

**Trucks without bailouts:
Equilibrium product characteristics
for commercial vehicles
JOB MARKET PAPER**

THOMAS WOLLMANN*

Latest draft: [click here](#). This draft: November 12, 2014.

Abstract

In differentiated product markets, the entry and exit of individual product models—rather than of firms—often serve as the main equilibrating force. Market structure changes that lead to high prices also tend to encourage entry, partially offsetting the effects. Thus, accurately predicting changes from a merger or bankruptcy should incorporate this behavior. I develop a model of equilibrium product characteristics in oligopoly and show how inequalities can identify the sunk costs of offering them. I apply these methods to an original dataset of all US commercial vehicles from 1987 to 2012 and assess their importance in the context of the \$85 billion bailout of the automotive industry in 2009. In the case where GM and Chrysler are liquidated rather than rescued, estimated sunk costs are low enough to induce product entry by rivals, and this has a dramatic effect on prices and purchases. Markup increases for the most affected products drop by two-thirds when I account for model-level entry and exit. The decrease in total output is also moderated by about one-half. Acquisitions are an alternative to liquidation. I show that although the policy choice, including who the acquiring firm is, matters a lot when model-level entry and exit are ignored, it matters little when they are accounted for.

*Harvard University, Department of Economics, Littauer Center, Cambridge, MA 02138 (email: wollmann@fas.harvard.edu). I am grateful to my advisors, Ariel Pakes, Juan Alcacer, Greg Lewis, and Dennis Yao, for their guidance and encouragement. This paper benefited from conversations with Richard Sweeney, whom I have worked closely with on related work, and Thomas Covert, James Lee, Robin Lee, Julie Mortimer, Mark Shepard, Michael Sinkinson, Haris Tabakovic, Che-Lin Su, and Ali Yurukoglu. I also thank Mike Landry of Span Alaska for his help with vehicle data and Steve Matsil, who was very generous in sharing his knowledge of commercial vehicle production and demand.

I Introduction

The response of firms to heterogeneous consumer preferences has improved considerably since Henry Ford famously remarked, “You can have the Model T in any color, so long as it’s black.”¹ In many industries, a second pattern emerged alongside the surge in differentiated production. Industries would come to be organized around a single, relatively stable set of firms but a rapidly evolving set of product offerings. The US automotive industry, for example, has witnessed virtually no firm-level entry or exit for thirty years despite wide variation in the number and nature of products over time. Outside of autos there are many other examples, ranging from aircraft to bicycles to cat food. In each case, it is the entry and exit of individual product models—not the Schumpeterian (1942) creation and destruction of firms—that serve to drive these markets into equilibrium. Yet despite rich evidence on the importance of accounting for entry and differentiation separately,² little is known about their combined impact. This paper examines how model-level entry and exit impacts policy analysis in differentiated product markets.

The theoretical motivation is simple. Entry and exit tend to work in the opposite direction of the price mechanism, so failing to account for them can overstate the impact to prices and purchases of a change in market structure. To illustrate, consider the acquisition or bankruptcy of a firm in oligopoly. Ignoring entry, this exit increases the market power of surviving firms, who raise markups and earn windfall profits. However, the prospect of high profits lures entrants. Even if startup costs are large enough to prohibit firm-level entry, the cost of adding or repositioning products for incumbents may be small enough to permit model-level entry. If the latter is true, firms will set lower prices that reflect a more crowded product space where consumers can more easily substitute. Thus, counterfactuals depend on accurately determining where and to what degree model-level entry will occur.

I provide methods for estimating the parameters governing model-level entry and exit decisions as well as computing counterfactuals that account for this behavior. Using these methods, I study the impact of the \$85B government bailout of the US automotive industry in late 2008 and early 2009. During this period, General Motors (“GM”) and Chrysler were headed for default.³ Whether to provide federal assistance was hotly contested and even became a Presidential campaign topic in 2008⁴ and 2012.⁵ For tractability, I narrow the scope of my assessment to the commercial vehicle segment of the auto industry, and then construct a original dataset of all product offerings between 1987 and 2012. Taken together, the panel data and methods allow me to ask, “What would have happened to output and prices had the government not rescued the automakers?” There are two obvious policy alternatives. The first is liquidation, an effective removal of the GM and Chrysler brands and products from the marketplace. The second is an acquisition by an existing, rival firm.

¹There are variations of this famous saying and this seems to be the most common. Ford (1922) recalled it slightly differently in his autobiography, but without a change in meaning.

²For early examples of the former see Borenstein (1989) and Bresnahan and Reiss (1991) and of the latter see Berry, Levinsohn, and Pakes (1995) and Nevo (2001).

³Isidore, Chris. “Bush announces auto rescue.” *CNNMoney* 19 Dec 2008.

⁴Romney, Mitt. “Let Detroit Go Bankrupt.” *New York Times* 18 Nov 2008.

⁵Rosche, Jedd. “Obama quotes Romney’s auto bailout op-ed in UAW speech.” *Politico* 28 Feb 2012.

To illustrate the range of outcomes, my counterfactuals consider a sale to Ford, which overlapped most heavily with GM and Chrysler in product space, and to Paccar, which overlapped the least.

The commercial vehicle market is representative of the auto industry overall and an ideal place to study this phenomenon. No firms have entered or exited in three decades but product offerings have changed frequently.⁶ Ownership is concentrated among about ten firms and is especially concentrated in sub-segments of the market, even though most firms have produced most product variants at some time. Furthermore, the highly modular production of commercial vehicles allows manufacturers to quickly swap parts and introduce new models, often within months of changing demand or competitive conditions. This creates a tight link between the incentives to adjust products and the actual adjustment decisions. Finally, their physical weight, tariff structure, and legal standing isolate commercial vehicles from the passenger vehicle market and foreign automotive markets.

I find that in the event of liquidation, profit increases are high enough relative to sunk costs to induce product entry among surviving firms, and that this has a dramatic impact on the market. Relative to the bailout, markups for the most affected products rise by over 70% when entry is ignored but only 18% when it is accounted for. At the same time, the probability of not purchasing a vehicle rises nearly 50% for the most affected buyers when entry is ignored, but only 14% when it is accounted for. The median impacted buyer and product, on the other hand, see virtually zero change. This reflects the heavy overlap of GM and Chrysler products and strong preference heterogeneity in the buyers. For this reason, the impact on total output is much more muted: it falls 7.6% when entry is ignored and 3.2% when it is accounted for.

Acquisition is another alternative. If entry is ignored, an acquisition by Ford, which overlaps heavily with GM and Chrysler in product space, looks qualitatively the same as a liquidation. In sharp contrast, an acquisition by Paccar, which does not overlap with GM or Chrysler, closely resembles the bailout. If model-level entry and exit are accounted for, however, all three counterfactual policies are essentially symmetric. Relative to the bailout, markups on the most affected products rise between 14% and 18% while total output falls between 2% and 3%. The policy choices—including the identity of the acquiring firm—appear to matter a lot when model-level entry and exit are ignored, but in fact matter little when this important equilibrating mechanism is accounted for.

These results require two methodological contributions. The first is a model of equilibrium product offerings that handles multi-product firms selling multi-attribute goods but remains tractable in applied settings. In this model, buyers vary in their preferences over characteristics, e.g. urban buyers on congested roads prefer short and maneuverable vehicles while long-distance freight haulers prefer large and rugged ones. Sellers face sunk costs to add and remove products as well

⁶In the passenger segment, the most recent entrants are Hyundai and Kia, who began selling in 1986 (Kia began by re-badging its exports and did not legally incorporate a standalone US brand until 1992). This ignores Tesla, who at the time of writing is still very small (but growing). In the commercial segment, the most recent were Hino and UD in 1984 and 1985, respectively, although Hyundai exported a small number of vehicles between 2000 and 2001 under the Bering badge before being absorbed into Daimler.

as marginal costs of production. In each year, firms choose which vehicles to offer in the first stage and, conditional on those offerings, choose what prices to charge in the second stage. First stage product entry and exit decisions, therefore, weigh sunk costs against changes in second stage profits. Simultaneous first period choices form a Nash equilibrium in product space, while second period choices form a Bertrand-Nash equilibrium in prices. Sunk costs can induce forward-looking behavior among the producers, but the dynamic programming problem in this market—and most richly-defined differentiated product markets—places an unreasonably large computational burden on firms. It requires taking expectations over billions or more of states and rules out learning by repeated play, but raises the question of what managers actually do in these situations. Interviews with commercial vehicle product designers and division managers suggest the use of hurdle rates to best approximate this process. This capital budgeting rule greatly simplifies the entry and exit decisions that firms make. Later in the paper, I provide evidence that the discounted profits firms actually earn are quite close to what is implied by hurdle rates.

The second is a method of identifying sunk costs in industries with rich time variation but only one geographic market. The first step in the process entails recovering the primitives governing demand and using the Bertrand-Nash pricing condition to recover marginal costs. This yields an estimate of the profits that firms would earn under any hypothetical set of product offerings. The second step constructs inequalities based on the fact that firms must find the observed product offerings—i.e. equilibrium choices—preferable to all other alternative sets of product offerings. The model is identified by combining these inequalities with exogenous shifts in the composition of buyers over time, which firms respond to by changing product offerings. My data reveal that, for example, the size of the construction industry is a strong predictor of vehicle offerings tailored to the preferences of those buyers. This response is true for other industries as well as for a law change that shifts demand. These factors provide variation that can be combined with revealed preference and the necessary conditions for a Nash equilibrium to identify sunk costs.

Relying on only necessary equilibrium conditions is important since the presence of multiple equilibria rules out a one-to-one map between the parameters and the outcomes. For example, for a guess of the parameters I may only be able to say that Firm A will enter if Firm B does not and vice versa (but cannot say which is more likely). The econometrician can rarely take a stand on which equilibria is played, so calculating a likelihood is impossible here. Nonetheless, this still provides information. For instance, I can penalize instances where both Firm A and B enter or where neither does. Inequalities have been used in this way before (Tamer (2003), Pakes, Porter, Ho, and Ishii (2014)), but combining them with exogenous demand shocks to identify parameters over a time series has not. Several plausibly exogenous ownership changes, which are another helpful feature of the market, also provide exogenous variation in product entry and exit incentives.

This paper contributes to a growing empirical literature on firm positioning. Prior work has largely relied on cross-sectional variation provided by multiple geographic markets (e.g. Mazzeo (2002), Seim (2006), Fan (2009), Draganska, Mazzeo, and Seim (2009)). These papers cleanly demonstrate the inherent tradeoffs in differentiation in many important settings; however, they

require that industries comprise many local, isolated markets. This has tended to restrict their application to either goods that are costly to transport (e.g. ice cream, which melts and spoils) or service-related products (e.g. hotels, retail, or news production). In contrast, commercial vehicle models—like many consumer and industrial goods—are produced in at most one or two locations but distributed nationwide. This rules out geographic variation on the supply side. Thus, identifying fixed or sunk costs necessarily requires looking over a time series. Recent work on microprocessors do precisely that, although it assumes that the possible set of product offerings evolves exogenously over time (Nosko (2014), Eizenberg (2014)). Such an approach may prove difficult to extend past industries that are not guided by plausibly exogenous technological variation (e.g. Moore’s Law). I instead rely on exogenous shifts in demand. This rests on only observably heterogeneous buyer preferences, which have already been shown critical to understanding purchasing and pricing patterns in differentiated product markets (Petrin (2002), Berry, Levinsohn, and Pakes (2005)).

Recent work on dynamic games has also estimated sunk costs of entry, exit, and repositioning. Recovering the primitives of these games was made easier with a two-step approach introduced by Hotz and Miller (1994), which in effect transferred the burden of solving these problems from the computation to the data (Pesendorfer and Schmidt-Dengler (2008), Aguirregabiria and Mira (2007), Pakes, Ostrovsky, and Berry (2007), Bajari, Benkard and Levin (2007)). Agents in these models fully internalize the future impact of their decisions. Applying these methods, however, has required a smaller action space than I consider here (e.g. Blonigen, Knittel, and Soderbery (2013), who notably study passenger vehicle refresh and scrapping decisions), often with large numbers of geographic markets (e.g. Collard-Wexler (2013), Ryan (2012) and Sweeting (2013)). Again, in most differentiated product markets, the number of potential product types is large while the number of geographic markets rarely exceeds one. Morales, Sheu, and Zahler (2014) provide an alternative that is robust to large choice sets, although limited to the case of monopolistic competition, so firms face only a single agent problem. They show that Euler equation perturbations yield inequalities that can flexibly identify fixed and sunk costs (although these would be ruled out if there was strategic interactions).

The rest of the paper is as follows. Section II describes the market and data. Section III shows descriptive evidence of how changes in demand impact changes in product offerings. Section IV presents the model and Section V discusses estimation. Section VI reports the results. Section VII describes the policy setting and computes counterfactual outcomes. Section IX concludes.

II Market Setting and Data

Overview

The commercial vehicle segment of the US automotive industry accounts for about 10% of total US automotive sales (WardsAuto (1986-2013)), which themselves account for 4% of gross domestic product.⁷ The segment comprises any on-road vehicle rated for over 10,000 lbs gross vehicle weight (defined below) and sold domestically.⁸ In terms of use and users, their scope is quite broad. For example, the market includes inner-city delivery vans, landscaping flatbeds, dump trucks, and highway tractor-trailers. In terms of capabilities, the high end of the segment has carried loads in excess of 250,000 lbs, such as oil rigs and turbine engines, and as cumbersome as an Airbus A320 fuselage, which was the case when the “Miracle of the Hudson” was hauled from New York City to the Smithsonian museum in Washington DC.⁹

Data

Three main sources of data are used in this paper. First, I compile a panel of all commercial vehicle models sold in the US from 1986 to 2012 from annual issues of *The Truck Blue Book*. Each observation includes the brand, model year, model name and number, and a host of product characteristics. These include price as well as detailed specifications related to the load capacity, cab, chassis, powertrain and drivetrain. Each guide contains data for the prior ten years, easing compilation, but unfortunately does not contain retail price data for the most recent model-year vehicles, so I proceed with MSRP. I convert all prices to 2005-equivalent dollars via the Consumer Price Index.

I merge the product characteristic data to unit sales data from the R.L. Polk & Co. *New Vehicle Registration Database*. The Polk data covers US vehicle sales Class 3 and above and is collected through state motor vehicle registrations so is believed to be extremely accurate. Observations are broken down by brand, model name and/or number, as well as gross vehicle weight rating class, fuel-type, and body description. All sales figures are compiled on a calendar year basis. In rare cases, a model was identified in the quantity data by a model number and in the product characteristics data by a model name, or vice versa, although I was able to resolve these conflicts using the *Official Commercial Truck Guide* published by the American Truck Division of the National Automotive

⁷Dept. of Commerce *Auto Industry Snapshot*. <http://selectusa.commerce.gov/industry-snapshots/automotive-industry-united-states>. Accessed 12/20/13.

⁸The “on-road” distinction here, or in US Dept. of Transportation terms “on-highway,” merely excludes irrelevant vehicles like one-seat “terminal tractors” and fifty-foot-tall mining trucks. I exclude motorhomes, buses, and “step” vans from the dataset because these are by-and-large manufactured by different firms and bought by different consumers than the vehicles in this dataset. Finally, I do not include low-entry cab-forward models. Although these are made by Mack/Volvo, and Peterbilt, who appear in this dataset, there is no variation in who produces them or their characteristics (with the exception of the short-lived Sterling Condor). Their exclusion should also not affect the demand system, since they are used almost exclusively in urban garbage collection.

⁹The “Miracle on the Hudson” refers to US Airways commercial flight 1549, which Capt. Chesley Sullenberger emergency landed on the Hudson River following a bird strike.

Dealers Association (NADA). In two cases, I resolved the issue by calling a dealer.

The third main data source is microdata on commercial vehicle purchases available through the US Census. This microdata was collected up to and including 2002 under a program known as the Truck Inventory Use Survey (TIUS) and later the Vehicle Inventory and Use Survey (VIUS). Every five years, the Census mailed approximately 130,000 owners of trucks and vans and asked them various questions about their vehicle, the use of their vehicle, and about the owners themselves (relating to their vehicle use). The response rate was approximately 80% and relatively stable over time. I observe the industry and state that the buyer operates in, whether the vehicle was acquired new or used, and the characteristics of the vehicle the buyer owns.

Three additional pieces of information complete the dataset. First, the US Census *County Business Patterns* contributes the number of US firms by industry, state and year, while the US Department of Transportation (“DOT”) *Highway Statistics* contributes urban and non-urban road mileage by state and year. Together, these provide an empirical distribution of buyer types that serve as the basis for taking simulation draws in the demand system and that determine changes in size of the total potential market for commercial vehicles. Last, the Bureau of Labor Statistics contributes product worker wages at the state and year level, which I match to product models based on their respective factory locations.

Product Characteristics

Commercial vehicles are conveniently summarized by a short list of product characteristics. The first and most important is gross vehicle weight rating (“GWR”), defined as the maximum load that may legally rest on the axles. GWR is the main means by which both the automotive industry and the Department of Transportation characterize vehicles “classes” of vehicles. The threshold of 10,000 lbs GWR provides a natural separation between the passenger and commercial segments. Below this threshold are cars, minivans, station wagons, nearly all pickup trucks, SUVs, and cargo vans; above it are what would typically be thought of as work trucks. In terms of use, the Census microdata show a very obvious distinction. More than 95% of vehicles below 10,000 lbs are designed for “personal use” while less than 5% above this threshold are. In terms of production, the physical design is also distinct. Vehicles below this cutoff usually feature unibody design, meaning the exterior skin provides primary support to the load. Vehicles above this cutoff feature “body-on-frame” design, meaning the load is supported by a ladder frame, onto which assemblers attach the axles, cab, power unit, and controls.¹⁰ That is, the driver sits inside the load-bearing structure in a passenger vehicle but sits on top of it in a commercial one. Body-on-frame design allows assemblers allows vehicle assemblers to quickly modify the characteristics of a vehicle, often in as little as a few months.

¹⁰The technical term is monocoque construction, meaning the exterior skin supports the load. Variants of this term, like semi-monocoque and unitary construction, more precisely describe modern passenger vehicles, but all are distinct from body-on-frame.

GWR determines the possible uses of a vehicle. Since carrying loads in excess of it is illegal and unsafe, and since it increases price, buyers purchase vehicles with the minimum GWR that safely covers their needs. Other characteristics, like the transmission and engine, relate quite closely (even though they may not map exactly).¹¹ Since firms design vehicles with the lowest cost—and therefore lowest weight-rated—parts that can safely haul the GWR load at highway speeds, this tight relationship between GWR and all other load-related attributes is not surprising. The only exception to the usefulness of GWR as a load capacity measure is at the top end, where vehicles are more likely “pulling” rather than “carrying” loads. I adjust GWR to the gross combination weight rating, which is the appropriate measure for these vehicles, although the measures are so correlated that in practice this adjustments only affects ten models in the panel.

The industry’s desire for compatible parts also means that while GWR is technically a continuous choice variable, it takes on only discrete values in practice. For example, over thirty vehicles in the data have a GWR of exactly 52,000 lbs. The small group of models that did not match exactly to any group were always close to some larger group, so the group’s GWR was substituted in for these models. After this adjustment, GWR takes up to 22 values.

The second characteristic is the style of cab, the portion of the vehicle that encloses the passengers, controls, and engine. Cabs come in three distinct varieties. The most popular cab type is the “conventional,” which is distinguished by its relatively long hood and placement of engine well ahead of the occupants. The conventional cab places the axle ahead of the driver, making for a smooth ride and spacious interior, but its long hood limits the amount of maneuverability in tight spaces. The cab-over-engine (“cabover”) type features a flat front and places the occupants directly over the front axle and engine. This makes for exceptional visibility and turning but an uncomfortable and less safe driving experience.¹² Although not particularly popular overall, they are common congested city environments. A compromise between the two is the compact-front-end, or what is commonly called the “van.” Vans push steering and controls as forward as possible, making for a short but slanted hood. They feature some benefit of both the cabover and conventional design, but are limited in engine size.¹³ A final characteristic specific to heavier conventional vehicles is a long-hood design. The long-conventional design provides maximal comfort and minimal noise to drivers making long or difficult hauls. The characteristic is sufficiently important to be listed separately in product guides (e.g. *The Truck Blue Book*).

Table I summarizes the product data. The top panel provides a summary of the data by year, while the bottom panel aggregates the data into equal nine-year periods. The most striking aspect of the data is that despite rich year-to-year variation in the number and type of products offered as well as units sold, the market exhibits few strong trends (with the exception of heavy cabover

¹¹Engines deserve discussion. Although each model typically carries a large number of engine options, the base option is usually the same across manufacturers for a given GWR and cab type. Product guides even include charts that relate one engine type to another across manufacturers, suggesting close comparability. Additionally, the market for large trucks (where engine quality matters most) is dominated by Cummins and Caterpillar, which are independent of any manufacturers in our data.

¹²Cab-over-engine vehicles place no distance between the drive and the vehicle in front of it, which creates the safety issue.

¹³For this reason, compact-front-end vehicles take only the four lowest GWR values in the data.

and compact-front-end vehicles, discussed at length later on). For example, although quantity swings considerably between years, the bottom panel shows that the size of the commercial vehicle market has grown little in the past 26 years. Relatedly, there is little movement in price over time (after a CPI adjustment to 2005 constant dollars). This suggests (but clearly does not strictly imply) production costs have not fallen much over time, which would not be surprising considering that modular production limits the extent to which large capital investments can automate the manufacturing process. GWR tells a similar story. Together these facts suggest stable preferences and slow long-term growth but large compositional changes.

Individual product types also move quite independently from one another over time. Although the number of models tends to peak during periods of high economic activity, it is clear this is driven almost completely by medium weight vehicles. The next section shows that their sensitivity to the construction industry—rather than the economy overall—drives this relationship. Other vehicle types do not exhibit this relationship. Nonetheless, model-level entry and exit is still quite frequent. In the case of heavy cabover vehicles, something else is clearly at work. In fact, this variation across product types and time is at the core of the identification of sunk costs.

[Table I about here.]

Buyer Attributes

Buyers differ by industry, driving environment, and legal climate. Industry matters because it determines the load size and driving distance, which in turn influence preferences over GWR and cab size. According to US Census microdata, over 94% of buyers belong to either the for-hire transportation (“freight”) industry, one of three construction-related industries, or the business and personal service industry. Driving environments vary by whether surrounding areas are urban or not, which I measure by dividing urban road mileage by total road mileage. The third buyer attribute relates to a series of vehicle length laws. Initially states regulated the entire vehicle length, from the front of the power unit to the rear of the load being carried. A series of federal legislative and judicial decisions, beginning with the Surface Transportation Assistance Act, mandated that only the length of the load being carried be regulated. This allowed the power unit, i.e. truck or tractor, to be as long as the driver desired. Rather than analyze each complicated decision, I take a simple count of the regulatory actions. In Section V, I discuss the law change in more depth and, in section VI, show the measure has good explanatory power for demand.

Table II highlights the tight link between buyer attributes and product characteristics. The left-most column in each row reports the average product characteristic conditional on purchase by any buyer, while the columns to the right report the average product characteristic conditional on purchase by a particular group of buyers. The first row shows how GWR varies by industry type. Delivery and service firms are purchasing most light commercial vehicles, construction-related firms are purchasing most medium weight units, and freight firms are purchasing the majority of heavy

units. It should not surprise readers that, for example, the florist’s van handles much lighter loads than the builder’s dump truck. The degree of heterogeneity may be of some surprise, although these purchase patterns are not atypical (and are certainly inline with the Petrin (2002) findings that relate family size to minivan purchase). The second row shows that shorter cab types—with better visibility and agility—are preferred by urban buyers. It is not a particularly popular choice overall, but mere observation of, for example, downtown Manhattan, reveals that it dominates congested areas. The third row shows that compact-front-end cabs are preferred choice by local delivery firms. In most areas, the noisy and bumpy cabover is an unnecessary choice, although the extra visibility and tighter turning of compact front end vehicles is a help to these drivers.

[Table II about here.]

Firms

This paper focuses on automotive assembly firms. It treats upstream and downstream operations as either independent or completely determined by assembly operations. There are several reasons for this. First, the assemblers are few in number but large in size, and serve as the central party to contracts between the disaggregated parts suppliers and geographically diverse dealerships. Unlike the carmakers, who typically build major components in-house or sign exclusive contracts with third-parties, commercial vehicle assemblers incorporate components from a host of suppliers. Whereas carmakers typically build many major components in-house, for example engines and axles, commercial builders rarely do. Cummins and Caterpillar, for example, account for a majority of heavy-duty truck engines but are not active in the market themselves. Also unlike the passenger vehicle segment, commercial dealers often carry competing brands.¹⁴ One exception is that firms with both commercial and passenger vehicle operations occasionally leverage their passenger vehicle dealerships to sell commercial units (which I account for in estimation).¹⁵ Second, assembly firms map directly to the “brand” that identifies the vehicle. Third, extending the analysis along the value chain is simply beyond the scope of the current data and methods.

US commercial vehicle production is also separate from its foreign market counterparts. The catalyst for this separation is a 25% import tariff on trucks imposed in 1963 and in effect today. The duty is part of Proclamation 3654 and ubiquitously referred to in the auto industry as the “Chicken Tax” due to the fact that President Johnson aimed it primarily at stemming poultry imports. It applies to all “truck” imports to the United States. Although what is meant by “truck” is hotly contested, all vehicles in this paper are covered. Together with the heavy weight and high shipping costs of commercial vehicles, imports and exports are below 3% of this market.

As of 2012, nine independent parent companies offered fourteen brands. General Motors (GM), Ford, and Chrysler, which owns Dodge, are American firms with large passenger segment operations

¹⁴For example, the closest commercial vehicle dealership to Cambridge, Massachusetts (as of 12/1/13) sells both Ford and International vehicles. The second closest sells Mack, Western Star, Isuzu, Volvo, Isuzu, and Peterbilt vehicles, all rivals.

¹⁵In particular, I allow the sunk costs of offering some characteristics to vary by firm. See section VII for details.

throughout the panel and collectively referred to as the “Big Three.”¹⁶ Volvo and Daimler are European firms, while Hino and Isuzu are Japanese firms. International, also known as Navistar, and Paccar are American firms without passenger segment operations. Table III reports In terms of 2012 market share, the largest brand is Ford while the largest parent firm is Freightliner, which includes the Mitsubishi-Fuso (hereafter “Fuso”) and Western Star brands.

[Table III about here.]

Some brands have changed owners, although the reasons behind these changes are plausibly exogenous to the US commercial vehicle market. European and Japanese commercial segment assemblers, as well as American passenger segment manufacturers, own subsidiaries in the domestic commercial vehicle market, although US commercial segment sales always comprise a small portion of total sales. Table IV presents 7 such changes over the panel. To illustrate with the first entry, Germany-based Daimler purchased Chrysler in 1998. The former owned the Mercedes brand while the latter owned the Dodge, Plymouth, Jeep, and Chrysler brands. At the time of the merger, only 1.6% of Chrysler’s sales were in the commercial vehicle market. It is then unlikely that this acquisition, as well as the others, were driven by concerns related to the assembly operations included in this panel.¹⁷

[Table IV about here.]

III Model

This section presents a two stage model that captures how firms endogenously adjust the set of products they offer to changing market conditions. In the first stage firms choose product offerings. In the second stage, firms set prices and consumers make purchase decisions. Figure I describes this timing.

Firms solve the problem by working backwards from the second stage: calculate the equilibrium profits that will likely accrue to them under any possible set of product offerings and then choose the products that maximize those profits. For this reason, I also begin with the second stage decisions.

¹⁶GM sells under the Chevrolet and GMC badges and although there is a distinction between these in the passenger segment, there is none in the commercial segment over my panel. I follow *The Truck Blue Book* by combining all GM models.

¹⁷The one exception is the acquisition of Renault by Volvo. Renault at the time held a controlling stake in Mack, which accounted for about one-quarter of their combined size. Still, both are based in Europe and mentions of the merger in the annual report tended to focus on the European market.

TIMING OF THE GAME.

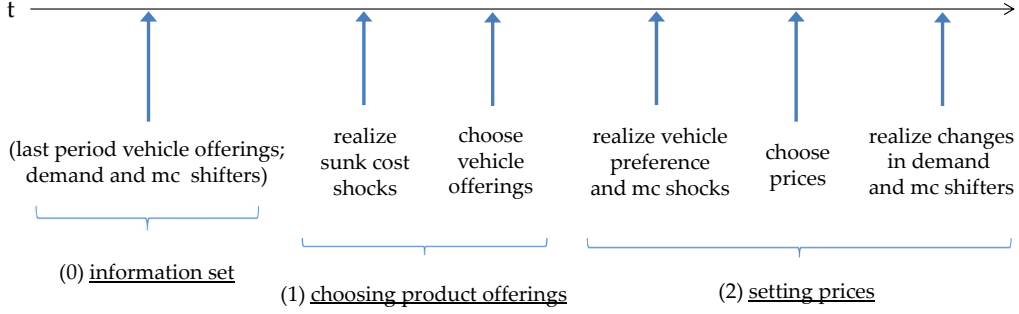


FIGURE I

Demand

Each buyer, r , decides whether to purchase a vehicle j from among J choices or the outside good so as to maximize utility. In the event of purchase, they derive utility from the vehicle based on an interaction between their attributes and the vehicle's characteristics. They also derive disutility from price. The total utility from product j is given by the following:

$$U_{r,j} = x_j (\beta_x + \beta_x^o z_r^o + \beta_x^u z_r^u) - p_j \beta_p + \xi_j + \epsilon_{r,j} \quad (1)$$

x_j denotes the vector of vehicle characteristics, excluding price. These include a constant, the gross weight rating, and dummies for the cab types and options, as well as an interaction between the gross weight rating and cab-over-engine ("cabover"). z_r^o and z_r^u denote these buyer attributes, which can be observed or unobserved by the econometrician. β_x denotes the mean taste for each product characteristic, while β_x^o and β_x^u are coefficients on the interaction of buyer attributes and product characteristics. β_p denotes the consumers distaste for price, p_j . For convenience, let β denote the vector of taste parameters, $(\beta_x, \beta_x^o, \beta_x^u, \beta_p)$, and z denote the vector of both unobservable and observable consumer attributes, (z^o, z^u) . ξ_j denotes a product-specific preference shock while $\epsilon_{r,j}$ denotes a preference shock specific to the choice and buyer. Buyers can also consume the outside good, whose mean utility is normalized to zero so that buyers receive only $\epsilon_{r,0}$. The distribution of ξ is only restricted to be i.i.d. across products. ϵ is distributed extreme value and is i.i.d. across products and buyers.

This specification yields the familiar logit choice probabilities for each consumer. After integrating out over the total number of simulated consumers, ns , I arrive at the market share for any product j :

$$s_j = \frac{1}{ns} \sum_r s_{r,j} = \frac{1}{ns} \sum_r \left(\frac{e^{x_j(\beta_x + \beta_x^o z_r^o + \beta_x^u z_r^u) - p_j \beta_p + \xi_j}}{1 + \sum_k e^{x_k(\beta_x + \beta_x^o z_r^o + \beta_x^u z_r^u) - p_k \beta_p + \xi_k}} \right) \quad (2)$$

Adding a time t subscript,¹⁸ we have that the product of market share and market size, M_t , yields total unit sales, $q_{j,t}$.

This setup assumes that static, unit demand closely approximates the actual purchasing decisions made and that buyers are price takers. In practice, many buyers do in fact own “fleets” of vehicles, although in most cases they are purchasing only one or two vehicles at a time.

Pricing

The second stage decision from the firm’s perspective is to set prices. Firms, f , offering a set of products $J_{f,t}$, choose prices to maximize profits, given by:

$$\Pi_{f,t} = \sum_{j \in J(f)} [p_{j,t} - mc_{j,t}] s(x_{j,t}, x_{-j,t}, p_t, z_t; \beta, \xi_t, mc_{j,t}) M_t \quad (3)$$

where $mc_{j,t}$ denotes the marginal costs of producing j at t .

This requires a first order condition of the profit function, rearranging terms, and taking other firms’ prices as fixed to arrive at Nash equilibrium prices given by:

$$p_{j,t}^* = mc_{j,t} - \frac{s_{j,t}}{\beta_p} \left[s_{j,t} - \frac{1}{ns} \sum_r \sum_{k \in J(f)} s_{r,j,t} s_{r,k,t} \right]^{-1} \quad (4)$$

Marginal costs are a parametric function of observable product characteristics, wages, time, parameter γ , and an unobserved factor specific to the product and time. That is, $mc_{j,t} = mc(x_{j,t}, w_{j,t}, t; \gamma, \omega_{j,t})$. For now, I will remain agnostic as to the functional form; later on I will show that since the demand parameters are recovered without information from the supply side, I can be somewhat non-parametric with respect to the exact relationship of marginal costs and its determinants.

Product Offerings

In the first period, firms choose product offerings, i.e. make model-level entry and exit decisions, with the understanding that their actions and their rivals’ actions will impact the second stage. They do not know ξ_t and ω_t but do know the distribution of the disturbances, (F_ξ, F_ω) , so they form an expectation over them to compute the hypothetical expected profits from any set of production offerings. Then the expected “variable” profits are

$$\pi(J_{f,t}, J_{-f,t}, z_t, w_t, t, p^*; \beta, \gamma, F_\xi, F_\omega) \equiv \int_{\xi, \omega'} \Pi_f(J_t, z_t, w_t, t, p^*; \beta, \gamma, \xi', \omega') dF_{\xi'} dF_{\omega'} \quad (5)$$

¹⁸The applied setting below considers annual decisions, so I use t and “year” interchangeably.

The principal decision firms face in the first stage is to weigh the added profits of introducing or continuing to offer existing product models against the sunk costs of doing so. To proceed, I need to take a stand on the nature of sunk costs,¹⁹ given below.

Sunk Entry/Exit Costs Assumption (I).

$$SC_{f,j,t} = x_j \tilde{\theta}_{f,x(j),t} \times [\{j \in J_{f,t}, j \notin J_{f,t-1}\} + \{j \notin J_{f,t}, j \in J_{f,t-1}\} \times \frac{1}{\lambda}]$$

The first braces term is an indicator function for products offered this year but not last year, whereas the second braces term is an indicator function for products not offered this year but offered last year. There are two features of sunk costs. First, they are linear in the observable product characteristics, although at this point can freely vary by product space, time, and firm. Second, the sunk cost of adding a model is a multiplicative scaling by λ of the sunk cost of retiring one.

Sunk costs can induce forward-looking behavior, but in differentiated product markets, the dynamic solution requires storage of and an expectation over billions or more of states. This computational burden is orders of magnitude too hard for firms in practice. A common rebuttal to this observation is that although agents do not appear to explicitly optimize, repeated play can nonetheless converge to equilibrium strategies²⁰—although learning is also out of the question here. When the product space is rich, fifty years of annual—or even daily decisions—would not provide anywhere close to meaningful convergence over the state space, even if only a recurrent class of states are considered.

This raises an important question as to what managers actually do. Survey data suggests managers cut computational corners.²¹ Solutions to these problems suggested by practitioners or research staff in private sector firms often suggest the same. For example, Jeff Alden, group manager of Manufacturing Systems Research at General Motors, wrote in *Operations Research*, “By far the most common planning procedure found in practice is to approximate the solution” (1992).²² To figure out as accurately as possible what is done in this setting, I interviewed engineers, designers, and veteran managers. A common thread ran through these interviews, the clearest of which was given by the former head of General Motors Commercial Division, who said:

“Each year we look at demand, what we offer, and what the competition is going to offer. We consider changing the lineup like adding a vehicle... We know who the customers would be, what we can charge, and the production costs—so we have the added margin.

¹⁹With sufficiently much data, one could imagine semi-parametric or even non-parametric identification of sunk costs along the lines of Matzkin (1992). This would eliminate the need for this assumption, but requires more data than is available in this setting.

²⁰Fudenberg and Levine (1998) explore how agents can learn through repeated play. Fershtman and Pakes (2012) provide an empirical model where simple updating converges to equilibrium strategies, and Doraszelski, Lewis, and Pakes (2014) estimate a related model in the context of electricity markets.

²¹Graham and Harvey (2002) found that CFOs are 2-3 times more likely to use contingent-free methods like the payback ratios and hurdle rates. With respect to adjusting for risk, Summers (1987) showed that 94% of surveyed Fortune 500 firms use the same discount rate across all projects, and that 23% used a discount rate above 19%.

²²This article provides a simplified single-agent problem and shows how “rolling horizon procedures” approximate the solution.

The margin over the investment gives a return on capital, and *we'll build it when it crosses some threshold* (emphasis added).”

Hurdle rates are interpreted here in the following way.

Capital Budgeting Assumption (II).

$$\{j \in J_{f,t}, j \notin J_{f,t-1}\} \Leftrightarrow \mathcal{E} \left[\frac{\pi(J_{f,t}, J_{-f,t}) - \pi(J_{f,t} \setminus j, J_{-f,t})}{x_j \theta_{f,t,x}} \mid \mathcal{I}_{f,t} \right] \geq \text{HurdleRate}$$

$$\{j \in J_{f,t}, j \in J_{f,t-1}\} \Leftrightarrow \mathcal{E} \left[\frac{\pi(J_{f,t}; J_{-f,t}) - \pi(J_{f,t} \setminus j, J_{-f,t})}{x_j \theta_{f,t,x} \times \frac{1}{\lambda}} \mid \mathcal{I}_{f,t} \right] \geq \text{HurdleRate}$$

The first inequality refers to products not offered by f at $t - 1$ while the second refers to products that are.

Hurdle rates are a straightforward rule-of-thumb. To illustrate, consider some product j that is not offered by f at $t - 1$. Introducing it at t would increase expected second-stage profits by \$20 and require a \$100 sunk costs. This action yields a 20% expected (static) return. Firms with a 19% hurdle rate would accept while firms with a 21% hurdle rate would not.

They also may capture a large share of what would be earned under more complex, fully dynamic strategies. In seeking to explain why “most firms do not make explicit use of real option techniques” and “projects are taken based on whether or not IRRs exceed arbitrarily high discount rates,” McDonald (2000) shows that at least in single-agent settings, hurdle rates do quite well. The reason is that because as option value increases, for example due to an increase in volatility, deviations from the optimal strategy are less costly. These hurdle rates are larger than the firms discount rate since these need to capture option value. Second, recall that the profits of adding products in any sub-segment of the market, or the market overall, are not predictably growing or shrinking over time. This presumably mitigates the impact of deterrence, since firms are not looking to move early and foreclose that action from rivals in the future. A companion paper to this one, Sweeney and Wollmann (2014), provide simulations that support the use of hurdle rates in environments like this one by comparing two-period strategies to multi-period dynamic strategies. Finally, I provide calculations showing that what firms actually earned is quite close to and centered around what they expected to earn.

IV Estimation

Demand

The estimation of demand closely resembles Petrin (2002). It minimizes a generalized method of moments objective function based on two sets of moments.²³ The first are constructed as follows. Firms do not know ξ_t when they choose product characteristics, so $\mathbb{E}[\xi_{j,t}|x_{j,t}] = 0$. On the other hand, firms do know ξ_t when they choose prices. Approximations to the optimal instruments are constructed from variables that shift either marginal costs or markups. Wages are a valid instrument in the former case. I match vehicle models to the areas in which they were assembled and proxy for the factory wage rate with a Bureau of Labor Statistics estimate of the production wage in that area. Production locations are unlikely to relate to current commercial vehicle market conditions since the decision to launch a new facility precedes production by several years and, once launched, the mix of products are rarely re-allocated across factories. The competitive conditions provides a valid instrument in the latter case. The markup on j produced by f is decreasing in the number of competing products that are close in characteristic space but increasing in the proportion of these owned by f . The timing of choices, again, guarantees that $\mathbb{E}[\xi_{j,t}|x_{-f,t}] = 0$ and $\mathbb{E}[\xi_{-j,f,t}|x_{j,t}] = 0$.²⁴ Observed shares are calculated by simply dividing the units sold by market size, the calculation of which is given in the Appendix.

The second set matches first and second moments from the microdata to model-predicted analogs of these moments. Specifically, I choose three subsets of moments to match for each buyer-product relationship: the mean and variance of the buyer attribute conditional on the product characteristic and the probability of purchase conditional on that buyer type. Formally these moments are

$$\begin{aligned} mm_1 &= \mathbb{E}(z_r | r \text{ buys } j) \\ mm_2 &= \mathbb{E}\left[(z_r)^2 - \mathbb{E}(z_r | r \text{ buys } j) | r \text{ buys } j\right] \\ mm_3 &= \mathbb{E}(z_r | r \text{ buys any } j) \end{aligned}$$

The micro-moments include buyer-product relationships between the following: each industry type and the GWR; each industry type and a constant; the delivery industry and the compact-front-end cab type; the general freight industry and the cabover; the general freight industry interacted with length laws and the cabover; the urban measure and the cabover.

²³In practice, I exploited a technique to speed up the parameter search given in Varadhan and Roland (2008). I then checked the values using the standard contracting mapping proposed in Berry, Levinsohn, and Pakes (1995).

²⁴Armstrong (2013) makes a strong case for including cost shifters, like production wages, in demand estimation. He shows that as the number of products grows large, markups converge to a constant. This leaves a model with only “BLP instruments,” i.e. those based on competing and own product characteristics, unidentified in the limit.

Marginal Costs

With demand estimates in hand, only the Nash pricing condition is needed to back out marginal costs. This merely requires rearranging the pricing equation given in the previous section:

$$\hat{m}c_{j,t} = p_{j,t} + \frac{s_{j,t}}{\hat{\beta}_p} \left[s_{j,t} - \frac{1}{ns} \sum_r \sum_{k \in J(f)} \hat{s}_{r,j,t} \hat{s}_{r,k,t} \right]^{-1} \quad (6)$$

Notice that $\hat{m}c_{j,t} = mc(p_{j,t}, s_{j,t}, \hat{\beta}_p, s_{r,j,t})$. That is, estimated marginal costs are a function of quantities available to the econometrician: prices and shares, which are observed in the data, as well as the price coefficient and the individual purchase probabilities, which are recovered from the demand system.

Much of the prior work has used this equation with an explicit functional form for marginal costs to add supply-side moments to the demand estimation. This is particularly helpful in pinning down the price coefficient, which can often be difficult to instrument for outside data. In contrast, I estimate demand without these assumptions and then confirm the implied price-cost margins are in line with what we see in audited financial data and that the elasticities are sensible. Presumably, rich microdata and an observable marginal cost shifter for marginal costs (wages) help a great deal here. The payoff lies in the fact that we can now analyze rather than assume the marginal cost shape with respect to the right-hand side variables. With enough data, one could be non-parametric here or provide a formal shape test. In practice, I provide graphical evidence suggesting that the log of marginal costs is linear in the continuous product characteristic. That is, I have that:

$$\ln(mc_{j,t}) = [x_{j,t}, w_{j,t}, t] \gamma + \omega_{j,t} \quad (7)$$

Rearranging terms and solving for the additively separable error term provides:

$$\omega_{j,t} = \exp \left(p_{j,t} - \frac{s_{j,t}}{\hat{\beta}_p} \left[s_{j,t} - \frac{1}{ns} \sum_r \sum_{k \in J_f} \hat{s}_{r,j,t} \hat{s}_{r,k,t} \right]^{-1} \right) - [x_{j,t}, w_{j,t}, t] \hat{\gamma} \quad (8)$$

γ is estimated via ordinary least squares or a weighted least squares, since $\omega_{j,t}$ is not known when the firms choose which products to offer. That is, we have $\mathbb{E}[\omega_{j,t} | x_{j,t}, w_{j,t}, t] = 0$. The last step is to plug back in for $\hat{\beta}$ and $\hat{\gamma}$ and recover an empirical distribution of ξ and ω , which provide \hat{F}_ξ and \hat{F}_ω . Together, these provide unbiased (but potentially measured with error) estimates of the second-stage payoffs firms would expect from offering any alternate set of products, $\pi(J_{f,t}, J_{-f,t}, z, w; \hat{\beta}, \hat{\gamma}, \hat{F}_\xi, \hat{F}_\omega)$.

Sunk Costs

Setup.

The estimation of fixed costs follows the logic of revealed preference. Firms were free to offer any alternative set of products to the ones that appear in the data (i.e. those that were chosen in equilibrium) but did not because these alternatives were less profitable. These alternatives provide intuitive upper and lower bounds on the parameters of interest: fixed costs could not be too low—or else firms would offer more products than appear in the data—and could not be too high—or else firms would offer less products than appear in the data. As a consequence of simultaneous moves, the necessary conditions for the Nash equilibrium provide that firms take rivals' decisions as fixed. Any unilateral deviation should be less profitable in expectation for the firm than the chosen product offerings.

To arrive at inequalities that are linear in observed and estimated quantities, the parameters of interest, and a set of disturbances, some algebra is necessary. Rewrite the negative of the initial sunk cost term multiplied by the hurdle rate, $-\tilde{\theta}_{f,x,t} \times \text{HurdleRate}$, as $\theta_{f,x,t}$. If the hurdle rate is unaffected by the policy change, then these terms need not be separately identified.²⁵ Henceforth, refer to $\theta_{f,x,t}$ as the “sunk cost” for convenience, rather than the term given in Assumptions I and II. Write these sunk costs as a mean and a deviation away from this mean which is firm, time, and product space specific. That is, $\theta_{f,t,x}^k = \theta^k + \nu_{f,t,x}^k$, where k refers to either the constant, gross weight rating, or dummy variable for cab-over-engine, compact-front-end, or long-option. Also re-write the expected profits based on the true parameters as the expected profits based on the econometrician's estimate plus an error, so that $\Delta\pi(\cdot; \beta, \gamma, F_\xi, F_\omega) = \Delta\pi(\cdot; \hat{\beta}, \hat{\gamma}, \hat{F}_\xi, \hat{F}_\omega) + \nu^\pi = \Delta\hat{\pi}(\cdot) + \nu^\pi$. Last, notice that any comparison of the equilibrium offerings, denoted $J_{f,t}$, to an alternative set, $J'_{f,t}$, provides an inequality. Three such deviations are particularly helpful: dropping a model, adding a model, and substituting one model for another. These yield

$$\mathcal{E} \left[\Delta\hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \setminus j; J_{-f,t}}^\pi - x_{j'} \left(\theta + \nu_{f,t,x_j} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) \middle| \mathcal{I}_{f,t} \right] \geq 0 \quad (9)$$

$$\mathcal{E} \left[\Delta\hat{\pi}(J_{f,t}, J_{f,t} \cup j', J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \cup j'; J_{-f,t}}^\pi + x_{j'} \left(\theta + \nu_{f,t,x_{j'}} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) \middle| \mathcal{I}_{f,t} \right] \geq 0 \quad (10)$$

$$\begin{aligned} \mathcal{E} \left[\Delta\hat{\pi}(J_{f,t}, J_{f,t} \setminus j \cup j', J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \setminus j \cup j'; J_{-f,t}}^\pi + x_j \left(\theta + \nu_{f,t,x_j} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) \right. \\ \left. - x_{j'} \left(\theta + \nu_{f,t,x_{j'}} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) \middle| \mathcal{I}_{f,t} \right] \geq 0 \end{aligned} \quad (11)$$

²⁵Firms, at least anecdotally, rarely change hurdle rates. For example, *The Economist* reports that “Shell [Oil Company] left its hurdle rates unchanged for two decades until it ‘nudged them down’ in 1997, and now intends to keep them at present levels for years to come.” That said, the failure of one or more major competitors could have a non-trivial impact on the risk, real or perceived, of operating in this market. This would increase the cost of capital and, in turn, the hurdle rate. The intuition for why this affects the analysis is that an increase in risk makes firms value the future less; future profits are less valuable relative to sunk costs paid in the current period, and entry is less likely.

It is easy to see how Equations (9)-(11) can provide bounds on the sunk cost parameters. The inequalities directly reflect the tradeoffs between changes in profits and changes in sunk costs in any product offering decision. Notice here that if the econometrician assumed all disturbance terms are zero, then it is sufficient to merely solve the system of linear inequalities that bound (θ, λ) . This assumes that the econometrician has the same information that the agents have as well as that agents have perfect information about what they will earn, conditional on their choices. This is both an unreasonable assumption and almost certainly not able to rationalize the data. Notice also that if the econometrician assumed all disturbance terms are unknown to the agents when decisions are made, then it is sufficient to merely minimize the violations of these inequalities. This is still too strong an assumption but will rationalize the data. Relaxing it, however, introduces a classic endogeneity problem: firms enter products when it is less costly to do so. The discussion below addresses this problem.

Disturbances.

Three assumptions will identify the sunk costs, which currently can freely vary across product space, firms, and years. The first assumption is that sunk cost disturbances are independently and identically distribution over product space and time.

Product and Time Disturbances (Assumption III).

$\nu_{f,t,x}$ is i.i.d. over x and t .

Independence over product space implies that we can re-write as the k -characteristic specific portion of disturbance $\nu_{f,t,x}$ as $\nu_{f,t}^k$, where $k \in \{con, GWR, COE, CFE, Long\}$. That is, conditional on (f, t) , knowing $\nu_{f,t}^k$ tells us nothing about $\nu_{f,t}^{k'}$ for $k \neq k'$.

The second assumption allows for observable heterogeneity in the firm specific portion of the sunk cost disturbances that are based on two important features of the commercial vehicle market. First, congestion and stringent length regulation have made the cabover vehicle ubiquitous in Asia, which may affect brands with their headquarters based in Japan. For example, they may find it cheaper to introduce cabover vehicles, find it more expensive to introduce non-cabover vehicles, or both. For this reason, the constant term and cabover term are allowed to be different for Japan-based brands. Second, the Big Three firms have large assembly operations in the passenger vehicle segment, which are lighter than commercial vehicles. Hence they may have an advantage in introducing light vehicles but a disadvantage in producing heavy vehicles, or both. Because of this, the constant term and GWR term are allowed to be different for Big Three brands. After accounting for these differences, however, I assume the remaining firm specific portion of the sunk cost disturbances is not known by the firms when they make their decisions. For notational ease, denote the portion of disturbances that are and are not in the agents' information sets when choices are made as $\nu_{2,\cdot}$ and $\nu_{1,\cdot}$, respectively.

Firm Disturbances (Assumption IV).

For the constant term, $\nu_{f,x,t}^{con} = \theta_{Big3}^{con} + \theta_{Japan}^{con} + \nu_{2,t,x}^{con} + \nu_{1,f,t,x}^{con}$. For gross weight rating,

$\nu_{f,x,t}^{GWR} = \theta_{Big3}^{GWR} + \nu_{2,t,x}^{GWR} + \nu_{1,f,t,x}^{GWR}$. For the cab-over-engine dummy variable, $\nu_{f,x,t}^{COE} = \theta_{Japan}^{COE} + \nu_{2,t,x}^{COE} + \nu_{1,f,t,x}^{COE}$. For $k \in \{CFE, Long\}$, $\nu_{f,x,t}^k = \nu_{2,t,x}^k + \nu_{1,f,t,x}^k$.

The third assumption states that profit disturbances are not known to the firms when they make their decisions.

Profit Disturbances (Assumption V).

$$\nu_{J_{f,t}, J'_{f,t}; J_{-f,t}}^\pi = \nu_{1, J_{f,t}, J'_{f,t}; J_{-f,t}}^\pi$$

The final assumption provides sufficient variation in the instruments to achieve a bound on each side of each parameter.

Disturbance Support (Assumption VI).

Let $z^{g,c} \subset Z$ be the set of demand shifters that comprise the largest share of demand for vehicles with GWR g and cab type c . $(\nu_{2,t}^{con}, \nu_{2,t}^{GWR}, \nu_{2,t}^{COE}, \nu_{2,t}^{CFE}, \nu_{2,t}^{Long})$ is bounded such that if $z_t^{g,c} - z_{t-1}^{g,c} = \operatorname{argmax}\{z_t^{g,c} - z_{t-1}^{g,c}\}$, then for all \tilde{J} where $j(g, c) \notin \tilde{J} \max_f \left\{ \mathcal{E} \left[\pi(\tilde{J}, \tilde{J} \cup j(g, c), J_{-f,t}) - \theta^{con} - \nu_{2,t}^{con} - \{c = k\}(\theta^c + \nu_{2,t}^c) - g(\theta^{GWR} + \nu_{2,t}^{GWR}) | \mathcal{I}_{f,t} \right] \right\} \leq 0$.

Taken together, these assumptions provide that the characteristic-specific portion of the sunk costs that vary over time to be observable by firms but not the econometrician. To illustrate, the cabover may be particularly expensive to introduce at time t , causing firms to add fewer cabovers to their product lineup in this period relative to other cab types. Formally, the product offering decisions, $J_{f,t}$, are selected on the ν_2 terms, and this will bias estimates if it were ignored. On the other hand, the assumptions provide that knowing the cabover-specific sunk cost disturbance term at t reveals nothing about the disturbance term specific to other characteristics at t or to the cabover at any time other than t . These assumptions also rationalize the data.²⁶

²⁶The ν_1 terms rationalize the data, i.e. ensure the model does not over-fit. This implies firms are surprised by the net profits they receive, and occasionally have ex post regret (but overall are still right on average). Technically, this would be satisfied if firms choose to offer $J_{f,t}^*$ but instead receive $J_{f,t}$, which the econometrician observes. For example, upon entering a product j with $x_j^{GWR^*}$ GWR, it receives $x_j^{GWR} = x_j^{GWR^*} + \text{error}_{j,f,t}$. A more detailed discussion that relates disturbance assumptions to the underlying data generating process is found in Pakes (2010).

Substituting in to the prior three inequalities with the assumptions yields

$$\begin{aligned}
& \mathcal{E} \left[\Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \setminus j; J_{-f,t}}^{\pi} | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[\left(\theta^{con} + \theta_{Big3}^{con} \{Big3\} + \theta_{Japan}^{con} \{Japan\} + \nu_{2,t,x_j}^{con} + \nu_{1,f,t,x_j}^{con} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[x_j^{GWR} \left(\theta^{GWR} + \theta_{Big3}^{GWR} \{Big3\} + \nu_{2,t,x_j}^{GWR} + \nu_{1,f,t,x_j}^{GWR} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[x_j^{COE} \left(\theta^{COE} + \theta_{Japan}^{COE} \{Japan\} + \nu_{2,t,x_j}^{COE} + \nu_{1,f,t,x_j}^{COE} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[x_j^{CFE} \left(\theta^{CFE} + \nu_{2,t,x_j}^{CFE} + \nu_{1,f,t,x_j}^{CFE} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[x_j^{Long} \left(\theta^{Long} + \nu_{2,t,x_j}^{Long} + \nu_{1,f,t,x_j}^{Long} \right) \left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \geq 0 \quad (12)
\end{aligned}$$

$$\begin{aligned}
& \mathcal{E} \left[\Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j', J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \cup j'; J_{-f,t}}^{\pi} | \mathcal{I}_{f,t} \right] \\
& - \mathcal{E} \left[\left(\theta^{con} + \theta_{Big3}^{con} \{Big3\} + \theta_{Japan}^{con} \{Japan\} + \nu_{2,t,x_{j'}}^{con} + \nu_{1,f,t,x_{j'}}^{con} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& - \mathcal{E} \left[x_{j'}^{GWR} \left(\theta^{GWR} + \theta_{Big3}^{GWR} \{Big3\} + \nu_{2,t,x_{j'}}^{GWR} + \nu_{1,f,t,x_{j'}}^{GWR} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& - \mathcal{E} \left[x_{j'}^{COE} \left(\theta^{COE} + \theta_{Japan}^{COE} \{Japan\} + \nu_{2,t,x_{j'}}^{COE} + \nu_{1,f,t,x_{j'}}^{COE} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& - \mathcal{E} \left[x_{j'}^{CFE} \left(\theta^{CFE} + \nu_{2,t,x_{j'}}^{CFE} + \nu_{1,f,t,x_{j'}}^{CFE} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \\
& - \mathcal{E} \left[x_{j'}^{Long} \left(\theta^{Long} + \nu_{2,t,x_{j'}}^{Long} + \nu_{1,f,t,x_{j'}}^{Long} \right) \left(\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) | \mathcal{I}_{f,t} \right] \geq 0 \quad (13)
\end{aligned}$$

$$\begin{aligned}
& \mathcal{E} \left[\Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j \cup j', J_{-f,t}) + \nu_{J_{f,t}, J_{f,t} \setminus j \cup j'; J_{-f,t}}^{\pi} | \mathcal{I}_{f,t} \right] \\
& + \mathcal{E} \left[\left(\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\} - \{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\} \right) \left(\theta^{con} + \theta_{Big3}^{con} \{Big3\} + \theta_{Japan}^{con} \{Japan\} + \nu_{2,t,x_j}^{con} - \nu_{1,f,t,x_{j'}}^{con} \right) \right] \\
& + \left(x_j^{GWR} (\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\}) - x_{j'}^{GWR} (\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\}) \right) \left(\theta^{GWR} + \theta_{Big3}^{GWR} \{Big3\} + \nu_{2,t,x_j}^{GWR} + \nu_{1,f,t,x_{j'}}^{GWR} \right) \\
& + \left(x_j^{COE} (\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\}) - x_{j'}^{COE} (\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\}) \right) \left(\theta^{COE} + \theta_{Japan}^{COE} \{Japan\} + \nu_{2,t,x_j}^{COE} + \nu_{1,f,t,x_{j'}}^{COE} \right) \\
& + \left(x_j^{CFE} (\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\}) - x_{j'}^{CFE} (\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\}) \right) \left(\theta^{CFE} + \nu_{2,t,x_j}^{CFE} + \nu_{1,f,t,x_{j'}}^{CFE} \right) \\
& + \left(x_j^{Long} (\{j \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j \in J_{f,t-1}\}) - x_{j'}^{Long} (\{j' \notin J_{f,t-1}\} + \frac{1}{\lambda} \{j' \in J_{f,t-1}\}) \right) \left(\theta^{Long} + \nu_{2,t,x_j}^{Long} + \nu_{1,f,t,x_{j'}}^{Long} \right) \Big| \mathcal{I}_{f,t} \Big] \geq 0 \quad (14)
\end{aligned}$$

Identification.

Equations (11)-(13) are now additively separable in quantities available to the econometrician, the disturbance terms, and the parameters of interest. Together with sufficient variation in a set of instruments discussed below, these equations provide upper and lower bounds on θ and λ . Formally, they yield the following proposition.

Proposition I. If Assumption I-IV hold, sunk costs are identified.

Proof. See Appendix.

The argument is summarized here. The first step is to bound (θ^{con}, λ) , conditional on θ^{GWR} . Construct four inequalities based on any firm at any time taking the following actions with respect to any GWR conventional cab vehicle: one more product than it did, add one less, drop one more, and drop one less. (Note that the conventional cab is chosen only as an example and because it is the cab type represented by the constant term.) Suppose each action was a possibility. This would create one inequality, or bound, on either side of the parameters for each firm, time, and GWR value. An expectation over those inequalities would average the additively separable disturbance terms across the data to their mean. Since the disturbances are independent of past offerings and independently distributed over time, they average out to their unconditional mean, which is by construction zero. This requires that each action was possible, a point revised below. The next step is to bound θ^{GWR} , conditional on (θ^{con}, λ) . Construct two inequalities based on the fact that any conventional cab vehicle that was added could have been designed with either one more unit or less unit of GWR. Suppose each firm in each period enters at least one conventional cab vehicle for each GWR. Once again, the additively separable disturbance terms average out to their unconditional mean, which are zero.

In the data, however, all firms do not add, drop, or even offer each GWR cab in each period. At least in the case of the conventional cab, however, one or more firms always offer at least one of each GWR. This means that in all t , at least one firm could drop one more product of each GWR. Since the part of sunk costs observable to the firm but not the econometrician does not vary by firm, a weighting scheme can be devised to average out the disturbance terms to zero. Nonetheless, there are still periods when *no firm* adds or drops a given GWR conventional cab vehicle. To omit these periods from the expectation would be to select on periods where $\nu_{2,t}$ are particularly high, preserving the classic endogeneity/selection problem. To be clear, what precisely is needed is an instrument that selects periods where a given vehicle type is almost surely added by some firm without selecting on the disturbance terms. In fact, the demand shifters satisfy *exactly* this condition. Recall that the microdata reveal that demand for each vehicle type can be tied to a set of exogenous demand shifter in Z . To condition on sufficiently large changes in this subset of shifters is to practically guarantee entry of this product type in this period. Similarly, to condition on a sufficiently large fall in is to practice guarantee product exit of this product type in this period. However, selecting on Z , or subsets or functions thereof, does not affect the distribution of the disturbances. Ultimately, this identifies $(\theta^{con}, \theta^{GWR}, \lambda)$. An analogous routine identifies the remaining parameters.

The objective function is a sum of squared violations of the inequality conditions, weighted by the inverse of the moments. Formally, let

$$Q_n(\theta) = \sum_k (\sigma_k(W_i, \theta)^{-1} (m_k(W_i, \theta))_-)^2 \quad (15)$$

where $(\cdot)_-$ denotes the negative portion of the quantity within the parenthesis, $m_k(W_i, \theta)$ denotes

moment k for data W_i , and $\sigma_k(W_i, \theta)$ represents the square root of the variance of that moment. A moment $m_k(W_i, \theta)$ is the expectation of the interaction of instrument k , given by $h_k(Z)$ and the left-hand side of the inequalities. Q is common in recent empirical and econometric work. Standard efficiency concerns suggest over-weighting the most informative moments by scaling with inverse variance. These do not, of course, impact the bias of the estimator, since $\mathbb{E}[m_k(W_i, \theta)] \geq 0$ is equivalent to $\mathbb{E}[\zeta_k m_k(W_i, \theta)] \geq 0$ as long as the scaling parameter ζ does not vary within moment k . It can also be seen as analogous to a generalized method of moments estimator where the correlation between the moments is ignored. The parameter search then merely satisfies $\hat{\Theta} = \arg \min_{\theta \in \Theta} \{Q_n(\theta)\}$.

Inference.

I construct sets in which the true sunk cost parameters will lie 95% of the time. Inference based on inequalities is somewhat less straightforward than inference based on equalities—e.g. generalized method of moments—because inequalities provide only one-sided restrictions. The most informative of these is the *least upper bound* and the *greatest lower bound*. As such, these bounds represent a minimum and maximum, respectively, of the moments, rather than an average. This rules out the use of the central limit theory to provide a direct formula for standard errors.

I follow Andrews and Soares (2010) in constructing these sets but with shifted rather than selected moments. The process is summarized below. For details, I refer to their paper (and for readers uninterested in this section, it can be skipped without a loss in continuity). They suggest inverting a test, as in Chernozhukov, Hong, and Tamer (2007). In short, begin with any suitable parameter guess. Compute the objective function. Find the mean and variance of the sample moments, which are normally distributed. Bootstrap from their underlying empirical distribution, computing the objective function each time. Accept points for which $Q_n(\theta)$ falls below the bootstrapped objective function evaluations 95% of the time. Otherwise, reject.

Andrews and Soares (2010) show this method by itself can have poor power properties and, in my calculations, it tended to produce wide confidence sets. The problem stems from cases where some moments are satisfied by a wide margin. Loosely speaking, these moments will cause the econometrician to accept any non-perverse parameter guesses. To illustrate, suppose that all but one moment pertaining to the upper bound of the sunk cost parameter is close to binding while the last moment is very slack. That is, all but one moment provides a moment close to $y - \bar{\theta}^{Enter} \geq 0$ but the last provides $1000 \times y - \bar{\theta}^{Enter} \geq 0$. The last moment is not very informative, and the lower bound should be near y . Yet, the last moment ensures that $Q(\theta)$ will be below the bootstrapped objective function values for most θ below $1000 \times y$.

One solution is to remove such moments from the estimation process, letting the data guide the selection. I begin with the inequality Equations (11)-(12). In particular, they suggest selecting moments where $\zeta_{n,k}(\theta) > 1$ where $\zeta_{n,k}(\theta) = \kappa_n^{-1} n^{1/2} \hat{D}_{n,k}^{1/2}(\theta) \bar{m}_{n,k}(\theta)$. $\zeta_{n,k}$ is the slackness measure of moment k computed over n observations. $\bar{m}_{n,k}(\theta)$ is sample moment k , and $\hat{D}_{n,k}(\theta)$ comprises the diagonal elements of the sample variance of the moments, i.e. $\hat{D}_n(\theta) = \text{Diag}(\hat{\Sigma}_n(\theta))$ with

$\hat{\Sigma}_n = n^{-1} \sum_i (m(W_i, \theta) - \bar{m}(W_i, \theta)) (m(W_i, \theta) - \bar{m}(W_i, \theta))'$. They suggest $\kappa = (\ln(n))^{1/2}$.²⁷ The intuition for this selection is to disregard moments that are sufficiently slack as measured by their standardized distance away from zero.

I test the null hypothesis for all possible values in the parameter space, Θ . To speed up the search, a large and sparse grid was started with and then iteratively made smaller and more granular. This was essential. The parameter space covers six dimensions, so even twenty points per parameter equates to nearly 11,400,000 value calculations. Each of these has a bootstrap of size 1,000, although the moments are linear so the calculations are fast. Some quantities can be precomputed and, of course, the computation can be run in parallel.²⁸

Unlike Andrews and Soares (2010), I shift rather than select the moments (as in Pakes, Porter, Ho, and Ishii (2014)). This requires adding $\zeta_{n,k}(\theta)$ back to the sample moments under both the null and alternative. For each bootstrap sample, I evaluate Q and construct a critical value $\hat{c}_n(\theta_0, 1 - \alpha)$ where $\alpha = 5\%$. I accept if the computed test statistic, $T_n(\theta) = Q(n^{1/2}\bar{m}_n(\theta), \hat{\Sigma}_n(\theta))$ is less than this critical value and reject otherwise. The true parameter vector θ_0 will lie in the space defined by the intersection of the estimated values of $[\underline{\theta}, \bar{\theta}]$ and $[\underline{\lambda}, \bar{\lambda}]$.

V Descriptive Evidence

Purchasing patterns from the microdata presented in Section III have already demonstrated strongly heterogeneous preferences among the buyers. What remains to be seen is whether firms are actually adjusting their product characteristics to changes in buyer composition. This would provide support for the notion that firms are solving a problem close in nature to the one presented in Section IV, but would also suggest a set of instruments to identify sunk costs. To illustrate, suppose the construction industry expands rapidly, as it did between 2005 and 2007, and then contracts steeply, as it did in the period subsequent to that. The microdata suggests this group prefers only a small subset of the vehicles produced, so firms adapting to market conditions should expand and contract product offerings tailored to their needs over the same periods. For estimation purposes, it is clear that the introduction of, say, four vehicles of this sort will imply lower sunk costs than three but higher sunk costs than five. Algebraically, the expansion in the construction industry can be presented by a change from z to z' . If j is a vehicle that is highly preferred by construction-related industry buyers, then $\pi(J_{f,t} \cup j, J_{f,t}, J_{-f,t}, z', \cdot) > \pi(J_{f,t} \cup j, J_{f,t}, J_{-f,t}, z, \cdot)$. If the difference between the left- and right-hand side expressions is large enough, f or one of its rivals is likely to enter.

²⁷For notational comparability, note that I use ζ and Q in place of ξ and $S_{(2)}$, which appear in Andrews and Soares (2010).

²⁸This search was run in MATLAB using the combined efforts of two Dell Precision 7500 terminals, each with two Intel Xeon 3.47GHz processors and 56GB of RAM. These are powerful machines (as of 2014). The alternative was to run these in parallel across a high powered cluster. Since the inference procedure consists only of affine transforms and sorts of the data, I preferred to run sub-spaces of Θ in a loop, with the data organized over observations, moments, and the bootstrap draws.

Encouraging reduced-form evidence that directly relates the instruments and outcomes is presented in Figures II-IV.

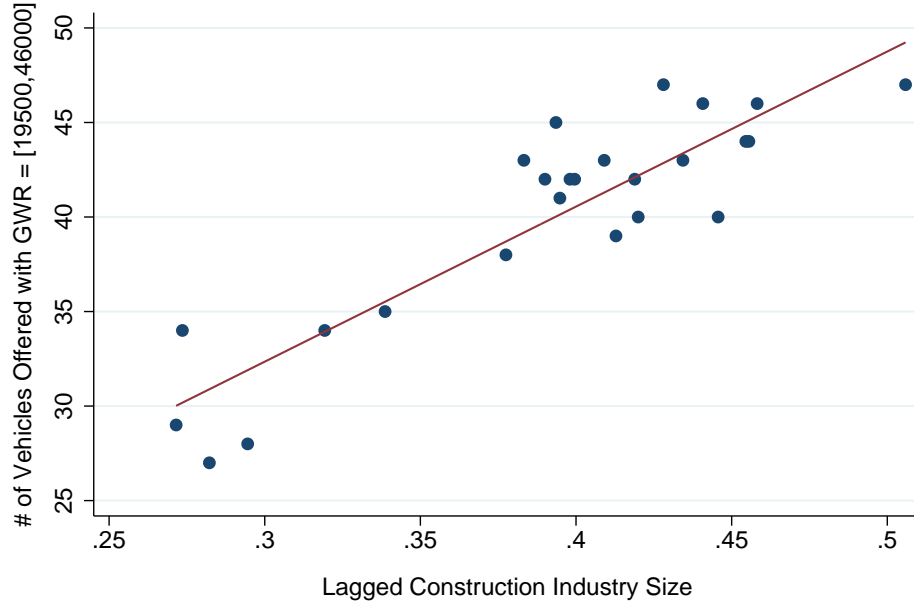


FIGURE II

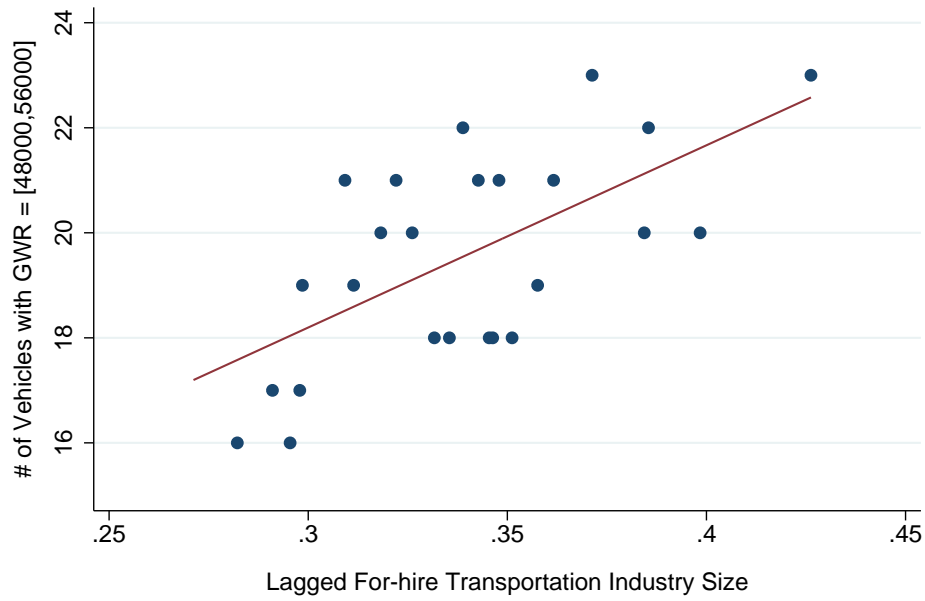


FIGURE III

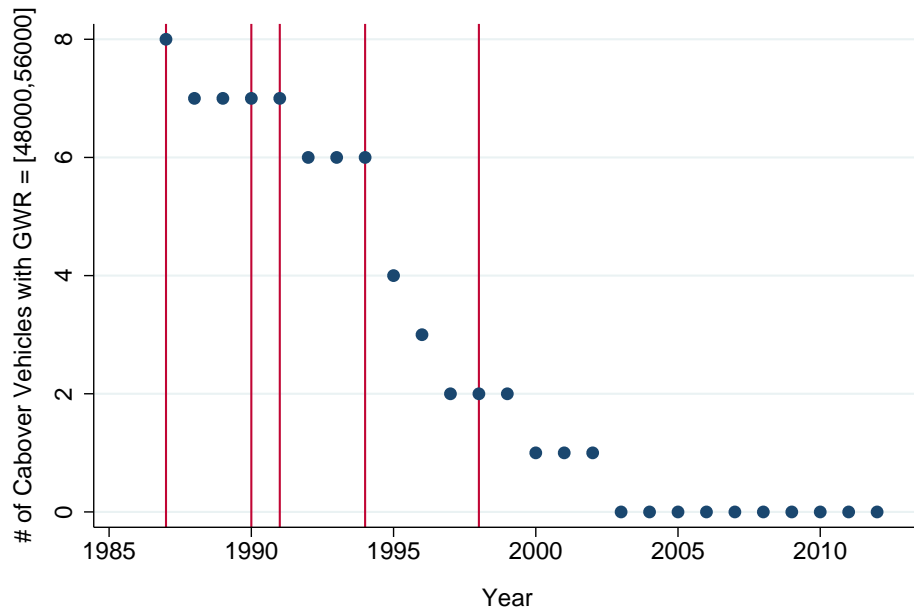


FIGURE IV

I follow the example above and begin with construction-related industries. These buyers account for under 40% of total purchases but over 80% of the sales of vehicles with a GWR between 19,500 and 40,000. Figure II shows that the number of offerings in this sub-segment is quite closely tied to the industry. In fact, this data was already presented in Table I, although its significance was not obvious at the time. For example, there is a steep increase in medium weight offerings from the early- to mid-2000s and an equally steep decline in the late-2000s.

The freight industries provide the second piece of evidence. These industries again account for only about 40% of total purchases but well over 90% of the sales of vehicles with a GWR above 48,000. Figure III shows the number of offerings in this segment is again closely linked to the industry.

The deregulation of cab length provides the final evidence. For the early and middle part of the 20th century, states independently regulated the use of their highways. One restriction states imposed was a limit on the combined length of vehicle and trailer. This affected freight companies, but had virtually no affect on local service or construction firms, which tend to carry exclusively small loads. Strict length laws advantaged the cab-over-engine relative to conventional vehicles since short cabs translated directly into larger loads and higher revenue. Regulations varied considerably and in many cases created blocks to interstate commerce in some regions of the country. Beginning with the Surface Transportation Assistance Act of 1982, the federal government began to standard the maximum legal load being carried and at the same time deregulated the length of the vehicle pulling it. Heavy cabover sales, once favored by the freight industry in some states, were crushed.

Since the process unraveled slowly and modeling its idiosyncrasies are beyond the scope of this paper, I construct an instrument for deregulation from a simple count of the relevant legislative and court decisions.

Figure IV shows the product response to this deregulation, with red lines indicating relevant legislative or court action on the length laws. Although not a one-to-one mapping, the impact on the heavy cabover is undeniable. The number of these vehicles dropped from eight at the beginning of the panel to zero by 2003.

VI Results

Demand

Table V reports the results from estimation of the demand system. Note that coefficients referencing the cab-over-engine, compact-front-end, and long-option are relative to the conventional cab, which is the omitted discrete category. All parameters are estimated very precisely (with the exception of the long-option) and this is especially true for interaction terms, whose precision is greatly aided by the microdata.²⁹ A few parameters deserve discussion. First, GWR is positive for all buyers. This is an important check of the model, since price is always increasing in GWR and a negative coefficient would imply these buyers should choose an alternate (lighter) means of transportation. Moreover, the industry interactions are ordered in precisely the same way as the microdata would suggest. Specialty freight buyers, often called “heavy haul” firms, value GWR the most. Business and personal service firms like bulk couriers and local delivery firms, represented by the constant (since it is the omitted industry type), value it the least.

Second, the cabover is disliked by the average buyer because of its cramped, bumpy, and noisy ride. These problems are exacerbated by heavy loads, which reflected in the negative interaction with GWR. Yet, the cabover is strongly preferred by urban buyers, who value its agility and visibility, and this is reflected in the large positive interaction with the road density measure. This squares with casual observations: the cabover is an uncommon site in the United States except for areas like downtown Manhattan or Chicago, where these vehicles are ubiquitous. The impact of the length regulation, which primarily bound freight carriers, is also evident here. The interaction of the freight industry buyer dummy with the cab-over-engine is positive, although the further interaction with length deregulation is very negative.

Buyers dislike price. The coefficient is precisely estimated but difficult to interpret in units of utility. I translate this into more easily understood measures in the following section.

²⁹I began with a large set of potential interactions and dropped those that were consistently neither statistically significant nor impactful. Micro-data summary statistics, however, gave a strong indication as to which interactions would ultimately be included.

Measuring Fit for Demand

The demand system implies a mean price elasticity of demand equal to 2.23. General Motors suggested an overall price elasticity of demand for passenger vehicles of 1.0 in Berry, Levinsohn, and Pakes (2004), so my estimates suggest the commercial vehicle market is more elastic. This is not surprising. Commercial vehicle buyers tend to be small businesses that are price-sensitive and unmoved much by styling, color, or brand prestige.

They also imply an average price-cost margin of 8.95%. Since eight out of nine parent companies are public entities with audited financial statements, we can simply compare against the reported figures. A few caveats are necessary. First, all of these firms operate in either other product markets or other geographic areas or both. Paccar is active in Australia, Isuzu and Hino are based in Japan, Volvo makes motorboat engines, Ford makes passenger vehicles, et cetera. Second, I estimate marginal costs while the firms report average costs. I cannot tell, for example, what proportion of the Selling, General, and Administrative (“SG&A”) line-item costs are sunk rather than marginal and what portion are related to fixed headquarters activities. I should see, however, that my reported price-cost margin falls between the gross margin, which does not include SG&A, and pre-tax operating margin, which does. These bounds are reported in Figure V . The red vertical line represents the average price-cost margin computed from the model, 8.95%, which falls in between the upper and lower bounds reported in audited financial statements (averaged over the period 2007 to 2012) for all but one firm. Although SG&A is a significant portion of gross margin, making for wide bounds, these nonetheless provide a sense of fit of the model overall.

Marginal Costs

Using demand estimates and data, I back out marginal costs and compare them to gross weight rating (the continuous characteristic) conditional on cab type. Figure VI provides two examples from representative “bust” and “boom” years. The left panel is from 2001 while the right panel is from 2007. In both panels, the x-axis measures GWR (in 0,000s of lbs) while the y-axis measures the log of marginal costs (where marginal costs are in \$0,000s). Red markers indicate cab-over-engine vehicles while blue markers indicate conventionals. The compact-front-end and long-option do not exhibit much variation in weight, so these were not included. Two things are apparent: the log of marginal costs are linear in GWR and that cab-over-engine vehicle marginal costs are a positive additive shift upwards from the conventional cab. Other years looked similar. This provides some confidence in assuming that the log of marginal costs is linear in the observable regressors.

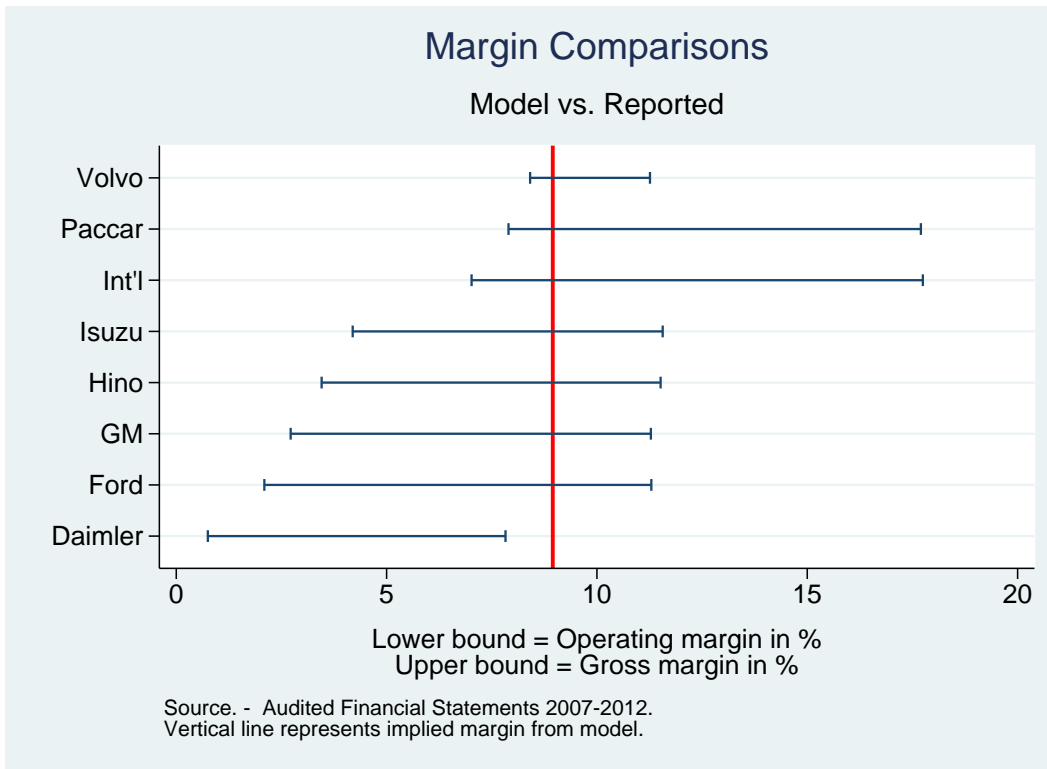


FIGURE V

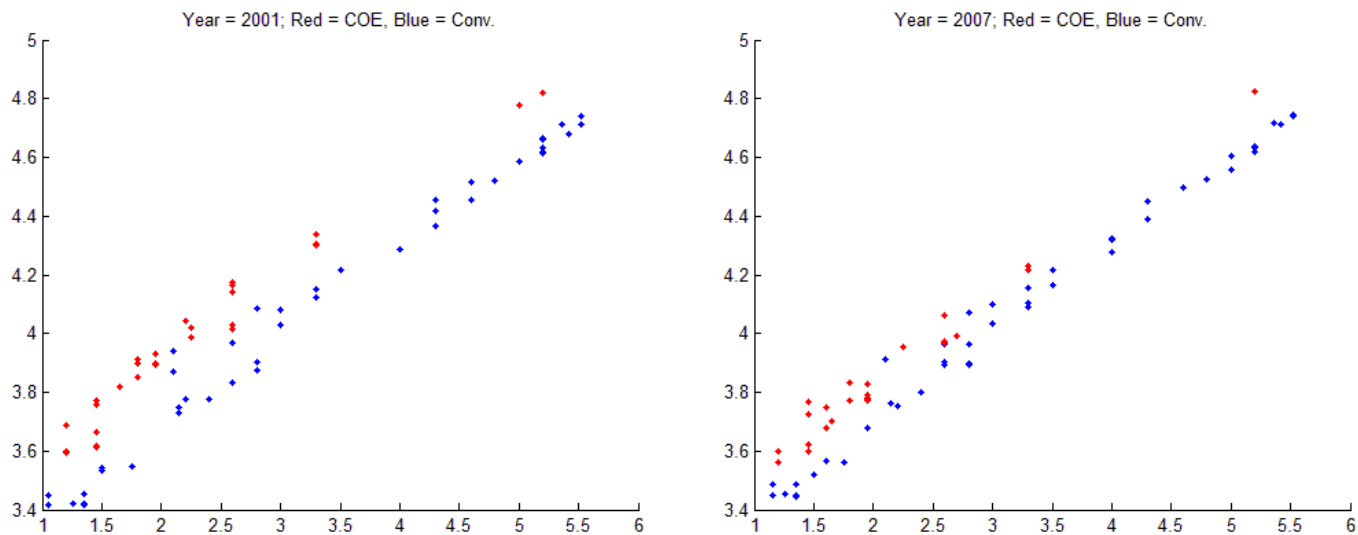


FIGURE VI

Table VI now reports the estimated marginal cost coefficients. All are precisely estimated. A 10,000 lb increase in GWR translates to about a 40% increase in marginal cost. Adding the long-option also contributes to higher marginal costs, as do higher wages. The time coefficient is small but positive and significant.

Sunk Costs

Table VII reports sunk cost estimates of model-level entry and exit. The constant term here refers to a zero-GWR conventional cab. Five things are important to check here. First, confidence intervals that contain the true parameter 95% of the time are both non-empty and reject zero. Second, they imply positive fixed costs for all product configurations I see in the data. Third, firms with their headquarters in Japan face lower sunk costs to introduce the cab-over-engine but higher sunk costs for all other models. The likely cause is design spillovers from the home market, where the cabover is ubiquitous. Fourth, the Big Three face relatively lower sunk costs to introduce low GWR vehicles but relatively higher sunk costs to introduce high GWR vehicles. The break-even point where the Big Three face essentially identical costs is around 33,000 lbs. Finally, the long option adds to the sunk cost of introduction. It is important to recover a non-trivial positive cost here, since this characteristic is a non-standard extension of the cabin and/or hood that requires the modification of parts and designs.

[Table VII about here.]

For perspective, Figure VII reports sunk costs across product space for the baseline firms, i.e. those that are neither headquartered in Japan nor a member of the Big Three. As an example, the sunk cost of introducing a conventional cab with a GWR of 33,000 and without the long option is about \$60 million. Sunk costs fall with GWR, which is surprising at first since more rugged vehicles seem to be more complicated to build and design. In practice, however, GWR affects assembly through the quality of the parts—mainly the strength of the steel chassis, the weight rating and number of axles, and durability of the transmission. These drive marginal costs, which, recall, increase nearly 40% for every 10,000 increase in GWR. Differences in sunk costs across GWR, however, may load on distribution, marketing, selling, and related expenses. Reaching buyers of low-GWR vehicles (e.g. bulk couriers, local delivery businesses, landscapers, and moving companies) may be more costly and challenging than reaching long-haul and heavy-haul freight carriers, who are simply more readily informed and interested in new models, and this could explain the difference.

For a similar perspective, Figure VIII presents sunk costs across the firm types, contrasting Big Three and Japan-headquartered firms against the rest. These illustrate the cost differences across GWR in the former case and cab-over-engine vehicles in the latter.

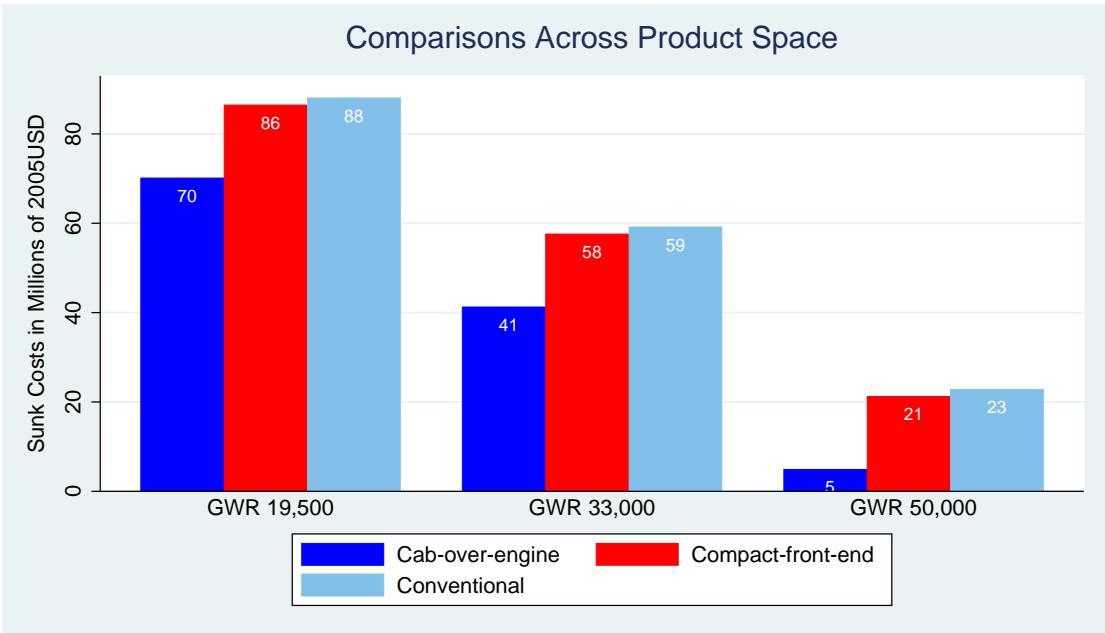


FIGURE VII

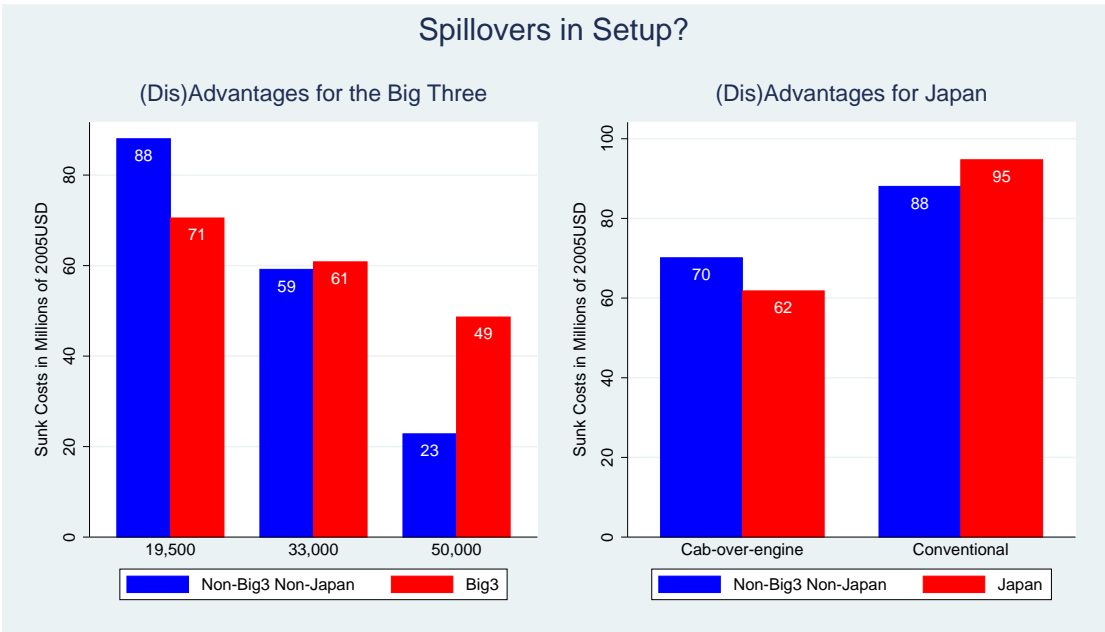


FIGURE VIII

Assessing the Capital Budgeting Rule

Sunk costs can be compared against profit estimates, conditional on how long firms kept products in the market for, to assess the capital budgeting decisions and implied hurdle rates. The estimation is based on Assumption II, which states that for a product j offered at t but not at $t - 1$, the agent's expected difference in profits, earned in perpetuity at the hurdle rate, weakly exceed the sunk costs. That is, $1/HurdleRate \times \mathcal{E} [\Delta\pi(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}; z_t, w_t, t) | \mathcal{I}_{f,t}]$ is greater than or equal to $\mathcal{E} [x_j \tilde{\theta}_{f,t,x} | \mathcal{I}_{f,t}]$. Realized profits, however, differ from this expectation. Firms do not receive profits in perpetuity, since they may choose to exit j in the future, although they will receive a scrap value when they do so. They also face different $J_{f,t}$, $J_{-f,t}$, z_t , and w_t in the future. Thus, the realized discounted cash flows from adding j that will be offered for T periods, which firms can then compare to sunk costs, are given by

$$\sum_{\tau=1}^T \frac{1}{HurdleRate^\tau} \times \Delta\pi(J_{f,\tau}, J_{f,\tau} \setminus j; J_{-f,\tau}; z_\tau, w_\tau, t) + \frac{1}{HurdleRate^T} \times \frac{1}{\lambda} \times x_j \tilde{\theta}_{f,t,x_j}$$

That is, in terms of realized cash flows, f adds j at t when

$$\underbrace{\sum_{\tau=1}^T \frac{1}{HurdleRate^\tau} \times \Delta\pi(J_{f,\tau}, J_{f,\tau} \setminus j; J_{-f,\tau}; z_\tau, w_\tau, t)}_{\text{Realized Sum of Discounted Profits}} \geq \underbrace{\tilde{\theta}_{f,t,x}}_{\text{Sunk costs}} - \underbrace{\frac{1}{HurdleRate^T} \times \frac{1}{\lambda} \times x_j \tilde{\theta}_{f,T,x_j}}_{\text{Realized discounted scrap value}}$$

Net present value of sunk costs

In the data, a close relationship between the left-hand side and right-hand side expressions, at reasonable hurdle rates, would imply firms earn what they expect to, on average, using the much simpler capital budgeting rule. Figure IX reports this ratio for three different hurdle rates: 13.26%, 15.26%, and 17.26%. Some products are still offered in the final year of the data so rather than deal with a more complicated truncation problem, I simply force these products to exit and recover the relevant scrap values.

Figure IX can reject the notion that firms earn radically different profits than they would expect to under the hurdle rate rule. At 15.26%, the distribution of realized ratios is centered close to 1.0. Since the truncation issue slightly understates the left-hand side variable, the true hurdle rate that sets these equal may be slightly higher. This would fall close to the midpoint of the 12% reported in the Poterba and Summers (1995) survey and the 19% reported in the Summers (1989) survey.

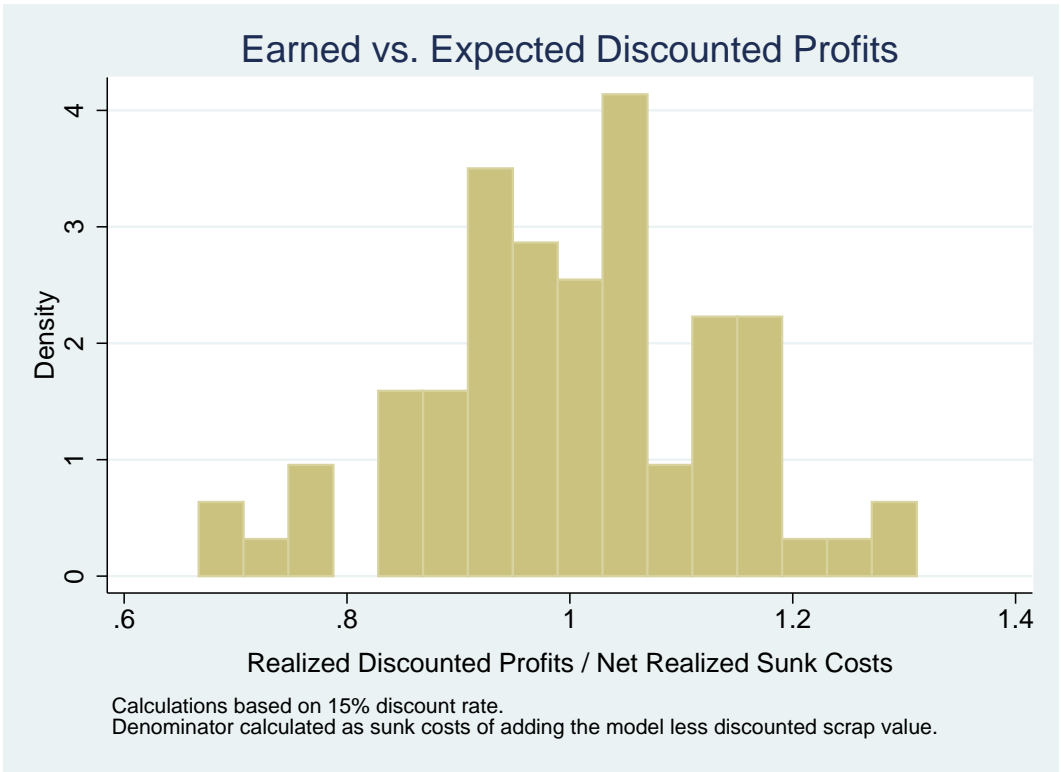


FIGURE IX

VII The Impact of the “Bailout”

Policy Setting

The \$85 billion of federal assistance to GM and Chrysler in 2009 constitutes the largest government bailout of a non-financial industry in modern history. Its causes are debated but there is mostly consensus on three factors: a global recession beginning in 2008 and prompting a trough in sales, a rise in fuel prices coupled with American manufacturers’ focus on trucks and SUVs, and legacy costs from pension and retiree healthcare benefits. By late 2008, there was an immediate fear that GM and Chrysler would default, prompting \$17.4 billion in assistance.³⁰ Shortly afterward in 2009, the federal government agreed on more funds, bringing the total to \$85 billion.

Whether to provide assistance was hotly contested, split partly along partisan lines, and even became a major US Presidential campaign issue. Republican Presidential Nominee Mitt Romney argued for a market-based solution in a November 2008 Op-Ed in the *The New York Times* titled “Let Detroit Go Bankrupt.” Later, in 2012, Barack Obama took credit for its apparent success, saying “I said we’re going to bet on American workers and the American auto industry, and it’s come surging back.”³¹ In a 2012 Op-Ed in the *Wall Street Journal*, Robert Barro pointed out that rhetoric supporting the bailout ex post ignored counterfactual policy outcomes. He wrote, “If GM had disappeared, its former workers and other inputs would not have sat around doing nothing. Another company—be it Toyota, Honda or Ford—would likely have taken over its operations.”³²

GM and Chrysler comprise nearly 15% of commercial purchases in 2009, so the bailout is relevant for this segment. Caveats are in order, however. First, debtor-in-possession financing was extremely scarce so the analysis below considers only complete liquidation or a rival firm’s takeover, even though more complex arrangements might have been likely. Second, moral hazard is ignored. Firms that expect support in the future may take riskier bets or under-expend effort in unprofitable states of the world. Third, the focus on the commercial segment clearly illustrates why model-level entry and exit matter and is suggestive of what could happen in other markets but is not a comprehensive automotive study. For example, it ignores general equilibrium effects that are probably small if failure is confined to the commercial segment but problematic if we extend this to the passenger segment, which is more than nine times larger. Fourth, changing an assembly line is presumably quick and painless relative to relocating labor. This may be fine for marginal expansions of product offerings but may prove challenging for a big shift in productive capacity. As a final note, financial distress is rarely random. Often it signals some underlying problems and that the efficient solution is shut down. I will assume—and am aided now by hindsight—that the issues of GM and Chrysler are mainly poor past decisions coupled with a very rare and deep capital drought.

³⁰See, for example, Wilcox, Gregory., “Big 3 detail plans for bailouts,” *Daily News* 3 Dec 2008., and Isidore, Chris., “Bush announces auto rescue,” *CNNMoney* 19 Dec 2008., and Carty, Sharon., “Automakers ask for billions in loans, credit; Two say they’re nearly out of money,” *USA Today* 3 Dec 2008.

³¹Rosche, Jedd. “Obama quotes Romney’s auto bailout op-ed in UAW speech.” *Politico* 28 Feb 2012.

³²Barro, Robert J. “Ryan and the Fundamental Economic Debate.” *The Wall Street Journal* 14 Aug 2012.

Alternate Policies

The following analysis compares the automotive bailout, i.e. federal support for GM and Chrysler that allowed them to continue operating as independent entities, against three alternate policies. One is liquidation, effectively a removal of the GM and Chrysler brands and products from the market. The other two are acquisitions: one by Ford, which overlapped heavily in product space with the troubled firms, and one by Paccar, which did not.

For each counterfactual policy, I compare the predictions of a model that allows for just prices to re-equilibrate against a model which allows for prices and product offerings to re-equilibrate. There are four broad findings. First, in the case of liquidation, sunk costs are low enough relative to profits to induce entry. Second, in this case, model-level entry and exit have a strong, moderating effect on the impact of liquidation. Finally, with respect to the acquisitions, although the identity of the acquiring firm matters a great deal when model-level entry and exit are ignored, it matters little when they are accounted for.

Computing Counterfactual Policy Outcomes

Assessing *what would have happened* in the event that GM and Chrysler were not rescued by the federal government requires recomputing the product offerings that would result from the change in the environment. In positioning games, multiple equilibria are the rule rather than the exception. Differentiated product markets with multiple attributes, as in the present setting, feature a large number of potential product offerings and a corresponding large number of potential equilibria. Lee and Pakes (2009) suggest a learning process that results in a distribution of equilibria played as well as potentially reducing the computational burden of this problem. The counterfactual policy outcomes reported below rest on a best response dynamic, which is computed as follows. Begin with product offerings from the prior period and a predetermined order of firm moves. The first firm chooses product offerings that are the best action conditional on what all other firms are currently offering. Their choice updates the product offerings that others observe when making their choices. The second, third, and so on, firms do the same. When the last firm in the ordering has chosen a best response, the order repeats. The process terminates when no firm has any profitable deviations. For each decision, the sunk costs used to compute the best response are based on moves from the 2009 product offerings, not the prior iteration of the learning process.

The rest point of this system is consistent with the necessary conditions that were used in the estimation of sunk costs, so the equilibrium selection process here is internally consistent with the model presented above. The policy analysis below proceeds with an ordering based on market share, with the largest share (“leader”) moving first. Details of the calculations are provided in the Appendix.³³

³³One alternative method uses random orderings, which will result in a distribution over the outcomes. This is in process.

Findings

It is helpful to start with an idea of what products GM and Chrysler offered in the year prior to the decision. Table VIII reports the offerings for 2009. GM offers twelve models while Dodge offers four. Both were operating in the lower one-half of the GWR distribution and produced all three types of cabs but did not feature the long/extended option. Several of these models overlap. All of the Chrysler models and over two-thirds of the GM models overlap with offerings by Ford, while none overlap with Paccar (not shown).

[Table VIII about here.]

Liquidation is assessed first. Table IX reports the impact on the most affected and median affected products, ordered by their respective changes in markups. In Panel A, model-level entry and exit are ignored. In this case, all three of the most affected products are owned by Ford, are conventional cab vehicles, and tend to be at the low end of the weight distribution. Markups on these models rise by over 60%. Despite higher prices, their market shares also expand, capturing a subset of buyers who find it difficult to substitute away from low-GWR conventional cabs as well as having a pressing need to purchase a vehicle. In Panel B, model-level entry and exit are accounted for. In stark contrast to the prior results, markups for the most affected products increase only around 10% to 20%. The impact to market share is more muted. These reflect increased sales driven by lower prices that are partially offset by substitution to newly introduced models. The sharp differences in Panel A and B are driven by model-level entry in precisely the places where markups increase the most.³⁴ In total seven products enter: two by Daimler, two by International, and one each by Paccar, Volvo, and Hino. Daimler enters direct competitors to the F250 and F450 while International enters a direct competitor to the F350. The other introductions are slightly more dispersed. No products exit.

[Table IX about here.]

The policies' distributional effects on buyers are strong. Table X reports these results. Liquidation has the biggest impact on business and personal service industry firms that reside in areas slightly more dense or urban than average. These firms increase their substitution to the outside good by up to almost 50% when entry is ignored but only 14% when it is accounted for. Unsurprisingly—and reassuringly—this is precisely the subset of firms that the microdata indicated are most likely to purchase low-GWR conventional vehicles. The median impacted buyer, meanwhile, experiences virtually no change. For this reason, the change in total output measured in levels is relatively muted overall. Total output falls 8% when model-level entry is ignored and about 3% when it is accounted for. Nonetheless, this still translates to a nearly 60% drop in the effects of liquidation.

[Table X about here.]

Acquisitions are assessed next. Figure X reports these results and compares them with liqui-

³⁴See the Appendix for more details on how these are computed.

dation. The left-hand side graph reports the increase in markups for the most affected products relative to the bailout. The y-axis measures level changes in percents (not to be confused with percentage changes). It shows that when model-level entry and exit are ignored, an acquisition by Ford would resemble liquidation while an acquisition by Paccar would resemble the bailout. It further suggests, however, that when model-level entry and exit are accounted for, it matters little which policy is chosen. Markups for the most affected products rise between 14% and 18%, regardless as to whether GM and Chrysler are liquidated or sold to a rival. The differences between the bailout and the alternate policies are much larger than between the alternate policies themselves. The elimination of two independent owners matters much more than what ultimately happens to the products owned by them, since the sunk costs are low enough that product portfolios will flexibly adjust ex post anyway. These are very different results than one would obtain under high sunk product entry costs. In this case, policymakers would care a great deal about transferring GM and Chrysler products to Paccar or another rival with whom the trouble firms did not previously overlap. The right-hand side graph shows the net number of product entries, i.e. the total number of individual models entered less the total number exited. An acquisition by Ford leads to seven total product introductions and five net product introductions. Two exits by Ford of duplicate models that it inherited make up the difference. An acquisition by Paccar leads to no entry and two exits of duplicate modes that it inherited. A high number of net entrants in the former two cases drive the large wedge between markups in the left-hand side graph, while the exits by Paccar push prices up slightly.

Figure XI compares the impacts of the policy changes to buyers by showing the change in the probability of not purchasing a vehicle. The y-axis measures level changes in percents. The left-hand side graph studies the most affected buyers. Predictably, these are tied closely to markups. There is one minor difference when model-level entry and exit are ignored. While markup changes are about even for a Ford acquisition relative to liquidation, purchasing changes are about one-third lower. This reflects idiosyncratic preferences of the buyers for the liquidated products. It is, of course, mechanical in the logit model but also captures the reality of differences in stylizing, dealer relationships and locations, and a host of other factors outside the scope of a tractable model. The right-hand side graph studies the mean affected buyers. As above, the impact of the policies is very slight and the sign of the effect of allowing model-level entry and exit is the same in all cases. Also, as above, the impact to the median affected buyers (not show) is virtually zero.

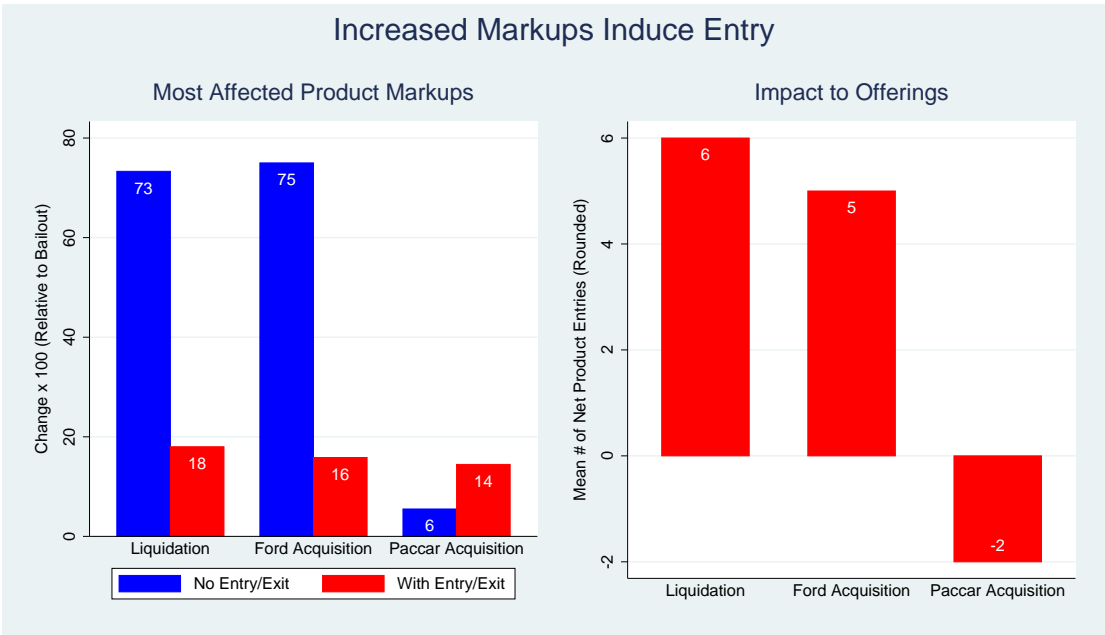


FIGURE X

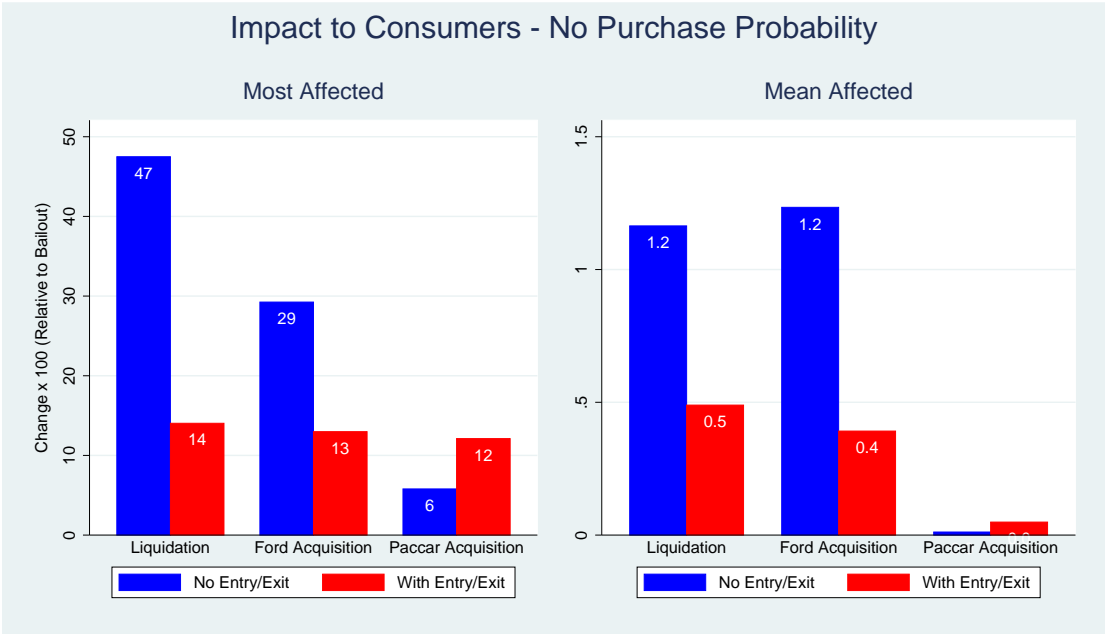


FIGURE XI

VIII Conclusions

In markets like the commercial vehicle segment of the US automotive industry, and the automotive industry in general, there has been virtually no firm-level entry and exit for decades and yet frequent product additions and removals. This suggests that even if the cost of entry to startup firms is prohibitively high, the cost to incumbents to enter and exit individual models is sufficiently low for the market to adjust to some policy changes. For this reason, the predictions that come out of models that ignore this fact and allow only prices to re-equilibrate will tend to overstate the impact of policy or market structure changes on markups, profits, and purchases.

Reduced-form evidence indicates that in the commercial vehicle market, producers respond to changes in the composition of demand with changes in the product offerings. For example, as the construction industries expands, the number of product offerings tailored to the preferences of construction buyers expands as well. This indicates that, in short, sunk costs of introducing these models cannot be so sufficiently high as to discourage this behavior. On the other hand, that some rather than all firms are entering these models indicates that these introductions are not free. These facts provide clear intuition for why the data can bound fixed costs, although to quantify those bounds and predict alternate policy outcomes, I needed a structural model.

Several challenges presented themselves. As with many large, economically important differentiated product markets, commercial vehicles are sold nationally but produced in at most one or two locations nationwide. This ruled out geographic variation on the supply side and, in turn, the use of estimation methods that have previously been exploited in the literature. I instead used the intuition provided by the reduced-form evidence above along with a rich twenty-seven year panel of product offerings to identify the sunk costs. This entailed taking a stand on how managers make potentially very complex decisions, for which I relied on practitioner interviews. To deal with multiplicity, which is the rule rather than the exception in positioning games, I relied only on the necessary—rather than sufficient—Nash equilibrium conditions. To deal with this complication in simulating counterfactual policy outcomes, I used a learning process based on best response dynamics. Taken together, they suggested that sunk costs were low enough to induce entry for policies where markups would have otherwise risen considerably. They demonstrated that in this market, when model-level entry and exit are ignored, it matters a lot which policy is chosen while when they are accounted for, it matters little. Nonetheless, the loss of two independent operating entities increases the overall concentration of ownership in this market. Markups for the most affected products rise slightly, although the impact to the mean and median affected products and buyers, ordered by markup and purchasing changes, respectively, was essentially zero.

REFERENCES

Aguirregabiria, V., and Mira, P. (2007): “Sequential Estimation of Dynamic Discrete Games.” *Econometrica*, 75:1, 1-53.

- Andrews, D., and G. Soares (2010): “Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection.” *Econometrica*, 78:1, 119–157.
- Bajari, P., Benkard, L., and Levin, J. (2007): “Estimating Dynamic Models of Imperfect Competition.” *Econometrica*, 75:5, 1331-1370.
- Berry, S. (1994): “Estimating Discrete-Choice Models of Product Differentiation.” *RAND Journal of Economics*, 25(2): 242-262.
- Berry, S., J. Levinsohn, and A. Pakes. (1995): “Automobile Prices in Market Equilibrium.” *Econometrica*, 63:4, 841-890.
- . (1995): “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market.” *Journal of Political Economy*, 12:1, 68-105.
- Blonigen, B., Knittel, C., and Soderbery, A. (2013): “Keeping It Fresh: Strategic Product Redesigns and Welfare.” Mimeo.
- . (2001): “Do Mergers Increase Product Variety? Evidence from Radio Broadcasting.” *Quarterly Journal of Economics*, 116:3, 1009-1025.
- Chernozhukov, V., H. Hong, and E. Tamer (2007): “Estimation and Confidence Regions for Parameter Sets in Econometric Models.” *Econometrica*, 75:5, 1243-1284.
- Collard-Wexler, A. (2013): “Demand Fluctuations in the Ready-Mix Concrete Industry.” *Econometrica*, 81:3, 1003–1037
- Draganska, M., M. Mazzeo, and K. Seim (2009): “Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions.” *Quantitative Marketing and Economics*, 7:2, 105-146.
- Doraszelski, U., Lewis, G., and Pakes, A. (2014): “Just starting out: Learning and price competition in a new market.” Mimeo.
- Eizenberg, A. (2014): “Upstream Innovation and Product Variety in the United States Home PC Market.” *The Review of Economic Studies*, 81:1003-45.
- Fan, Y. (2013): “Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market.” *American Economic Review*, 103(5): 1598-1628.
- Fershtman, C., and Pakes A. (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work.” *Quarterly Journal of Economics*. 127(4): 1611-1661.
- Ford, H. (1922): *My Life and Work*. Garden City, New York, USA: Garden City Publishing Company, Inc.
- Fudenberg, D., and Levine, D. (1998): *The Theory of Learning in Games*. Cambridge: The MIT Press.
- Haile, P., and Tamer, E. (2003): “Inference With an Incomplete Model of English auctions.” *Journal of Political Economy*, 111:1, 1-51.

- Hotz, J., and Miller, R. (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice.” *The Review of Economic Studies*, 61:2, 265-289.
- Ishii, J. (2008): “Compatibility, Competition, and Investment in Network Industries: ATM Networks in the Banking Industry.” Mimeo.
- Matzkin, R.L. (1992): “Nonparametric and Distribution Free Estimation of the Binary Choice and the Threshold Crossing Models.” *Econometrica*, 60:2, 239-270.
- Mazzeo, M. (2002): “Product Choice and Oligopoly Market Structure.” *RAND Journal of Economics*, 33:2, 221-242.
- Morales, E., Sheu, G., and Zahler, A. (2014): “Gravity and Extended Gravity: Using Moment Inequalities to Estimate a Model of Export Entry.” NBER Working Paper 19916.
- Nevo, A. (2001): “Measuring Market Power in the Ready-to-eat Cereal Industry.” *Econometrica*, 69:2, 307-342.
- Nosko, C. (2010): “Competition and Quality Choice in the CPU Market.” Mimeo.
- Pakes, A., Ostrovsky, M., and Berry, S. (2007): “Simple estimators for the parameters of discrete dynamic games (with entry/exit examples).” *RAND Journal of Economics*, 38:2, 373-399.
- Pakes, A., Porter, J., Ho, K., and Ishii, J. (2014): “Moment Inequalities and Their Application.” *Econometrica*, forthcoming.
- Pesendorfer, M., Schmidt-Dengler, P. (2008): “Asymptotic Least Squares Estimators for Dynamic Games.” *The Review of Economic Studies*, 75, 901-928.
- Petrin, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan.” *Journal of Political Economy*, 110:4, 705-729.
- Ryan, S. (2012): “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica*, 80:3, 1019-1062.
- Schumpeter, J. (1942): *Capitalism, Socialism and Democracy*. London: Routledge.
- Summers, L. (1987): “Investment incentives and the discounting of depreciation allowances.” In M. Feldstein (ed.), *The effects of taxation on capital accumulation*. University of Chicago Press: Chicago.
- Poterba, J., and Summers, L. (1995): “A CEO Survey of U.S. Companies’ Time Horizons and Hurdle Rates.” *MITSloan Management Review*. 37:1 Fall, 27-59.
- Sweeting, A. (2013): “Dynamic Product Repositioning in Differentiated Product Markets: The Case of Format Switching in the Commercial Radio Industry.” *Econometrica*, 81:5, 1763-1803.
- Tamer, E. (2003): “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria.” *The Review of Economic Studies*, 70:1, 147-165.
- Varadhan, R., and Roland, C. (2008): “Simple and Globally Convergent Methods for Accelerating

the Convergence of Any EM Algorithm.” *Scandinavian Journal of Statistics* 35:2, 335-353.

Ward’s Communications (1986-2013): *Ward’s Automotive Yearbook*, Southfield: Penton.

FIGURES PRESENTED IN BODY. TABLES PRESENTED BELOW.

TABLE I
SUMMARY STATISTICS

	All Years		9-YEAR GROUPS		
	Min	Max	1986-1994	1995-2003	2004-2012
<u>Count of Offerings</u>					
Conventional, light-medium	8	10	8.0	7.9	8.3
Conventional, medium	21	35	21.3	24.8	27.1
Conventional, heavy	8	14	8.3	10.8	9.6
Conventional, long option	6	12	8.3	9.7	10.8
Cab-over-engine, light-medium	7	13	12.0	11.8	9.8
Cab-over-engine, medium	3	18	12.7	15.0	9.2
Cab-over-engine, heavy	0	8	6.0	1.8	0.0
Compact-front-end, all GWR	5	9	5.0	5.7	7.2
All vehicles	70	97	81.7	87.3	82.0
<u>Market Outcomes</u>					
Price, CPI-adjusted	\$65,958	\$71,334	\$64,510	\$68,012	\$67,644
Quantity, 000s of units	193	658	467	502	494
<u>Buyer Composition</u>					
Freight-related	27.1%	42.6%	30.9%	35.1%	34.7%
Construction-related	27.2%	50.6%	37.6%	41.3%	38.8%
Bus. & Pers. Service-related	13.3%	44.4%	31.6%	24.4%	26.5%

TABLE II
MEAN CHARACTERISTICS CONDITIONAL ON BUYER TYPE AND PURCHASE

EXPECTED [CHARACTERISTIC BUYER TYPE, PURCHASE]							
	All Industries	Bus. & Pers. Service	Contractor	General Construction	Heavy Building	General Freight	Specialty Freight
GWR	31,818 <i>(16,612)</i>	12,700 <i>(5,263)</i>	21,495 <i>(6,191)</i>	31,494 <i>(4,193)</i>	43,462 <i>(5,216)</i>	51,616 <i>(1,465)</i>	54,277 <i>(1,205)</i>
Cabover	All Buyers 8%	+1.5 σ Road Density 56%					
Compact- front-end	All Industries 14%	Bus. & Pers. Service 46%					

TABLE III
DESCRIPTIVE STATISTICS BY BRAND

Brand	Parent	Models Per Year	Snapshot in 2012		
			Mean GWR	Mean Price	Market Share
Dodge	Chrysler	2.8	13,139	28,389	5.56%
Ford	"	12.8	21,555	39,680	12.50%
Freightliner	Daimler	8.7	39,994	85,529	12.50%
Fuso	Daimler	4.6	20,454	45,729	4.17%
GM	"	11.6	17,193	33,545	6.94%
Hino	"	3.3	22,044	47,464	4.17%
International	"	13.6	39,974	85,130	16.67%
Isuzu	"	5.1	18,554	40,192	6.94%
Kenworth	Paccar	4.9	49,848	118,498	6.94%
Mack	Volvo	2.5	48,647	110,053	2.78%
Peterbilt	Paccar	6.1	50,457	118,524	11.11%
UD	Volvo	3.9	19,228	43,686	4.17%
Volvo	"	4.0	47,128	107,163	4.17%
Western Star	Daimler	1.1	52,000	131,854	1.39%

Note. Brands taken as of 2012.

TABLE IV
OWNERSHIP CHANGES

Parent	Action	Target	Year	Target's US CV Sales / Target's Total Sales
Daimler	acquisition	Chrysler	1998	1.6%
Daimler	acquisition	Hyundia Truck	2001	3.0%
Volvo	acquisition	Mack (Renault)	2001	24.1%
Daimler	spinoff	Hyundia Truck	2001	2.5%
Daimler	acquisition	Mitsu. Fuso	2004	4.2%
Daimler	spinoff	Chrysler	2006	1.3%
Volvo	acquisition	Nissan Diesel	2006	6.1%

TABLE V
DEMAND ESTIMATION

	Mean β		Interactions		
	Estimate	SE	Buyer Type	Estimate	SE
GWR (0,000s lbs.)	0.60***	0.13	SPECIALTY FREIGHT INDUSTRY	21.88***	5.78
			GENERAL FREIGHT INDUSTRY	17.90***	5.35
			HEAVY BUILDING INDUSTRY	10.70***	2.54
			GEN. CONSTRUCTION INDUSTRY	6.68**	2.26
			CONTRACTOR INDUSTRY	3.61***	1.00
Cab-over-engine	-1.85**	0.62	URBAN	13.96***	2.62
			GENERAL FREIGHT INDUSTRY	15.93***	4.43
			GENERAL FREIGHT INDUSTRY X LAW	-4.25**	1.83
Compact-front-end	-24.58**	8.37	LOCAL DELIVERY INDUSTRY	29.67***	7.84
Long-conventional	-0.56	0.39	FREIGHT INDUSTRY X N(0,1)	10.49***	2.05
Cabover X GWR	-3.92***	0.94	---		
Constant	-18.13***	5.28	SPECIALITY FREIGHT INDUSTRY	-89.37***	14.45
			GENERAL FREIGHT INDUSTRY	-67.57***	15.36
			HEAVY BUILDING INDUSTRY	-23.54***	7.71
			GENERAL CONSTRUCTION INDUSTRY	-11.07**	4.87
			CONTRACTOR INDUSTRY	-4.33**	1.86
			LOCAL DELIVERY INDUSTRY	-5.44***	1.21
			N(0,1)	20.74**	7.31
Price (\$000s)	-0.35***	0.06	---		

Note. - Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

TABLE VI
MARGINAL COST ESTIMATES

Product Characteristic	ln(mc)	
	Estimate	SE
Cab Over Engine	0.176***	0.037
Compact Front End	-0.187***	0.035
Long Option	0.073***	0.024
GWR (0,000s lbs)	0.392***	0.081
Hourly Wage (\$)	0.018***	0.004
Time	0.009**	0.004
constant	2.033***	0.391

No. of Observations: 1928

Note. - Robust standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

TABLE VII
SUNK COST ESTIMATION

Characteristic	Parameter	Point	95% Conf. Interval*
		Estimate	
constant	θ^0	[\$129.73]	[\$117.30 , \$144.06]
Cab Over Engine	θ^{COE}	[-\$17.92]	[-\$16.17 , -\$20.58]
Compact Front End	θ^{CFE}	[-\$1.59]	[-\$1.31 , -\$1.75]
Long Option	θ^{Long}	[\$33.89]	[\$30.25 , \$38.70]
GWR (0,000s lbs.)	θ^{GWR}	[-\$21.38]	[-\$19.05 , -\$25.10]
constant X BigThree	θ^{Big3}	[-\$45.22]	[-\$40.77 , -\$51.82]
constant X Japan	θ^{Japan}	[\$6.69]	[\$5.50 , \$8.12]
Cab Over Engine X Japan	$\theta^{COE,Japan}$	[-\$15.00]	[-\$13.46 , -\$17.90]
GWR X BigThree	$\theta^{GWR,Big3}$	[\$14.21]	[\$12.46 , \$16.22]
Scaling for Exit	$1/\lambda$	[-0.386]	[-0.278 , -0.426]

Note. All figures in millions of constant 2005USD.

* Probability that the true parameter θ lies in this space.

TABLE VIII
2009 GM AND CHRYSLER PRODUCT OFFERINGS

Parent	Brand	Cab Type	Long/Ext. Option	GWR
Chrysler	Dodge	Conventional	No	12,000
Chrysler	Dodge	Conventional	No	13,500
Chrysler	Dodge	Conventional	No	15,300
Chrysler	Dodge	Conventional	No	19,500
GM	GM	Compact Front End	No	11,000
GM	GM	Conventional	No	12,000
GM	GM	Conventional	No	13,500
GM	GM	Compact Front End	No	14,000
GM	GM	Cab Over Engine	No	14,700
GM	GM	Conventional	No	15,300
GM	GM	Cab Over Engine	No	19,500
GM	GM	Conventional	No	19,500
GM	GM	Conventional	No	21,500
GM	GM	Cab Over Engine	No	22,400
GM	GM	Cab Over Engine	No	27,100
GM	GM	Conventional	No	28,000

TABLE IX
COUNTERFACTUAL: GM & CHRYSLER LIQUIDATED
 EFFECTS OF MODEL-LEVEL ENTRY ON PRODUCTS

---- PANEL A: JUST PRICES RE-EQUILIBRATE ----

Rank	GWR	Cab Type	Brand	Model	Markup Per Model		Market Share
					Level Difference	% Chg.	% Chg.
1	12,000	Conventional	Ford	F250	\$2,288	73%	39%
2	13,500	Conventional	Ford	F350	\$2,168	70%	34%
3	15,300	Conventional	Ford	F450	\$1,889	63%	27%
Median	35,000	Conventional	Sterling	Acterra33	\$90	3%	7%

---- PANEL B: PRICES AND PRODUCT OFFERINGS RE-EQUILIBRATE ----

Rank	GWR	Cab Type	Brand	Model	Markup Per Model		Market Share
					Level Difference	% Chg.	% Chg.
1	12,000	Conventional	Ford	F250	\$453	18%	23%
2	13,500	Conventional	Ford	F350	\$420	14%	21%
3	15,300	Conventional	Fuso	FM330	\$114	4%	8%
Median	33,000	Conventional	Peterbilt	386	\$0	0%	0%

TABLE X
COUNTERFACTUAL: GM & CHRYSLER LIQUIDATED
 EFFECTS OF MODEL-LEVEL ENTRY ON BUYERS

Measure	Level changes		Level Difference	% Difference
	<i>Without Entry/Exit</i>	<i>With Entry/Exit</i>		
	(1)	(2)		
<u>No-Purchase Probability</u>				
Most Affected Buyer	-43.2%	-14.0%	-29.2%	-67.5%
Mean Buyer	-1.2%	-0.5%	-0.7%	-57.9%
Median Buyer	0.0%	0.0%	0.0%	-30.8%
<u>Output</u>				
Total market	-7.6%	-3.2%	-4.4%	-57.9%

Most Affected Buyer Type: Industry = Bus. & Personal Service, Urban Measure = -.36 σ below mean

DATA AND COMPUTATIONAL APPENDIX

Empirical Distribution of Buyers.

To construct the empirical distribution of buyers, I match state-level observations on the road density measure (from the US Department of Transportation *Highway Statistics*) to state-level observations on industry (from the US Census *County Business Patterns*). Weights in the distribution are based on the number of employees rather than establishments or companies since it is the employees that operate vehicles, not the fictitious legal entities that employ them. This data provides that, for example, the New York State-based construction industry accounts for 0.26% of all employment in the buyer industries for 2011 and that this group of potential buyers face surrounding roads that are classified 75% of urban (based on total road mileage).

Industries vary in the portion of employees that will operate vehicles. This is presumably quite high in the freight transportation industry but low in, for example, construction. I scale industry weights in the empirical distribution to match the average vehicle ownership in the microdata. For example, if freight firms account for 30% of commercial vehicle owners in the microdata but 15% of buyers in the US Census data, the weight of commercial vehicle buyers in the empirical distribution is doubled (technically, it would be scaled by $2.43 \approx \frac{.3}{1-.3} \times \frac{1-.15}{.15}$ were the other industries to scale proportionately and in the opposite direction).

Buyers also vary in unobservable ways and these are drawn from independent standard normal distributions.

Market Size Calculations.

The market size for each year, M_t , is constructed as a product of a mean market size over the panel and a scaling factor for each year. First, to compute the mean market size over the panel, write the total units sold in t as q_t such that $q_t = \sum_j q_{j,t}$. Set mean market size \bar{M} to a level such that the average “inside share” across the years equals $\frac{\frac{1}{T} \sum_t q_t}{\bar{M}}$. Second, to compute the scaling factor, write the scaled industry-time specific employment levels (described in the Appendix section immediately above) as $y_{I,t}$ at t and let $y_t = \sum_I y_{I,t}$. Set the scaling factor for t so that the change in market size is proportional to the change in y_t . That is, set the scaling factor \tilde{M}_t such that $\tilde{M}_t = \min\{q_t\} + (q_t - \min\{q_t\}) \frac{\max\{q_t\} - \min\{q_t\}}{\max\{y_t\} - \min\{y_t\}}$, which yields $M_t = \bar{M} \times \tilde{M}_t$.

Equilibrium Selection Details.

This section provides computational details on selecting out equilibria using a learning process based on best response dynamics. The order of decisions is based on 2009 market share, aggregated to the parent level, which is where the decisions are made. With GM and Chrysler eliminated from the set of individual owners, there are seven remaining parent companies. The ordering of their market shares is as follows: Ford (26%), Daimler (18%), International (17%), Paccar (14%), Volvo (7%), Isuzu (3%), Hino (1%). Best response dynamics simply begin with the set of product offerings inherited from the previous period and then cycle through the firms, updating the set of

product offerings following each choice. The cycle terminates when no profitable deviations can be found.

Candidates for a best response action by a firm f is any possible J_f and is of the size $2^{\mathbb{J}}$, where \mathbb{J} is the total number of distinct product types. Large portions of the potential action space will never be a (conditional) best response and can be ruled out *ex ante*, speeding up computations considerably. First, four of the product types are heavy cab-over-engine vehicles. The impact of length deregulation has reduced demand for the vehicles to effectively zero in 2010, conditional on any other product being offered over 33,000 lbs GWR. Recalling Figure V, not a single vehicle of this type has been offered in the entire six year span preceding 2010. This alone reduces the number of potential actions by any firm to approximately 500,000. Second, recall firms headquartered in Japan introduce the conventional cab at a higher sunk cost than their rivals. For this reason, it turns out that these firms do not find it profitable to add any conventional cab vehicle over 33,000 lbs GWR conditional on at least three competing firms offering conventional cab vehicles over 40,000 lbs GWR. Third, recall the Big Three firms find it cheaper to introduce lighter vehicles and more expensive to introduce heavier vehicles relative to their rivals. For this reason, it also turns out that Ford (the only remaining Big Three firm) does not find it profitable to offer any cab over 48,000 lbs GWR conditional on at least three competing firms offering conventional cab vehicles over 48,000 lbs GWR. In my experience to date, Volvo, Paccar, International, and Daimler never completely vacate the heavy conventional cab sub-segment, so these restrictions are always in effect.

The fourth restriction employs different logic, which is based on the number of offerings. If a firm finds it more profitable to offer n products than $n + 1$ products, then it will always find it more profitable to offer n than $n + k$, $k > 1$.

PROOF APPENDIX

For clarity, I prove Proposition I with two lemmas.

Lemma IA. λ is identified, while θ^{con} is jointly identified with θ^{GWR} .

Proof. Let $J_{f,t}$ be the observed, i.e. equilibrium, product offerings by f at t . For notational simplicity, temporarily ignore observations specific to Big Three and Japanese firms and consider only conventional cab vehicles without the long/extended hood option, recalling that $\nu_{2,t}^{GWR} | \nu_{2,t}^k \sim \nu_{2,t}^{GWR}$ where $k \in \{con, COE, CFE, Long\}$ by Assumption III. For further notational simplicity, let the G of discrete values of x^{GWR} be denoted g . Define two instruments as

$$h^1(f, g, t) = \left[\frac{1}{F} \sum_f \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \in J_{f,g,t-1}\} \right]^{-1} \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \in J_{f,g,t-1}\}$$

$$h^2(f, g, t) = \left[\frac{1}{F} \sum_f \{j_{f,g,t} \notin J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\} \right]^{-1} \{j_{f,g,t} \notin J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\}$$

The first instrument takes a positive value when a product model stays in the market, weighted inversely by the number of products with GWR g at t that stay in. The second instrument is constructed analogously, but refers to product models that stay out. Next, interact h^1 with Equation 12 and take an expectation over $T \times G \times F$ observations. If there is at least one GWR g conventional cab offered at t , which is true in the data, then this yields

$$\begin{aligned} & \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 \left(\Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \theta^{con} - \nu_{2,t}^{con} - \nu_{1,j,t}^{con} \right. \\ & \qquad \qquad \qquad \left. - g(\theta^{GWR} + \nu_{2,t}^{GWR} + \nu_{1,j,t}^{GWR}) + \nu_{1,J_{f,t}, J_{f,t} \setminus j; J_{-f,t}}^\pi \right) \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 (\Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \theta^{con} - \theta^{GWR} g - \nu_{2,t}^{con} - \nu_{2,t}^{GWR} g) \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \frac{1}{T} \frac{1}{G} \sum_t \sum_g \left[\theta^{con} + \theta^{GWR} g + \nu_{2,t}^{con} + \nu_{2,t}^{GWR} g \right] \left[\frac{1}{F} \sum_f h_{f,g,t}^1 \right] \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \frac{1}{T} \frac{1}{G} \sum_t \sum_g \left[\theta^{con} + \theta^{GWR} g + \nu_{2,t}^{con} + \nu_{2,t}^{GWR} g \right] \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \theta^{con} - \frac{1}{T} \sum_t \left[\nu_{2,t}^{con} - \left(\frac{1}{G} \sum_g g \right) \theta^{GWR} + \left(\frac{1}{G} \sum_g g \right) \nu_{2,t}^{GWR} \right] \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t}) - \theta^{con} - \bar{g} \theta^{GWR} \\ &\geq 0 \end{aligned}$$

The first equality holds because ν_1 terms were not known when product offerings were chosen, so they are distributed independent of the choices made (as well as prior choices, demand shifters, etc) by Assumption IV. This means that $\mathbb{E}[h_{f,g,t}^1 \nu_{1,f,j,t}^{GWR} | J_{f,t}, J_{-f,t}, h^1] = \mathbb{E}[h_{f,g,t}^1] \mathbb{E}[\nu_{1,f,j,t}^{GWR} | J_{f,t}, J_{-f,t}]$ and since $\mathbb{E}[\nu_{1,f,j,t} | J_{f,t}, J_{-f,t}] = \mathbb{E}[\nu_{1,f,g,t}^{GWR}] = 0$, then $\frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h^1 \nu_{1,j,t}^{GWR} \xrightarrow{p} 0$. The second inequality holds because θ^{con} , $\theta^{GWR} g$, $\nu_{2,t}^{con}$, nor $\nu_{2,t}^{GWR} g$ are firm specific, so the final bracketed term can be multiplicatively separated from the others. The third equality holds since, as stated above, we have at least one GWR g conventional cab offered at every t , and so $\frac{1}{F} \sum_f h_{f,g,t}^1 = 1$. The fourth equality holds since no remaining sunk cost terms are firm specific, since no term is GWR specific except g itself, and because neither θ^{con} nor θ^{GWR} are time specific. The fifth equality holds since $\nu_{2,t}$ is i.i.d. over time, indexed only by time, and unconditionally mean zero by construction. That is, $\frac{1}{T} \sum_t \nu_{2,t}^k \xrightarrow{p} \mathbb{E}[\nu_{2,t}] = 0$, where $k \in \{con, GWR\}$. This provides an upper bound on $\theta^{con} - \bar{g} \theta^{GWR}$, where $\bar{g} = \frac{1}{G} \sum_g g$ is a known constant. For a lower bound, interact h^2 with Equation (13) and go through an analogous procedure. Together these yield

$$\theta^{con} + \theta^{GWR} \bar{g} \geq -\frac{1}{TGF} \sum_t \sum_g \sum_f h^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j; J_{-f,t})$$

$$\lambda(\theta^{con} + \theta^{GWR} \bar{g}) \leq \frac{1}{TGF} \sum_t \sum_g \sum_f h^2 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j; J_{-f,t})$$

These inequalities are based on models that were not dropped and not added, respectively. Next we consider the case where firms could, data permitting, drop one more and add one more product, respectively. In this case, however, there are some (t, g) -couples for which no conventional cab vehicle is added or dropped by any firm. The solution is to construct an instrument h that selects sub-segments of the characteristic space and periods of the data that are particularly unattractive or attractive to drop or add models in, respectively. These choices must be made agnostic with respect to what the firms actually chose. One strategy is the following: for each GWR g conventional cab vehicle, determine the subset z of Z that most impacts demand, and then select periods where $z_t^g - z_{t-1}^g$ took either a sufficiently large or small value. These are given by $(z_t^g - z_{t-1}^g)^*$ and $(z_t^g - z_{t-1}^g)^{**}$, respectively. These instruments would formally be constructed as

$$h_{f,g,t}^3 = \left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^* \right\}$$

$$\times \left[\frac{1}{F} \sum_f \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\} \right]^{-1} \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\}$$

$$h_{f,g,t}^4 = \left\{ (z_t^g - z_{t-1}^g) < (z_t^g - z_{t-1}^g)^{**} \right\}$$

$$\times \left[\frac{1}{F} \sum_f \{j_{f,g,t} \notin J_{f,t}\} \{j_{f,g,t} \in J_{f,g,t-1}\} \right]^{-1} \{j_{f,g,t} \notin J_{f,t}\} \{j_{f,g,t} \in J_{f,g,t-1}\}$$

Interact h^3 with Equation (12) and take an expectation over $T \times G \times F$ observations. Notice, for example, in h^3 that when $\left\{ (z_t^g - z_{t-1}^g) > (z_t - z_{t-1})^\star \right\}$ takes a non-zero value, we expect $\{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\}$ to also take a non-zero value for some f . This is true in the data. This yields

$$\begin{aligned} & \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^3 \left(\Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j; J_{-f,t}) + \nu_{1,J_{f,t}, J_{f,t} \setminus j; J_{-f,t}}^\pi \right. \\ & \quad \left. + \lambda \left[\theta^{con} + \nu_{2,t}^{con} + \nu_{1,j,t}^{con} + g(\theta^{GWR} + \nu_{2,t}^{GWR} + \nu_{1,j,t}^{GWR}) \right] \right) \\ &= \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j; J_{-f,t}) + \lambda \left[\theta^{con} + \bar{g} \theta^{GWR} \right] \left[\frac{1}{TG} \sum_t \sum_g \left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} \right] \\ &\geq 0 \end{aligned}$$

The algebra is analogous to the procedure above, with one important exception. The $\nu_{2,t}$ terms are independent of Z or functions thereof by Assumption III, so that

$$\begin{aligned} & \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} \nu_{2,t}^{con} | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \\ &= \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \mathbb{E}_t [\nu_{2,t}^{con}] \\ &= \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \times 0 \\ &= 0 \end{aligned}$$

and

$$\begin{aligned} & \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} \nu_{2,t}^{GWR} g | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \\ &= \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \bar{g} \mathbb{E}_t [\nu_{2,t}^{GWR}] \\ &= \mathbb{E}_{g,t} \left[\left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^\star \right\} | z_t^g, z_{t-1}^g, (z_t^g - z_{t-1}^g)^\star \right] \bar{g} \times 0 \\ &= 0 \end{aligned}$$

Similarly, interact h^4 with Equation (14) and take the same expectation. Together these yield

$$\begin{aligned} \lambda &\geq - \frac{\sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t})}{\left[\theta^{con} + \theta^{GWR} \bar{g} \right]} \geq - \frac{\sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t})}{\sum_t \sum_g \sum_f h_{f,g,t}^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j, J_{-f,t})} = \underline{\lambda} \\ \lambda &\leq \frac{\sum_t \sum_g \sum_f h_{f,g,t}^2 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j, J_{-f,t})}{\left[\theta^{con} + \theta^{GWR} \bar{g} \right]} \leq \frac{\sum_t \sum_g \sum_f h_{f,g,t}^2 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j, J_{-f,t})}{\sum_t \sum_g \sum_f h_{f,g,t}^4 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j, J_{-f,t})} = \bar{\lambda} \end{aligned}$$

Lemma IB. θ^{GWR} is jointly identified with λ .

Proof. Let $J'_{f,t}(j) = J_{f,t} \setminus j \cup j'(j)$ where $x_{j'}^{GWR} = x_j^{GWR} + 1$ and let $J''_{f,t}(j) = J_{f,t} \setminus j \cup j''(j)$ where $x_{j''}^{GWR} = x_j^{GWR} - 1$. Carry through all notational simplifications (e.g. with respect to only considering conventional cab vehicles produced by non-Japanese non-Big Three firms). Define instruments as follows

$$h_{f,g,t}^5 = \left\{ (z_t^g - z_{t-1}^g) > (z_t^g - z_{t-1}^g)^* \right\} \\ \times \left[\frac{1}{F} \sum_f \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\} \right]^{-1} \{j_{f,g,t} \in J_{f,t}\} \{j_{f,g,t} \notin J_{f,g,t-1}\}$$

Interact the instrument h^3 with Equation (14), where $J'_{f,t}(j)$ is the alternative product set, and take an additional summation over the firms, which yields

$$\frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^3 (\Delta \hat{\pi}(J_{f,t}, J'_{f,t}(j); J_{-f,t}) - \lambda [\theta^{GWR} + \nu_{2,t}^{GWR} + \nu_{1,j,t}^{GWR}] + \nu_{1,J_{f,t}, J'_{f,t}; J_{-f,t}}^\pi) \\ = \frac{1}{T} \frac{1}{G} \frac{1}{F} \sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J'_{f,t}(j); J_{-f,t}) - \lambda \theta^{GWR} \\ \geq 0$$

The expectation of disturbance terms that are known and not known to the agents when they make their decisions average out to zero by arguments identical to the ones above. Analogously, interact the instrument with h^3 using $J''_{f,t}(j)$ as the alternate product set and take an expectation. Together these yield

$$\lambda \theta^{GWR} \geq -\frac{1}{TGF} \sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J'_{f,t}, J_{-f,t}) \\ \lambda \theta^{GWR} \leq \frac{1}{TGF} \sum_t \sum_g \sum_f h_{f,g,t}^3 \Delta \hat{\pi}(J_{f,t}, J''_{f,t}, J_{-f,t})$$

Proposition I. If Assumptions I-IV hold, the model is identified.

Proof. Lemma IA provides $[\underline{\lambda}, \bar{\lambda}]$. Combine this with Lemma IB to yield

$$\overline{\theta^{GWR}} = \theta^{GWR} \leq \frac{[TGF]^{-1} \sum_t \sum_g \sum_f h^3 \Delta \hat{\pi}(J_{f,t}, J'_{f,t}(j); J_{-f,t})}{\bar{\lambda}}$$

and

$$\underline{\theta^{GWR}} = \theta^{GWR} \geq -\frac{[TGF]^{-1} \sum_t \sum_g \sum_f h^3 \Delta \hat{\pi}(J_{f,t}, J''_{f,t}(j); J_{-f,t})}{\underline{\lambda}}$$

Combine this with Lemma IA to yield

$$\underline{\theta^{con}} = \theta^{con} \geq -[TGF]^{-1} \sum_t \sum_g \sum_f h^2 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \cup j; J_{-f,t}) - \underline{\theta^{GWR} \bar{g}}$$

$$\overline{\theta^{con}} = \theta^{con} \leq \frac{[TGF]^{-1} \sum_t \sum_g \sum_f h^1 \Delta \hat{\pi}(J_{f,t}, J_{f,t} \setminus j; J_{-f,t}) - \overline{\theta^{GWR} \bar{g}}}{\bar{\lambda}}$$

There are two sets of remaining parameters. To identify the sunk cost parameters for the non-conventional cab types, i.e. the cab-over-engine, follow an identical procedure to the one identifying θ^{con} . To identify the interaction terms (specific to Japanese and Big Three firms), repeat the entirety of the above exercise, conditioning only on Japanese and Big Three firms and replacing the non-interaction parameters with the upper and lower bound estimates of the parameters (where appropriate).

End Proof.