hat{z}^{r\prime}_{jt}\hat{q}^e_{jt}\hat{Z}\hat{W}^{+}\Big)\hat{W}^{3/4}\hat{g}_m.\end{aligned}$$

The second line in the definition of this influence function captures the added variability from estimation of the best linear predictor $\tilde{q}$ and relies on the definitions $\hat{q}^e_{jt} = \mathfrak{q}_{jt} - \hat{\tilde{q}}_{jt}$, $\hat{z}^r_{jt} = z_{jt} - w_{jt}'(w'w)^{-1}w'z$, and $\hat{Z} = (\frac{1}{n}\sum_{j,t}\hat{z}^r_{jt}\hat{z}^{r\prime}_{jt})^{-1}$. This variance estimator is transparent and easy to implement. Adjustments to $\hat{\psi}_{mjt}$ and/or $\hat{V}^{\mathrm{RV}}_{\ell k}$ can also accommodate initial demand estimation and clustering.

The F-statistic in the current context is a joint test statistic for the two hypotheses: H^{AR}_{0m} : $\pi_m = 0$ in the equations $\omega_{mjt} = z^{e\prime}_{jt}\pi_m + e_{mjt}$ for $m = 1, 2$. Formulaically, it is

$$F = \big(1 - \hat{\rho}^2\big)\frac{n}{2(d_z - 1)}\frac{\hat{\sigma}_2^2\hat{g}_1'\hat{W}\hat{g}_1 + \hat{\sigma}_1^2\hat{g}_2'\hat{W}\hat{g}_2 - 2\hat{\sigma}_{12}\hat{g}_1'\hat{W}\hat{g}_2}{\hat{\sigma}_1^2\hat{\sigma}_2^2 - \hat{\sigma}_{12}^2},$$

where

$$\hat{\rho}^2 = \frac{\big(\hat{\sigma}_1^2 - \hat{\sigma}_2^2\big)^2}{\big(\hat{\sigma}_1^2 + \hat{\sigma}_2^2\big)^2 - 4\hat{\sigma}_{12}^2},\ \hat{\sigma}_m^2 = \frac{\mathrm{trace}\big(\hat{V}^{\mathrm{AR}}_{mm}\hat{W}^{+}\big)}{d_z - 1},\ \hat{\sigma}_{12} = \frac{\mathrm{trace}\big(\hat{V}^{\mathrm{AR}}_{12}\hat{W}^{+}\big)}{d_z - 1},$$

$\hat{V}^{\mathrm{AR}}_{\ell k} = n^{-1}\sum_{j,t}\hat{\phi}_{\ell jt}\hat{\phi}_{kjt}'$ where $\hat{\phi}_{mjt} = \hat{W}\hat{z}^e_{jt}\big(\hat{\omega}_{mjt} - \hat{z}^{e\prime}_{jt}\hat{\pi}_m\big)$, and $\hat{\pi}_m = \hat{W}\hat{g}_m$ is a (reduced rank) OLS estimator of π_m.

Assumptions The following assumptions list direct analogues of Assumptions 1–4 in Duarte et al. (2024). The only key new requirement is in Assumption 5, part (ii), which ensures that the instruments are strong for pinning down the economies of scale. We refer to Duarte et al. (2024) for further discussion.

Assumption 4. z_{jt} is a vector of d_z excluded instruments, so that $E[z_{jt}\omega_{0jt}] = 0$.

Assumption 5. (i) $\{p_{jt}, \Delta_{0jt}, \Delta_{1jt}, \Delta_{2jt}, z_{jt}, \mathrm{w}_{jt}, q^p_{jt}, \omega_{0jt}\}_{j,t}$ are jointly iid with the observables being $p_{jt}, \Delta_{1jt}, \Delta_{2jt}, z_{jt}, \mathrm{w}_{jt}, q^p_{jt}$; (ii) $E\left[(\Delta_{1jt}-\Delta_{2jt})^2\right]$ is positive, $E\left[(z'_{jt}, \mathrm{w}'_{jt})'(z'_{jt}, \mathrm{w}'_{jt})\right]$ is positive definite, and $E\left[(z'_{jt}, \mathrm{w}'_{jt})'(\mathfrak{q}_{jt}, \mathrm{w}'_{jt})\right]$ has full rank; (iii) the entries of $\mathfrak{c}_{1jt}$, $\mathfrak{c}_{2jt}$, z_{jt}, w_{jt}, $\mathfrak{q}_{jt}$, ω_{1jt}, and ω_{2jt} have finite fourth moments.

Assumption 6. The error term in Equation (4), e_{mjt}, is homoskedastic, i.e., $E[e^2_{mjt} z^e_{jt} z^{e\prime}_{jt}] = \sigma^2_m E[z^e_{jt} z^{e\prime}_{jt}]$ with $\sigma^2_m > 0$ for $m \in \{1, 2\}$ and $E[e_{1jt} e_{2jt} z^e_{jt} z^{e\prime}_{jt}] = \sigma_{12} E[z^e_{jt} z^{e\prime}_{jt}]$ with $\sigma^2_{12} < \sigma^2_1 \sigma^2_2$.

Assumption 7. Define $\pi_m = W g_m$. For both $m = 1$ and $m = 2$,

$$\pi_m = q_m/\sqrt{n} \qquad \text{for some finite vector } q_m.$$

Establishing Remark 4 To establish the guidance provided in Remark 4, we prove the following result, which is a direct analog of Proposition 4 in Duarte et al. (2024) with the only difference being that $d_z - 1$ replaces d_z:

Proposition 2. *Suppose Assumptions 4–7 hold. Then*

$$(i) \qquad \begin{pmatrix} |T^{\mathrm{RV}}| \\ F \end{pmatrix} \xrightarrow{d} \begin{pmatrix} |\Psi'_- \Psi_+| / \left(\|\Psi_-\|^2 + \|\Psi_+\|^2 + 2\rho \Psi'_- \Psi_+ \right)^{1/2} \\ \left(\|\Psi_-\|^2 + \|\Psi_+\|^2 - 2\rho \Psi'_- \Psi_+ \right) / (2d_z) \end{pmatrix}$$

where $\hat{\rho}^2 \xrightarrow{p} \rho^2$ *and*

$$\begin{pmatrix} \Psi_- \\ \Psi_+ \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_- \mathbf{e}_1 \\ \mu_+ \mathbf{e}_1 \end{pmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \otimes I_{d_z - 1} \right),$$

$$(ii) \qquad H_0^{\mathrm{RV}} \text{ holds if and only if } \mu_- = 0,$$

$$(iii) \qquad H_{0,1}^{\mathrm{AR}} \text{ and } H_{0,2}^{\mathrm{AR}} \text{ holds if and only if } \mu_+ = 0,$$

$$(iv) \qquad 0 \le \mu_- \le \mu_+.$$

To establish Proposition 2, we establish the following lemma, which generalizes Lemma A.1 in Duarte et al. (2024) to the setting of this paper. Given this lemma, Proposition 2 follows by the use of the same proof as for Proposition 4 in Duarte et al. (2024).

To state the following lemma and give a formulation of σ^2_{RV}, we introduce population versions of $\hat{\psi}_{mjt}$ and $\hat{\phi}_{mjt}$ along with notation for their variances. Let

$$\begin{aligned} \psi_{mjt} = {} & W^{1/2} z^e_{jt} \omega_{mjt} - \tfrac{1}{2} W^{3/4} z^e_{jt} z^{e\prime}_{jt} W^{3/4} g_m - \tfrac{1}{2} W^{1/2} g_m \\ & + \tfrac{1}{2} W^{3/4} \left(W^+ Z z^r_{jt} q^e_{jt} \lambda'_q + \lambda_q z'_{jt} q^e_{jt} Z W^+ \right) W^{3/4} g_m. \end{aligned}$$

and $\phi_{mjt} = W z^e_{jt} e_{mjt}$ where $q^e_{jt} = \mathfrak{q}_{jt} - \tilde{q}_{jt}$, $z^r_{jt} = z_{jt} - w'_{jt} E[w'w]^{-1} E[w'z]$, $Z = E[z^r_{jt} z^{r\prime}_{jt}]^{-1}$,

and

$$\begin{bmatrix} \lambda_q' \\ \Lambda_{\mathrm{w}}' \end{bmatrix} = E\big[(\tilde{q}, \mathrm{w})'(\tilde{q}, \mathrm{w})\big]^{-1} E\big[(\tilde{q}, \mathrm{w})' z\big].$$

Also, let $V_{\ell k}^{\mathrm{RV}} = E[\psi_{\ell jt}\psi_{kjt}']$, $V_{\ell k}^{\mathrm{AR}} = E[\phi_{\ell jt}\phi_{kjt}']$, and $V^{\mathrm{RV}} = E[(\psi_{1jt}', \psi_{2jt}')'(\psi_{1jt}', \psi_{2jt}')]$, which is a matrix with V_{11}^{RV}, V_{12}^{RV}, and V_{22}^{RV} as its entries. Finally,

$$\sigma_{\mathrm{RV}}^2 = 4\Big[g_1' W^{1/2} V_{11}^{\mathrm{RV}} W^{1/2} g_1 + g_2' W^{1/2} V_{22}^{\mathrm{RV}} W^{1/2} g_2 - 2 g_1' W^{1/2} V_{12}^{\mathrm{RV}} W^{1/2} g_2\Big].$$

Lemma B.1. *Suppose Assumptions 4 and 5 hold. For $\ell, k \in \{1, 2\}$, we have*

$$(i)\quad \sqrt{n}\begin{pmatrix} \hat{W}^{1/2}\hat{g}_1 - W^{1/2} g_1 \\ \hat{W}^{1/2}\hat{g}_2 - W^{1/2} g_2 \end{pmatrix} \xrightarrow{d} N\big(0, V^{\mathrm{RV}}\big), \qquad (ii)\quad \hat{V}_{\ell k}^{\mathrm{RV}} \xrightarrow{p} V_{\ell k}^{\mathrm{RV}},$$

$$(iii)\quad \sqrt{n}\,(\hat{\pi}_m - \pi_m) \xrightarrow{d} N\big(0, V_m^{\mathrm{AR}}\big), \qquad (iv)\quad \hat{V}_m^{\mathrm{AR}} \xrightarrow{p} V_m^{\mathrm{AR}}.$$

Remark 8. When the RV test is not degenerate, i.e., when $\sigma_{\mathrm{RV}}^2 > 0$, it follows from parts (i), (ii), and a first order Taylor approximation that $T^{\mathrm{RV}} \xrightarrow{d} N(0,1)$ under H_0^{RV} so that the RV test is asymptotically valid. Details of this step can be found in Rivers and Vuong (2002); Hall and Pelletier (2011) and are omitted. When non-degeneracy fails to hold, a first-order Taylor approximation does not capture the behavior of T^{RV}.

Proof of Lemma B.1. To prove their analog of this lemma, Duarte et al. (2024) proceeds in three steps to establish (i) and (ii) before commenting on the minor modifications in the argument that establish (iii) and (iv). Here we comment on the modifications needed due to the introduction of economies of scale. Step (1) shows that $\frac{1}{\sqrt{n}}\sum_{j,t}(\psi_{1jt}', \psi_{2jt}')' \xrightarrow{d} N(0, V^{\mathrm{RV}})$ and $\check{V}_{\ell k}^{\mathrm{RV}} := \frac{1}{n}\sum_{j,t}\psi_{\ell jt}\psi_{kjt}' \xrightarrow{p} V_{\ell k}^{\mathrm{RV}}$ for $\ell, k \in \{1, 2\}$, and this step require no new arguments. Step (2) establishes that $\sqrt{n}\big(\hat{W}^{1/2}\hat{g}_m - W^{1/2} g_m\big) - \frac{1}{\sqrt{n}}\sum_{j,t}\psi_{mjt} = o_p(1)$ for $m \in \{1, 2\}$, and step (3) proofs that $\mathrm{trace}\big((\hat{V}_{\ell k}^{\mathrm{RV}} - \check{V}_{\ell k}^{\mathrm{RV}})'(\hat{V}_{\ell k}^{\mathrm{RV}} - \check{V}_{\ell k}^{\mathrm{RV}})\big) = o_p(1)$ for $\ell, k \in \{1, 2\}$. In the approximations of the last two steps, the only new complication to handle is the estimation of the best linear predictor $\tilde{q}$ and the economies of scale parameter γ_m.

As in Remark 2 of Duarte et al. (2024), we have that $g_m = n^{-1}\hat{z}^{e\prime}\hat{\omega}_m = n^{-1} z^{e\prime}\hat{\omega}_m$ leading to the approximation

$$n^{-1}(\hat{z}^{e\prime}\hat{\omega}_m - z^{e\prime}\omega_m) = n^{-1} z^{e\prime}(\hat{\omega}_m - \omega_m) = \underbrace{n^{-1} z^{e\prime}(\mathfrak{q}, \mathrm{w})}_{=O_p(n^{-1/2})} \underbrace{\begin{pmatrix} \gamma_m - \hat{\gamma}_m \\ \tau_m - \hat{\tau}_m \end{pmatrix}}_{=O_p(n^{-1/2})} = O_p(n^{-1}).$$

For $n^{-1}\hat{z}^{e\prime}\hat{z}^{e}$, we instead have

$$n^{-1}(\hat{z}^{e\prime}\hat{z}^{e} - z^{e\prime}z^{e}) = n^{-1}z^{e\prime}(\hat{z}^{e} - z^{e}) + n^{-1}(\hat{z}^{e} - z^{e})'z^{e} + n^{-1}(\hat{z}^{e} - z^{e})'(\hat{z}^{e} - z^{e}).$$

Here we can write $\hat{z}^{e} - z^{e}$ as

$$\hat{z}^{e} - z^{e} = (\tilde{q}, \mathrm{w})\begin{pmatrix}\lambda_q' \\ \Lambda_{\mathrm{w}}'\end{pmatrix} - (\hat{\tilde{q}}, \mathrm{w})\begin{pmatrix}\hat{\lambda}_q' \\ \hat{\Lambda}_{\mathrm{w}}'\end{pmatrix} = (\tilde{q}, \mathrm{w})\begin{pmatrix}\lambda_q' - \hat{\lambda}_q' \\ \Lambda_{\mathrm{w}}' - \hat{\Lambda}_{\mathrm{w}}'\end{pmatrix} - \big((\hat{\tilde{q}} - \tilde{q})\hat{\lambda}_q', \mathbf{0}\big)$$

where in turn $\hat{\tilde{q}} - \tilde{q} = z(\hat{\zeta}_z - \zeta_z) + \mathrm{w}(\hat{\zeta}_{\mathrm{w}} - \zeta_{\mathrm{w}})$. Since $n^{-1}z^{e\prime}(\tilde{q}, w) = O_p(n^{-1/2})$ and $n^{-1}z^{e\prime}z = W^{+} + O_p(n^{-1/2})$, standard arguments (among which a key one is $\hat{\zeta}_z - \zeta_z = n^{-1}Zz^{r\prime}q^{e} + O_p(n^{-1})$) imply that

$$n^{-1}\hat{z}^{e\prime}\hat{z}^{e} = n^{-1}z^{e\prime}z^{e} - n^{-1}W^{+}Zz^{r\prime}q^{e}\lambda_q' - n^{-1}\lambda_q q^{e\prime}z^{r}ZW^{+} + O_p(n^{-1}).$$

Thus, it follows by the same steps as in the proof of the analog lemma in Duarte et al. (2024), that

$$\hat{W}^{1/2}\hat{g}_m - W^{1/2}g_m = \tfrac{1}{n}\sum\nolimits_{j,t}\psi_{mjt} + O_p\big(n^{-1}\big).$$

Tracing the proof in Duarte et al. (2024), the key step is to establish that $n^{-1}\sum_{j,t}\|\hat{\psi}_{mjt} - \psi_{mjt}\|^2 = o_p(1)$, which follows by arguments used in the first half of this proof. □

Appendix C Monte Carlo Simulations

We now illustrate the performance of the procedure developed in Section 4 of the paper through Monte Carlo simulations.

Setup: We simulate data for 100,000 markets using `PyBLP` (Conlon and Gortmaker, 2020). In each market, two single-product firms compete. Consumer utility follows a logit specification with a price coefficient α and three product characteristics (including a constant), with coefficients $\beta = [1, 2, 1]$. For each product-market pair, we draw two characteristics independently from a uniform distribution $U[0, 3]$.

On the supply side, marginal costs include two observed shifters (also drawn from $U[0, 3]$) and a constant, with coefficients $\tau = [3, 0.5, 1.5]$. When present, economies of scale enter through a parameter $\gamma = -0.5$ that multiplies quantity. The unobserved demand (ξ_{jt}) and cost (ω_{jt}) shocks follow a bivariate normal distribution with unit variances and correlation 0.9, which is the default in `PyBLP`. We select demand and cost parameters to generate

reasonable elasticities and outside good shares. In all simulations, the true model is Bertrand competition, while we attempt to falsify Cournot competition.

Results: Table 10 reports results from six simulation experiments that illustrate our theoretical findings on instrument relevance. For each experiment, we report both the RV test statistic T_{RV} and the F-statistic that diagnoses instrument strength, computed with the package `pyRVtest` (Duarte et al., 2022).

TABLE 10: Illustrating Instrument Relevance

Experiments	(1)	(2)	(3)	(4)	(5)	(6)
DGP and Instruments						
Econ. of Scale (γ)	0	-0.5	-0.5	-0.5	-0.5	-0.5
Instrument	$w_{-j}^{(1)}$	$w_{-j}^{(1)}$	$w_{-j}^{(1)}, w_{-j}^{(2)}$	$w_{-j}^{(1)}, x_j^{(1)}$	$w_{-j}^{(1)}, x_{-j}^{(1)}$	$x_j^{(1)}, x_{-j}^{(1)}$
Panel A: Bertrand vs. Cournot						
T^{RV}	-2.5	-0.0	-0.4	-3.2	-0.2	-5.0
	$***$			$***$		$***$
F	2,894.2	0.0	0.2	1,443.6	0.0	4,189.4
	$\dagger\dagger\dagger$ $\wedge\wedge\wedge$			$\dagger\dagger\dagger$ $\wedge\wedge\wedge$		$\dagger\dagger\dagger$ $\wedge\wedge\wedge$
Panel B: Model-implied Estimated Economies of Scale $(\hat{\gamma}_B, \hat{\gamma}_C)$						
$\hat{\gamma}_B$	–	-0.6	-0.8	-0.5	-0.7	-0.5
	–	(0.46)	(0.32)	(0.02)	(0.23)	(0.02)
$\hat{\gamma}_C$	–	2.3	2.3	-0.7	2.2	-0.7
	–	(0.49)	(0.33)	(0.02)	(0.24)	(0.02)

The table reports, for each experiment 1-6, the RV test statistics T^{RV} and effective F-statistic in panel A, and estimated economies of scale parameters in panel B. For panel A, a negative RV test statistic suggests a better fit of the true Bertrand model. The symbol $***$ indicates rejection of the null of equal fit 0.01 confidence level. The symbols $\dagger\dagger\dagger$ and $\wedge\wedge\wedge$ indicated that F is above the appropriate critical values for worst-case size below 0.075, and best-case power above 0.95, respectively. Both T^{RV} and the F-statistics account for market-level clustering.

In Simulation 1 we consider constant marginal costs ($\gamma = 0$) and use a single rival cost shifter as an instrument. With a large F-statistic and significant RV test statistic ($T_{RV} = -2.5$), the results confirm that a researcher can falsify a wrong model of conduct with a single relevant instrument when marginal costs are constant (Dearing et al., 2024). However, once we introduce economies of scale ($\gamma = -0.5$) in Simulation 2 while maintaining the single rival cost shifter instrument, both statistics drop to near zero, indicating a complete failure of falsification as predicted by Remark 5.

In line with Remark 6, adding a second rival cost shifter in Simulation 3 fails to resolve this problem. These instruments are not economically distinct, and yield an F-statistic of only 0.2, indicating they remain weak for testing. The breakthrough comes in Simulation 4, where we pair a rival cost shifter with an own product characteristic - these are economically distinct instruments and available in standard datasets, in line with Remark 7. These

instruments successfully falsify the wrong model even with non-constant costs, yielding both a significant test statistic ($T_{RV} = -3.2$) and strong instruments ($F = 1,443.6$). This finding illustrates how instruments that affect the environment through different economic channels can overcome the challenges posed by non-constant marginal costs.

The final two specifications further validate our theoretical framework by exploring alternative instrument pairs, further illustrating Remarks 6 and 7 . Simulation 5 combines a rival cost shifter with a rival product characteristic, but these instruments fail to be economically distinct and provide no power for testing. In contrast, Simulation 6 pairs own and rival product characteristics, achieving strong power as diagnosed by the effective F-statistic ($F = 4,189.4$) and clear rejection of the wrong model ($T_{RV} = -5.0$).

Panel B reports the estimated economies of scale parameters (γ) under both Bertrand and Cournot specifications across our experiments. These estimates highlight how instrument choice affects our ability to separately identify cost structure and conduct. With weak instruments (columns 2, 3, and 5), the estimates vary substantially between models. In contrast, when using economically distinct instruments (columns 4 and 6), we obtain more stable and similar estimates across specifications.

Appendix D Additional Results, Counterfactuals, and Robustness

D.1 Additional Results

Full Parameter Estimates from Marginal Cost Regressions: Table 3 reports the economies os scale parameters obtained from the regression in Equation 3 for the five models of conduct we consider in the car industry. In Table 11, we report all parameter estimates from these regressions.

TABLE 11: Full Implied Marginal Cost Regression Results

			Mixed Models		
	Bertrand	Cournot	Asian Bertrand	US Bertrand	Europe Cournot
Constant	−6.658	−7.270	−6.795	−6.563	−6.618
	(1.015)	(1.042)	(1.010)	(1.001)	(1.009)
RXR	0.047	0.050	0.046	0.046	0.047
	(0.037)	(0.037)	(0.037)	(0.037)	(0.037)
t	−0.124	−0.110	−0.115	−0.124	−0.125
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
t^2	0.002	0.002	0.002	0.002	0.002
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
log(height)	−0.884	−0.949	−0.903	−0.909	−0.882
	(0.142)	(0.146)	(0.142)	(0.142)	(0.142)
log(footprint)	0.020	0.011	0.009	0.005	0.018
	(0.152)	(0.155)	(0.151)	(0.152)	(0.152)
log(horsepower)	0.437	0.469	0.449	0.446	0.437
	(0.047)	(0.049)	(0.048)	(0.047)	(0.047)
log(MPG)	0.310	0.312	0.313	0.303	0.309
	(0.074)	(0.075)	(0.073)	(0.073)	(0.073)
log(curbweight)	1.353	1.414	1.361	1.364	1.352
	(0.121)	(0.122)	(0.119)	(0.121)	(0.121)
log(# trims)	0.016	0.012	0.013	0.014	0.016
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Release Year	−0.130	−0.130	−0.130	−0.128	−0.130
	(0.028)	(0.029)	(0.028)	(0.028)	(0.028)
SUV	0.043	0.040	0.045	0.046	0.043
	(0.024)	(0.025)	(0.024)	(0.024)	(0.024)
Truck	−0.156	−0.204	−0.179	−0.150	−0.157
	(0.038)	(0.040)	(0.038)	(0.037)	(0.038)
Van	−0.121	−0.145	−0.128	−0.114	−0.121
	(0.041)	(0.042)	(0.041)	(0.041)	(0.041)
PHEV/EV	0.216	0.229	0.222	0.224	0.216
	(0.045)	(0.045)	(0.045)	(0.045)	(0.045)
Sport	0.094	0.099	0.097	0.092	0.095
	(0.027)	(0.027)	(0.027)	(0.026)	(0.026)
Years Since Design	−0.003	−0.003	−0.003	−0.003	−0.003
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
$\log(q_{jt}^p)$	−0.119	−0.114	−0.115	−0.116	−0.118
	(0.032)	(0.033)	(0.032)	(0.032)	(0.032)

Table reports full results of estimating Equation (3) via 2SLS under different models of conduct. Table 3 reports the economies of scale parameters from these regressions. Specifications include manufacturer fixed effects. SE are clustered by car model.

Counterfactual Results under Bertrand with Constant Marginal Cost: In Table 12, we contrast the main results of our counterfactuals predicted under our preferred model of conduct (Cournot with economies of scale) against counterfactual results predicted under the standard model in the literature, Bertrand competition with constant marginal cost. The results show the importance of accounting for cost and conduct, of which we highlight three differences.

First, the trade literature has estimated pass-through effects of tariffs in the first Trump administration to be around one. Under a model of Bertrand competition with constant marginal cost, we find that prices of foreign-produced cars increase by 19.5% in C1 and 22.7% in C3, both of which correspond to a pass-through of the port cost of approximately one (21.25%). Instead, Cournot with economies of scale predicts pass-through of the port cost above one.

Second, the price effects for US models in C1 further underscores the importance of learning conduct for predicting trade policy. Our preferred model predicts a price reduction while the Bertrand model with constant marginal cost predicts a price increase. While differences in both the model of conduct and the functional form of cost give rise to discrepancies in the predicted price effects, the differences in cost play a larger role in our application.

Third, the differences in equilibrium effects give rise to differences in predicted welfare and employment effects. In C3, the Bertrand model with constant marginal costs understates the total effect on net surplus by over $6 billion. It also predicts less growth in employment arising from the stacked tariffs.

TABLE 12: Counterfactual Results under Alternative Models

	Cournot, Econ. of Scale			Bertrand, Constant MC		
	C1	C2	C3	C1	C2	C3
Price Effects (percent)						
All Models	9.05	12.96	14.56	8.03	12.29	13.49
US Models	−0.85	6.52	6.39	0.47	7.40	7.45
Non-US Models	24.21	22.83	27.07	19.59	19.79	22.72
Surplus Effects (billions $)						
Consumer Surplus	−25.91	−56.89	−59.27	−30.50	−59.16	−62.28
Manufacturer Profits	−5.51	−16.31	−17.13	−6.03	−13.25	−13.98
Tax Revenue - car imports	13.47	17.93	15.57	17.88	22.01	19.65
Tax Revenue - parts imports		23.65	25.64		25.42	27.96
Total Net Surplus	−17.95	−31.62	−35.19	−18.65	−24.98	−28.64
Employment Effects (thousands of jobs)						
Car Assembly	68.7	32.65	40.28	52.32	23.79	30.12
Parts Manufacturing	38.3	3.44	3.33	24.67	−8.21	−9.74

Table reports effects of tariffs in Counterfactuals 1, 2, and 3 predicted by our preferred model (Cournot with economies of scale) and the standard model in the literature (Bertrand with constant marginal cost). As there is no uncertainty in cost function under the constant marginal cost assumption, we suppress standard errors for Cournot with economies of scale (which were reported in Section 6).

D.2 Additional Counterfactuals

Which Car Models Would Exit?: While our main analysis examines medium-run effects holding the product set fixed, manufacturers facing substantial profit declines might eventually discontinue certain models or relocate production. Table 13 examines this possibility by identifying models at risk of discontinuation under tariff scenario C3.

To assess how car manufacturers' product portfolios could be affected, we focus on two popular segments (SUVs and sedans), and identify segment-specific profit thresholds based on the lowest variable profit levels observed in the pre-tariff market—$16.3 million for sedans (represented by the Fiat 500) and $21.3 million for SUVs (represented by the Mitsubishi Outlander PHEV).[36] These figures are consistent with other estimates in the literature: Sabal (2025) quantifies median market entry costs for automobiles at $8-15 million. Models falling significantly below these thresholds post-tariff would be candidates for exit or production relocation.

[36] We exclude from this analysis vehicles with less than 2,500 sales in 2018—these may not have been on sale for the full year (e.g., because discontinued), or correspond to very small niches.

TABLE 13: Models at Risk Under Counterfactual C3

Car Model	Profit (millions $)	Assembly	US Content (%)	Profit Change (%)
Panel A: Sedan Segment				
Honda Clarity PHEV	7.1	Japan	0	-94
Kia Stinger	8.6	Korea	0	-93
Ford Fusion PHEV	15.0	Mexico	30	-71
Nissan 370z	7.3	Japan	0	-70
Kia Cadenza	11.6	Korea	0	-62
Subaru BRZ	10.3	Japan	0	-56
Fiat 500	8.9	Mexico	19	-46
Fiat 124 Spider	11.6	Japan	0	-46
Mitsubishi Lancer	11.3	Japan	1	-37
Panel B: SUV Segment				
Toyota Land Cruiser	3.1	Japan	5	-90
Mitsubishi Outlander PHEV	7.0	Japan	0	-67
Kia Niro PHEV	16.0	Korea	2	-34

The table reports models with the largest profit declines under the C3 tariff scenario, showing variable profits under C3, assembly location, US parts content percentage, and percentage change in profits relative to baseline.

In total, nine sedan models and three SUV models fall below the profit thresholds after the tariffs are imposed. Table 13 reports those and their respective profit decline under C3. For sedans, several models face severe profit reductions: the Honda Clarity PHEV (94% decline), Kia Stinger (93% decline), Ford Fusion PHEV (71% decline) and Nissan 370z (70% decline). In the SUV segment, the model predict a large profit reduction for the Toyota Land Cruiser (90% decline). These reductions would likely trigger production decisions beyond price adjustments.

The assembly location and parts content data reveal clear patterns. Most severely affected models are assembled in Asia (primarily Japan and Korea) with minimal US parts content. For instance, the Honda Clarity PHEV, Kia Stinger, and Mitsubishi models all show 0% US content. Instead, Mexican-assembled models have higher US content (30% for the Ford Fusion PHEV and 19% for the Fiat 500), making them more vulnerable to the reciprocal tariffs in C3. Beyond assembly location, the results highlight that some specialized, low-volume models and electrified vehicles would be at risk. Eight of the nine most-affected models are either niche performance vehicles (370z, Stinger, 124 Spider), more premium offerings (Land Cruiser), or plug-in hybrid electric vehicles (Clarity PHEV, Fusion PHEV, Outlander PHEV, Niro PHEV).

Reshoring Incentives Under Tariffs: The stacked and reciprocal tariffs in C3 create complex incentives for manufacturers to adjust their supply chains. To explore these dynamics, we simulate a counterfactual scenario where each of the top-selling models increases

its US/Canadian parts content by 10 percentage points after the implementation of C3 tariffs, holding all other factors constant. This allows us to isolate the strategic incentives for reshoring parts production on a model-by-model basis.

It is important to note that this exercise should be interpreted as an enhanced back-of-the-envelope calculation rather than a full evaluation of parts sourcing decisions—this is not a margin of decision that is directly captured by our model. Thus, we quantify only the potential benefit of reshoring in terms of decreasing the tariff bill, while abstracting from the costs and constraints that might make domestic sourcing more expensive or even infeasible. In reality, manufacturers source parts globally for numerous reasons beyond cost, including access to specialized expertise, quality considerations, capacity constraints, and technological advantages. Our analysis thus quantifies only one side of the trade-offs involved in the decision to reshore parts production, as it does not account for the potentially substantial costs that led manufacturers to establish global supply chains in the first place.

TABLE 14: Effects of Increasing US Content by 10 Percentage Points Under C3

Car Model	Price Effect(%)		Share Effect(%)		Profit Effect(%)		US Content	Assembly
	C3	Change	C3	Change	C3	Change	(%)	
Ford F-Series	6.6	−1.1	−3.3	5.4	−1.2	7.1	53	US
Chevrolet Silverado	5.9	−1.5	3.6	10.5	2.2	12.0	46	US
Ram Pickup	6.0	−1.2	1.1	7.9	1.7	10.2	56	US
Toyota RAV4	29.0	2.2	−62.6	−4.0	−62.9	−3.9	35	J
Nissan Rogue	10.1	−1.7	8.7	8.8	13.1	10.2	25	US
Honda CR-V	13.1	−1.9	−13.2	8.8	−6.9	10.1	23	US
Toyota Camry	3.6	−1.7	18.6	9.6	20.1	10.3	65	US
Chevrolet Equinox	39.5	5.2	−77.8	−5.2	−78.5	−4.9	45	CN
Honda Civic	7.3	−1.3	7.5	5.4	12.6	6.0	48	US
Honda Accord	5.4	−1.9	5.8	11.1	10.6	12.1	65	US

The table reports effects for the top 10 models by 2018 sales volume. For each model, we show the C3 tariff effect on price, market share, and profits (in percentage terms), followed by the change in percentage points when that specific model increases its US parts content by 10 percentage points. Assembly locations are US except for Toyota RAV4 (Japan) and Chevrolet Equinox (Canada).

Table 14 presents the results of this counterfactual for the top-selling models in 2018. Several patterns emerge that highlight the complex incentives created by stacked and reciprocal tariffs. First, increasing US parts content generally leads to lower price decreases across most models, as indicated by the negative values in the price change column. For example, the Honda CR-V would see a 1.9 percentage point reduction in its price increase under C3, while the Toyota Camry would experience a 1.7 percentage point smaller price increase.

However, the Toyota RAV4 and Chevrolet Equinox are exceptions. Despite being among the most impacted models under C3 (with a 29% and 39.5% price increase, and a 62.9%

and 78.5% profit decline, respectively), increasing their US content would counterintuitively worsen their situation, with price effects increasing by an additional 2.2 and 5.2 percentage points. This stems from the models' foreign assembly location, where the reciprocal tariffs create perverse incentives that penalize reshoring of production parts.

For market shares, the reshoring effects are generally positive for US-assembled vehicles, with the Honda Accord showing the largest improvement (11.1 percentage points). These share gains translate into profit gains, with all US-assembled models seeing profit enhancements ranging from 6 to 12.1 percentage points.

These findings highlight three insights about reshoring incentives under a tariff regime like C3. First, the benefits of increasing domestic content are not monotonic in initial US content levels. For instance, the Honda Accord (65% US content) receives a larger profit benefit (12.1 percentage points) from increasing its domestic content than the Honda Civic (48% US content) with only a 6 percentage point improvement. Second, assembly location alters the reshoring calculus. For the Toyota RAV4 assembled in Japan and Chevrolet Equinox assembled in Canada, increasing US content actually amplifies its competitive disadvantage in the C3 scenario.

In summary, while reshoring parts production could mitigate some negative tariff impacts for certain models, the benefits are unevenly distributed and sometimes counterintuitive. A more complete analysis would need to weigh these tariff avoidance benefits against the fundamental economic reasons that led to global sourcing in the first place. These findings show how the complexity of global value chains can produce unexpected outcomes when disrupted by stacked tariff policies, highlighting the importance of considering these nuanced effects when evaluating trade interventions.

D.3 Robustness

Alternative Cost Specifications: In Table 15, we explore the sensitivity of our economies of scale estimates to both functional form assumptions and the definition of the total production at which economies of scales accrue. Unless noted, all regressions contain the full set of exogeneous cost shifters and fixed effects in Table 11 and we rely on 2SLS estimation using the instruments discussed in Section 5. In our main specification in Equation 3, economies of scale accrue across units of a car model produced in a given country and we adopt a log-log functional form for marginal cost. Column 1 reproduces our IV estimates and Column 2 reports the OLS estimates: endogeneity attenuates the economies of scale, as expected. Column 3 drops the SUV time trend from the regression while in Column 4, we adopt a log-linear specification of marginal cost, including both total production and its square (the implied economies of scale from these estimates is -0.19).

Finally, in Column 5, we allow economies of scale to accrue across the country-level pro-

duction of all car models using the same platform. A prominent feature of the car industry is that car models in the same segment produced by the same manufacturer often share the same engineering platform, meaning that they use the same mechanical underpinnings, such as engine and transmission (Van Biesebroeck, 2003). Therefore, it is possible that economies of scale and scope accrue at the platform level. The Markline data contains information on car platforms, and we are able to construct platform data for around half of the car models in our data.

Across specifications that we estimate via 2SLS, economies of scale estimates range from -0.114 to -0.19, suggesting our main estimates are fairly robust to the functional form of Equation 3 and the definition of q_t^p.

TABLE 15: Robustness of Cournot Implied Economies of Scale to Alternative Functional Forms

	Dependent variable: $\log c_C$				
	Main Spec	OLS	Drop SUV Trend	Log-Linear Quad	Platform Level
	(1)	(2)	(3)	(5)	(8)
$\log(q_{jt}^p)$	−0.114	−0.038	−0.176		
	(0.033)	(0.007)	(0.042)		
q_{jt}^p (100 ths)				−0.299	
				(0.152)	
$(q_{jt}^p)^2$ (100 ths)				0.034	
				(0.028)	
$\log(\Sigma_{j\in\text{platform}} q_{jt}^p)$					−0.160
					(0.050)

Table reports economy of scale estimates from log marginal cost regression in Equation 3 under Cournot model of conduct. All regressions include the exogenous cost shifters and fixed effects reported in Table 11. Column 1 reproduces the main results in Table 3 - total production at the model-country level enters in logs. Column 2 reports OLS results obtained without instruments. Column 3 removes the SUV time trend. In Column 4, total production and its square enter in levels. Column 5 models total production at the platform-model level. Standard errors clustered by car model in parentheses. 3,929 model-year observations.

Alternative Tariff Assumptions: Here, we explore the robustness of our results in Counterfactual 3 to two maintained assumptions when computing our tariff counterfactuals. First, when imposing tariffs on foreign car manufacturers, we modeled the port cost of a vehicle as 80% of its retail price (in line with Goldberg, 1995). Alternatively, Coşar et al. (2018) estimate the port cost to be 68% of the implied marginal cost. Column 3 of Table 16 reports results in Counterfactual 3 under this alternative assumption.

TABLE 16: Robustness of Counterfactual 3 Results to Alternative Tariff Assumptions

			Parts Tariff Pass-Thru =	
Tariff Effects	Main	Port Cost = 0.68 c_m	0.2	1.2
Price Effects (percent)				
All Models	14.56	11.33	10.11	15.71
	(0.99)	(0.88)	(0.80)	(1.06)
US Models	6.39	6.58	0.55	7.88
	(0.10)	(0.19)	(0.34)	(0.17)
Non-US Models	27.07	18.60	24.74	27.70
	(2.51)	(1.98)	(2.51)	(2.50)
Surplus Effects (billions $)				
Consumer Surplus	−59.27	−52.50	−33.30	−65.15
	(0.57)	(0.73)	(2.45)	(0.52)
Manufacturer Profits	−17.13	−16.09	−7.98	−19.19
	(0.80)	(0.49)	(1.10)	(0.78)
Tax Revenue - car imports	15.57	14.93	13.90	16.02
	(1.39)	(0.96)	(1.55)	(1.32)
Tax Revenue - parts imports	25.64	24.80	6.04	29.62
	(0.18)	(0.26)	(0.11)	(0.29)
Total Net Surplus	−35.19	−28.87	−21.34	−38.69
	(1.10)	(1.58)	(2.31)	(1.31)
Employment Effects (thousands of jobs)				
Car Assembly	40.28	22.78	62.19	35.11
	(3.93)	(3.38)	(8.00)	(3.45)
Parts Manufacturing	3.33	−7.00	31.24	−3.35
	(5.26)	(4.09)	(10.31)	(4.91)

Table reports effects of stacked and reciprocal tariffs in Counterfactual 3 under alternative tariff assumptions. Column 2 reproduces our main results while Column 3 assumes the port cost of a vehicle is 68% of implied marginal cost. Columns 3 and 4 assume the pass-through rate of parts tariffs to input prices for car manufacturers is 0.2 and 1.2 respectively. Standard errors in parentheses.

Second, when levying tariffs on car parts, we assumed that the pass-through of these tariffs to the input prices paid by car manufacturers was one. Columns 4 and 5 of Table 16 report C3 results for alternative assumptions on the pass-through rate (0.2 in Column 4 and 1.2 in Column 5). The overall results hold under these alternative assumptions: stacked and reciprocal tariffs increase overall car prices and reduce consumer and producer surplus. Overall, total net surplus is lower with tariffs and the employment effects imply that the surplus reduction per job created is less than the average manufacturing wage in these industries.