Online Appendix

TRUCKS WITHOUT BAILOUTS: EQUILIBRIUM PRODUCT CHARACTERISTICS FOR COMMERCIAL VEHICLES [2017]

by Thomas Wollmann

Section 1 describes aspects of the data construction not covered in the text. Section 2 provides computational details. Section 3 presents supplementary results, e.g. tables and figures that test robustness.

1 Data

Empirical buyer distribution

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 a higher proportion of commercial vehicle owners employees in the US Census data (it is roughly 2x as large), the weight of commercial vehicle buyers in the empirical distribution is doubled.

Market size

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 $\tilde{M}_t = min\{q_t\} + (q_t - min\{q_t\}) \frac{max\{qt\} - min\{q_t\}}{max\{y_t\} - min\{y_t\}}$, which yields $M_t = \bar{M} \times \tilde{M}_t$.

Ownership changes

There are seven ownership changes across the panel. Table A.I reports these transactions. All are plausibly driven by factors outside the US commercial vehicle market.

[Table A.I about here.]

Daimler's acquisition and divestiture of Chrysler was driven by passenger vehicle market concerns. Daimler's acquisition of the Bering-branded Hyundai commercial vehicle unit was a result of the late-1990s financial crisis, in which the Hyundai parent sold assets. Volvo's acquisition of Renault was driven by European market concerns, but a controlling stake in Mack went along with it. Volvo's acquisition of Nissan Diesel and Daimler's acquisition of Mitsubishi Fuso were clearly driven by Asian market concerns.

Laws affecting vehicle length

Up to the early 1980s, individual states regulated the total length of vehicles, from the back bumper of the trailer to the front bumper of the cab. This advantaged cab-over-engine ("cabover") commercial vehicles, since the shorter cabs allowed for more revenue-generating trailer space. However, the varying lengths across states created impediments and—across some state lines—a complete block on truck-borne interstate commerce. Beginning with the Surface Transportation Assistant Act of 1982, a series of legislative and court decisions ultimately harmonized the length of commercial vehicles across states and completely relaxed the length of cabs, so that the only relevant length for commercial vehicles is that of the object being pulled, e.g. a trailer, platform, et cetera.

This process was drawn out and complicated, so I take a simple cumulative count over the panel to construct the length law index. Table A.II reports a brief summary of the process.

[Table A.II about here.]

2 Computation details

Demand estimation

Demand estimation generally follows the procedure first proposed in Berry, Levinsohn, and Pakes [1995]. It does not, though, incorporate supply side moments, i.e. those provided by taking first order conditions of the second stage profits with respect to prices [Nevo, 2001]. It also incorporates microdata to form "micromoments", which aid in understanding the relationship between the buyer attributes and product characteristics [Petrin, 2002] (considered also in [Berry, Levinsohn, and Pakes, 2004]).

The procedure minimizes a GMM objective function, which is equal to a weighted average of a set of moments. The first set are formed by multiplying a set of instruments by the vector of ξ's. The instruments comprise the following: each of the product characteristics excluding price, the sum of own-firm offerings for the relevant year, the sum of other firms' offerings for the relevant year, the product-factory-region matched wage, and the year less 1986 (so that this component of the instrument equals 1 in 1987, 2 in 1988, and so forth). The second set of moments comprise the difference between the empirical model's predicted micromoments and the empirical micromoments provided by the Census data. For each relevant characteristic-attribute relationship, there are three sets of moments based on the following: the probability of purchase conditional on the buyer attribute, the expected product characteristic value conditional on the buyer attribute, and the variance in the product characteristic conditional on the buyer attribute.[1] The five characteristic-attribute relationships I consider are between GWR and an indicator for whether the buyer was in the freight industry, heavy building industry, general construction industry, and contractor industry, respectively, and between the urban density measure and an indicator for cabover.

I begin the minimization process at a guess of the (β_x^o, β_x^u) parameters and iterate. Each guess of those parameters provides the vector of δ's, which in turn provides the vector of ξ's, as well as the micromoments. The program initially weights the moments equally. After convergence, I repeat the process, but compute the objective function by weighting the moments by their inverse variance. I report the final parameter values as well as the robust standard errors, taking care to also include error derived from the sampling procedure.

Three aspects of the final implementation drastically improved speed over initial runs. As opposed to the contraction mapping proposed in Berry, Levinsohn, and Pakes [1995], I relied on an alternative, much faster procedure proposed by Reynaerts, Varadhan, and Nash [2012] [2] (and then check that it provides identical δ values, which it does). Also, I pre-computed several matrices, e.g. the product of product characteristics and buyer attributes, the product of the product characteristics squared and the buyer attributes (needed for the micromoments), the expected value of a buyer attributes (also needed for the micromoments), etc. Of narrower interest, I re-wrote the algorithm, which was originally run in MatLab, for Julia, which is a relatively new, "high-level high-performance dynamic programming language for

[1] In practice, I estimate transformations of these moments to speed computation. For example, in place of the first moment, equal to $\mathbb{E}[Y_{r,j}|z_r^o]$, I use $\mathbb{E}[Y_{r,j}z_r^o]$, which provides the same information (conditional on the fact that z_r^o is data rather than a prediction of the model).

[2] I thank Mark Shepard for this suggestion.

technical computing." This improved the speed of the program considerably.[3]

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 one potential selection procedure to make this problem tractable, based on a best response dynamic, which is computed as follows:

1. Begin with product offerings from the current period and a predetermined order of firm moves.
2. The first firm in the order best responds to what all other firms offered in the current period.
3. The second firm best responds analogously to the first, but replaces the first firm's offerings with the first firm's best response.
4. The third firm best responds analogously to the second, but replaces the first firm's and second firm's offerings with their updated best responses.
5. This process repeats for the fourth firm, fifth firm, *et cetera*, until the order is complete.
6. When the last firm has best responded, the process returns to the first firm. It terminates only when all firms have cycled through the process without any profitable deviations.

Note that when firms are calculating how to respond, they compute sunk costs based on moves away from the current period's offerings, not their prior best responses.

At a rest point of this procedure, the necessary conditions that were used in the estimation of sunk costs hold. That is, it provides a simultaneous move Nash equilibrium in product space and is internally consistent with the model presented in the body of the paper. I based the order of choices on market share, with the largest share firm ("leader") moving first.

Two comments are necessary with respect to the procedure. First, while it is natural to consider an ordering based on firm market share, a more robust approach would consider randomly selected orderings. The program is computationally very costly as it is, so I did not consider this. Informally, however, I should add that my experience running a large number of *ad hoc* permutations—with the sunk costs computed using the midpoint of the (θ, λ) confidence intervals and no disturbances—is that this will change the identity of the firms entering products but would not change the main results. The reason is that, regardless

[3] Julia better handles loop commands. If one is careful, Julia can execute a program that relies on loops at computational times that are on the order of an identical program in Matlab that uses "vectorized" code, which forces users into complicated workarounds, including painfully populating and re-populating sparse block-diagonal matrices.

of the policy choice, the final rest point of this program will almost surely not entail large markup increases (and the large output drops and consumer losses associated with them); at least some firms will deviate and enter, and even a few entries can "discipline" prices.

Nonetheless, Section 3 below provides one obvious, feasible robustness test. It replicates the main counterfactual exercise but completely reverses the order of the firms. In line with the intuition directly above, there are no substantive differences between the two approaches.

Second, computational costs also prohibit considering all possible product deviations, although I believe this results in no meaningful (or any) loss of generality. One restriction imposes that all offerings over 40,000 GWR are left untouched. Chrysler's, Ford's, and GM's largest vehicles are 19,500, 28,000 and 40,000 GWR, respectively, and GM and Chrysler do not substantively compete with PACCAR-owned brands. Any competition is far away from the 40,000 GWR boundary I impose (and besides this, I check that the counterfactuals do not change profit incentives above this point). Another restriction imposes that no firm considers more than three entry or three exit decisions simultaneously. After considerable experimentation, it is clear that no firm would respond to any of the counterfactuals with more changes than that. Entry and exit combinations, however, are of course allowed so that the total number of product changes are greater than three. Another restriction imposes that no brand owned by an acquiring or acquired parent can enter more than one product. For instance, if Ford acquires GM and Chrysler, then their incentives to enter products fall (while the incentives to exit increase), so one need not consider more than one product entry point here (in fact, one need not consider any, but I did so to be conservative). A related restriction imposes that no brand owned by a non-acquiring parent firm or by a non-acquired parent firm can exit more than one product. If rival brands merge, incentives to enter products increase (while incentives to exit products fall). The final restriction imposes that if GM and Chrysler are liquidated, no firm can exit more than one product. Under liquidation, all firms' incentives to enter increase.

Note that these restrictions limit the generality of the results if commercial vehicle manufacturing entailed *fixed* rather than *sunk* cost investments. Standard fixed cost models require a per-period payment that does not depend on what the firms offered last period, so there are no adjustment costs. They conveniently avoid dynamic considerations. However, without further restrictions, they typically lead to a large number of re-positioning decisions each period, even if the characteristic changes are small. For example, a fixed cost model would predict frequent shifts back and forth between, say, 16,000 and 18,000 lb offerings, as manufacturers tweak the portfolio to adjust to small shifts in the buyer composition. Sunk costs prohibit this: this would require the costly removal of one product and introduction of another. (In line with the economic model, the data do not indicate many small shifts in the characteristics of the products offered by commercial vehicle manufacturers.)

Inference procedure

The inference procedure follows Andrews and Soares [2010], which involves inverting a test of the null of any (many) feasible parameter vectors. A summary of each step is provided below. For a detailed step-by-step implementation, please see their Section 4.2.

1. Choose a feasible parameter vector (θ'_f, λ'). This will test the null hypothesis that the true vector is this vector.
2. Compute a test statistic, here equal to Q_n at (θ'_f, λ').
3. Draw a large number of bootstrap samples R.
4. For each sample, compute the objective function for each at (θ'_f, λ').
5. For an α% confidence interval, compute a critical value at the $(1 - \alpha\%)$ quantile of the distribution of bootstrapped objective functions.
6. Accept any (θ'_f, λ') that lies above this critical value and reject the rest.
7. Repeat for all feasible parameter vectors.

Note that although re-centering the bootstrap moments may be intuitively appealing and result in a test for each (θ'_f, λ') of the right (asymptotic) size, the test has poor power. In the event that some moments bind by a wide margin, i.e. they are satisfied for a very large number of parameter vectors, re-centering can produce very wide intervals—a problem that, in the author's experience, is a practical issue in many applied problems. To illustrate a simple example, suppose there are a large number of moments, half of which provide that $1 - \theta \geq 0$ and half of which provide that $10 - \theta \geq 0$. The upper bound of the confidence intervals should be around one (if the variance in the underlying moments is not "too large"). Yet, for a test with re-centered bootstraps, the second half of the moments almost guarantee that a θ of, say, 5 will be inside the confidence interval, since the second half of the moments evaluated at 5 are zero. Thus, although the second half of the moments are uninformative, conditional on the first half of the moments, they are very influential in determining critical values. The solution is to discard these "slack" moments that are satisfied by a wide margin. Note that this was not a problem for Chernozhukov, Hong, and Tamer [2007], who are interested in confidence intervals over the identified set rather than the true parameter. I shift rather than drop these moments in a way that makes them (continuously) more likely to bind.[4] The amount they are shifted by is equal to $\sqrt{ln(n_k)}^{-1}\sqrt{n}\hat{\sigma}_k\bar{m}_k$ where n_k denotes the number of observations for moment k.

I test the null hypothesis at a large number of possible parameter values. The parameter space is high dimensional, so this can be computationally intensive. Two techniques speed up this process. One is to

[4] As in [Pakes, Porter, Ho, and Ishii, 2015]

begin with a large, sparse grid of values and then iteratively make the grid more granular. Another is to parallelize the procedure, spreading each evaluation of the null across processing units.

3 Supplementary results

Product offerings and markup changes

In the body of the paper, the first two panels of Figure IV report that, in the event of an acquisition of GM and Chrysler by Ford and in the case where one does not account for entry and exit, areas of the product space where GM and Chrysler are most represented (as a proportion of total offerings) correspond to sharp increases in markups. Figure A.I provides more detail. In particular, it provides the count of GM and Chrysler products as well as the count of GM and Chrysler products as a proportion of all offerings. The quotient of these two provide the proportion that appears in Figure IV. For the convenience of the reader and the sake of comparison, Figure A.I also replicates those first two panels from Figure IV.

[Figure A.I about here.]

Also in the body of the paper, the final two panels Figure IV show that, in the event of an acquisition of GM and Chrysler by Ford, predicted markup changes are much smaller when one allows for product entry and exit. To facilitate visual comparison between the markup changes between the second and fourth panel, i.e. those that occur when one does and does not allow for product entry and exit, the same y-axis scale is maintained. Since markup changes are very small in the fourth panel, ascertaining the actual variation across GWR is difficult. Figure A.II rescales the y-axis (from a range of zero to sixty to a range of zero to only ten). Again, for the convenience of the reader and the sake of comparison, Figure A.II also replicates the first three panels in Figure IV.

[Figure A.II about here.]

Sensitivity of findings to reversal of moves in first stage

The counterfactual exercise in the body of the paper reflects an ordering of moves by the firms based on their market share. Ideally I would consider a large number of permutations of this order, although each run of the counterfactual exercise is computationally very costly. Instead, as a basic test of sensitivity, I replicate the counterfactual exercise but completely reverse the order of the moves.

Table A.III reports these results, which show that the main findings are insensitive to the reordering. For example, allowing for product entry and exit again sharply decreases markups and increases output.

It also decreases compensating variation 76% in the event of an acquisition by Ford (relative to 81% in the baseline results), 78% in the event of liquidation (relative to 88% in the baseline results), and leaves it essentially unchanged in the event of an acquisition by PACCAR (the same as the baseline results).

[Table A.III about here.]

Sensitivity of findings to where parameters lie within the bounds

For the counterfactual exercise in the body of the paper, sunk costs are drawn from a distribution centered at the midpoint of the confidence intervals. Moment *inequalities* used to recover the sunk cost parameters provide only bounds, and the distributions of the true parameters are not necessarily symmetrically distributed around the intervals' midpoints. Thus the midpoints may not be the best guess. If the true parameters lie to the left of these points, i.e. are more negative, then product entry is less likely and product exit is more likely. (In the event they lie to the right, then I understate the effects of entry and exit. In the interest of brevity, I do not further address this case.) Testing robustness with respect to the sunk cost parameters one-by-one is ideal but too burdensome, so I instead shift *all* sunk cost parameters left and re-run the counterfactual analysis. To be precise, I use parameter values that are at the one-quarter of the way through the set identified by the tightest bounds. To illustrate, if a parameter's tightest lower bound is -4 and the tightest upper bound is zero, then mean sunk costs are determined using a parameter value equal to -3. This results in sunk costs that are 11.3% higher in absolute value terms than those used for the baseline results.

Table A.IV reports these results, which show that the main findings are mostly insensitive to the change. Allowing for product entry and exit decreases markups and increases output. It also decreases compensating variation 58% in the event of an acquisition by Ford, 69% in the event of liquidation, and leaves it essentially unchanged in the event of an acquisition by PACCAR.

[Table A.IV about here.]

In line with intuition, though, the outcomes are all worse from the point of view of the buyers. The deviations from the baseline results are small in magnitude but directionally apparent, especially when the final three columns here are placed side-by-side with those in Table IV.

Sensitivity of counterfactual results to increased hurdle rates

One concern is that the policy choice could directly affect discount rates in a way that would limit product entry and promote product exit among surviving, rival firms. That is, observed product entry and exit decisions reflect hurdle rates that, in turn, are determined by the required return on capital of investors in the commercial vehicle industry. These investors placed weight on the likelihood of federal government assistance in the face of a tail event like the Great Recession. That weight would diminish going forward if the federal government did not act—especially in the case where default would be so chaotic as to induce complete liquidation. If discount rates, and by extension hurdle rates, increase, the value of future cash flows is less important relative to sunk costs, which are paid up front, so product entry is less attractive. (Symmetrically, the value of future cash flows is less important relative to *scrap values*, which are received front, so product *exit* is *more* attractive.) Thus, to assess this concern, I replicate the counterfactual exercise but assume higher hurdle rates.

This sensitivity analysis assumes that hurdle rates would increase proportionally with the cost of debt and that the cost of debt will increase commensurately with a full "letter grade" drop in credit rating. (Accurately forecasting how much hurdle rates would change by policy choice is beyond the scope of this evaluation, so this approximation provides, at the very least, a well-understood benchmark for the effect of the policy choice.) Commercial vehicle manufacturers have historically received credit ratings from the major agencies—Moody's, Standard & Poor's, and Fitch—slightly below but usually above investment grade. For example, at the time of writing, the average rating over the last decade for Isuzu is about a 'BBB,' Volvo an 'A,' International a 'BB,' PACCAR an 'A' or 'AA,' Daimler an 'A.' To forecast the proportional change in the cost of debt, I use data from The Federal Reserve Bank of St. Louis, which provides the effective yield of long term US corporate bonds of various credit ratings from December 31, 1996 to the end of the panel, December 31, 2012.[5] The mean proportional change in effective yield of receiving a letter grade drop, i.e. three "notch" downgrade, from 'AA' to 'A,' 'A' to 'BBB,' and 'BBB' to 'BB' is 17.96%.

Table A.V reports the results of this exercise. Despite the nearly 20% increase in sunk costs of entry and exit in absolute value terms, the results are qualitatively the same. Allowing for product entry and exit sharply reduces markups and compensating variation and increases output—the same general findings as Table IV reported. For example, the predicted increases in average markups fall by 63.5%, and the predicted decreases in total output fall by 43.8%. Compensating variation falls by 43.9%. All seven rows, however, do report worse outcomes from the consumers' standpoint relative to the baseline results. This again squares with intuition. Higher sunk costs in absolute value terms reduce the likelihood of product entry and increase the likelihood of product exit, so naturally these changes mute the effects of product entry and exit.

[Table A.V about here.]

[5]Federal Reserve Bank of St. Louis: Effective yields of 'AA,' 'A,' 'BBB,' and 'BB' bonds, retrieved from FRED, Federal Reserve Bank of St. Louis, website: https://fred.stlouisfed.org/categories/32348, April 2, 2017.

Sensitivity of counterfactual results to market size

The counterfactual exercise in the body of the paper is based on market conditions in 2009. Commercial vehicle manufacturers faced lower demand in this period than the years leading up to it. To be precise, the market size—the approximate number of total prospective buyers at this point in time—was 7% lower than in the preceding five years and 19% lower than its highest point, which was realized two years earlier. If demand were higher, then the likelihood of product entry is higher, since firms would weigh the same sunk costs against larger profit increases that result from adding models. By similar reasoning, the likelihood of product exit is lower. Thus, the results reported in the rightmost three columns of Table IV in the body of the paper should slightly understate the effects of accounting for product entry and exit. To confirm this intuition and assess the magnitude of the understatement, I replicate the counterfactual exercise but calculate product entry and exit decisions using the average market size over the prior five years.

Table A.VI reports these results. In magnitude, they are overall quite close to the baseline results in the paper. In line with intuition, though, all seven rows in each of the three rightmost columns report better outcomes from the consumers' standpoint relative to the baseline results reported by Table IV in the body of the paper.

[Table A.VI about here.]

As the final three columns show, the increased likelihood of entry and decreased likelihood of exit resulting from 7% more prospective buyers roughly compensate for the change in market structure, which results from either the elimination of the GM and Chrysler brands or their acquisition by Ford or PACCAR. (Market size is, of course, not a policy choice, so I merely point this out as a coincidence.) When product entry and exit are allowed, average markups across products change by less than 0.4% in either direction, and total output changes by less than 0.2% in either direction. Compensating variation is very near zero but in fact negative. From the consumers' point of view, the former effect more than compensates for the latter, albeit very modestly so.

Sunk costs by Gross Weight Rating

Table III reports sunk cost estimates parameter-by-parameter. Figure A.III reports average sunk costs estimates by GWR. These collapse the sunk cost differences across cab types into a single measure. This facilitates a comparisons between sunk costs and the counterfactual outcomes—in particular, net entry—that are reported in Figures, IV, A.I, and A.II.

[Figure A.III about here.]

Collapsed by GWR, sunk costs of entry vary from approximately $4 million to $18 million. The coefficient on GWR is positive in Table III, so sunk costs of entry are declining in absolute value terms in that characteristic. The long/extended cab option comprises a meaningful share of vehicles at 52,000 to 56,000 GWR. Since the coefficient on this characteristic is large and negative in Table III, sunk cost rises sharply there. Cabover comprise most of the offerings at 14,500 and 18,000 GWR. Since the coefficient on this characteristic is positive in Table III, sunk costs fall sharply at that point.

References

Donald W. K. Andrews and Gustavo Soares. Inference for parameters defined by moment inequalities using generalized moment selection. Econometrica, 78(1):119–157, 2010. ISSN 1468-0262.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. Econometrica: Journal of the Econometric Society, pages 841–890, 1995.

Steven Berry, James Levinsohn, and Ariel Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. Journal of Political Economy, 112(1):68, 2004.

Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. Econometrica, 75(5):1243–1284, 2007.

Robin S Lee and Ariel Pakes. Multiple equilibria and selection by learning in an applied setting. Economics Letters, 104(1):13–16, 2009.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. Econometrica, 69(2):307–342, 2001.

Ariel Pakes, Jack Porter, Kate Ho, and Joy Ishii. Moment inequalities and their application. Econometrica, 83:315–334, 2015.

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

Jo Reynaerts, Ravi Varadhan, and John C Nash. Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping. 2012.

Thomas Wollmann. Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. 2017.

Figures

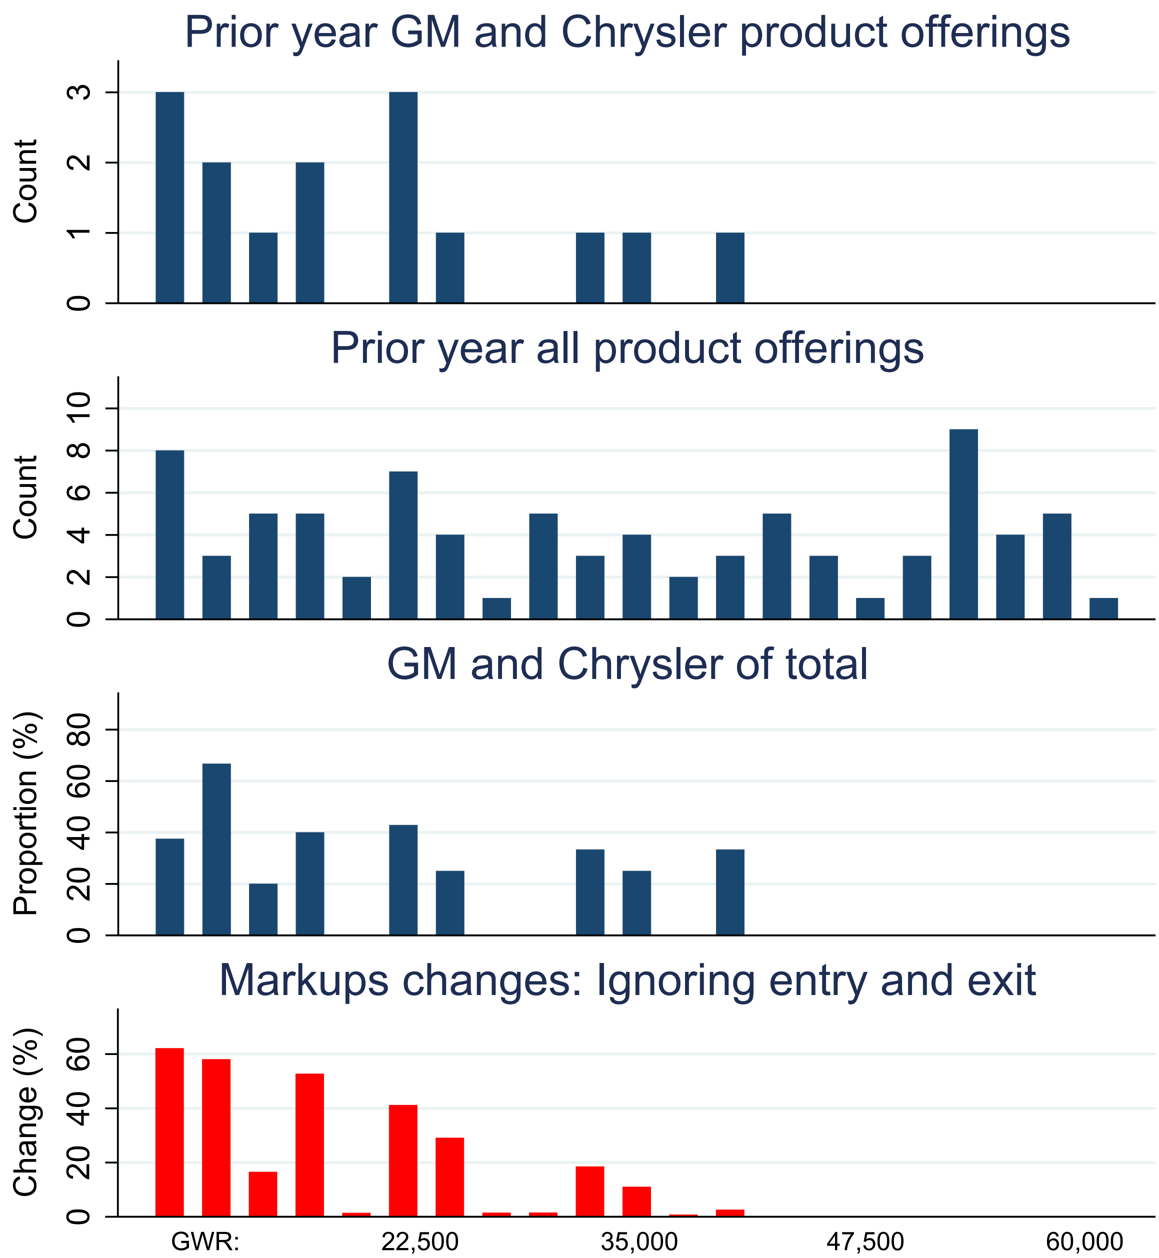

Figure A.I: *Relationship of GM and Chrysler products to total products and markup changes*

Note.– This relates to Figure IV in the body of the paper. The x-axis is the GWR (likely the single most important characteristic). The y-axis in the first panel is a combined count of the number of offerings by GM and Chrysler. The y-axis in the second panel is a count of the number of offerings by any firm. The y-axis in the third panel is the proportion of all products that are offered by GM and Chrysler. For example, the first bar has a high of 37.5, equal to 3 GM and Chrysler products (first panel) divided by 8 total products (second panel). The y-axis in the fourth panel is the percentage change in dollar markup, i.e. $(p_{Ford\ Acquisition} - mc)/(p_{Bailout} - mc)$, *for the most affected product in each GWR bin (if the economic model does not allow for product repositioning). Thus, the final two panels merely replicate the first two panels of Figure IV in the body of the paper.*

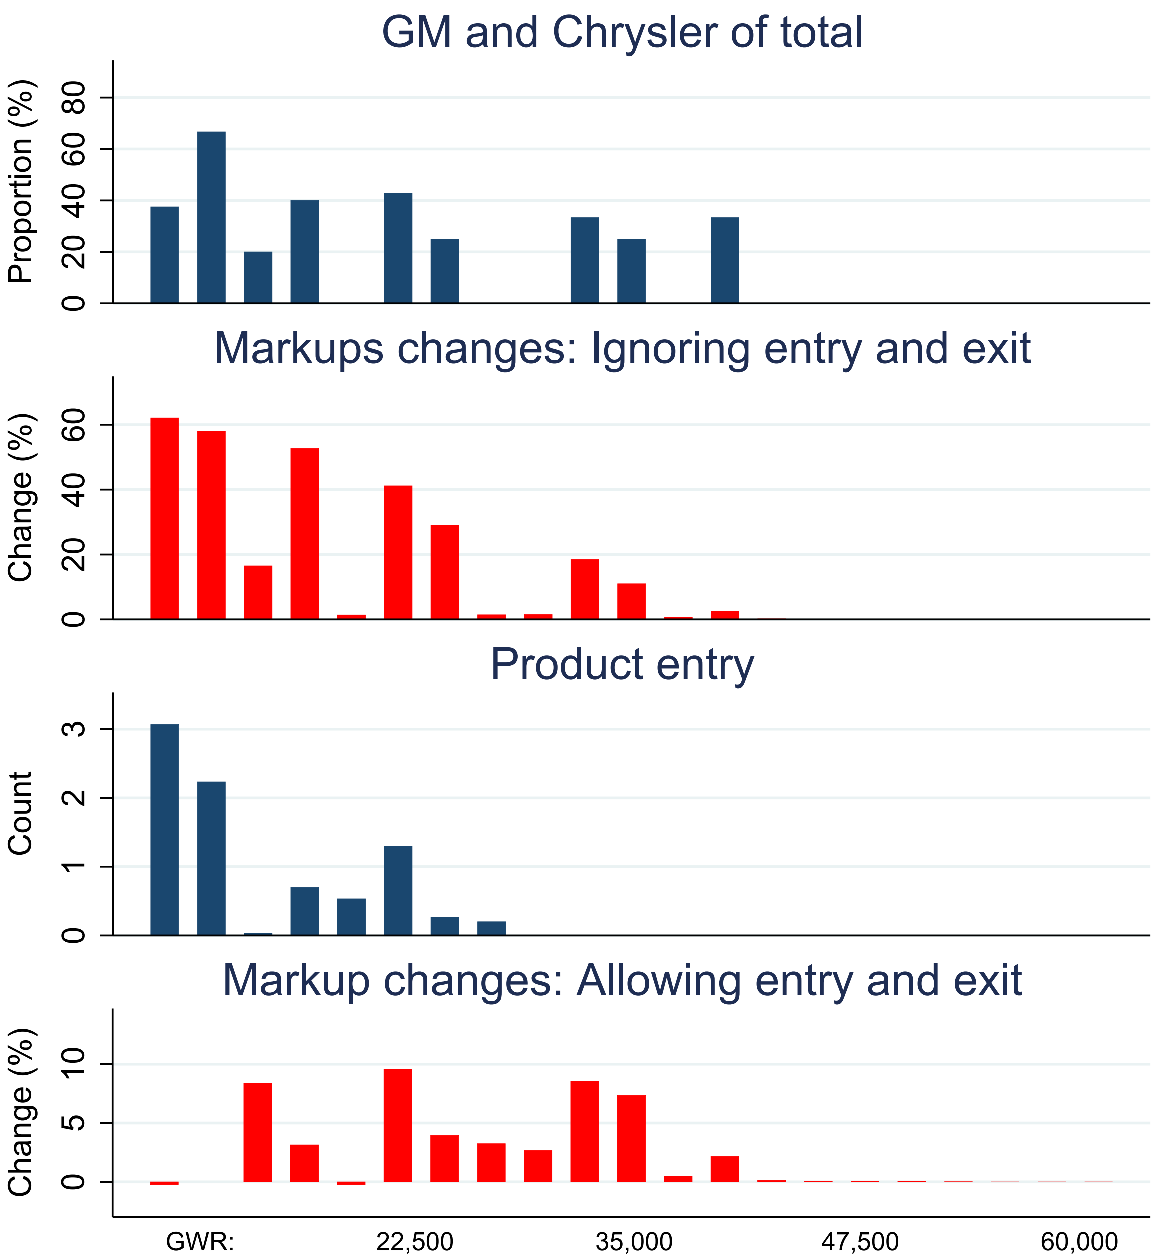

Figure A.II: *Products offerings, entry, and markups*

Note.– This paper merely replicates Figure IV in the body of the paper, but rescales the y-axis of the in the fourth panel to ease comparisons across product space. Thus, the x-axis is the GWR (likely the single most important characteristic). The y-axis in the first panel is the proportion of all products that are offered by GM and Chrysler. The y-axis in the second panel is percentage change in dollar markup, i.e. $(p_{Ford\ Acquisition} - mc)/(p_{Bailout} - mc)$*, for the most affected product in each GWR bin (if the economic model does not allow for product repositioning). The y-axis in the third panel is a combined count of the net number of product entries. The y-axis in the fourth panel is analogous to the second (but the economic model allows for product repositioning).*

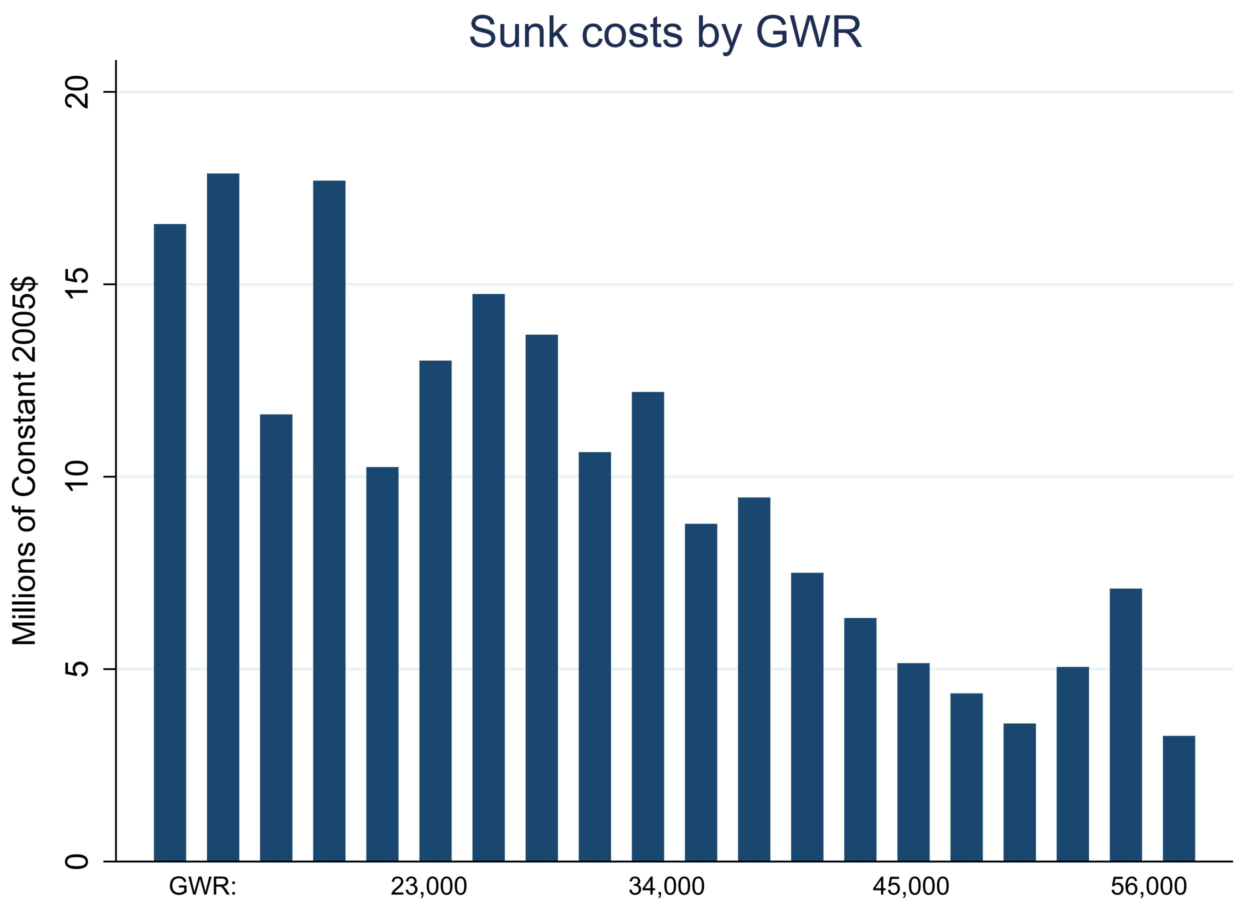

Figure A.III: *Relationship of GM and Chrysler products to total products and markup changes*

Note.– The x-axis is the GWR (likely the single most important characteristic). The y-axis measures sunk costs of entry in millions of constant 2005 US dollars. To facilitate a comparison to earlier tables, the averages are taken by GWR over all offerings in 2009.

Tables

Table A.I: *Ownership changes (FOR ONLINE PUBLICATION)*

Year	Parent	Action	Target	US CV Share
1997	Daimler	Acquisition	Chrysler	1.60%
2001	Daimler	Acquisition	Hyundai Truck	3.00%
2001	Volvo	Acquisition	Mack (Renault)	24.10%
2004	Daimler	Acquisition	Mitsubishi Fuso	4.20%
2006	Daimler	Spinoff	Chrysler	3–4%
2006	Volvo	Acquisition	Nissan Diesel	6.10%

Note.– "US CV Share" denotes the proportion of total target firm revenues that are attributed to the commercial vehicle segment of the domestic automotive industry. The remaining revenues are either attributed to the non-commercial, i.e. passenger vehicle segment of the US automotive industry, which is primarily the case with Chrysler, or foreign commercial markets, which is the case for the remaining cases. For the Chrysler divestiture, this figure is imprecise, since detailed trailing-period financial statements are not available, but in any case is low, i.e. $\leq 5\%$.

Table A.II: *Legislative and court decisions pertaining to cab-over-engine (FOR ONLINE PUBLICATION)*

Year	Title	Summary
1982	Surface Transportation Assistance Act (STAA)	Improves uniformity of truck safety and weight ("TS&W") regulation by, for example, mandating minimums on the interstate highway system, including a 48' minimum on trailers.
1991	Intermodal Surface Transportation Efficiency Act (effect. 1992)	Extends STAA "National Network" to include all Federal aid Primary System roads; also provides "grandfather" exemptions to Western states for extra long trailer configurations, e.g. "Western Doubles."
1995	National Highway Designation Act	Expands interstate system (it is the last major change to it), which now includes just 4% of road mileage but 40% of highway traffic and 75% of heavy truck traffic.
1995	ICC Termination Act (Effec. 1996)	Transfers licensing and other functions to the FHA and STB, i.e. (federal) USDOT agencies, and further removes states ability to restrict length.
2003	EPA Emissions decree	Required compliance date for 2004 medium–heavy–duty emissions standards. Not per se a length law, but additional drag of cab–over–engine vehicles made long–term compliance with chassis–specific emissions laws impossible.

Table A.III: *Counterfactual outcomes when order of first stage moves are reversed*

	Product entry and exit ignored			Product entry and exit allowed		
	Ford Acquisition	PACCAR Acquisition	Liquidation	Ford Acquisition	PACCAR Acquisition	Liquidation
Markups						
Most affected model	62.1%	23.0%	26.9%	35.0%	15.7%	7.5%
Most affected vehicle type	53.2%	13.0%	32.3%	27.5%	8.4%	5.1%
Market	10.9%	3.0%	4.0%	4.4%	1.6%	0.1%
Output						
Most affected model	-34.4%	-18.0%	NR	9.5%	12.0%	15.6%
Most affected vehicle type	-30.0%	-8.6%	-100.0%	-71.1%	-58.7%	-63.7%
Market	-5.1%	-1.4%	-11.2%	-1.5%	-1.6%	-2.7%
Comp. Var. (in MMs of 2005$)	119.0	33.0	253.0	28.0	33.0	56.0

Note.– This table replicates the results of Table IV in the body of the paper; however, it reverses the order of the moves of the firms when they make product entry and exit decisions. The original ordering was based on market share. "NR" denotes "not relevant."

Table A.IV: *Counterfactual outcomes using more negative sunk cost parameters*

	Product entry and exit ignored			Product entry and exit allowed		
	Ford Acquisition	PACCAR Acquisition	Liquidation	Ford Acquisition	PACCAR Acquisition	Liquidation
Markups						
Most affected model	62.1%	23.0%	26.9%	16.6%	15.4%	6.3%
Most affected vehicle type	53.2%	13.0%	32.3%	9.7%	7.8%	7.2%
Market	10.9%	3.0%	4.0%	1.0%	0.6%	0.6%
Output						
Most affected model	-34.4%	-18.0%	NR	11.4%	15.3%	12.2%
Most affected vehicle type	-30.0%	-8.6%	-100.0%	-59.1%	-55.6%	-63.0%
Market	-5.1%	-1.4%	-11.2%	-2.3%	-1.6%	-3.6%
Comp. Var. (in MMs of 2005$)	119.0	33.0	253.0	50.0	36.0	78.0

Note.– This table replicates results of Table IV in the body of the paper; however, it shifts the distribution of sunk costs to the left. That is, rather than using the midpoint of the confidence intervals to compute the mean sunk costs, it uses parameters values that are at the one-quarter of the way through the set identified by the tightest bounds. These are 11.3% higher in absolute value terms on average than those used for the baseline results. "NR" denotes "not relevant."

Table A.V: *Counterfactual outcomes when hurdle rates are increased*

	Product entry and exit ignored	Product entry exit allowed	
	Baseline results	Baseline results	Higher hurdle rates
Markups			
Most affected model	26.9%	6.2%	8.5%
Most affected vehicle type	32.3%	5.3%	11.3%
Market	4.0%	-0.0%	1.5%
Output			
Most affected model	0.0%	NR	24.0%
Most affected vehicle type	-100.0%	-64.7%	-62.4%
Market	-11.2%	-1.6%	-6.3%
Comp. Var. (in MMs of 2005$)	253.0	28.0	142.0

Note.– Columns 1 and 2 in this table replicate columns 3 and 6 in Table IV in the body of the paper. Column 3 in this table replicates column 3 in Table IV in the body of the paper, but assumes that hurdle rates—and, as a result, sunk costs of entry and exit—are 17.96% higher in absolute value terms. This increase approximates the result of a full letter grade drop in the firms' credit ratings. "NR" denotes "not relevant."

Table A.VI: *Counterfactual outcomes using largest market size*

	Product entry and exit ignored†			Product entry and exit allowed		
	Ford Acquisition	PACCAR Acquisition	Liquidation	Ford Acquisition	PACCAR Acquisition	Liquidation
Markups						
Most affected model	62.1%	23.0%	26.9%	16.2%	15.3%	6.2%
Most affected vehicle type	53.2%	13.0%	32.3%	9.1%	7.5%	5.0%
Market	10.9%	3.0%	4.0%	0.4%	0.2%	-0.3%
Output						
Most affected model	-34.4%	-18.0%	NR	1.8%	10.9%	7.0%
Most affected vehicle type	-30.0%	-8.6%	-100.0%	-63.1%	-54.9%	-63.9%
Market	-5.1%	-1.4%	-11.2%	0.2%	-0.2%	-0.1%
Comp. Var. (in MMs of 2005$)‡	119.0	33.0	253.0	-15.0	-4.0	-7.0

Note.– This table replicates results of Table IV in the body of the paper; however, it substitutes for the market size in 2009 with the average market size of the prior five years when determining product entry and exit decisions. "NR" denotes "not relevant."
† *Since these columns ignore product entry and exit, and since the industry is assumed free of capacity constraints throughout the paper, changing market size does not affect these figures.*
‡ *To facilitate comparisons with other tables, the market size used to compute the level of compensating variation is the same as in the baseline results.*