Online Appendix to:
"Mergers, Foreign Competition, and Jobs: Evidence from the U.S. Appliance Industry"

Felix Montag

This version: March 25, 2026

I Appendix to Section 2: Data Set Construction

I.A Product market data set

This appendix provides further details on the construction of the product market data set.

Product data. As described in Section 2, I define a washing machine product by the combination of brand, retailer, and washer type: front-loader, regular top-loader (with an agitator), or high-efficiency top-loader (without an agitator). These characteristics primarily differentiate products.

Figure A.1 illustrates the key difference between front-loaders, which are loaded from the front, and top-loaders, which are loaded from above. Front-loaders can be stacked (e.g., placing a dryer on top), are more water- and energy-efficient, offer better cleaning performance, and typically cost more. Top-loaders cannot be stacked; however, among top-loaders, a critical distinction is the presence of an agitator, as shown in Figure A.2. Top-loaders without agitators, known as high-efficiency top-loaders, perform better than regular top-loaders in terms of efficiency and cleaning, but not as well as front-loaders.[41]

Within each market (defined nationally at the yearly level), I group survey responses based on these three characteristics.[42] This yields 2,939 products from 2005 to 2015. Under this definition, many products have small market shares, often based on a single household's response. Additionally, responses sometimes lack brand information. Thus, I exclude products labeled as "Other Brands" or "Store Brand/Generic," and any product with a market volume share below 0.01 percent. This reduces the data set to 1,590 products, which still represent between 97.3 and 99.0 percent of total washing machine sales volume in the *TraQline* data, ensuring minimal selection bias.

[41] See, for example, McCabe (2016) for a detailed comparison of washer types.

[42] I classify Maytag products as Whirlpool products in the first quarter of 2006, before the official acquisition, to avoid artificially inflating the number of washing machine products that year. Given that merger discussions were public from mid-2005 onward, it is reasonable to assume minimal competitive interaction between Maytag and Whirlpool by early 2006.

Figure A.1: Difference between front-loader and top-loader

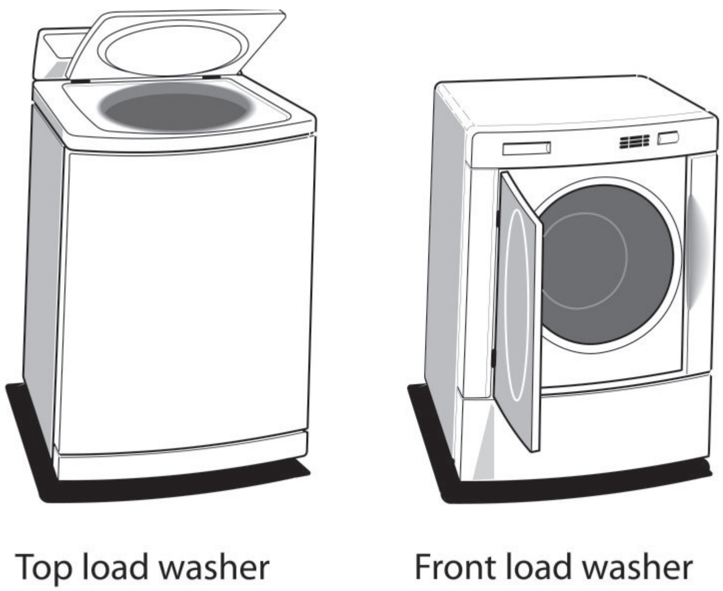

Figure A.2: Top-loaders with and without agitator

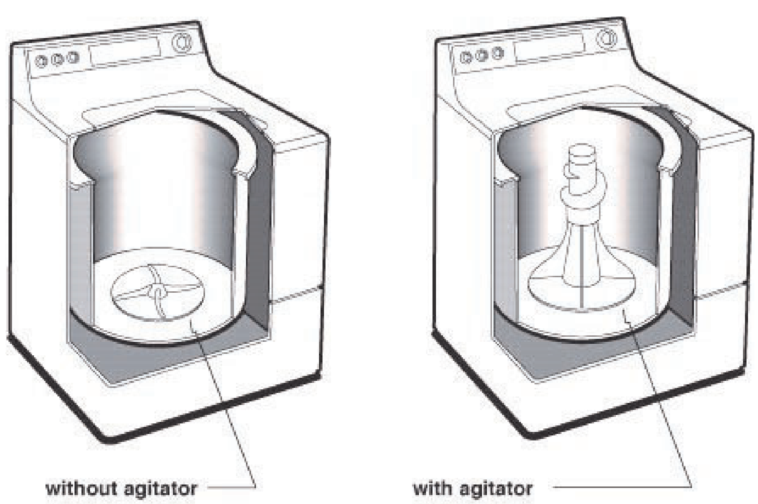

For characteristics collected only from a random subset of *TraQline* respondents (whether a washer is part of a stacked pair; has a stainless steel, white, or other-colored exterior; is Energy Star certified; includes additional noise insulation or a child lockout; and the number of special programs), I use the within-group average response.

Household income. The CPS data include exact household incomes, whereas the *TraQline* data only report income ranges. To estimate the relationship between price sensitivity and household income using a single parameter, I assign exact income values to each household based on the empirical income distribution and reported ranges, proceeding as follows:

1. Compute midpoints of the nonoverlapping income ranges reported for each respondent.
2. For each year, fit a log-normal distribution to the set of midpoint values.
3. Draw 1,000,000 random incomes from this fitted log-normal distribution.
4. Allocate each drawn income to the corresponding reported income bucket.
5. For each household, randomly draw (with replacement) an income value from the set of incomes matching its income bucket.

I.B Plant locations and plant location weights

Plant locations. Constructing the dataset on washing machine manufacturing plants involves three steps. First, I identify plant locations from annual reports, news articles, and transcripts from the USITC's antidumping hearings on washing machine imports from Mexico and South Korea.

For LG and Samsung, pre-2012 production locations primarily rely on the USITC investigation. For 2012–2015, I use firm-level washer imports based on PIERS shipping data, as documented in Flaaen, Hortaçsu, and Tintelnot (2020). For Electrolux, Maytag, and Whirlpool, plant locations come largely from their annual reports. Since General Electric's annual reports lack detailed appliance plant information, I combine USITC documents and news reports to identify its plant locations.

Second, to confirm which plants supplied the U.S. market, I cross-check these locations against USITC data on front- and top-loading washer imports by country, excluding plants that cannot plausibly produce substantial volumes for the U.S. market.

Finally, I validate that each significant source country's washing machine exports to the U.S. correspond plausibly to the identified plants. Figure A.3 illustrates the evolution of annual U.S. imports of front- and top-loaders by country. Over half of the front-loaders sold in the U.S. during this period were imported. Initially, Germany was the primary exporter, reflecting Whirlpool's Schorndorf plant, which closed in 2012. Before 2012, LG and Samsung imported primarily from South Korea and Mexico. Following antidumping duties in 2012, imports from these countries declined, shifting LG and Samsung production to China (see Flaaen, Hortaçsu, and Tintelnot (2020) for details). In contrast, significant top-loader imports into the U.S. began only after 2011, primarily from China, driven by LG and Samsung's increased market presence.

Figure A.3: Washer imports to the United States by source country

(a) Front-loaders **(b)** Top-loaders

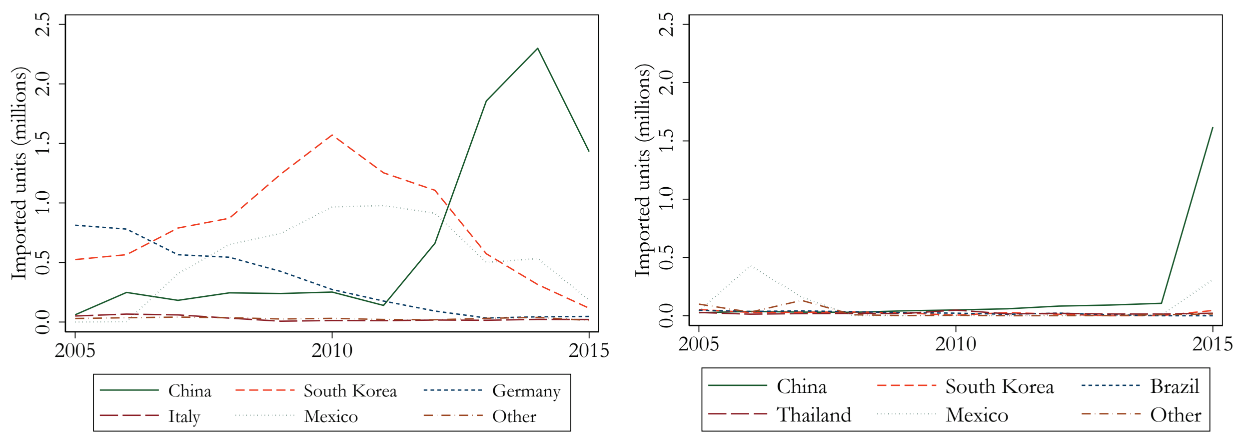

Notes: The left panel plots annual U.S. import volumes of front-loader washers (HS8450110080, HS8450200080, HS8450200090) by source country. The right panel plots annual U.S. import volumes of top-loader washers (HS8450110040, HS8450200040) by source country. Both panels show the top six source countries, grouping all others into "Other".

For reference, according to Appliance Portrait (2006), approximately 9.3 million washers were sold in the U.S. in 2005. Based on the *TraQline* data, about one-third of these were front-loaders, with the remainder being top-loaders. The front-loader share rose to over 40 percent by 2010 and then declined to approximately 25 percent by 2015. Although significant numbers of front-loaders were imported throughout the period, most top-loaders were produced domestically.

Figures A.4, A.5, and A.6 display washer plant locations for manufacturers holding more than 3 percent of the U.S. market volume in any year during the sample period.

Plant location weights. Table A.1 reports the weights used to compute product-level average real exchange rates, indicating the share of each product sourced from different countries annually. Weights are based on identified plant locations, aggregate USITC import data, and firm-level washer imports (2012–2015) from PIERS data reported in Flaaen, Hortaçsu, and Tintelnot (2020).

Figure A.4: Washer plants manufacturing for the U.S. market, 2007

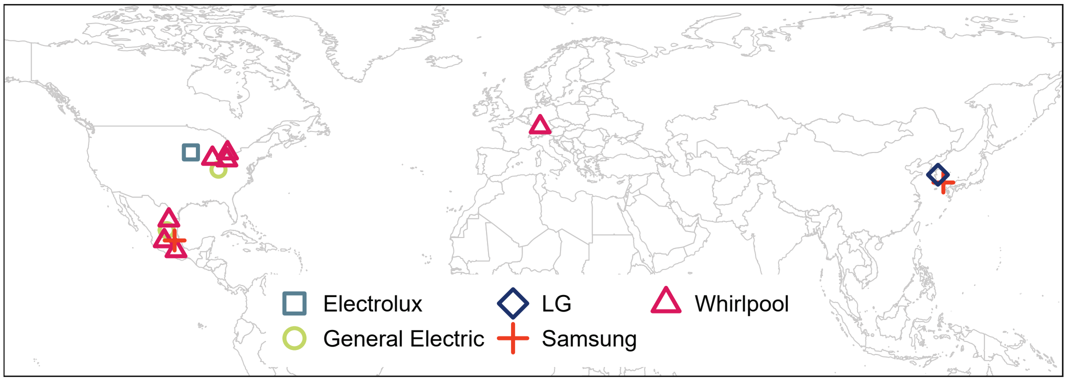

Notes: The map displays all plants manufacturing washing machines for the U.S. market in 2007, for manufacturers with a market share exceeding 3 percent in any year of the sample period.

Figure A.5: Washer plants manufacturing for the U.S. market, 2009

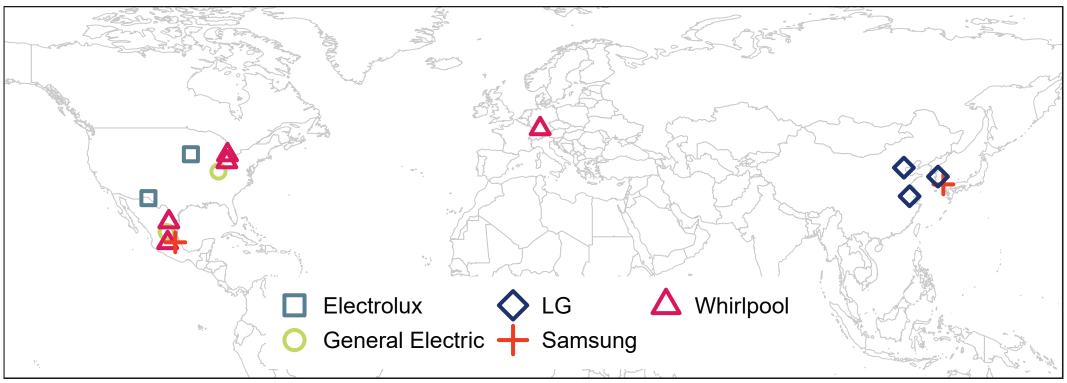

Notes: The map displays all plants manufacturing washing machines for the U.S. market in 2009, for manufacturers with a market share exceeding 3 percent in any year of the sample period.

Figure A.6: Washer plants manufacturing for the U.S. market, 2011

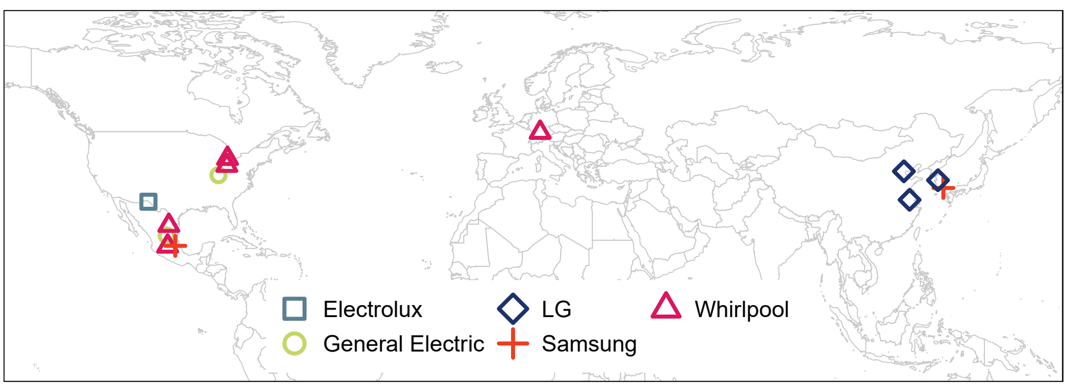

Notes: The map displays all plants manufacturing washing machines for the U.S. market in 2011, for manufacturers with a market share exceeding 3 percent in any year of the sample period.

Table A.1: Plant location weights

Owner	Brand	Product	Years	China	Germany	Mexico	South Korea	USA
Electrolux	All brands	Front-Loader	2005–2007					1
Electrolux	All brands	Front-Loader	2008–2015			1		
Electrolux	All brands	Top-Loader	2005–2010					1
Electrolux	All brands	Top-Loader	2011–2015			1		
General Electric	All brands	Front-Loader	2005–2012			1		
General Electric	All brands	Front-Loader	2013–2015					1
General Electric	All brands	Top-Loader	2005–2015					1
Whirlpool	Roper	Front-Loader	2005–2007					1
Whirlpool	All other WP brands	Front-Loader	2005–2007		1			
Whirlpool	All other WP brands	Front-Loader	2008–2010		0.5	0.5		
Whirlpool	All other WP brands	Front-Loader	2011–2012		0.33	0.33		0.33
Whirlpool	All brands	Front-Loader	2013–2015					1
Whirlpool	Admiral, Amana, Maytag	Front-Loader	2007–2010			1		
Whirlpool	Admiral, Amana, Maytag	Front-Loader	2011–2012			0.5		0.5
Whirlpool	All brands	Top-Loader	2005–2015					1
LG	All brands	Front-Loader	2005–2012				1	
LG	All brands	Front-Loader	2013	0.67			0.33	
LG	All brands	Front-Loader	2014–2015	1				
LG	All brands	Top-Loader	2005–2007				1	
LG	All brands	Top-Loader	2008–2015	1				
Samsung	All brands	Front-Loader	2005–2011			0.5	0.5	
Samsung	All brands	Front-Loader	2012	0.33		0.33	0.33	
Samsung	All brands	Front-Loader	2013–2015	1				
Samsung	All brands	Top-Loader	2005–2011				1	
Samsung	All brands	Top-Loader	2012–2015	1				
Maytag	All brands	Front-Loader	2005–2006					1
Maytag	All brands	Top-Loader	2005–2006					1

Plant-level data on output and employment. Data on employment and output at Whirlpool's Clyde, Ohio, and Schorndorf, Germany plants (producing for the U.S. market), and Radomsko, Poland (producing for Europe), are sourced from news reports. Data on Whirlpool's Amiens, France, and Poprad, Slovakia plants are from case studies by Ferencikova (2002) and Rubens, Ferencikova, and Bardy (2019), who cooperated with Whirlpool Slovakia to obtain data on employment and output. Data for BSH plants in Berlin, Germany, and Alcalá, Spain, originate from historical internal documents obtained from the BSH company archive.

I.C Details on the instrumental variable for price

Figure A.7 illustrates the evolution and source of variation in the deflated average real exchange rate. The left panel plots the average deflated RER across all production locations by manufacturer. This average is computed using the country-level RER of each manufacturer's plants, plant-level production shares, and sales volume weights, all deflated to 2012 U.S. dollars. Although this plot aggregates across products of a manufacturer, significant variation remains over time. The right panel separates the average RER into top-loaders and front-loaders for Whirlpool and Maytag products, highlighting within-manufacturer variation.[43] For example, while Maytag and Whirlpool top-loaders were consistently produced in the U.S., front-loaders were sourced from multiple locations, including Mexico and Germany.

The substantial variation in RER over time aligns with anecdotal evidence on the strategic importance of production location costs. Maytag's financial struggles before the merger were partly attributed to high domestic production costs and limited offshoring.[44] Similarly, Electrolux implemented a global cost-reduction initiative in 2004, aiming to relocate over half of its production to lower-cost countries by 2009 (Electrolux, 2007). These examples underscore the relevance of international production decisions in driving costs in the appliance market, illustrating the source of variation used in the instrument: fluctuations in RER between the U.S. and specific production locations and shifts in plant locations.

[43] Maytag products include Admiral, Amana, Magic Chef, and Maytag brands; Whirlpool products include all other Whirlpool-owned brands.

[44] Maytag's 2004 annual report explicitly highlighted these pressures: "Globalization of manufacturing is allowing companies to reduce costs by reaching around the world farther, faster and cheaper than ever before. It's no longer a trend we can watch with interest but a reality to which we are responding" (Maytag, 2005, p. 3).

Figure A.7: Average real exchange rate over time

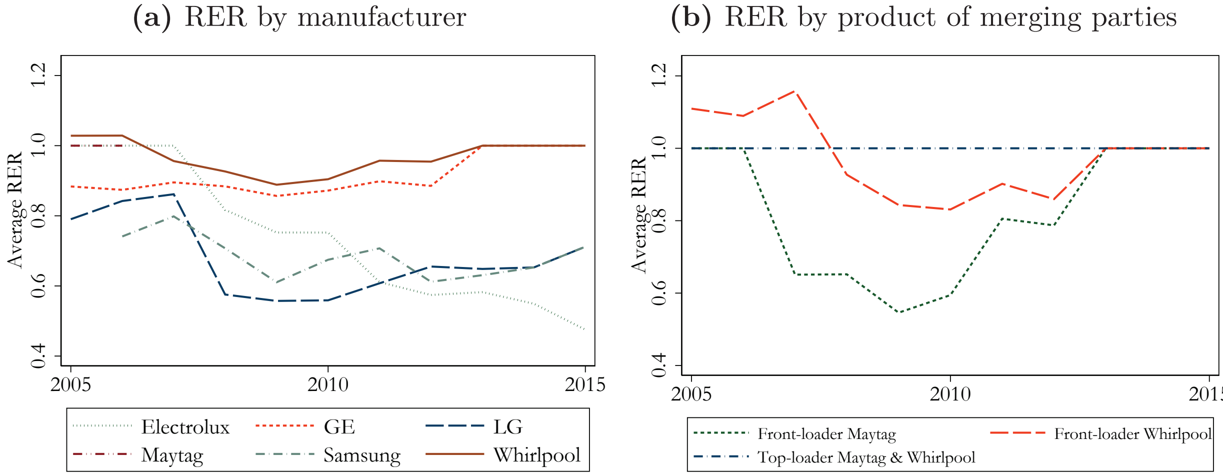

Notes: The left panel plots the average deflated RER of production locations by manufacturer over time for manufacturers with at least a 3 percent market share in any year. The right panel plots this RER separately for Maytag and Whirlpool products, based on their respective production locations and plant-level shares. Maytag products include Admiral, Amana, Magic Chef, and Maytag brands; Whirlpool products cover all other Whirlpool-owned brands.

II Appendix to Section 3: Descriptive Results

II.A Additional descriptive results on price effects of the merger

To ensure comparability of my descriptive results to the analysis in Ashenfelter, Hosken, and Weinberg (2013) and to provide evidence on dryers, I replicate the event study using freestanding ranges in the U.S. as a control market.[45] This serves as a suitable control if washer and dryer prices would have evolved similarly to freestanding ranges absent the merger.

Figure A.8 displays the estimated event study from Equation 1 using freestanding ranges as the control. Two main insights emerge: first, washers and dryers exhibit similar price dynamics; second, washer and dryer prices do not increase relative to ranges.

To further understand the merger's effects, I separately estimate price differences around the merger for Maytag and Whirlpool products using freestanding ranges as the control. Specifically, I estimate:

$$log(p_{it}) = \alpha_1 \text{Maytag}_{it} \times \text{post}_t + \alpha_2 \text{Whirlpool}_{it} \times \text{post}_t + \beta x_{it} + \tau_i + \gamma_t + \epsilon_{it}\,. \tag{19}$$

Here, α_1 captures the average price effect for Maytag products and α_2 for Whirlpool products.

[45] Ashenfelter, Hosken, and Weinberg (2013) use ranges, cooktops, ovens, and freezers as control group.

Figure A.8: Price effects of the merger for washers and dryers

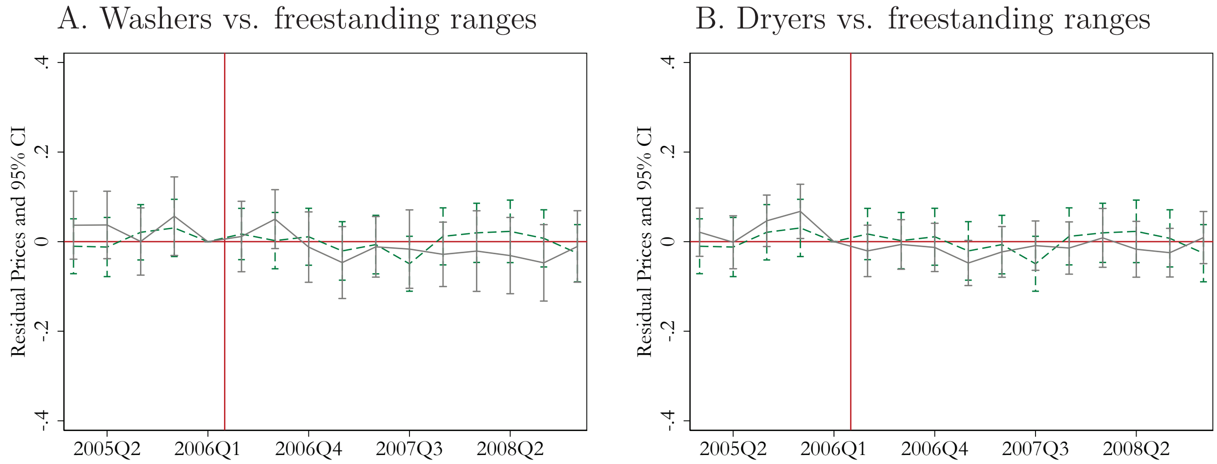

Notes: The figure shows the residualized logarithm of prices for Maytag and Whirlpool washers and dryers in the United States compared to freestanding ranges from other manufacturers, unaffected by the merger. The solid line depicts the residualized price evolution for washers and dryers; the dashed line represents the control group (freestanding ranges). The merger occurred at the end of 2006Q1, normalized to zero. Confidence bounds are at the 95 percent level, with standard errors clustered at the model level.

Following Ashenfelter, Hosken, and Weinberg (2013), I restrict the observation period from 2005Q2 to 2008Q3, noting that I use quarterly instead of monthly data.

Table A.2 summarizes these price changes. Columns (1) and (4) pool Maytag and Whirlpool products and indicate no average price increase for washers or dryers. Columns (2) and (5) report separate estimates for Maytag and Whirlpool products. Columns (3) and (6) replicate these estimates with more granular product fixed effects instead of brand fixed effects. The results consistently suggest no price increase and possibly price decreases for Maytag and Whirlpool washers and dryers compared to freestanding ranges.

Comparing washers to ranges, Ashenfelter, Hosken, and Weinberg (2013) report price declines for old and new Whirlpool washers and no price changes for old and new Maytag washers. While direct reconciliation is limited by differences in data sources and the fact that *TraQline* does not allow me to control for product age to distinguish between old and new products, my results align broadly with their findings.

For dryers, Ashenfelter, Hosken, and Weinberg (2013) report price increases for new Whirlpool dryers, decreases for old Whirlpool dryers, no change for old Maytag dryers, and a modest increase for new Maytag dryers. Depending on the relative importance of new vs. old dryers, these results may align with each other.

However, interpreting these findings as causal evidence that Whirlpool's acquisition of Maytag did not affect laundry product prices would likely be erroneous. Comparisons of U.S. washer prices to prices in other similar international markets, unaffected by the merger, strongly suggest that U.S. prices would have fallen in the absence of the merger. Therefore,

Table A.2: Price evolution for laundry products relative to freestanding ranges

	Washers vs. ranges			Dryers vs. ranges		
	(1)	(2)	(3)	(4)	(5)	(6)
Merging parties × post	-0.020			-0.038		
	[-0.080, 0.040]			[-0.089, 0.012]		
Maytag × post		-0.024	-0.037*		-0.036	-0.018
		[-0.076, 0.029]	[-0.078, 0.003]		[-0.099, 0.026]	[-0.068, 0.032]
Whirlpool × post		-0.017	-0.014		-0.040	-0.002
		[-0.094, 0.059]	[-0.050, 0.023]		[-0.098, 0.017]	[-0.055, 0.051]
Characteristics controls	Yes	Yes	Yes	Yes	Yes	Yes
Quarter × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Brand fixed effects	Yes	Yes	No	Yes	Yes	No
Product fixed effects	No	No	Yes	No	No	Yes
Observations	2114	2114	1972	5011	5011	4636

Notes: Columns (1) to (3) compare the logarithm of prices for washers and freestanding ranges. Columns (4) to (6) compare the logarithm of prices for dryers and freestanding ranges. The differences in observations in Columns (3) and (6) from the preceding columns are due to the iterative dropping of singleton observations in the clustering of standard errors. Ninety-five percent confidence intervals are reported in parentheses. Standard errors are clustered at brand level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

freestanding ranges may be influenced by distinct industry dynamics and thus may not be an appropriate control for washers and dryers in this context.

Since pre-treatment price data begin only in 2005, just prior to the merger, one might question how appropriate washer prices in the European Union are as a control for U.S. washer prices. Although a specific consumer price index (CPI) for washing machines is unavailable, comparing the CPIs for major appliances in the U.S. and household appliances in the European Union offers insight into broader trends.

Figure A.9 illustrates this comparison. Consistent with the descriptive evidence in Section 3 and anecdotal reports, deflated appliance prices declined in both the U.S. and EU prior to Whirlpool's acquisition of Maytag. Beginning in 2005, when merger negotiations commenced, this downward trend halted in the U.S., whereas prices continued to decline in the EU. After 2010, deflated prices in the U.S. resumed their decline.

II.B Additional descriptive results on product entry

Figure A.10 illustrates the combined U.S. market share of LG and Samsung for washers, dryers, and freestanding ranges. Between 2005 and 2008, LG and Samsung significantly expanded their share for washers and dryers—from approximately 2 percent to about 12 percent. In contrast, their combined market share for freestanding ranges remained low, reaching only about 2 percent by the end of 2008.

Figure A.9: Consumer price index for major appliances in the U.S. and the EU

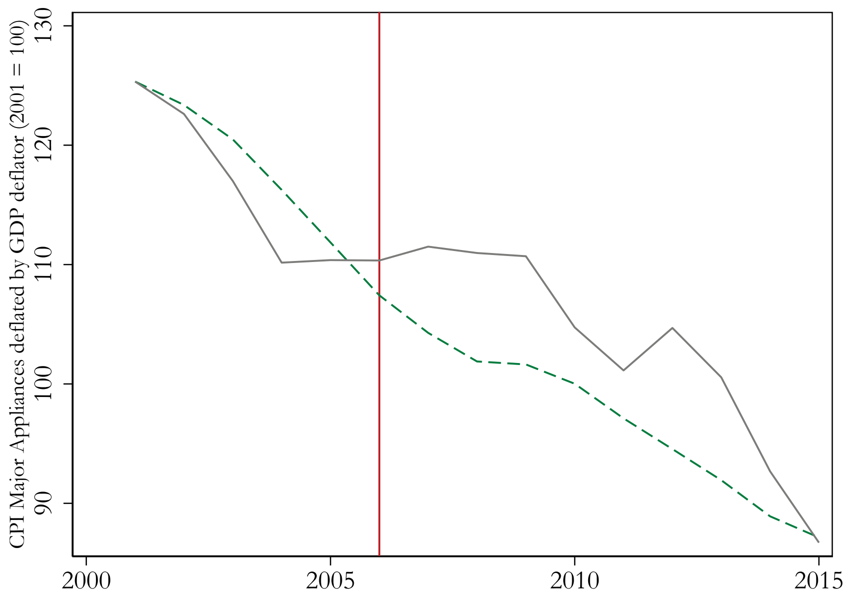

Notes: The figure compares consumer price indices for major appliances in the U.S. and the European Union from 2001 to 2015, deflated by the U.S. GDP deflator. Both indices are normalized to 100 in 2001. The solid line represents the United States, and the dashed line represents the European Union.

Figure A.10: Combined LG and Samsung market share by appliance

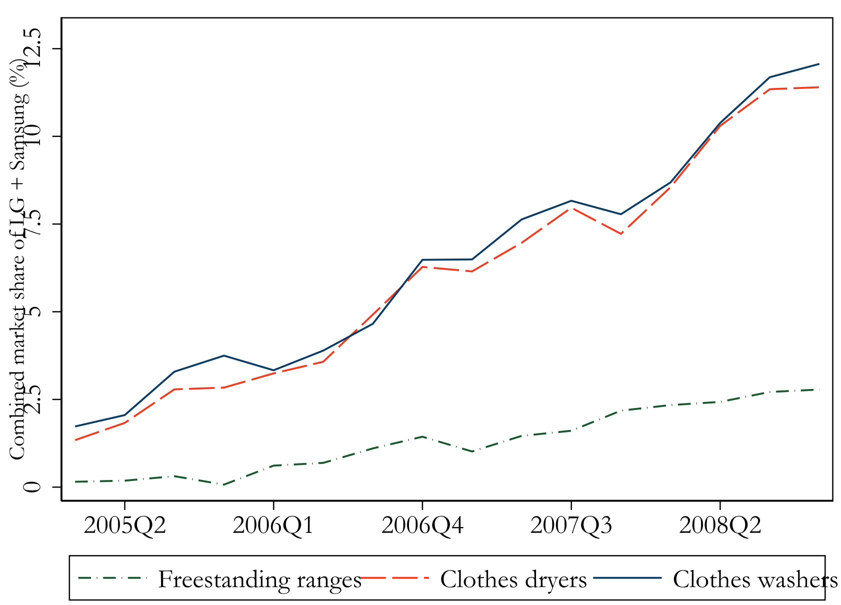

Notes: The figure showss the combined market share (volume sales) evolution in the U.S. and key European control markets for 2005 to 2008.

II.C Additional descriptive results on mass layoffs and new jobs

Figure A.11 presents event study plots illustrating the local labor market effects of mass layoffs and new job creation following the merger. The outcomes analyzed are county-level unemployment, manufacturing employment, and manufacturing wages. The key outcomes exhibiting significant treatment effects in the main analysis—specifically, unemployment and manufacturing employment following mass layoffs, and manufacturing employment following new job creation—do not display pre-treatment trends. This strengthens the robustness of the main findings.

Figure A.11: Effects of mass layoffs and new jobs over time

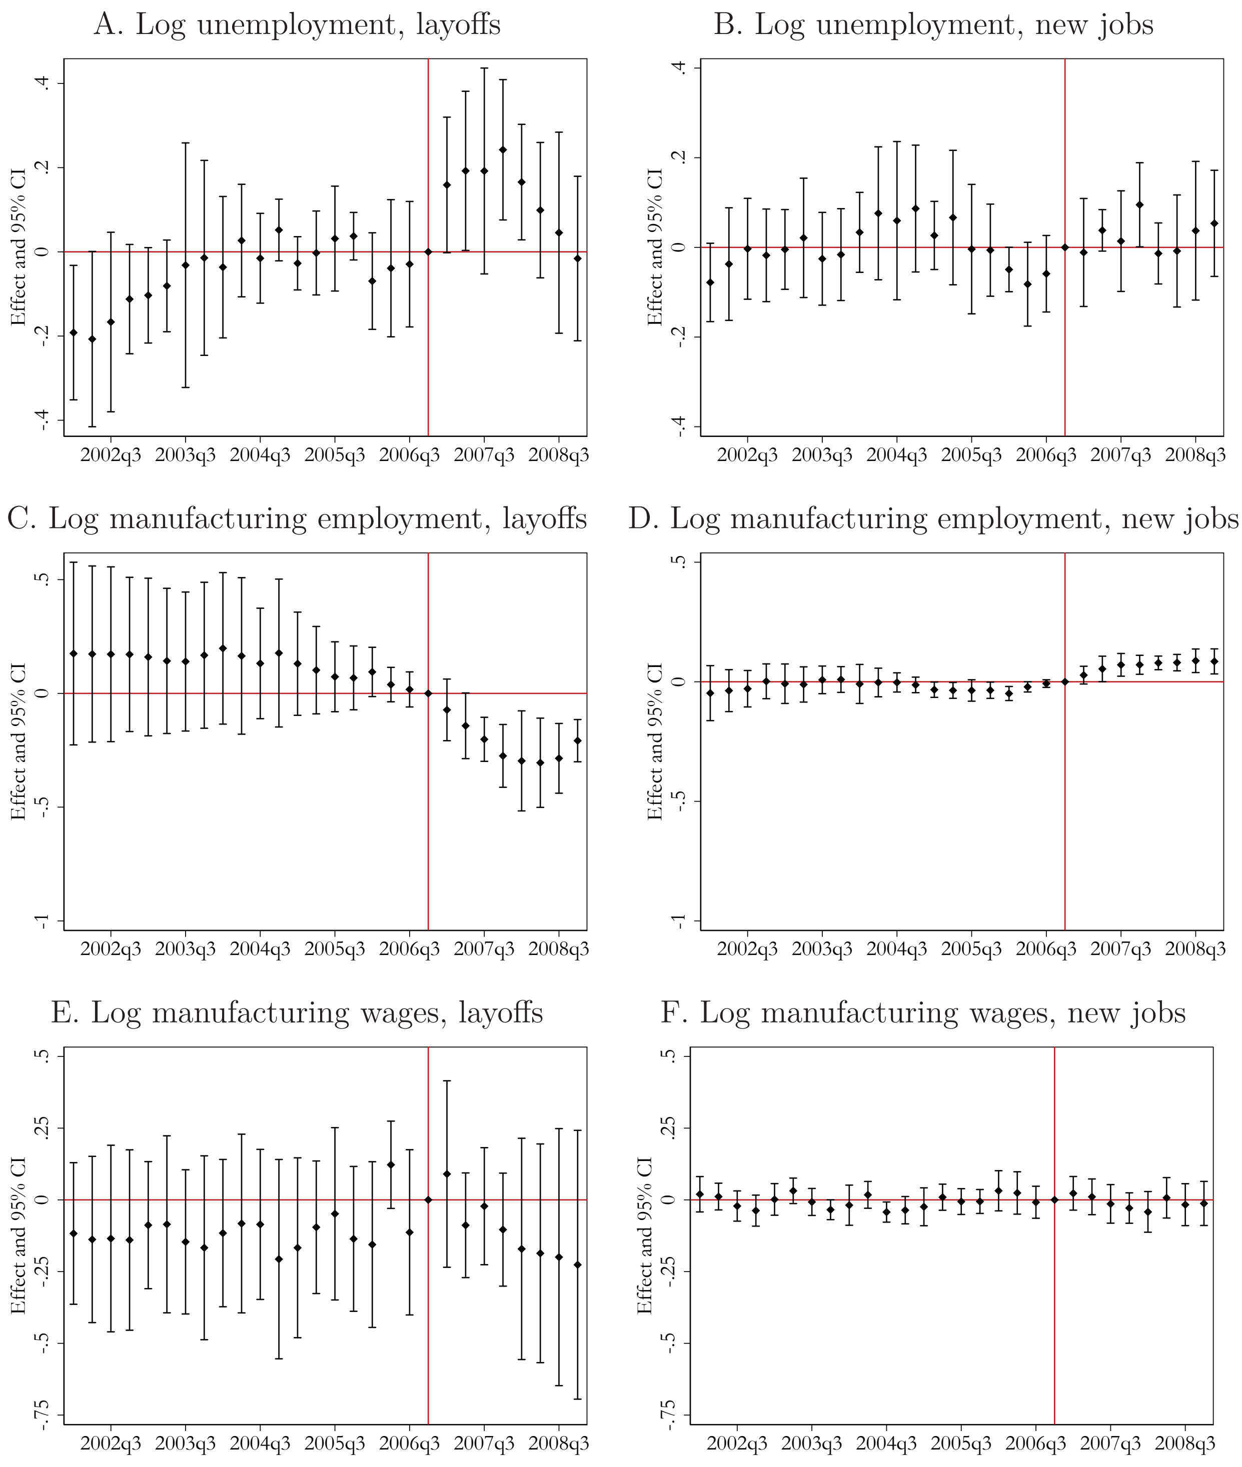

Notes: The figure shows the quarterly changes in labor market outcomes for counties experiencing mass layoffs or job creation at Whirlpool plants relative to matched control counties. Panels A, C, and E depict effects from the closure of Maytag manufacturing plants and headquarters; Panels B, D, and F depict effects from new job creation at existing Whirlpool plants. Standard errors are clustered at the county level.

III Appendix to Section 4: Model Details

III.A Plant-level output and employment

To test for economies of scale in labor, I specify a Cobb-Douglas relationship between plant-level output and employment.[46] Specifically, I assume that plant-level output ($q_{c(j)t}$) relates to employment ($L_{c(j)t}$) via:

$$q_{c(j)t} = AL^{\zeta}_{c(j)t}.$$

I estimate the parameters by regressing the log of annual plant-level output on the log of plant-level manufacturing employment:

$$\ln(q_{c(j)t}) = \ln(A) + \zeta \ln(L_{c(j)t}) + \eta_{c(j)t}, \tag{20}$$

where $\eta_{c(j)t}$ represents an error term that occurs because firms hire workers before observing transitory shocks, thus knowing only $E\left[q_{jt}\right]$, not q_{jt} itself.

Table A.3 shows the estimates. Although the sample size is small, I cannot reject the hypothesis that the elasticity of output with respect to labor is equal to one. This finding indicates constant returns to scale in labor.

[46] Nesting a Cobb-Douglas production function for labor within a Leontief production function is similar to the approach taken by Loecker and Scott (2024) who model the production of beer.

Table A.3: Relationship between annual plant-level output and employment

	(1)	(2)
Constant	6.40***	6.23***
	[4.61, 8.18]	[4.28, 8.18]
Logarithm of plant-level employment	1.11***	1.15***
	[0.84, 1.38]	[0.87, 1.43]
Restrict to 2000s	No	Yes
Observations	25	12

Notes: The dependent variable is the logarithm of plant-level output. Column (2) restricts observations to the years 2000–2009. 95% confidence intervals are reported in brackets. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

IV Appendix to Section 5: Estimation Details

IV.A Estimating product characteristics for potential products

Potential products include both the products brand owners actually offer (active products) and those they could offer but do not (inactive products). Active products are observed directly in the data. Estimating characteristics of inactive products is more complex.

The analysis focuses on products firms are already technologically capable of producing. Thus, if a firm does not sell front-loading washers, these are not potential products. For example, if Maytag sells regular top-loaders under its Amana brand at Best Buy and Lowe's, similar Amana top-loaders at other major retailers are potential but inactive products.[47]

Characteristics of products can vary slightly across retailers. For instance, Amana top-loaders sold at Best Buy might differ somewhat from those sold at Lowe's. To attribute characteristics to an inactive product—such as Amana regular top-loaders potentially sold at Sears—I must decide whether to use characteristics from the Amana top-loaders actually sold at Best Buy or Lowe's.

Whenever a particular combination of brand and key characteristic exists at two or more retailers, I use the most similar retailer, using the following ordering:

- **Sears**: Home Depot, Lowe's, Best Buy, H.H. Gregg, Others
- **Home Depot**: Lowe's, Sears, Best Buy, H.H. Gregg, Others
- **Lowe's**: Home Depot, Best Buy, Sears, H.H. Gregg, Others

[47] Major retailers are Best Buy, H.H. Gregg, Home Depot, Lowe's, and Sears.

- **Best Buy**: Lowe's, H.H. Gregg, Home Depot, Sears, Others
- **H.H. Gregg**: Best Buy, Lowe's, Home Depot, Sears, Others

IV.B Details on demand estimation

Following Berry, Levinsohn, and Pakes (2004), I estimate demand in two stages. First, I solve for the nonlinear parameters $\theta_2 = (\alpha, \kappa_\alpha, \sigma^{\mathrm{FL}})$ and corresponding mean utilities δ, obtaining $\hat{\theta}_2$ and $\hat{\delta}$. Second, I recover the linear taste parameters β. Throughout, I implement best practices as described by Conlon and Gortmaker (2020). For simplicity, the market index t is omitted; all expressions refer to a single market and are averaged subsequently.

The estimation of the nonlinear parameters and the mean utilities proceeds in two iterative steps. In the inner loop, I search for the mean utilities given a guess of the nonlinear parameters. In the outer loop, I search for the nonlinear parameters that minimize the objective function, solving the inner loop at each step.

To estimate the mean utilities δ, I follow Berry (1994) and invert the market share function $s_j(\delta_j;\theta)$ to obtain $\delta_j(s_j^n, s_j(\delta_j;\theta))$, where s_j^n denotes the market shares observed in the data and $s_j(\delta_j;\theta)$ denotes the simulated market shares implied by the model and the parameter vector θ.[48] Second, I use the fixed-point formulation from Berry, Levinsohn, and Pakes (1995) to estimate $\hat{\delta}_j$. I use the SQUAREM described in Reynaerts, Varadha, and Nash (2012) to accelerate the convergence of the fixed-point iterations. As convergence is not guaranteed, whenever it fails, I revert to the contraction mapping in Berry, Levinsohn, and Pakes (1995), which has guaranteed convergence. Finally, I speed up the inversion of market shares by using the reformulation of the contraction mapping in terms of consumer-specific choice probabilities for the outside option, described by Brunner et al. (2020).

To approximate model-predicted market shares from Equation 12, I integrate numerically by Monte Carlo simulation, drawing 2,000 households per market from the joint empirical distribution of demographics (from CPS data) and standard-normal distributed unobserved taste shocks (via scrambled Halton draws as in Owen, 2017).

Estimation relies on three moment conditions:

1. Income-price covariance: Matches the covariance between the first-choice washer's price and average income of households purchasing it:

$$m_1(\theta) = \sum_j \frac{n_j}{n} p_j \left\{ \frac{1}{n_j} \sum_{i:y_i^1=j} \iota_i - E\left[\iota \mid y^1 = j, \theta\right] \right\}, \tag{21}$$

[48] Note that $s_j(\delta_j;\theta)$ also depends on the product and household characteristics, which I omit to simplify notation.

where n_j is the number of buyers of product j, y_i^1 is household i's chosen product, p_j is product j's price, and ι_i denotes household i's income. Figure A.12 illustrates this relationship.

Figure A.12 shows a scatter plot of the relationship between household income and price.

Figure A.12: Correlation of average purchaser household income and price by product

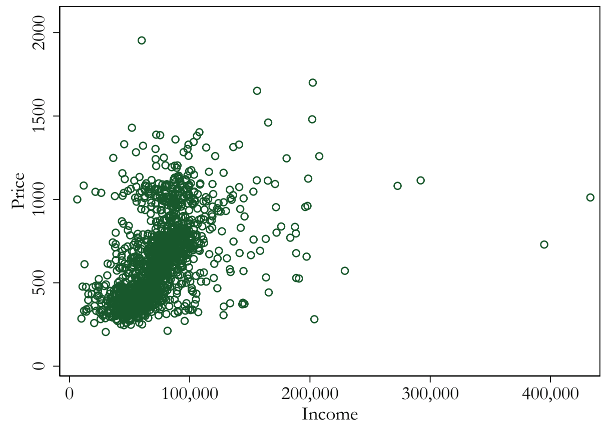

Notes: The plot shows the average annual income of households purchasing a particular washer on the x-axis and the average price of that washer on the y-axis. Each point is a product in a particular year.

2. Front-loader substitution pattern: Matches the covariance between choosing a front-loader and the share of front-loaders in the second-choice brand:

$$m_2(\theta) = \sum_j \frac{n_j}{n} x_j^{FL} \sum_{b' \neq b_j} x_{b'}^{FL} \left\{ \frac{n_{jb'}}{n_j} - E\left[\mathbb{1}(b^2 = b') \mid y^1 = j, \theta \right] \right\}, \tag{22}$$

where b_j denotes brand of product j, b^2 is the second-choice brand, and x_j^{FL} ($x_{b'}^{FL}$) indicates product (brand) front-loader shares.

The third moment condition is the orthogonality moment $m_3(\theta) = E\left[RER_{c(j)t}\xi_{jt}\right]$, which is based on the exclusion restriction of the price IV described in Section 5.

3. Instrumental variable exclusion restriction: The orthogonality condition using the real exchange rate as a price instrument:

$$m_3(\theta) = E\left[RER_{c(j)t}\xi_{jt}\right].$$

Stacking these moment conditions, I estimate θ_2 using the method of simulated moments (MSM):

$$\hat{\theta}_{2,MSM} = \arg\min_{\theta_2} \hat{m}(\theta_2)'\hat{m}(\theta_2). \tag{23}$$

After obtaining nonlinear parameters and mean utilities, I estimate the linear parameters β from:

$$\hat{\delta}_j + exp(\hat{\alpha} + \hat{\kappa}_\alpha \iota_i)p_j = x_j\beta + \xi_j, \tag{24}$$

assuming independence between observed nonprice product characteristics and unobserved quality ξ_j.

IV.B.1 Market size and share of the outside good

To estimate total market size, I assume that every seventh household considers purchasing a washer each year. According to Consumer Reports, the average life expectancy of a washer was ten years in 2009. Many households consider replacements before the end of a washer's life expectancy to benefit from new features. Some households consider purchases over multiple years, while those who recently purchased are unlikely to buy again soon. Therefore, a plausible range for annual market size is between one-fifth and one-tenth of all households. My results are robust to this assumption.

To translate estimates into total profits, consumer welfare, and fixed cost bounds in dollar terms, I scale the estimates by the annual market size for washers in the United States. Two alternative estimation methods yield similar results around the merger period. First, assuming one-seventh of total U.S. households are potential buyers each year; second, dividing annual total washer shipments reported by Appliance Portrait by the inside good's market share. Both approaches indicate a total U.S. market size of approximately 15 million households.

IV.C Speeding up the computation of expected profits

The estimation of fixed costs and the heuristic entry algorithm both require computing expected profits across numerous product portfolios. Since this step is repeated frequently, computational efficiency is critical. Here, I briefly outline key optimizations implemented to accelerate computations.

Computing equilibrium prices. Each draw of second-stage marginal cost and demand shocks e_{jt} requires recalculating equilibrium prices for all active products. Given that I approximate expected profits for each product portfolio using 500 draws, equilibrium prices must also be recomputed 500 times per portfolio. Efficient computation of these equilibrium prices is thus essential. Moreover, not all numerical methods for recalculating equilibrium prices reliably converge. Morrow and Skerlos (2011) evaluate several numerical methods for solving Nash–Bertrand equilibrium prices. They find that Newton methods are reliable but computationally slow, whereas fixed-point iteration using the standard BLP markup equation can be slow and sometimes fail to converge. Instead, they propose an alternative markup formulation—the ζ-markup equation—which is both fast and reliable. I therefore use fixed-point iteration based on the ζ-markup equation to compute equilibrium prices.

Drawing $\boldsymbol{e_{jt}}$. The heuristic portfolio choice algorithm repeatedly compares expected profits across product portfolios differing by at most a single product. Consequently, it revisits the same portfolios multiple times. A crucial efficiency gain arises from using the same set of e_{jt} draws for each product across different portfolio evaluations. Economically, this assumption is justified since firms' expectations about demand and cost shocks for a given product should not vary depending on the presence of other products. Computationally, this allows storing the expected profit estimates for each portfolio. When the algorithm revisits a portfolio, it reuses previously computed expected profits, thus eliminating redundant calculations of equilibrium prices and expected profits.

V Appendix to Section 6: Further Parameter Estimates

V.A Demand estimation

Table A.4: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)
	First stage	Logit OLS	Logit IV	Mixed logit
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Real exchange rate	2.033***			
	(0.365)			
Price ('00 2012 $)		-0.164**	-0.351**	
		(0.062)	(0.178)	
Front-loader	0.195	0.358	0.343	-0.730***
	(0.208)	(0.244)	(0.219)	(0.104)
Korean	-0.584***	1.569***	1.528***	1.522***
front-loader	(0.187)	(0.349)	(0.345)	(0.210)
Fisher & Paykel	Paykel	-4.536***	-1.455***	-2.705**
front-loader	(0.322)	(0.480)	(0.796)	(1.060)
European high-end	0.089	1.192***	1.246*	1.261

continued

Table A.4: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)
	First stage	Logit OLS	Logit IV	Mixed logit
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
front-loader	(1.268)	(0.438)	(0.640)	(0.892)
Agitator	-2.513***	0.540**	0.071	-0.135
	(0.276)	(0.252)	(0.481)	(0.088)
Stacked pair	0.490*	-0.147	-0.053	-0.028
	(0.279)	(0.149)	(0.183)	(0.158)
Stainless steel exterior	0.471	0.009	0.112	0.116
	(0.602)	(0.270)	(0.330)	(0.265)
White exterior	-0.285	0.624***	0.574***	0.541***
	(0.360)	(0.101)	(0.122)	(0.148)
Energy Star	0.019	0.092	0.097	0.115
	(0.182)	(0.126)	(0.133)	(0.129)
Extra noise insulation	0.397*	0.312**	0.387**	0.418***
	(0.207)	(0.120)	(0.153)	(0.101)
Number of special programs	0.008	0.052	0.054	0.051
	(0.058)	(0.039)	(0.045)	(0.033)
Child lockout	-0.075	0.200	0.181	0.185
	(0.163)	(0.167)	(0.169)	(0.122)
Repair rate	-2.384	1.627	1.200	1.091
	(3.151)	(2.957)	(2.778)	(1.799)
Total advertising expenditure	-0.006	0.003	0.002	0.001
	(0.005)	(0.002)	(0.002)	(0.005)
Retailer Best Buy	-0.098	-1.062***	-1.080***	-1.088***
	(0.085)	(0.307)	(0.306)	(0.094)
Retailer H.H. Gregg	-0.368***	-1.963***	-2.032***	-2.060***
	(0.120)	(0.299)	(0.278)	(0.099)
Retailer Home Depot	-0.161	-0.765**	-0.795**	-0.806***
	(0.106)	(0.321)	(0.320)	(0.096)
Retailer Lowe's	-0.179**	-0.334	-0.365	-0.381***
	(0.090)	(0.231)	(0.223)	(0.086)

continued

Table A.4: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)
	First stage	Logit OLS	Logit IV	Mixed logit
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Retailer Sears	0.014	-0.435	-0.431	-0.429***
	(0.114)	(0.445)	(0.442)	(0.101)
Brand FE	*Yes*	*Yes*	*Yes*	*Yes*
Brand time trends	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,590	1,586	1,590	1,590
Kleibergen–Paap F-statistic	31.041			
Own-price elasticity		-0.964	-2.058	-2.542

Notes: Column (1) reports the first-stage regression results of prices on the real exchange rate. Column (2) presents estimates from the simple logit model without instrumentation. Column (3) shows estimates from the simple logit using the RER as an instrument for price. Column (4) displays results from the mixed logit model described in Section 4. Standard errors are clustered at the brand level. Own-price elasticities of residual demand are computed at the product level and averaged across products, weighting by sales volume. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VI Appendix to Section 7: Details on Welfare Effects

VI.A Technical details on the portfolio choice algorithm

After initializing the algorithm, product portfolios are optimized in two loops to identify a one-step equilibrium. In the *inner loop*, a player computes the expected change in firm-level profits from adding each inactive product individually and removing each active product individually.[49] If there is a profitable one-step deviation, the player updates the product portfolio accordingly and repeats this until no profitable one-step deviations remain. In the *outer loop*, I sequentially repeat this process for each player.

In practice, the computational burden is reduced by optimizing portfolios at the brand, rather than firm, level. This requires checking fewer potential deviations per portfolio adjustment. Although I fully take into account how a portfolio adjustment impacts the firm's

[49] As realized demand and supply shocks are unobservable for potential products, I estimate expected profits based on 500 demand and supply residual draws per product.

expected profit (and not just that of the brand), a potential drawback is that for products from two brands of the same firm that are close substitutes, the entry order could affect the portfolio choice. However, firms typically segment products by brand, such that products within a brand are closer substitutes, making such cases unlikely.

To further reduce computational burden, I consider only one-step deviations, ignoring multistep deviations. Evaluating multistep deviations would be computationally infeasible in this context.[50] Given that washing machines are generally substitutes, adding multiple products simultaneously is unlikely to be profitable if adding each individually is not. However, profitable multistep deviations involving simultaneous additions and removals could exist, but checking these exhaustively is impractical. Furthermore, it may not be desirable to consider multistep deviations with many different portfolio adjustments simultaneously, since making complex portfolio adjustments is also more difficult for firms in practice.

Finally, because fixed costs are only partially identified, I implement the portfolio choice algorithm repeatedly using 50 fixed-cost draws. While the estimation of fixed-cost bounds places no restrictions on the distribution within brands, evaluating counterfactuals requires an explicit assumption. In the spirit of Wollmann (2018), I set the mean fixed cost for each brand at the midpoint of its 95-percent confidence bounds and draw the idiosyncratic product- and market-specific fixed-cost shocks, υ_{jt}, from a normal distribution with mean zero and a standard deviation equal to 25 percent of the width of these bounds. In all counterfactuals, I report 95-percent confidence sets for welfare effects across fixed-cost draws.

VI.B Additional results

Table A.5 repeats the analysis from Table 8, but without constraining the shipping cost parameter to the point estimate from the main sample in bootstrapped samples. As expected, this results in wider confidence intervals in Column (2). The results in Column (4) remain unaffected, as the endogenous product portfolio algorithm is estimated using only the main sample.

Table A.6 presents results from the various merger simulations comparing Whirlpool's acquisition of Maytag to a scenario with a standalone Maytag. This table extends the analysis in Table 7 by additionally accounting for scenarios that include offshoring efficiencies.

Table A.7 shows the results for the different merger simulations when comparing an acquisition of Maytag by Haier to a standalone Maytag. As expected, without cost changes, the product market implications of these two scenarios are the same. With offshoring cost

[50]For example, brands can have up to 15 potential products, implying $2^{15} = 32,768$ candidate deviations per brand iteration.

Table A.5: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier (bootstrapping shipping cost parameter)

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Cost pass-through				
Whirlpool relocation	–	86%	–	85%
	–	[-75%, 247%]	–	[84%, 87%]
Haier relocation	–	85%	–	84%
	–	[68%, 103%]	–	[83%, 85%]
Prices and consumer welfare				
Average price	3.2%	4.6%	4.1%	5.5%
	[1.9%, 4.5%]	[-0.5%, 9.7%]	[3.5%, 4.7%]	[4.7%, 6.2%]
Consumer welfare	$-166M	$-215M	$-222M	$-271M
	[$-213M, $-119M]	[$-380M, $-49M]	[$-264M, $-179M]	[$-316M, $-225M]
Employment				
Domestic jobs maintained	566	571	754	797
	[400, 733]	[387, 755]	[559, 949]	[612, 982]
Industry profits				
Variable profits	$81M	$55M	$79M	$47M
	[$58M, $104M]	[$-53M, $163M]	[$54M, $105M]	[$19M, $74M]
Total profits	$81M	$55M	$106M	$83M
	[$58M, $104M]	[$-53M, $163M]	[$80M, $132M]	[$56M, $111M]
Maytag + Whirlpool profits				
Maytag cost change	–	5.9%	–	4.7%
	–	[-17.9%, 29.7%]	–	[3.9%, 5.5%]
Variable profits	$23M	$-16M	$4M	$-38M
	[$5M, $41M]	[$-168M, $136M]	[$-37M, $44M]	[$-81M, $6M]
Total profits	$23M	$-16M	$40M	$11M
	[$5M, $41M]	[$-168M, $136M]	[$24M, $57M]	[$-5M, $26M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$293k	$376k	$299k	$344k
	[$209k, $377k]	[$93k, $660k]	[$203k, $396k]	[$250k, $438k]
Consumer welfare + total industry profits	$150k	$280k	$157k	$238k
	[$90k, $211k]	[$201k, $761k]	[$100k, $213k]	[$172k, $304k]

Notes: The first two columns compare the effect of Maytag acquisitions by Whirlpool vs. Haier without product portfolio adjustments. The final two columns account for endogenous portfolio adjustments. Columns (1) and (3) do not allow for acquisition-related cost changes. Columns (2) and (4) account for labor and shipping cost changes due to offshoring. Ninety-five percent confidence intervals for the first two columns are computed from 200 residual bootstrap draws. Confidence sets for final two columns are based on 50 fixed cost draws for each potential product.

Table A.6: Simulated effects of Maytag acquisitions by Whirlpool vs. no acquirer

	Pre-acquisition product portfolios		Endogenous product portfolios	
Cost adjustments:	None	Offshoring	None	Offshoring
Cost pass-through				
Whirlpool relocation	–	86%	–	85%
	–	[76%, 96%]	–	[84%, 87%]
Haier relocation	–	–	–	–
	–	–	–	–
Prices and consumer welfare				
Average price	3.2%	3.2%	4.1%	3.6%
	[1.9%, 4.6%]	[1.8%, 4.6%]	[3.5%, 4.7%]	[3.0%, 4.2%]
Consumer welfare	$-167M	$-156M	$-224M	$-175M
	[$-215M, $-119M]	[$-207M, $-106M]	[$-267M, $-182M]	[$-218M, $-131M]
Employment				
Domestic jobs maintained	-735	-759	-1009	-951
	[-869, -601]	[-900, -618]	[-1188, -829]	[-1120, -781]
Industry profits				
Variable profits	$82M	$90M	$81M	$105M
	[$58M, $105M]	[$57M, $123M]	[$55M, $106M]	[$79M, $131M]
Total profits	$82M	$90M	$107M	$132M
	[$58M, $105M]	[$57M, $123M]	[$81M, $133M]	[$107M, $157M]
Maytag + Whirlpool profits				
Maytag cost change	–	-3.9%	–	-7.0%
	–	[-7.4%, -0.3%]	–	[-7.6%, -6.4%]
Variable profits	$23M	$42M	$4M	$53M
	[$5M, $41M]	[$8M, $76M]	[$-36M, $45M]	[$11M, $95M]
Total profits	$23M	$42M	$41M	$90M
	[$5M, $41M]	[$8M, $76M]	[$24M, $58M]	[$74M, $106M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	–	–	–	–
	–	–	–	–
Consumer welfare +	–	–	–	–
total industry profits	–	–	–	–

Notes: The first two columns compare the effect of a Maytag acquisition by Whirlpool vs. the outcomes in the absence of an acquisition. The final two columns account for endogenous portfolio adjustments. Columns (1) and (3) do not allow for acquisition-related cost changes. Columns (2) and (4) account for labor and shipping cost changes due to offshoring. Ninety-five percent confidence intervals for the first two columns are computed from 200 residual bootstrap draws. Confidence sets for final two columns are based on 50 fixed cost draws for each potential product.

changes, U.S. consumers benefit from an acquisition of Maytag by Haier, whereas U.S. workers lose out.

Table A.8 repeats the analysis in Table 8 drawing fixed cost from a distribution around the 25^{th} percentile of the distance between the lower- and upper-bound brand-level fixed costs, instead of its midpoint. The results are similar to the baseline results. All of the point estimates are within the confidence sets of the main estimation results in Table 8.

Table A.9 repeats the analysis in Table 8 drawing fixed cost from a distribution around the 75^{th} percentile of the distance between the lower- and upper-bound brand-level fixed costs, instead of its midpoint. The results are similar to the baseline results. All of the point estimates are within the confidence sets of the main estimation results in Table 8.

Table A.7: Simulated effects of Maytag acquisitions by Haier vs. no acquirer

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Cost pass-through				
Whirlpool relocation	– –	85% [72%, 98%]	– –	84% [83%, 85%]
Haier relocation	– –	– –	– –	– –
Prices and consumer welfare				
Average price	0.0% [-0.1%, 0.1%]	-1.4% [-3.5%, 0.7%]	0.1% [0.1%, 0.1%]	-1.8% [-2.2%, -1.3%]
Consumer welfare	$-1M [$-3M, $1M]	$58M [$-17M, $133M]	$-3M [$-3M, $-2M]	$96M [$73M, $120M]
Employment				
Domestic jobs maintained	-1302 [-1579, -1025]	-1330 [-1609, -1050]	-1763 [-1869, -1657]	-1748 [-1879, -1617]
Industry profits				
Variable profits	$1M [$-1M, $2M]	$35M [$-15M, $86M]	$2M [$1M, $2M]	$58M [$45M, $72M]
Total profits	$1M [$-1M, $2M]	$35M [$-15M, $86M]	$2M [$1M, $2M]	$48M [$35M, $62M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-9.7% [-21.9%, 2.4%]	– –	-11.7% [-12.1%, -11.3%]
Variable profits	$0M [$-1M, $1M]	$74M [$-12M, $159M]	$1M [$0M, $1M]	$110M [$83M, $137M]
Total profits	$0M [$-1M, $1M]	$74M [$-12M, $159M]	$1M [$0M, $1M]	$98M [$90M, $105M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	– –	– –	– –	– –
Consumer welfare + total industry profits	– –	– –	– –	– –

Notes: The first two columns compare the effect of a Haier acquisition by Whirlpool vs. the outcomes in the absence of an acquisition. The final two columns account for endogenous portfolio adjustments. Columns (1) and (3) do not allow for acquisition-related cost changes. Columns (2) and (4) account for labor and shipping cost changes due to offshoring. Ninety-five percent confidence intervals for the first two columns are computed from 200 residual bootstrap draws. Confidence sets for final two columns are based on 50 fixed cost draws for each potential product.

Table A.8: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier, fixed cost draws around 25^{th} percentile

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Cost pass-through				
Whirlpool relocation	– –	86% [76%, 96%]	– –	86% [86%, 87%]
Haier relocation	– –	85% [72%, 98%]	– –	84% [84%, 84%]
Prices and consumer welfare				
Average price	3.2% [1.9%, 4.5%]	4.6% [2.0%, 7.3%]	4.6% [4.0%, 5.1%]	6.1% [5.4%, 6.7%]
Consumer welfare	$-166M [$-213M, $-119M]	$-215M [$-296M, $-133M]	$-243M [$-267M, $-218M]	$-288M [$-314M, $-262M]
Employment				
Domestic jobs maintained	566 [400, 733]	571 [407, 734]	679 [520, 839]	720 [575, 865]
Industry profits				
Variable profits	$81M [$58M, $104M]	$55M [$2M, $107M]	$102M [$85M, $118M]	$71M [$54M, $89M]
Total profits	$81M [$58M, $104M]	$55M [$2M, $107M]	$118M [$104M, $132M]	$91M [$77M, $104M]
Maytag + Whirlpool profits				
Maytag cost change	– –	5.9% [-4.3%, 16.0%]	– –	5.3% [4.9%, 5.8%]
Variable profits	$23M [$5M, $41M]	$-16M [$-82M, $51M]	$26M [$0M, $51M]	$-12M [$-39M, $15M]
Total profits	$23M [$5M, $41M]	$-16M [$-82M, $51M]	$46M [$38M, $54M]	$12M [$4M, $20M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$293k [$209k, $377k]	$376k [$229k, $524k]	$362k [$265k, $460k]	$404k [$317k, $490k]
Consumer welfare + total industry profits	$150k [$90k, $211k]	$280k [$59k, $502k]	$186k [$133k, $239k]	$277k [$217k, $337k]

Notes: The first two columns compare the effect of a Maytag acquisition by Whirlpool with that of one by Haier without product portfolio adjustments. The final two columns account for endogenous portfolio adjustments. Columns (1) and (3) do not allow for acquisition-related cost changes. Columns (2) and (4) account for labor and shipping cost changes due to offshoring. Ninety-five percent confidence intervals for the first two columns are computed from 200 residual bootstrap draws. Confidence sets for final two columns are based on 50 fixed cost draws for each potential product.

Table A.9: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier, fixed cost draws around 75^{th} percentile

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Cost pass-through				
Whirlpool relocation	– –	86% [76%, 96%]	– –	85% [83%, 87%]
Haier relocation	– –	85% [72%, 98%]	– –	84% [83%, 85%]
Prices and consumer welfare				
Average price	3.2% [1.9%, 4.5%]	4.6% [2.0%, 7.3%]	3.6% [2.6%, 4.6%]	4.9% [3.9%, 5.9%]
Consumer welfare	$-166M [$-213M, $-119M]	$-215M [$-296M, $-133M]	$-213M [$-277M, $-150M]	$-253M [$-324M, $-183M]
Employment				
Domestic jobs maintained	566 [400, 733]	571 [407, 734]	976 [666, 1287]	1028 [721, 1334]
Industry profits				
Variable profits	$81M [$58M, $104M]	$55M [$2M, $107M]	$45M [$-1M, $91M]	$17M [$-35M, $69M]
Total profits	$81M [$58M, $104M]	$55M [$2M, $107M]	$96M [$62M, $130M]	$73M [$35M, $112M]
Maytag + Whirlpool profits				
Maytag cost change	– –	5.9% [-4.3%, 16.0%]	– –	4.7% [3.9%, 5.5%]
Variable profits	$23M [$5M, $41M]	$-16M [$-82M, $51M]	$-36M [$-109M, $38M]	$-69M [$-152M, $13M]
Total profits	$23M [$5M, $41M]	$-16M [$-82M, $51M]	$34M [$12M, $56M]	$8M [$-15M, $31M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$293k [$209k, $377k]	$376k [$229k, $524k]	$226k [$111k, $340k]	$254k [$131k, $376k]
Consumer welfare + total industry profits	$150k [$90k, $211k]	$280k [$59k, $502k]	$125k [$52k, $198k]	$180k [$97k, $263k]

Notes: The first two columns compare the effect of a Maytag acquisition by Whirlpool with that of one by Haier without product portfolio adjustments. The final two columns account for endogenous portfolio adjustments. Columns (1) and (3) do not allow for acquisition-related cost changes. Columns (2) and (4) account for labor and shipping cost changes due to offshoring. Ninety-five percent confidence intervals for the first two columns are computed from 200 residual bootstrap draws. Confidence sets for final two columns are based on 50 fixed cost draws for each potential product.

References for Online Appendix

Appliance Portrait. 2006. 29^{th} *Annual Portrait of the U.S. Appliance Industry.* Dana Chase.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2013. "The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool." *American Economic Journal: Economic Policy* 5 (1): 239–61.

Berry, Steven. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics* 25 (2): 242–262.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

———. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): 68–105.

Brunner, Daniel, Florian Heiss, André Romahn, and Constantin Weiser. 2020. "Reliable Estimation of Random Coefficient Logit Demand Models." *mimeo.*

Conlon, Christopher, and Jeff Gortmaker. 2020. "Best practices for differentiated products demand estimation with PyBLP." *The RAND Journal of Economics* 51 (4): 1108–1161.

Electrolux. 2007. *Annual Report 2006.* https://www.electroluxgroup.com/en/annual-reports-27641/?decade=2000-2009.

Ferencikova, Sonia. 2002. "Transition at Whirlpool–Tatramat: From joint venture to acquisition." *Transnational Corporations* 1 (1): 69–98.

Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot. 2020. "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines." *American Economic Review* 110 (7): 2103–27.

Loecker, Jan De, and Paul T. Scott. 2024. "Markup Estimation using Production and Demand Data: An Application to the US Brewing Industry." *Review of Economic Studies.*

Maytag. 2005. *2004 Annual Report.* Newton, IA: Author.

McCabe, Liam. 2016. "Should You Get a Front-Load or Top-Load Washing Machine?" *New York Times,* https://www.nytimes.com/wirecutter/blog/should-you-get-a-front-or-top-load-washing-machine/.

Morrow, W. Ross, and Steven J. Skerlos. 2011. "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research* 59 (2): 328–345.

Owen, Art B. 2017. *A randomized Halton algorithm in R.*

Reynaerts, Jo, R Varadha, and JC Nash. 2012. "Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping." *KU Leuven Working Papers of VIVES - Research Centre for Regional Economics 482522.*

Rubens, Arthur, Sonia Ferencikova, and Roland Bardy. 2019. “How to transcend from a local subsidiary into the global leadership position of your industry: the case of Whirlpool Slovakia.” *Transnational Corporations Review* 2 (11): 120–132.

Wollmann, Thomas G. 2018. “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles.” *American Economic Review* 108 (6): 1364–1406.