

Mergers, Foreign Competition, and Jobs: Evidence from the U.S. Appliance Industry

Felix Montag
Tuck School at Dartmouth*

This version: January 3, 2023

Abstract

Policy choices often entail trade-offs between workers and consumers. I assess how foreign competition changes the consumer welfare and domestic employment effects of a merger. I construct a model accounting for demand responses, endogenous product portfolios, and employment. I apply this model to the acquisition of Maytag by Whirlpool in the household appliance industry. I compare the observed acquisition to one with a foreign buyer. While a Whirlpool acquisition decreased consumer welfare by \$250 million, it led to 1,300 fewer domestic jobs lost. Jobs need to be worth above \$220,000 annually for domestic employment effects to offset consumer harm.

JEL classification: F61, L13, L40

*Email: felix.montag@tuck.dartmouth.edu. Tuck School of Business at Dartmouth College. Affiliate Fellow, Stigler Center, Booth School of Business at the University of Chicago. I am deeply grateful to my advisors Monika Schnitzer and Ali Yurukoglu for their patient guidance, constant encouragement and invaluable advice. I also thank Alessandra Allocca, Antonin Bergeaud, Laura Grigolon, Anna Gumpert, Ying Fan, Mitsuru Igami, Torsten Jaccard, Helena Perrone, Claudia Steinwender, Jose Vasquez, Fabian Waldinger, Lulu Wang, Martin Watzinger and seminar participants at Dartmouth, UBC Sauder, Bristol, HEC Paris, LMU Munich, Mannheim, IESE, DICE, Oslo, MaCCI Annual Conference, CEPR Applied IO, EDGE Jamboree, MaCCI Summer Institute, the Austrian BWB, the UK CMA, the Stigler Center at Chicago Booth, ULB ECARES, and CREST for helpful comments and discussions. Financial support by the German Research Foundation (DFG) through CRC/TRR190 (Project number 280092119) and the International Doctoral Program “Evidence-Based Economics” of the Elite Network of Bavaria is gratefully acknowledged.

1 Introduction

Mergers can have important employment effects. For some mergers between firms, this can create a trade-off between inducing competition and securing jobs.¹ While competition enforcers tend to focus exclusively on evaluating the consumer outcomes associated with potential mergers, there is little understanding of the employment effects, and a lack of studies evaluating this trade-off explicitly.² Such analysis is particularly relevant when the firms in question face competition from abroad, since blocking a merger could lead to aggregate employment loss for the jurisdiction in question.³

In this paper, I specify and estimate a structural model of demand and supply to analyze how foreign competition changes the effect of a merger on consumer welfare and employment. I incorporate two key margins of firm decision-making in response to a merger: adjustments of prices and changes to the product portfolio. To account for the effects of product entry and exit on consumers and employment, I embed the consumer demand model into an endogenous product choice model, where the demand for domestic labor depends on production locations and equilibrium product market quantities. The key methodological innovation is to construct a model that can be used to empirically study the trade-off between workers and consumers. Using this model, I estimate the average value of a domestic job such that gains to workers offset losses to consumers.

To illustrate the trade-off between workers and consumers, I use this model to study the landmark merger between Maytag and Whirlpool in the U.S. appliance market. I simulate the consumer welfare and employment effects of two acquisitions: the observed acquisition of Maytag by Whirlpool, as well as a hypothetical purchase of Maytag by the alternative buyer at the time, the Chinese company Haier. Production locations are exogenous.⁴ For the Whirlpool acquisition, I use the observed post-merger production locations. For the Haier acquisition, I assume that Haier completely offshores Maytag’s production to China.⁵ Whilst there was no threat to consumers from a Haier acquisition, because it had no prior

¹Examples include the proposed mergers between PSA and Fiat Chrysler (FCA), T-Mobile and Sprint, or Albertson and Kroger (see Noble, 2021, Legere, 2019, and Kroger, 2022).

²There is a nascent literature on labor market power in merger analysis (e.g., Prager and Schmitt, 2021, Shapiro, 2019 or Marinescu and Hovenkamp, 2019). This literature is concerned with overlaps between the merging parties in local labor markets. Instead, I focus on employment effects when no such overlaps exist.

³Some jurisdictions incorporate public interest considerations into merger control (see OECD, 2016). These do not exist in the European Union or the United States.

⁴Relatedly, Igami (2018) studies the dynamics leading to offshoring in the hard disk drive industry.

⁵This is based on anecdotal evidence around the Haier bid, described in further detail in Section 2.

presence in the U.S., its offshoring plans would have put more Maytag jobs in peril.

Three main findings emerge: first, a Whirlpool acquisition of Maytag is always worse for consumers. The merging parties increase prices of clothes washers by more than five percent without inducing substantial rival product entry, leading to an annual loss in consumer welfare of \$250 million. Second, a Whirlpool purchase leads to the maintenance of around 1,300 more U.S. jobs than after a purchase by Haier. Finally, I find that each additional job maintained by Whirlpool (relative to a Haier acquisition) must be worth around \$220,000 per year to offset the consumer harm.

While I find that the trade-off between workers and consumers in this application strongly depends on an alternative buyer with different offshoring plans, this is not necessary for this trade-off to arise.⁶ With sufficient information on the set of counterfactual plant locations, the methodology outlined in this paper can be applied to any merger – or policy choice – where there is a trade-off between workers and consumers.

For the empirical analysis, I construct a comprehensive data set of the U.S. residential laundry machines market between 2005 and 2015. The core of the data comes from *TraQline*, a representative survey of 600,000 households per year. On the production side, I hand-collect product-level data on the locations of plants manufacturing for the U.S. market.

I descriptively document trends around the time of Whirlpool’s acquisition using a difference-in-differences identification strategy. I show that although there is a significant increase in concentration within the U.S. market for clothes washers and dryers, there is no appreciable change in prices as compared to the control market. New product introductions by foreign rivals occur but also in the control market. Finally, I show a rise in local unemployment, decrease in employment, and decrease in average wages of the employed for those counties which experienced the closing of a Maytag plant.⁷

The descriptive evidence raises several questions: did rival product entry prevent price increases despite a strong increase in concentration or do other developments unrelated to the merger lead to an underestimation of the price effects? Was rival product entry induced by the merger? If consumers were harmed by the Whirlpool acquisition, could this harm be offset by benefits to U.S. workers? Answering these questions requires a model. For simplicity, I focus on washers from hereon after.

The model features manufacturers and consumers. Manufacturers choose their product

⁶For example, Stellantis CEO Carlos Tavares argued that the merger between automakers FCA and PSA into Stellantis increases capacity utilization, making the merged firm more competitive and preventing otherwise necessary job cuts (see Noble, 2021.)

⁷This is in line with recent evidence showing that the presence of multinational firms affects the wages of workers at other firms (see Card et al., 2018, Alfaro-Ureña, Manelici, and Vasquez, 2021, or Setzler and Tintelnot, 2021). Furthermore, Jacobson, LaLonde, and Sullivan (1993) show that workers separating from distressed firms suffer long-term earnings losses and that these depend on local labor market conditions.

portfolios and prices. Consumers make purchase decisions. The model is set up as a two-stage game. At the beginning of the game, each manufacturer is endowed with a set of potential products that it is technologically capable of producing. Each product is associated with an exogenous set of characteristics, a production location, and a marginal cost of production. In the first stage, each firm chooses which potential product to introduce into the market, at a per product fixed and sunk entry cost.⁸ Next, marginal cost and demand shocks are realized. In the second stage, firms set prices and consumers make purchases. I model consumer demand using a static random coefficients discrete choice model, where the price sensitivity of consumers depends on income and some consumers have an unobserved taste for front-loading clothes washers. Finally, the number of manufacturing jobs is determined. This is linear in the quantities of the product market equilibrium.⁹ Whether a job is created domestically or abroad depends on the exogenous production location for each product.

On the demand side, the estimation is in the spirit of Berry, Levinsohn, and Pakes (2004). Informally, the non-linear demand parameters are identified by the correlation between household income and purchase prices and the correlation between the characteristics of the first and second choice products. I construct a cost shifter based on the production location of each product and the real exchange rate (RER) between the production location and the U.S. This cost shifter is then used as an instrumental variable for price, which is exogenous to product-level demand conditions (see Goldberg and Verboven, 2001 or Grieco, Murry, and Yurukoglu, 2021).

On the supply side, I estimate the product-level marginal costs that rationalize the data assuming differentiated Bertrand-Nash competition (see Nevo, 2001). A growing literature is concerned with estimating bounds on the fixed costs of introducing a new product into the market using moment inequalities (see Pakes et al., 2015). Intuitively, the fixed and sunk cost of adding a product that was introduced to the market can at most be the change in the expected variable profit due to the product. Similarly, the fixed and sunk cost of adding a product that is part of the set of potential products but is not introduced to the market must be at least as high as the change in variable profit due to the removal of that product. Methodologically, the estimation of fixed cost bounds is closest to Eizenberg (2014). Finally, I combine evidence on the number of clothes washers that a manufacturing worker produces per year with the product-level plant locations to estimate how different product market equilibria affect the demand for domestic manufacturing workers.

I encounter several empirical challenges. A first challenge is to identify the set of

⁸Since I only observe product-level entry but no firm-level entry around the time of the merger, I focus on endogenous product choices and abstract away from firm entry.

⁹Wages are determined outside the model. They affect the demand for manufacturing workers through their effect on marginal costs and the product market equilibrium.

products that multi-product firms can introduce.¹⁰ Studying an unconditionally cleared merger allows me to overcome this. Draganska, Mazzeo, and Seim (2009) and Fan and Yang (2021) exploit cross-sectional variation in market structure to estimate the set of potential products. In my setting, this is infeasible because product portfolio decisions are national.¹¹

A second empirical challenge is the multiplicity of equilibria when simulating counterfactual entry. Due to the large number of products, computing all potential equilibria is computationally infeasible. Instead, I follow a literature that uses heuristic learning algorithms to determine equilibrium entry (e.g., Lee and Pakes, 2009, Wollmann, 2018 or Fan and Yang, 2020). Each player optimizes her portfolio sequentially, taking the choices of rivals as given, until there is no profitable one-step deviation.

The higher number of domestic jobs maintained after a purchase of Maytag by Whirlpool as compared to a purchase by Haier could plausibly offset consumer harm. For clothes washers, an average job would need to be worth around \$220,000 per year and around \$130,000 per year if we consider the sum of consumer welfare and industry profits to be the relevant welfare metric. If we also consider other appliance categories, where the overlap between the parties, and thus the harm to consumers, was lower, these offsetting values can be expected to decrease. In comparison, Setzler and Tintelnot (2021) find that the total wage bill in a local labor market increases by around \$113,000 per year for each job created by a foreign multinational firm. This does not include any other benefits of employment, which further increase the value of a job. Furthermore, welfare effects are unequally distributed. Whereas losses to consumers are spread out across the country, employment effects are concentrated in a few local labor markets. These findings relate to a literature that quantifies the differential effects of trade liberalization (see Jaravel and Sager, 2020) and restrictions (see Hufbauer and Lowry, 2012 or Flaaen, Hortaçsu, and Tintelnot, 2020) on workers and consumers. Among these estimates, I find the lowest job values necessary to offset consumer harm.

Finally, I contribute novel evidence to how endogenous product portfolio choices change the consumer welfare effects of mergers.¹² I find that even for an actual merger that was marginally cleared because of an entry defense, endogenous portfolio adjustments increase

¹⁰An earlier literature on endogenous product entry focuses on single-product firms with discrete product types (e.g., Mazzeo, 2002 or Seim, 2006).

¹¹Eizenberg (2014) analyzes a market without cross-sectional variation in entry. He estimates the set of potential products based on existing product lines and technologies. This works in his context, as he studies how the removal of a frontier technology affects the presence of older products. This is not a viable strategy to study the introduction of new products.

¹²A related literature (e.g., Werden and Froeb, 1998; Li et al., Forthcoming; and Ciliberto, Murry, and Tamer, 2021) studies mergers and static entry for single-product firms. Garrido (2020) studies dynamic product entry decisions by multi-product firms assuming nested logit demand. Fan (2013) studies product repositioning after mergers. Several papers study the effect of mergers on entry and product variety for radio stations (e.g., Berry and Waldfogel, 2001; Sweeting, 2010; and Jeziorski, 2015).

the harm to consumers. This is because foreign entry is mostly independent of the merger, whereas the merger leads to fewer products offered by the merging parties. Existing studies mostly consider hypothetical changes in concentration and find mixed results.¹³ Under certain conditions, Caradonna, Miller, and Sheu (2021) show that without marginal cost efficiencies product portfolio adjustments can never be profitable for the parties and also fully offset consumer harm.

The remainder of the paper is structured as follows: The next section discusses the details of the case and describes the data. Section 3 presents the descriptive evidence, Section 4 outlines the industry model, Section 5 sketches the estimation strategy, Section 6 presents the results, Section 7 describes the welfare effects, and Section 8 concludes.

2 Institutional Setting and Data

In the mid-2000s, around 90% of clothes washers and dryers sold in the U.S. were produced by the domestic manufacturers Whirlpool, Maytag, and General Electric, which also predominantly produced in the United States. Whereas the Swedish Electrolux produced all of its laundry machines for the U.S. market in the U.S., LG and Samsung were looking to enter the market using production facilities in Mexico and South Korea.

2.1 The acquisition of Maytag by Whirlpool

Prior to its acquisition by Whirlpool, Maytag had been struggling financially for several years. Although the company had already cut costs by reducing its workforce by 20 percent, in 2004 it continued to struggle with cost pressure, a further decline in revenues, and posted a net loss (Maytag, 2005). In May 2005, its management agreed to be bought by a group of private investors for \$1.13 billion (Barboza, 2005). In June 2005, the Chinese household appliance manufacturer Haier made a competing bid of \$1.3 billion. One month later, Maytag’s biggest manufacturing rival in the U.S. appliance market, Whirlpool, outbid Haier with an offer of \$1.4 billion. Haier withdrew its bid and in March 2006 Whirlpool acquired Maytag after an unconditional merger clearance by the Department of Justice (DoJ).

Haier’s bid came at a time when the Chinese government pushed its large companies to make foreign acquisitions to get access to foreign markets for its manufactured goods, particularly in the European Union and the United States.¹⁴ Since Chinese acquirers were

¹³Fan and Yang (2020) find that endogenous product adjustments exacerbate negative consumer welfare effects, whereas Wollmann (2018) finds the opposite. Fan and Yang (2021) show that product portfolio adjustments exacerbate negative merger effects in small markets and reduce consumer harm in larger markets.

¹⁴This was part of China’s “Go Out Policy”, promoting Chinese investments abroad (Goodman and

met with resistance, these acquisitions often targeted well-known brand names slipping into decline. This made it more likely that their bid would be accepted and also helped overcome the resistance of consumers towards Chinese brands in the product market.¹⁵ With its weak financial performance and its strong brand portfolio, Maytag fit the bill. Haier, who previously had negligible sales in the U.S. appliance market, planned to use Maytag’s brands, repair network and distribution channels, whilst offshoring production to its existing plants in China (Goodman and White, 2005).

Against this backdrop, Whirlpool’s bid for Maytag could be seen as fending off a foreign takeover. The main caveat, however, was that Whirlpool and Maytag were close competitors in the product market for several major appliance categories. In its investigation of the acquisition, the DoJ focused on residential clothes washers and dryers. For the manufacturing of laundry products, this was a merger from four to three, where Whirlpool and Maytag were the largest and second largest manufacturers in the U.S. market. Sears, with its Kenmore brand, was another large brand owner in the laundry market; however it did not manufacture any products itself but purchased them from original equipment manufacturers (OEMs) instead. The DoJ concluded that despite the high market shares of the merging parties, they would not be successful in raising prices because “LG, Samsung, and other foreign manufacturers could increase their imports into the U.S.” (Department of Justice, 2006). It therefore unconditionally cleared the acquisition. Baker and Shapiro (2008a) called this decision “[...] a highly visible instance of underenforcement” and Baker and Shapiro (2008b) described it as “fueling the perception that the Justice Department has adopted a very lax merger enforcement policy [...]”. They conclude that in this case the DoJ was willing to accept entry and expansion arguments in a highly concentrated merger case, although entrants had thus far only achieved relatively low market shares.

2.2 The data

To analyze the implications of the Maytag acquisition, I construct a comprehensive data set on the U.S. market for residential laundry products between 2005 and 2015.

2.2.1 Sales, products, and households

The centerpiece of the data comes from *TraQline*. This is a data set well-known across the appliance industry and is used by major retailers and major brands in the industry as a

White, 2005).

¹⁵A famous example is the 2005 acquisition of I.B.M.’s personal computer division by Lenovo.

source for market insights.¹⁶ In every quarter, a representative sample of around 150,000 U.S. households is asked about appliance purchases. The survey is a repeated cross-section and in total around 600,000 households are surveyed every year. The data spans the years 2005 until 2015. For each respondent, *TraQline* records the number of appliances bought, the price, a detailed set of product characteristics (e.g., the brand or whether a product is Energy Star certified), other brands that the household considered buying, the retailer at which the appliance was bought, as well as a detailed set of household demographics. The data includes information for clothes washers and dryers, as well as for freestanding ranges. I aggregate products at the national level, because product entry is determined for each major retailer at the national level. I also aggregate responses at the yearly level.

Although *TraQline* records detailed characteristic information, respondents are not asked to provide the exact model specification of the appliance they purchased. I therefore use brand, retailer and key characteristics information to aggregate appliance purchases into products. Most brand owners use different brands to cluster their product offering according to the consumers that they target.¹⁷ Thus, the brand of a product already captures much of the variation in, otherwise unobserved, product differentiation. Certain key product characteristics need to be reported by all survey respondents. For clothes washers, this includes whether a clothes washer is a regular top-loader (with an agitator), a high-efficiency top-loader (without an agitator) or a front-loader. Finally, I further refine the product definition by using information on the retailer at which the product is sold. Different retailers serve different customers. If a brand and key characteristics combination (e.g., a Whirlpool high-efficiency top-loading washing machine) is sold at both Sears and Best Buy, these products may still slightly differ in other characteristics.¹⁸ To capture all of these sources in observed and unobserved characteristics variation, I define a product as a brand, retailer and key characteristics combination.

To reduce the reporting burden, other characteristics only need to be reported by a random subsample of respondents. Households that are selected to answer the more detailed characteristics questions do not have the possibility to opt-out, ruling out any selection problems. For clothes washers, these more detailed characteristics include whether it has

¹⁶The only other comparable source of data on volume and value sales in the appliance industry is a, now discontinued, retailer panel by the NPD Group, which was the basis of the analysis by Ashenfelter, Hosken, and Weinberg (2013). To the best of my knowledge, the key difference between the data sets is that the retailer panel does not include any sales from Sears, which, at the time, was the largest U.S. retailer for household appliances and accounted for an important share of Maytag and Whirlpool sales.

¹⁷In its 2007 Annual Report, Whirlpool describes what each of its brands represents and what type of consumers it targets. Amana, for example, is described as stylish and affordable, whereas KitchenAid should stand for quality and craftsmanship, Whirlpool for innovation and Maytag for reliability.

¹⁸For retailers, I distinguish between Best Buy, H. H. Gregg, Home Depot, Lowe's, Sears, and all others. The latter group pre-dominantly includes smaller, regional retailers.

a child lockout, the number of special programs, whether it is a stacked pair or whether it has additional noise insulation. For each product, I calculate the average value of these characteristics among the subsample of respondents.

I enrich the *TraQline* product data set with two additional product characteristics: the brand repair rate and brand-level advertising expenditures.

The brand repair rates come from Consumer Reports. There exists a report on clothes washers at least once a year, which includes an overview of brand-level repair rates. This data is based on responses to the Annual Product Reliability Survey conducted by the Consumer Reports National Research Center for more than 100,000 clothes washers. I digitize this information to create a measure of the perceived product reliability of a brand over time.

Annual information on advertising expenditures comes from Kantar AdSpender between 2005 and 2015. I use the total advertising expenditure of a brand across media channels to capture variation in brand reputation over time. Benkard, Yurukoglu, and Zhang (2021) use this data set to track brand ownership over time.

The *TraQline* data set only includes household demographics for respondents that purchase an appliance. To identify how household income affects the sensitivity to prices in the demand estimation, I need data on the unconditional distribution of income among the population of all households. For this, I draw a random sample of households from the IPUMS Current Population Survey (CPS). This data set includes rich demographic information for a representative household sample for every year in the analysis period.

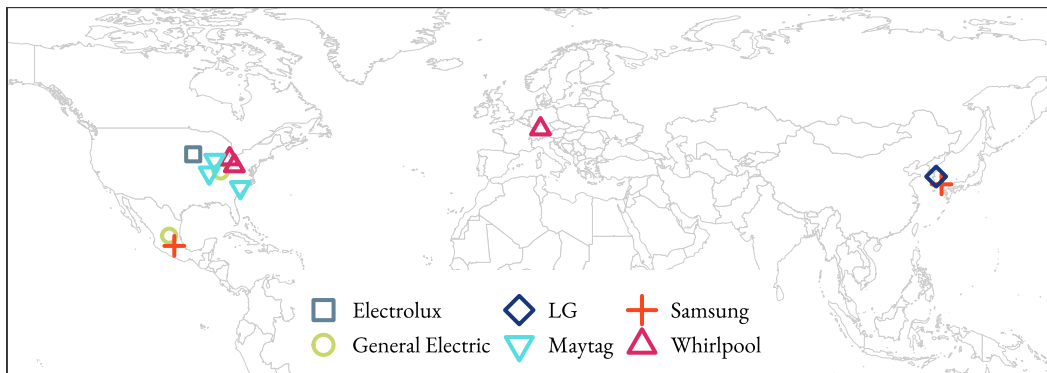
2.3 Production locations and an instrumental variable for price

The core of the supply-side data consists of hand-collected locations of plants manufacturing clothes washers for the U.S. market at the product level. This data set serves two purposes. First, it allows constructing a product-level instrumental variable for prices based on differences in production costs. Second, it allows simulating how the number of U.S. clothes washer manufacturing jobs changes between counterfactual scenarios.

Figure 1 shows the plant locations of major clothes washer manufacturers in 2005. To construct the panel of production locations, I collect production locations for all manufacturers with a market share of more than 3 percent in any year between 2005 and 2015. These are Electrolux, General Electric, LG, Maytag, Samsung, and Whirlpool.

Plant locations are based on the investigation by the U.S. International Trade Commission (USITC) into imports of large residential clothes washers from Mexico and South Korea, firm-level clothes washer imports based on the PIERS data set, which uses bill of landing documents and is reported in Flaaen, Hortaçsu, and Tintelnot (2020), annual reports, as

Figure 1: Clothes washer plants manufacturing for the U.S. market, 2005



Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2005 by manufacturers with a market share of more than 3 percent in any year in the sample. The Appendix includes a map for 2007 in Figure A.4, for 2009 in Figure A.5 and for 2011 in Figure A.6.

well as news reports. Details on the construction of this data set and the weights on the share of the product produced in each production location are described in Appendix I.B.

To explain the need for an instrumental variable for price and how I construct one, let us briefly jump ahead to the estimation of clothes washer demand as part of the structural model. As is well-known, there can be unobserved demand shocks that simultaneously affect prices and quantities. To get an unbiased estimate of the reaction of quantities to price changes, I need an instrument for price that is unrelated to unobserved demand shocks (exogeneity) and has a sufficiently strong effect on prices (relevance).

An ideal instrument is a variable that captures differences in product-level marginal costs and is unrelated to demand. I use the product-level weighted average real exchange rate (RER) between the U.S. and the countries in which the production of the product is located. The RER comes from the Penn World Table. Product-level plant weights are constructed as described above.

I use the RER based on consumption expenditures. This is calculated by dividing the consumption of households at nominal prices by the the same consumption using the U.S. price level in 2005 and then multiplying this by the nominal exchange rate between the local currency and the U.S. dollar (Feenstra, Inklaar, and Timmer, 2015). It therefore consists of differences in the relative price levels and serves as a proxy for the local wage level, as well as fluctuations in the nominal exchange rate. These should affect marginal costs, but should be unrelated to transitory demand shocks in the U.S. that have different effects on different products. Figure A.7 in Appendix shows the evolution of the average RER over time and illustrates the source of the variation.

2.4 Labor market data

To analyze the local labor market effects of plant closures, I use wage and employment data from the U.S. Bureau of Labor Statistics (BLS). The Quarterly Census of Employment and Wages (QCEW) collects quarterly employment and wage data at the county level as reported by employers. I use the wages per employee, disaggregated by county and industry. These wages include total compensation, bonuses, stock options, severance payments, the cash value of meals and lodging, tips, and other gratuities. I annualize these wages for ease of interpretability. The Local Area Unemployment Statistics (LAUS) aggregates data from state-level workforce agencies. It includes monthly information on the number of employed and unemployed individuals for every U.S. county.

3 Descriptive Evidence

Before diving into the theoretical model, I document trends around the Maytag acquisition in terms of concentration, prices, product entry, and appliance manufacturing employment.

3.1 Changes in concentration

Table 1 shows the evolution of brand owner shares around the time of the Maytag acquisition. Prior to the merger, Whirlpool and Maytag were the largest and third largest brand owners for laundry products in the U.S. market. Since Sears does not manufacture any appliances itself, Whirlpool and Maytag were also the largest and second largest laundry product manufacturers. In contrast, Haier had no significant market shares in either product market.

The pre-merger Herfindahl-Hirschman Index (HHI) and the increase in the HHI because of the merger based on pre-merger market shares indicate that the transaction led to a strong increase in concentration.¹⁹ According to the U.S. horizontal merger guidelines, the acquisition therefore potentially raises significant competitive concerns.²⁰

The evolution of market shares from just after the merger in 2007 to 2009 shows that although some rivals gained market shares and the HHI gradually declined (as compared to the post-merger HHI based on pre-merger market shares), the increase in concentration due to the merger remains substantial and persistent.

¹⁹The HHI is calculated as the sum of squared market shares using whole percentages.

²⁰The U.S. horizontal merger guidelines identify mergers with a pre-merger HHI between 1,500 and 2,500 and an increase in the HHI by more than 100 as potentially raising significant competitive concerns.

Table 1: Volume share by brand owner (%)

	Clothes washers			Clothes dryers		
	2005	2007	2009	2005	2007	2009
Whirlpool	25	44	42	27	44	42
Maytag	23			21		
Sears	25	20	18	25	21	19
General Electric	14	17	16	15	17	16
Electrolux	7	6	6	7	6	5
LG	3	7	10	2	6	10
Samsung	0	1	5	0	1	5
HHI	2,048	2,729	2,506	2,072	2,784	2,507
Δ HHI	1,149			1,124		

Notes: The table shows the market share in terms of volume sales by brand owners for clothes washers and clothes dryers pre-merger (2005) and post-merger (2007 and 2009). The HHI is calculated as the sum of squared market shares using whole percentages. The increase in the HHI is based on pre-merger market shares.

3.2 Evolution of prices

I next turn to the descriptive evolution of prices around the time of the acquisition. Ashenfelter, Hosken, and Weinberg (2013) compare the evolution of Maytag and Whirlpool product prices for appliance categories with a large increase in concentration to categories with low increases in concentration. Since I use a different data source, I repeat the descriptive price analysis. In particular, the *NPD* data used by Ashenfelter, Hosken, and Weinberg (2013) only includes product sales at a subset of retailers (e.g., omitting sales at Sears), which could lead to systematically different results.

As a comparison appliance category, I use freestanding ranges.²¹ This is an appropriate control group if, in the absence of the merger, prices would have evolved similarly in the treatment and the control groups.

I estimate the price effects around the time of the merger separately for Maytag and Whirlpool products. To do this, I estimate the parameters of the following model for washers (treatment) and freestanding ranges (control) and for dryers (treatment) and freestanding ranges (control)

$$\log(p_{it}) = \alpha_1 \text{Maytag}_{it} \times \text{post}_t + \alpha_2 \text{Whirlpool}_{it} \times \text{post}_t + \beta x_{it} + \tau_i + \gamma_t + \epsilon_{it}. \quad (1)$$

²¹Ashenfelter, Hosken, and Weinberg (2013) use ranges, cooktops, ovens and freezers as a comparison.

Table 2: Reduced form price effects of the Maytag acquisition

	Washers vs. ranges			Dryers vs. ranges		
	(1)	(2)	(3)	(4)	(5)	(6)
Merging parties \times post	-0.030 [-0.076, 0.016]			-0.017 [-0.081, 0.046]		
Maytag \times post		-0.049** [-0.097, -0.001]	-0.026 [-0.070, 0.018]		-0.043 [-0.097, 0.011]	-0.015 [-0.063, 0.032]
Whirlpool \times post		-0.016 [-0.077, 0.045]	-0.006 [-0.036, 0.023]		0.007 [-0.048, 0.062]	0.028 [-0.018, 0.075]
Characteristics controls	Yes	Yes	Yes	Yes	Yes	Yes
Quarter \times 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	3599	3599	3280	4088	4088	3739

Notes: Columns (1) to (3) compare the logarithm of prices for clothes washers and freestanding ranges. Columns (4) to (6) compare the logarithm of prices for clothes dryers and freestanding ranges. Differences in observations in columns (3) and (6) as compared to preceding columns are due to the iterative dropping of singleton observations when clustering standard errors. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the brand level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The parameters of interest are α_1 , which captures the average price increase for Maytag products and α_2 , which captures the average price increase for Whirlpool products.

Table 2 includes the estimates of the reduced form effects of the Maytag acquisition on the logarithm of prices. Columns (1) and (4) include estimates from a regression where I pool Maytag and Whirlpool products together and estimate a joint price effect. These results suggest that there is no large price increase for clothes washers or dryers. Based on the 95% confidence intervals, I reject price increases of more than 1.6 percent for clothes washers and 4.6 percent for dryers.

In Columns (2) and (5), I disaggregate this by Maytag and Whirlpool products. Based on the 95% confidence intervals, I reject large price increases for Maytag products in both categories. For Whirlpool products, the point estimates are just below (washers) and just above (dryers) zero, however, the width of the confidence intervals do not allow me to reject price changes of between -7.7 and $+4.5$ percent for clothes washers and -4.8 and $+6.2$ percent for clothes dryers. In Columns (3) and (6) I repeat the previous analysis, however swapping brand fixed effects for more granular product fixed effects. This leads to a smaller price decrease for merging party products after the merger, but decreases are still found for Maytag clothes washers and dryers and Whirlpool washers.

A causal interpretation of these results could lead to two conclusions: First, the acquisition of Maytag by Whirlpool at most mildly increased prices for laundry products. Second, the acquisition similarly affected clothes washers and dryers.

Irrespective of whether these findings are causal, they are only partially in agreement with the findings by Ashenfelter, Hosken, and Weinberg (2013). In line with their results, I

do not find any reduced form evidence for clothes washer price increases. In contrast to their results, I also do not find any reduced form evidence for large price increases for dryers.²² Given the very similar evolution of market shares and prices for washers and dryers, it seems plausible to expect similar price effects of the merger for both categories.

In any event, the estimated price effects from the reduced form regressions should be interpreted with great caution. As previously described, a causal interpretation of these results requires that prices for laundry products would evolve similarly to prices for freestanding ranges in the absence of the merger. For example, as noted by Ashenfelter, Hosken, and Weinberg (2013), product entry by LG and Samsung in the market for clothes washers may confound the reduced form estimates of the price effects of the merger. These entries may or may not be related to the merger.

Finally, the regression analysis does not treat products differently depending on their relative importance in the marketplace (i.e., their market share). Thus, if price changes are not homogeneous across all products, the estimated price changes may strongly be influenced by many products with relatively low market shares. If these are products that most consumers do not consider in any case, this may not be the most informative estimate to assess the price effects experienced by consumers.

3.3 Product entry

Rival product entry could affect the estimated price effects of the merger in two distinct ways: First, if the merger leads to merger-specific product entry, this can increase competition and decrease prices. Second, if there is merger-independent product entry by rivals around the time of the merger, this could increase competition and reduce prices.

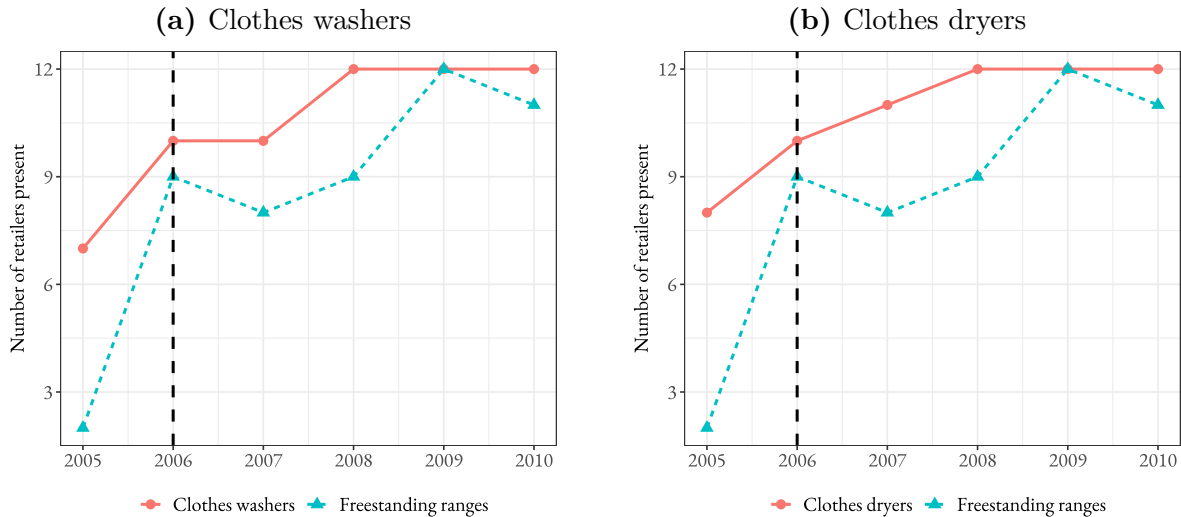
I therefore assess whether product entry by LG and Samsung occurred in the U.S. laundry market and whether this was different to entry patterns for freestanding ranges.

Figure 2 shows the evolution of the retailer presence by LG and Samsung for clothes washers, dryers, and freestanding ranges. Since I distinguish between five major retailers and “other retailers”, the sum of retailers carrying LG and Samsung can at most be twelve. Two trends emerge: First, the number of retailers carrying LG and Samsung laundry products increases around the time of the merger. By 2008, all major retailers carry LG and Samsung clothes washers and dryers. Second, there is a strong and persistent increase in the number

²²Compared to ranges, they find an increase in prices for Maytag dryers newly introduced after the merger of 3 percent and of 14 percent for Whirlpool dryers newly introduced after the merger. They also find that the acquisition did not change prices of old Maytag dryers and reduced prices of old Whirlpool dryers by 6 percent. Unfortunately, the data does not allow me to identify when a product was first introduced to the market and so I cannot make this additional decomposition.

of retailers carrying LG and Samsung freestanding ranges. Growth is stronger, as it starts from a very low level, however full retailer coverage is only temporarily reached in 2009.

Figure 2: Retailer presence LG and Samsung by product category



Notes: The solid red lines show the sum of retailers that carry clothes washers (left) or dryers (right) by LG and Samsung summed together. The dashed blue line shows the sum of retailers that carry freestanding ranges by LG and Samsung.

These results suggest that product entry occurred but was not necessarily merger-specific. Indeed, if we believe that merger-independent entry for laundry products is similar to the observed product entry for freestanding ranges, we would expect to observe product entry by LG and Samsung also in the absence of the Whirlpool acquisition.

3.4 Labor market effects of plant closures

The analysis so far focused on the product market effects of the acquisition. Different acquisitions may also entail different changes to employment. For those to enter the overall welfare effects, appliance manufacturing jobs need to matter for local labor markets. In the following, I assess how Maytag plant closures by Whirlpool post-acquisition affected employment, unemployment, and wages of the employed in affected counties.

Although Whirlpool maintained some of Maytag's manufacturing plants (e.g., in Amana, Iowa, or Cleveland, Tennessee), shortly after the acquisition it shut down appliance manufacturing plants in Searcy, Arkansas (700 manufacturing jobs) and Herrin, Illinois (1,000 manufacturing jobs), as well as manufacturing and headquarter operations in Newton, Iowa (1,000 manufacturing and 1,800 corporate jobs). At the same time, Whirlpool announced adding 1,500 jobs at two existing plants in Ohio.

Table 3: Reduced form labor market effects of plant and HQ closures

	Unemployment (persons)		Employment (persons)		Wages (\$)	
	(1)	(2)	(3)	(4)	(5)	(6)
Plant & HQ closure $\times \mathbb{1}$ (year = 2007)	163***		-1140***		-2472***	
	[151, 176]		[-1343,-937]		[-2666, -2278]	
Plant & HQ closure $\times \mathbb{1}$ (year = 2008)	291***		-1716***		-6508***	
	[263, 319]		[-1983,-1449]		[-6740, -6275]	
Plant closure $\times \mathbb{1}$ (year = 2007)		257		-288*		-329
		[-189, 704]		[-597,21]		[-1590, 931]
Plant closure $\times \mathbb{1}$ (year = 2008)		8		-336**		-400
		[-545, 561]		[-639,-33]		[-1815, 1014]
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,752	8,352	4,752	8,448	1,584	2,816
Mean outcome in treated counties	1,130	1,123	11,840	13,815	34,404	25,524

Notes: Columns (1) and (2) compare the absolute number of unemployed persons in treated counties to all other counties in the same state. Columns (3) and (4) compare the absolute number of employed persons in treated counties to all other counties in the same state. Columns (5) and (6) compare the average annualized gross wage of employed persons in treated counties to all other counties in the same state. Columns (1), (3) and (5) compare Jasper County (county of Newton) to all other counties in Iowa. Columns (2), (4) and (6) compare White County (Searcy) and Williamson County (Herrin) to all other counties in Arkansas and Illinois. Cook County (county of Chicago), is omitted from any analyses involving Illinois. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To quantitatively assess the local labor market effects of changes in Maytag employment, I estimate the parameters of the following regression model

$$\text{outcome}_{it} = \alpha_1 \mathbb{1}(\text{year}_t = 2007) \times \Delta \text{jobs}_i + \alpha_2 \mathbb{1}(\text{year}_t = 2008) \times \Delta \text{jobs}_i + \tau_i + \gamma_t + \epsilon_{it}, \quad (2)$$

where outcome_{it} is the number of employed persons, unemployed persons, or the average wage of employed persons in a particular county i and time period t , Δjobs_i is an indicator variable equal to one if a particular county is affected by job cuts or newly created jobs by the merging parties, τ_i are county fixed effects and γ_t are time fixed effects.

I group counties into three different treatment groups and estimate separate regressions. The first treatment group is Jasper County, in which there was a shut down of manufacturing and corporate operations. The second treatment group consists of White County and Williamson County, in which only manufacturing plants were shut down. The third group consists of Marion County and Sandusky County, where Whirlpool created new jobs.

Table 3 summarizes the regression estimates for the elimination of jobs. Column (1) reports the effects on unemployment in Jasper County. I find that there is a statistically and economically significant increase in unemployment. The effect is persistent throughout the observation period, but is small in magnitude (around 300 persons in 2008) compared to the number of Maytag jobs lost (1,000 manufacturing and 1,800 corporate jobs). This however only tells part of the story, as it masks other shifts into non-employment, such as

Table 4: Reduced form labor market effects of new jobs

	Unemployment (persons)	Employment (persons)	Wages (\$)
	(1)	(2)	(3)
New jobs \times 1 (year = 2007)	-33 [-178,112]	358 [-151,867]	-88 [-412,237]
New jobs \times 1 (year = 2008)	-230** [-458,-2]	656 [-169,1480]	-271 [-1299,758]
County fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Observations	4,224	4,224	1,408
Mean outcome in treated counties	2,067	27,006	32,452

Notes: Column (1) compares the absolute number of unemployed persons in Marion County (Marion) and Sandusky County (Clyde) to all other counties in Ohio. Column (2) compares the absolute number of employed persons in Marion County and Sandusky County to all other counties in Ohio. Column (3) compares the average annualized gross wage of employed persons in Marion County and Sandusky County to all other counties in Marion County and Sandusky County. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

early retirements, exits into education, as well as out-migration. The results in Column (3) show that the number of employed persons in Jasper County as compared to before the closing of operations declined by around 1,700. Finally, Column (5) shows the effect on annualized average wages of employed persons. Again, there are large and statistically significant decreases in average wages.

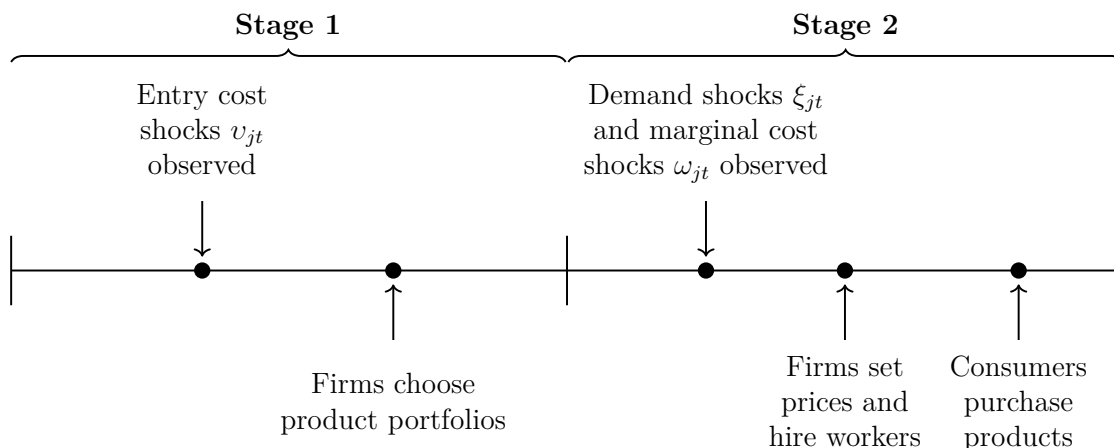
Columns (2) and (4) include the effects on unemployment and employment of only shutting down plants without the effects of closing the HQ. There is an economically meaningful but statistically noisy increase in unemployment. This effect however appears to only be transitory and disappears after a year. There is a more robust and persistent decrease in employment in affected counties of around 300 persons. Since the affected plants in the two treatment counties employed 700 and 1,000 persons respectively, this suggests that around a third to half of the jobs were permanently lost and led to out-migration or other shifts into non-employment, beyond unemployment. The results in Column (6) suggest that there is a small decrease in average wages, which is imprecisely estimated.

Table 4 shows the effect of relocating 1,500 new jobs to two existing Whirlpool plants in two different counties in Ohio. On average, this is equivalent to 750 new jobs per affected county. The results in Columns (1) and (2) suggest that these jobs led to a significant reduction in unemployment in 2008 and an increase in employment. Wages do not increase, suggesting that these jobs do not lead to positive wage pressure on the local labor market.

4 The Model

Three observations emerge from the preceding analysis. First, entry played a crucial role in the product market, so understanding entry is necessary to assess the effects of the merger. Second, although there were many changes in the product portfolios of existing firms, there was no entry by a new firm. The focus thus lies on endogenous portfolio choices and I abstract from firm-level entry. Third, there are frictions in local labor markets and important differences in the production locations of manufacturers. Where products are produced and by whom affects the welfare effects of the merger.

The model features manufacturers and consumers. Manufacturers choose their product portfolios and prices. Consumers make purchase decisions. The model proceeds in two stages. In the first stage, firms are endowed with a set of potential products that they are technologically capable of producing and their production locations. They observe product-level shocks to entry costs and decide which products to offer. At this stage, firms do not observe transitory demand and marginal cost shocks and only form expectations about them. In the second stage, these shocks realize and are observed by firms, upon which they set prices. Finally, households observe the products on offer and their characteristics, including prices, and make purchase decisions. The number of domestic production jobs depends on equilibrium quantities in the product market and the location of production.



I solve this game backwards by searching for the Subgame Perfect Nash Equilibria (SPNE) of the game.²³ To estimate the parameters of the game, I require the existence of a SPNE but not its uniqueness.

²³Whenever cost or demand shocks are observed by market participants, they remain unobserved by the econometrician.

4.1 Demand model

Demand is a household-level discrete choice between different clothes washers. The demand model is based on the empirical discrete choice demand literature following Berry, Levinsohn, and Pakes (1995) and Berry, Levinsohn, and Pakes (2004). Every year, a household makes the choice between different clothes washers on offer in the market as well as not purchasing a clothes washer, i.e., choosing the outside good. This can be thought of as keeping the clothes washer already owned by the household or not owning a clothes washer at all.

The utility of household i from buying clothes washer j in year t can be written as

$$u_{ijt} = x_{jt}\beta + \sigma^{FL}\nu_{it}^{FL}x_{jt}^{FL} - (\alpha + \kappa_{\alpha} \min(\$400k, z_{it}))p_{jt} + \xi_{jt} + \epsilon_{ijt}. \quad (3)$$

The vector x_{jt} includes non-price product characteristics, such as whether a clothes washer can be loaded from the front, whether it is Energy Star certified, or the number of special programs it includes. It also includes indicator variables for the brand and retailer at which the clothes washer was purchased, as well as year fixed effects and brand time trends.²⁴ p_{jt} is the price of a clothes washer j at time t . I denote the set of products among which households can choose at time t as J_t .

Average tastes for price and non-price characteristics are captured by α and β respectively. x_{jt}^{FL} is an indicator variable for whether a particular clothes washer is a front-loader. ν_{it}^{FL} is an i.i.d. draw from a standard normal distribution and represents a household-specific unobserved taste shock for front-loaders. z_{it} is the income of household i at time t . Household incomes are capped at \$400,000, as this avoids positive price coefficients for households with very high incomes which can arise when income enters the price coefficient linearly. Incomes beyond this threshold have negligible effects on the estimated demand parameters in practice. σ^{FL} measures the dispersion in taste for front-loaders between households. κ_{α} captures how the sensitivity to prices varies with household income.

The remaining part of the utility function consists of an unobservable component constant across households, ξ_{jt} , as well as an idiosyncratic household-specific unobservable, ϵ_{ijt} . ξ_{jt} includes quality differences not captured by the product characteristics and fixed effects, as well as transitory demand shocks that vary between products but are common across households. Finally, ϵ_{ijt} is an i.i.d. draw from a type I extreme value (Gumbel) distribution.

²⁴The full list of product characteristics are the price, the brand repair rate, the total advertising expenditure at the brand level, as well as indicator variables for whether a clothes washer is a front-loader, a Korean front-loader, a front-loader by Fisher & Paykel, a high-end European front-loader (i.e., Asko, Bosch, or Miele), has an agitator, is part of a stacked pair, has a stainless steel exterior, has a white exterior, is Energy Star certified, has additional noise insulation, has a child lockout. Finally, it includes retailer, brand and year fixed effects, as well as linear brand time trends.

To simplify notation, I separate utility into the mean utility δ_{jt} and the household-specific deviation $\mu_{ijt} + \epsilon_{ijt}$. The mean utility includes all utility components that are constant across households. I also define a vector $\theta = (\theta_1, \theta_2)$ which contains all the parameters of the demand model. Let $\theta_1 = (\alpha, \beta)$ contain all linear parameters of the model and $\theta_2 = (\sigma, \kappa)$ all nonlinear parameters. Since I can only identify utilities up to an affine transformation, I normalize the mean utility of the outside good to zero, after which the utility of a household for the outside good reduces to ϵ_{i0t} .

The distributional assumptions on the household-specific unobservable allow deriving the familiar logit choice probabilities from this specification. By integrating over the joint distribution of household demographics $P_D(D)$ and the joint distribution of unobserved taste shocks $P_\nu(\nu)$, the model-predicted market share of product j in market t becomes

$$s_{jt} = \int \frac{\exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_{k \in J_t} \exp(\delta_{kt} + \mu_{ikt})} P_D(D) P_\nu(\nu). \quad (4)$$

4.2 Second stage: pricing

In the second stage, firms observe demand and marginal cost shocks and subsequently set prices. Each firm f chooses prices for the set of products it offers, J_{ft} , to maximize its variable profits, given by

$$VP_{ft} = \sum_{j \in J_{ft}} (p_{jt} - mc_{jt}) s_{jt} M_t, \quad (5)$$

where p_{jt} denotes the price of j at t , mc_{jt} its marginal costs and M_t denotes the total market size. Firms set prices by maximizing variable profits. For each product j , the equilibrium price must satisfy

$$p_{jt} = mc_{jt} - [(\nabla_p s \bullet \Lambda)^{-1} s]_{jt}, \quad (6)$$

where Λ is the ownership matrix and $\nabla_p s$ is the matrix of partial derivatives of market shares with respect to prices.²⁵

Marginal costs can be decomposed into several components. In particular, the inverse hyperbolic sine of marginal costs depends on product- and market-specific components in the following way

$$\operatorname{arcsinh}(mc_{jt}) = [x_{jt}, ic_{jt}] \gamma + \omega_{jt}, \quad (7)$$

where ic_{jt} is a vector of input costs, γ captures how product characteristics and input costs affect marginal costs and ω_{jt} is a transitory product-level unobserved marginal cost shock

²⁵The ownership matrix contains information on whether two products are offered by the same firm and so cross-price effects matter for the optimal pricing decision of firm f .

that is realized and observed by the firms in the pricing stage.²⁶

4.3 First stage: entry

In the first stage, firms decide which products to offer. At the outset each firm is endowed with a set of potential products it can offer in market t , \mathcal{J}_{ft} . This can be thought of as the set of products that it is technologically capable of producing. It includes products that it sells already at a different retailer or in a different market and minor adjustments to existing products which it could perform in the short-term. It does not include products for which a firm would need to develop entirely new capabilities (e.g., launching its first front-loader).

Introducing a product into the market comes at a fixed and sunk cost. This includes costs related to the final development of a product (e.g., a particular front-loader), marketing or retailer investments. Empirically, I analyze markets at the yearly level. At the same time, Ashenfelter, Hosken, and Weinberg (2013) show that the volumes of a particular model rapidly decline after twelve months. It therefore seems plausible that the fixed and sunk cost of introducing a product at a retailer in a particular year is independent of the product portfolio in previous years.

The fixed cost of introducing a new product can be decomposed into a brand- and market-specific component F_{bt} and a mean-zero idiosyncratic product- and market-specific fixed cost shock v_{jt} . Thus $F_{jt} = F_{bt} + v_{jt}$ and $E[v_{jt}|j \in \mathcal{J}_{ft}] = 0$. Before deciding on its product portfolio, a firm observes the fixed cost shocks related to all products it could potentially add. It does not yet however observe the second stage marginal cost and demand shocks which I summarize as $e_{jt} = (\xi_{jt}, \omega_{jt})$. Instead, it chooses a product portfolio by trading off expected variable profits and the sum of fixed costs of different products. More specifically, it solves the following maximization problem:

$$\max_{J_{ft} \subseteq \mathcal{J}_{ft}} \{ \Pi = E[VP(p)|J_{ft}] - \sum_{j \in J_{ft}} F_{jt} \}. \quad (8)$$

Since choosing an optimal product portfolio is a discrete choice, the first order conditions of this profit maximization only hold with inequality.

²⁶The inverse hyperbolic sine is a transformation that approximates the natural logarithm. Its advantage is that zero is part of its definition area and it returns real numbers for negative inputs. See Bellemare and Wichman (2020) for more details.

4.4 Demand for domestic workers

Let us now turn to the employment side. The aim of this exercise is to model how the number of U.S. manufacturing jobs changes if we hold production locations and the production technology fixed. I therefore do not model demand and supply in the labor market itself. That is not to say that the number of U.S. clothes washer manufacturing jobs would not change if, for example, wages increased. This would be reflected in the marginal costs of a clothes washer and thus affect equilibrium prices and quantities in the product market.

I assume that firms make longer-term decisions on where to produce which product outside of the model. The share of each product that is produced in the U.S. is therefore exogenously given. Similarly, the production technology $G(\cdot)$ is fixed and the number of manufacturing workers required is linear in the number of clothes washers. The demand for domestic clothes washer manufacturing workers by firm f therefore is

$$LD_{ft} = \sum_{j \in J_{ft}} G(q_{jt}) \times \text{domestic}_{jt}. \quad (9)$$

5 Estimation

In this section, I describe how to estimate the parameters of the model. As for the model description I proceed in reverse-order, beginning with the demand parameters.

5.1 Demand

The estimation of the demand parameters is similar to Berry, Levinsohn, and Pakes (2004). In a first step, I estimate the non-linear parameters of the utility function, σ^{FL} and κ_{α} . I identify these parameters by matching simulated moments to their analogues in the data. Informally, we can think of the data moments as identifying the structural parameters of their simulated equivalent.

The first data moment is based on the correlation between the clothes washer bought being a front-loader and the average share of front-loaders among the second-choice brand. Respondents to the *TraQline* survey are only asked which other brands they considered buying but not which exact model. Some brands carry both front-loaders and top-loaders. However, the share of front-loading clothes washers differs greatly between brands. Furthermore, the correlation between whether the first choice is a front-loader and the share of front-loaders among the second choice brand is important, with a correlation coefficient of 0.4. This suggests that there is a strong unobserved taste for front-loaders among some households, which can affect substitution patterns.

The second data moment is based on the correlation between the household income and the price of a clothes washer bought. The correlation coefficient between the two is 0.5.²⁷ On average, the higher the income of a household, the higher the price of a clothes washer bought. This suggests that high income households are less sensitive to prices.

To estimate the linear parameters of the utility function, α and β , I first estimate the vector of mean utilities, δ , by matching simulated market shares for each product to observed market shares. To do so, I introduce an additional assumption:

Assumption 1. $E[e_{jt}|X_{jt}, F_{jt}] = 0$ for each $j \in \mathcal{J}_t$.

This means that the second stage demand and marginal cost shocks are independent of the non-price product characteristics and the fixed costs of introducing a product. As explained by Eizenberg (2014), this is slightly stronger than the assumption that e_{jt} is realized after products are chosen, since it also means that firms cannot predict e_{jt} . This assumption nevertheless seems reasonable, as firms may still predict future costs and demand as they relate to observable characteristics, which I can control for. It only means that firms cannot predict unobservable transitory marginal cost and demand shocks.

Since prices can be adjusted frequently, they are likely correlated with ξ_{jt} . As explained in Section 2, I use an instrumental variable based on the production location and the real exchange rate, which affects costs but is otherwise unrelated to demand.

For the linear and non-linear demand parameters, standard errors are clustered at the brand level using the residual bootstrap.

5.2 Marginal costs and fixed cost bounds

I compute marginal costs for each product by inverting the first order conditions of each firm's profit maximization problem. Under the model assumptions described above, the data are rationalized by a unique marginal cost and markup for each product.

The entry model in Section 4.3 only provides inequality conditions for profitable entry. It is hence not possible to point identify entry costs. Instead, I resort to partial identification and seek to estimate bounds on the identified set of fixed entry costs for every brand.

To estimate bounds on the fixed costs of adding a product, I need to determine the set of potential products of each firm. I refer to all products that a firm could have added as the potential products, to the potential products that it actually added as the active products, and to the potential products that it chose not to add as the inactive products. Recall that the set of potential products of firm f in market t is denoted as \mathcal{J}_{ft} and the set of active products as J_{ft} . I denote the set of inactive products of firm f as \tilde{J}_{ft} .

²⁷Figure A.8 shows a scatter plot of the relationship between the household income and price.

The set of active products are those products that we observe in the data. Before determining the set of inactive products, it is worth remembering that the goal is to estimate the fixed costs of adding or removing a product that is part of the set of products a firm is technologically capable of producing. Thus, if a firm does not have any front-loader among its active products, I do not consider that it could have added a such a washer in that year. Instead, I exploit the fact that I can distinguish sales at the retailer level and that appliance brand owners introduce different products at different retailers. For any active product (e.g., a front-loader by KitchenAid sold at Sears), all versions of the product that I do not observe in the data (e.g., a front-loader by KitchenAid sold at another retailer) is an inactive product. I therefore capture the fixed costs related to marketing, getting retail floor space for an additional product or customizing the product for the clientele of a particular retailer, but not of developing new technologies. This is appropriate in this case, since I am interested in estimating how the incentives to make portfolio adjustments change for existing players with already developed product portfolios.

The estimation of the bounds on fixed costs resembles the procedure described by Eizenberg (2014). If the product entry that I observe is a pure strategy SPNE, then no firm can profitably deviate unilaterally from this equilibrium. This means that no firm can increase its expected profits by unilaterally adding inactive products or removing active products. To estimate bounds on the fixed costs of adding a product, I exploit a subset of the equilibrium conditions, namely that no firm has a profitable one-step deviation.²⁸

Let us denote the equilibrium product portfolio (i.e., the set of active products) of firm f at time t as J_{ft}^* . For each active product j that a firm chooses to introduce in equilibrium, an upper-bound on the fixed cost of introducing the product is the expected incremental profit of offering that product holding other products fixed. That is,

$$F_{jt} \leq E_e[VP_{ft}(J_{ft}^*) - VP_{ft}(J_{ft}^* - \mathbf{1}_{ft}^j)] \equiv \bar{F}_{jt}, \quad (10)$$

where \bar{F}_{jt} is the upper-bound on fixed costs of adding product j at time t .

For each inactive product, a lower-bound on the fixed cost is the expected incremental profit of offering that product holding other products fixed. That is,

$$F_{jt} \geq E_e[VP_{ft}(J_{ft}^* + \mathbf{1}_{ft}^j) - VP_{ft}(J_{ft}^*)] \equiv \underline{F}_{jt}, \quad (11)$$

where \underline{F}_{jt} is the lower-bound on fixed costs of adding product j at time t .

²⁸In principle, I could add further restrictions on fixed cost bounds due to the lack of profitable multi-step deviations. In practice, restrictions based on multi-step deviations may be difficult to use, since the additional inequalities would include idiosyncratic fixed cost shocks v_{jt} for each product.

These two conditions allow estimating the upper-bound on fixed costs of active products and the lower-bound on fixed costs of inactive products. I estimate the expected incremental variable profits using 500 draws from the joint empirical distribution of the demand and marginal cost shocks e_{jt} . Ultimately, I am interested in bounds on the brand-level average fixed costs in market t , F_{bt} . Constructing the upper-bound on F_{bt} only based on active products and the lower-bound based on inactive products is inadmissible, since product portfolio decisions are not independent of v_{jt} , i.e. $E[v_{jt}|j \in J_{ft}] \neq 0$. Recall, however, that $E[v_{jt}|j \in \mathcal{J}_{ft}] = 0$, which means that the product-level fixed cost shock has mean zero conditional on products being part of the set of potential products. This means that if I can estimate a lower-bound on the fixed costs of adding active products and an upper-bound on the fixed costs of adding inactive products, I can get an unbiased estimate of bounds on the set of brand-level average fixed costs F_{bt} .

To fill the missing bounds, I follow the approach proposed by Eizenberg (2014). The details of the estimation procedure are described in Appendix II.D. I follow arguments by Imbens and Manski (2004) and Eizenberg (2014) to construct confidence sets for F_{bt} . These are based on one-sided intervals around the point-identified upper- and lower-bound of the fixed cost interval. I estimate the inputs into the confidence sets using the same bootstrapped samples as for the demand estimation to account for variance in the estimation of θ and simulation error in the estimation of variable profits.

5.3 U.S. employment

Estimating the equilibrium number of U.S. clothes washer manufacturing jobs under different scenarios requires the overall number of employees necessary to manufacture clothes washers in each scenario, as well as the corresponding production locations.

Recall that I assume that the number of employees necessary for the manufacturing process is directly proportional to the number of clothes washers sold. To simplify estimation, I assume that the production technology is linear and constant across products and manufacturers. I use information on the number of employees and clothes washer production from annual reports and news articles to calibrate the annual average number of clothes washers a manufacturing worker produces. I combine this number with the equilibrium quantity sold of each product, to estimate how many manufacturing jobs are necessary globally.²⁹

The second step is to estimate the share of clothes washers produced in the United States. As described in greater detail in Section 2.3, I construct a granular data set that contains product-level information on the production location of clothes washers produced

²⁹I describe this calibration in more detail in the Appendix Section II.E.

for the U.S. market. The equilibrium number of U.S. clothes washer manufacturing jobs is the share of global manufacturing jobs multiplied by the share of a product’s U.S. production.

6 Estimation Results

6.1 Demand

Table 5 includes the demand estimates. Column (1) reports the first-stage results, where I regress the endogenous price variable on the instrumental variable (IV) for price, which is the real exchange rate, and include full controls. The results indicate that an increase in the RER by a full unit leads to an increase in clothes washer prices by \$191. The F-statistic is approximately 23, suggesting that the IV is relevant.

Column (2) includes the reduced form estimates after regressing the outcome variable (the average utility that consumers get from purchasing clothes washer j at time t , δ_{jt}) on the instrument. The higher the RER, the lower the purchasing utility for a consumer. In Columns (3) and (4), I report the price coefficient for the simple logit demand model using OLS and the IV, respectively. By accounting for the endogeneity of prices, the average product-level own-price elasticity of residual demand changes from -0.96 to -2.42 . Finally, I report the price effects for the full mixed logit model using IV in Column (5). The results suggest that there are significantly heterogeneous but correlated preferences across households. As expected, households with a higher household income are less price sensitive. Furthermore, households that purchase front-loaders also have an above average unobserved preference for other front-loaders. Accounting for these effects, I estimate that the average own-price elasticity of residual demand for clothes washers further reduces to -3.26 .³⁰

6.2 Marginal cost

Across brands and time, I find that average marginal costs are around \$410 and range between close to zero and around \$1,500. The average Lerner Index is 40 percent.

Figure 3 shows the evolution of marginal costs and the Lerner Index by brand owner over time. The timing of the merger coincides with an industry-wide fall in marginal costs. As we do not expect marginal cost efficiencies of the merger at rival firms, these decreases in marginal costs cannot all be merger-specific. They should thus be incorporated in a counterfactual estimation of prices but for the merger. At the same time profit margins

³⁰These elasticity estimates are comparable in magnitude to results by Houde (2018), who finds short-term own-price elasticities of residual demand for refrigerators of between -5.41 and -4.15 , depending on household income and using weekly data.

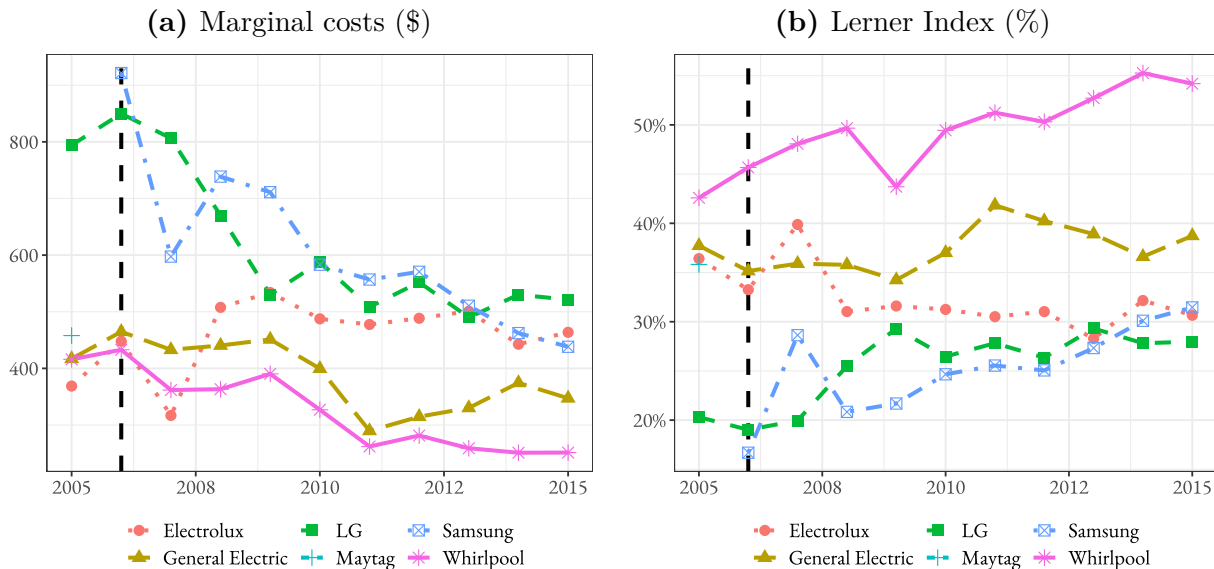
Table 5: Demand estimates

	(1)	(2)	(3)	(4)	(5)
	First-stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
<i>Dependent variable:</i>	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	
<i>Linear parameters</i>					
Real exchange rate	1.909*** (0.398)	-0.787** (0.358)			
Price ('00 2012 \$)			-0.164** (0.062)	-0.412** (0.202)	-0.637*** (0.024)
<i>Non-linear parameters</i>					
Income effect κ_α					0.089*** (0.011)
Unobserved taste σ^{FL}					2.495*** (0.017)
Characteristics	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Retailer FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Brand FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Brand time trends	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	1,590	1,590	1,586	1,590	1,590
Kleibergen-Paap F-statistic	22.979				
Avg. own-price elasticity			-0.964	-2.416	-3.301

Notes: Column (1) presents results for the first stage regression of prices on the real exchange rate. Column (2) includes reduced form estimates for the simple logit model. Column (3) reports demand estimates for the simple logit without a price instrument. Column (4) presents demand estimates for the simple logit model using the RER as an IV for price. Column (5) shows demand estimates for the mixed logit model presented in Section 4 and using the RER as an IV. I estimate κ_α , σ^{FL} , and $\hat{\delta}_{jt}$ using simulated method of moments and the linear parameters using linear IV regression. Standard errors are clustered at the brand level. The own-price elasticity of residual demand is computed at the product level and the average is calculated by weighting products according to their sales volume. Estimates for non-price characteristics are reported in Table A.2. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

increased.³¹ This is also true for Whirlpool, although its Lerner Index after 2006 also includes Maytag, which had a significantly lower Lerner Index than Whirlpool pre-merger.

Figure 3: Evolution of marginal cost and Lerner Index by brand owner



Notes: The plots show the evolution of marginal costs (left) and the Lerner Index (markup over price; right) by brand owner over time. The vertical line shows the time of the Maytag acquisition by Whirlpool. From 2006 onwards, Whirlpool also includes former Maytag products.

6.3 Fixed cost bounds

Finally, I estimate bounds on the fixed and sunk costs of product entry at the brand-level. Before interpreting these results, it is worth remembering that a product is defined as the combination of a brand, a retailer and major clothes washer characteristics (i.e., the distinction between front-loaders, regular top-loaders and high-efficiency top-loaders). Thus, the fixed cost sets that I estimate should be thought of as the cost of adding a product category (brand and major characteristic combination) at a particular retailer.

Table 6 describes the 95 percent confidence sets on the fixed costs of adding new products. I find that the range of plausible fixed costs to add products involves higher values for brands with large market shares (e.g., Maytag or Whirlpool) than brands with lower market shares (e.g., KitchenAid, Hotpoint or Westinghouse). This could be because the former are only offered at a retailer if this involves a full range of clothes washers within that product category, requiring more floor space as well as higher marketing expenditures.

³¹Large movements in the Lerner Index for Samsung between 2006 and 2007 should be interpreted with caution, since these are based on relatively few Samsung products at the time.

Table 6: Brand-level fixed costs of adding a product (\$M)

Brand owner	Brand	95 % confidence sets
Maytag	Admiral	[1.8, 2.8]
	Amana	[1.7, 3.8]
	Maytag	[7.5, 33.5]
Whirlpool	KitchenAid	[0.4, 1.9]
	Roper	[1.1, 4.9]
	Whirlpool	[6.5, 35.2]
General Electric	General Electric	[2.4, 21.0]
	Hotpoint	[0.4, 1.7]
Electrolux	Frigidaire	[2.4, 8.7]
	Westinghouse	[0.4, 1.9]
LG	LG	[2.0, 16.6]
Samsung	Samsung	[3.1, 7.8]

Notes: Brand-level fixed costs of adding or removing a product are based on all active and potential products in 2005 and 2007.

7 Welfare Effects of the Whirlpool Acquisition

Using the parameter estimates, I can simulate and compare the welfare effects of alternative Maytag acquirers. Since Haier had close to no presence in the U.S. laundry market prior to the merger, without marginal cost efficiencies, an acquisition by Haier is approximately equivalent to keeping a standalone Maytag in the product market.

To estimate how U.S. employment differs between the acquisition scenarios, I assume that Haier would offshore all Maytag jobs to China, whereas I use the observed post-merger production locations by Whirlpool after its acquisition of Maytag. The latter is necessary to also account for the partial offshoring of former Maytag manufacturing jobs by Whirlpool. Without doing so, I would overestimate the number of jobs maintained by Whirlpool.

7.1 Players and potential products

Endogenizing portfolio choices requires deciding who can change portfolios and which products they could potentially introduce. I define players as the producers with a volume share of more than 3 percent in any year. They can adjust product portfolios, whereas I keep product portfolios fixed for smaller competitors.³²

The set of potential products of each player consist of its potential products in 2005 (pre-merger) and 2007 (post-merger).³³ Since I observe the acquisition scenario with the

³²This includes Electrolux, General Electric, LG, Maytag, Samsung, and Whirlpool.

³³Whenever different versions of the same product exist for 2005 and 2007, I choose the 2007 version.

highest increase in market power, these observed sets of potential products should be a good approximation of the actual set of potential products. The higher the increase in market power, the lower the intensity of competition becomes and so the higher the incentives for rivals to add new products. Thus, any product that was not added by rivals after the merger is also unlikely to be added without the merger.

Finally, I fix the products of players at retailers not part of the five major retailers. This results in 135 potential products and 69 exogenously active products (products of non-players and products of players at smaller retailers).

7.2 Portfolio choice algorithm

A well-known feature of product entry games is that there can be many potential equilibrium product portfolios. One way of identifying the set of potential equilibria is to estimate the expected variable profits for all possible product entry combinations and then check whether there are any combinations of product entry costs contained in the fixed cost confidence sets that make these product portfolios a SPNE of the entry game.³⁴ In this case, this is computationally infeasible at this time, since there are 2^{135} candidate equilibria. Instead, I leverage specificities of the case at hand to construct a heuristic portfolio choice algorithm. This algorithm is most closely related to the heuristic algorithm by Fan and Yang (2020).

First, I recognize that although firms incur the fixed cost of adding a potential product to their active portfolio every year, they do not start in a vacuum. More specifically, if there are multiple equilibria of the post-merger entry game it seems plausible to assume that equilibria closer in product space to the pre-merger product portfolios are more likely to be realized. Thus, I initialize the portfolio choice algorithm at the pre-merger equilibrium.

In a nutshell, I allow players to iteratively optimize their portfolios by checking whether they have a profitable one-step deviation from their prior portfolio (i.e., removing an active product or adding an inactive product). If so, the player will make that adjustment. If not, I move on to the next player. This is done until I reach a rest point and no player has a profitable one-step deviation left. In practice, I can considerably reduce the computational burden by optimizing product portfolios brand-by-brand instead of firm-by-firm. Since I only set-identify fixed costs, I repeat this procedure for 50 different fixed cost draws and report 95 percent confidence sets for the welfare effects across draws.³⁵

³⁴This is the approach taken by Eizenberg (2014) in a setting where there are four brands and four product types. After adding some additional restrictions, he ends up with $2^9 = 512$ candidate equilibria.

³⁵Further technical details and a discussion of these details can be found in Appendix IV.A.

Table 7: Number of products offered by each firm in different acquisition scenarios

<i>Acquirer:</i>	Endogenous portfolio adjustments		2005	2007
	None / Haier	Whirlpool		
Maytag	22.4 [20.0, 24.9]	19.5 [17.1, 21.8]	21	23
Whirlpool	25.8 [22.4, 29.1]	23.4 [19.9, 27.0]	27	27
LG + Samsung	9.2 [7.1, 11.4]	9.7 [7.6, 11.8]	5	15
Electrolux + GE	25.6 [20.7, 30.5]	27.4 [23.5, 31.3]	34	38
Total industry	102.1 [96.1, 108.0]	99.0 [93.3, 104.6]	106	128

Notes: The first two columns include the 95% confidence sets on the number of products carried by each brand owner depending on who acquires Maytag. The final two columns show the observed number of products in 2005 and 2007. Confidence sets are based on 50 fixed cost draws and 500 demand and supply residual draws. Maytag includes all products marketed under the brands owned by Maytag pre-acquisition.

7.3 Product portfolio choices

Table 7 summarizes the number of products that firms choose to offer under different acquisition scenarios. By comparing the endogenous portfolio adjustments in different scenarios to the observed number of products in 2005 and 2007, I can disentangle merger-induced and merger-independent portfolio adjustments around the time of the merger.

By comparing the first two columns, I find that the Whirlpool acquisition induced a removal of products by the merging parties, a small expansion in product portfolios by rivals, and a reduction in the number of products on offer overall. Next, let us compare the merger-induced portfolio changes to the observed portfolio changes from 2005 to 2007. For the merging parties and their rivals, I observe more products in 2007 as compared to 2005 than predicted by the simulations. This suggests that general market trends independent of the merger led to an expansion in product portfolios. Reassuringly, rivals expand their portfolios more than the merging parties in the simulations and the observed outcomes.

7.4 Effects on consumers, workers, and firms

Table 8 summarizes the effect of different acquirers for Maytag on consumers, workers, and firms. All results are differences in the outcome between an acquisition by Whirlpool versus an acquisition by Haier. Since Haier has close to no presence in the U.S. clothes

Table 8: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier

	No adjustments	Independent adjustments	Endogenous adjustments
<i>Prices and consumer welfare</i>			
Average price	2.5% [1.2%, 3.9%]	2.8% [1.4%, 4.1%]	3.0% [1.4%, 4.7%]
Consumer welfare	\$-129M [\$-225M, \$-34M]	\$-153M [\$-282M, \$-24M]	\$-246M [\$-298M, \$-194M]
<i>Overall industry</i>			
Variable profits	\$65M [\$11M, \$119M]	\$80M [\$-22M, \$181M]	\$54M [\$16M, \$92M]
Total profits	\$65M [\$11M, \$119M]	\$80M [\$-22M, \$181M]	\$108M [\$83M, \$134M]
<i>Maytag + Whirlpool</i>			
Variable profits	\$15M [\$-39M, \$70M]	\$18M [\$-84M, \$120M]	\$-32M [\$-89M, \$25M]
Total profits	\$15M [\$-39M, \$70M]	\$18M [\$-84M, \$120M]	\$36M [\$26M, \$46M]
<i>Employment</i>			
Domestic jobs maintained	1234 [240, 2228]	1121 [421, 1820]	1258 [1101, 1414]
<i>Offsetting job values</i>			
Consumer welfare only	\$105k [\$99k, \$111k]	\$136k [\$104k, \$169k]	\$221k [\$145k, \$297k]
Consumer welfare + total industry profits	\$52k [\$47k, \$57k]	\$65k [\$60k, \$70k]	\$126k [\$71k, \$180k]

Notes: The first column compares the effect of Maytag acquisitions by Whirlpool vs. Haier without product portfolio adjustments. The second column shows the same comparison for merger-independent portfolio adjustments and the third column for endogenous portfolio adjustments. 95% confidence intervals for the first two columns are computed using 100 residual bootstrap draws. Confidence sets for the third column are based on 50 fixed cost draws for each potential product and 500 residual draws.

washer market pre-merger, without efficiencies an acquisition by Haier is equivalent to no acquisition. All outcomes unrelated to employment can thus also be interpreted as the effect of a Whirlpool acquisition as compared to no acquisition.

The previous results show that the merger induced portfolio changes and that not all observed portfolio changes are merger-induced. I nevertheless estimate the welfare effects under the assumptions of no portfolio adjustments and fully merger-independent portfolio adjustments, to illustrate the importance of endogenizing the product portfolio choice. For this, I consider the 2005 portfolios for the former scenario and the 2007 portfolios for the latter scenario.

Without any portfolio adjustments, prices after a Whirlpool acquisition increase by 2.5 percent and consumer welfare decreases by \$129 million annually. Total industry profits, as well as the profits of the merging parties, increase, however the increase in profits cannot offset the loss in consumer welfare.

If all portfolio adjustments are fully independent of the merger, the predicted price increases and consumer welfare losses are similar in magnitude to the scenario without product portfolio adjustments. Although there is an expansion of the overall product portfolio between 2005 and 2007, particularly by rivals, these are not sufficiently close substitutes to constrain the merging parties and prevent consumer harm.

With fully endogenous portfolio adjustments, the price increases of an acquisition by Whirlpool are modestly higher price than with fixed product portfolios. However, the decrease in consumer welfare is considerably higher than without endogenous portfolio adjustments. This is because I predict only a modest merger-induced expansion in product portfolios by rivals and a larger merger-induced reduction in the portfolio by Maytag and Whirlpool.

Total profits at the industry-level, as well as for the merging parties, increase most if firms can fully adjust their product portfolios after the merger. The large differences in total profits between scenarios with and without adjustments show that there are strong incentives for firms to re-optimize their product portfolios. Variable profits do not necessarily increase, as firms are willing to remove products and forego variable profits if they can save fixed costs of carrying the particular product variety.


Overall, a Whirlpool acquisition is always worse for consumers than an acquisition by Haier. This is because it leads to a larger increase in market power, only modestly more entry by rivals, and significantly fewer products by the merging parties. This does not change if we consider the sum of consumer welfare and total profits as the relevant outcome.

It is therefore worth taking a closer look at the domestic employment effects of the two potential acquisitions. Differences in the number of domestic jobs can arise from two sources: A relocation effect (due to different plant relocations between acquirers) and a reallocation effect (due to a reallocation of sales between sellers producing in different locations).

Across all scenarios, an acquisition of Maytag by Whirlpool maintains significantly more jobs in the U.S. than an acquisition by Haier. Figure 4 decomposes the employment effects into the relocation and the reallocation effects. With endogenous portfolio adjustments, most of the differences in U.S. employment come from a difference in Maytag jobs offshored by the acquirer after a potential acquisition. That is, although Whirlpool offshored some Maytag jobs after its acquisition, Haier would have offshored even more jobs. The reallocation effect partially offsets this effect. This is because after an acquisition of Maytag by Whirlpool, the merging parties increase their prices and more sales are diverted to rivals that produce abroad. Thus, the relocation and the reallocation effects go in opposite directions, but the relocation effect is quantitatively more important in this particular case.

Finally, I combine the product market and employment effects and estimate the average

Figure 4: Decomposition of employment effects



Notes: The histograms decompose the employment effects into a relocation and a reallocation effect. The former is related to differences in U.S. employment between Maytag acquisitions by Whirlpool and Haier because of different plant relocation plans. The latter is related to differences in U.S. employment between the acquisitions due to a reallocation of sales between sellers that produce in different locations, as well as the outside good.

job value necessary such that the domestic jobs saved by a Whirlpool acquisition offset domestic product market losses. First, I consider the offsetting job value necessary if we only consider consumer welfare to be relevant in the product market. With endogenous portfolio adjustments, the average offsetting value of each additional job is approximately \$221,000 per year. Second, I estimate the offsetting job value if we consider the sum of consumer welfare and industry profits to be relevant. In this case, the offsetting value of each additional job is approximately \$126,000 per year.

On their own, these figures are not sufficient to assess whether employment effects could have plausibly offset the product market effects. Ideally, we would like to know what is the value of these clothes washer manufacturing jobs to the local labor markets. The descriptive analysis in Section 3 indicates that these jobs are valuable, as workers that are made redundant do not quickly find equivalent jobs, but the exact value remains unclear. Simply looking at average wages is insufficient too, as these jobs likely have a multiplier effect. Two comparisons are helpful for reference: First, we can look at estimates in the literature of the value of jobs created by foreign multinationals to local labor markets. Second, we can compare the necessary offsetting value to that of trade policies.

Setzler and Tintelnot (2021) study the direct and indirect local labor market effects of a job created by a foreign multinational firm. They find that an additional foreign multinational job increases the total wage bill in a local labor market by \$113,000 per year. This

includes wages for workers coming from non-employment, as well as the direct effect of the foreign multinational wage premium on employees previously employed at domestic firms. It also includes wage gains for employed workers at domestic firms, as well as wages at newly created domestic jobs. Since all of the clothes washer manufacturers are foreign or domestic multinationals, it seems plausible that this is a good proxy for the value of the jobs considered in the appliance industry.³⁶

An alternative is to look at the average job values necessary such that the gains to domestic workers of past sector-specific tariffs imposed by the U.S. outweigh the losses to domestic consumers. Hufbauer and Lowry (2012) estimate that for the 2011 safeguard tariffs on tire imports from China, each job would have had to be worth at least \$900,000 per year to offset the losses to U.S. consumers. Even more relevant, Flaaen, Hortaçsu, and Tintelnot (2020) estimate that for the additional domestic jobs created by the 2018 U.S. global safeguard tariffs on clothes washer imports to offset losses to consumers, each of these jobs would have had to be worth at least 817,000 per year.

There are many other positive effects related to an increase in the availability of jobs, that go beyond an increase in wages. Bearing this in mind, I consider the increase in the total wage bill by \$113,000 per year as a lower-bound estimate of the value of a U.S. appliance manufacturing job to the U.S. economy. This is at the lower end of the necessary job values to offset losses in consumer welfare for clothes washers with endogenous portfolio adjustments. At the same time, the offsetting job value for other appliance categories in which Maytag was active (and where there is less of an overlap with Whirlpool) is most likely lower. The offsetting job value of the overall merger (i.e., also taking into account other product categories than clothes washers) is therefore going to be lower than the estimates for clothes washers, which can be taken as an upper-bound estimate.

Overall, I therefore cannot exclude that the gains to domestic workers of a Whirlpool acquisition as compared to an acquisition by Haier offset the losses to consumers. If we consider both, consumer welfare and total industry profits to be relevant, this further decreases the necessary offsetting job value.

7.5 Unequal distribution of welfare effects

So far, the analysis focused on how consumers and workers are affected overall by the two alternative acquisitions. However, an important dimension are the distributional effects of the policy alternatives.

³⁶I do not need to distinguish between local and national employment effects, since I consider each manufacturing job not created domestically to be created abroad.

For the U.S. economy as a whole, 1,000 additional clothes washer manufacturing jobs do not have any significant effect on employment or wages. However, as I show in Section 3, the closure of a manufacturing plant can decrease wages and employment at the county-level even two years after the plant closure. As illustrated in Figure 1, clothes washer manufacturing plants are concentrated in a few counties in Illinois, Iowa, Michigan, Ohio, and South Carolina. Although most local labor markets are unaffected by the potential acquisitions, some are strongly affected.

Whereas the effects on consumers are distributed across the country, the employment effects of the alternative potential acquisitions are geographically highly concentrated. This has implications for optimal policy. If households have diminishing marginal utility of income and employment effects are not concentrated among the very wealthy, then an acquisition by Whirlpool may be better for the sum of workers and consumers, even if the increase in the total domestic wage bill as compared to a Haier acquisition is lower than the relative consumer welfare loss. Furthermore, other non-wage considerations related to job loss, such as mental or physical health, can improve the domestic welfare effects of a Whirlpool as compared to a Haier acquisition.

Finally, political considerations cannot be neglected completely. Whereas a loss in consumer welfare in the clothes washer market by \$20 is unlikely to affect how voters cast their ballot, direct and indirect employment effects can. Thus, facilitating an acquisition of Maytag by Whirlpool as opposed to Haier can be politically more attractive.

8 Conclusion

In this paper, I propose a framework to quantify the trade-off that policy choices, such as whether to clear a merger, can have for workers and consumers. To consider how mergers change the incentives to launch new products, I allow firms to endogenously adjust their product portfolios. To account for employment effects, I model how the product market equilibrium affects the number of workers required to manufacture the product.

Empirically, I focus on how foreign competition changes the consumer welfare and domestic employment effects of a merger. I find that the Maytag acquisition by Whirlpool led to a substantial decrease in consumer welfare, modestly induced rivals to add new products, and that this merger-induced rival product entry was unable to mitigate consumer harm.

A Whirlpool acquisition of Maytag leads to more consumer harm than an acquisition of Maytag by Haier, but the latter leads to a larger decrease in domestic employment. I therefore calculate the average value of a job necessary for domestic employment effects to offset the consumer harm. A comparison to estimates by Setzler and Tintelnot (2021) on

the local labor market effects of new jobs by multinational firms leads me to conclude that the smaller losses to domestic workers after a Whirlpool acquisition is of similar magnitude than the additional consumer harm. Overall, I cannot exclude that an acquisition of Maytag by Whirlpool leads to higher domestic welfare than an alternative acquisition by Haier.

This has important implications for policy. Since the employment effects of a product market merger can be of first order importance, these should not be ignored in merger analysis. Blocking acquisitions that could lead to offshoring or allowing anti-competitive mergers that could preserve domestic jobs may still not be optimal. Instead, the framework laid out in this paper could be used to identify mergers in which employment effects are of first order importance. Whilst the merger decision could still be taken based on the consumer welfare standard, this would identify cases where there may be a need for complementary labor market policies.

References

- Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P. Vasquez. 2021. “The Effects of Multinationals on Workers: Evidence from Costa Rican Microdata.” *mimeo*.
- Appliance Portrait. 2006. *29th 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.
- Baker, Jonathan B., and Carl Shapiro. 2008a. “Detecting and Reversing the Decline in Horizontal Merger Enforcement.” *Antitrust* 22 (3): 29–36.
- . 2008b. “Reinvigorating Horizontal Merger Enforcement,” edited by Robert Pitofsky. Oxford University Press.
- Barboza, David. 2005. “From China, A New Bid for Maytag and Status.” *New York Times* (June). <https://www.nytimes.com/2005/06/22/business/worldbusiness/from-china-a-new-bid-for-maytag-and-status.html>.
- Bellemare, Marc F., and Casey J. Wichman. 2020. “Elasticities and the Inverse Hyperbolic Sine Transformation.” *Oxford Bulletin of Economics and Statistics* 82 (1): 50–61.
- Benkard, C. Lanier, Ali Yurukoglu, and Anthony Lee Zhang. 2021. “Concentration in Product Markets.” *NBER Working Paper 28745*.
- 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.
- Berry, Steven T., and Joel Waldfogel. 2001. “Do Mergers Increase Product Variety? Evidence from Radio Broadcasting*.” *The Quarterly Journal of Economics* 116 (3): 1009–1025.
- Brunner, Daniel, Florian Heiss, André Romahn, and Constantin Weiser. 2020. “Reliable Estimation of Random Coefficient Logit Demand Models.” *mimeo*.
- Caradonna, Peter, Nathan H. Miller, and Gloria Sheu. 2021. “Mergers, Entry, and Consumer Welfare.” *Georgetown McDonough School of Business Research Paper 3537135*.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics* 36 (1): 13–17.

- Ciliberto, Federico, Charles Murry, and Elie Tamer. 2021. “Market Structure and Competition in Airline Markets.” *Journal of Political Economy* 129 (11): 2995–3038.
- Conlon, Christopher, and Jeff Gortmaker. 2020. “Best practices for differentiated products demand estimation with PyBLP.” *The RAND Journal of Economics* 51 (4): 1108–1161.
- Department of Justice. 2006. *Statement on the Closing of its Investigation of Whirlpool’s Acquisition of Maytag*. Washington, DC: Author. https://www.justice.gov/archive/atr/public/press_releases/2006/215326.htm.
- Draganska, Michaela, Michael J. Mazzeo, and Katja Seim. 2009. “Beyond plain vanilla: Modeling joint product assortment and pricing decisions.” *Quantitative Marketing and Economics* 7 (2): 105–146.
- Eizenberg, Alon. 2014. “Upstream Innovation and Product Variety in the U.S. Home PC Market.” *The Review of Economic Studies* 81 (3): 1003–1045.
- Electrolux. 2007. *Annual Report 2006*. <https://www.electroluxgroup.com/en/annual-reports-27641/?decade=2000-2009>.
- Fan, Ying. 2013. “Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market.” *The American Economic Review* 103 (5): 1598–1628.
- Fan, Ying, and Chenyu Yang. 2020. “Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market.” *American Economic Journal: Microeconomics* 12 (2): 99–134.
- . 2021. “Merger, Product Variety and Firm Entry: the Retail Craft Beer Market in California.” *mimeo*.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. “The Next Generation of the Penn World Table.” *American Economic Review* 105 (10): 3150–82.
- 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.
- Garrido, Francisco. 2020. “Mergers Between Multi-Product Firms With Endogenous Variety: Theory and an Application to the RTE-Cereal Industry.” *mimeo*.
- Goldberg, Pinelopi Koujianou, and Frank Verboven. 2001. “The Evolution of Price Dispersion in the European Car Market.” *The Review of Economic Studies* 68 (4): 811–848.
- Goodman, Peter, and Ben White. 2005. “Haier Withdraws Maytag Bid.” *Washington Post* (July). <https://www.washingtonpost.com/archive/business/2005/07/20/haier-withdraws-maytag-bid/f05586dd-097d-41b8-8cee-99c3bd17a72b/>.
- Grieco, Paul L.E., Charles Murry, and Ali Yurukoglu. 2021. “The Evolution of Market Power in the Us Auto Industry.” *NBER Working Paper 29013*.

- Houde, Sébastien. 2018. “How consumers respond to product certification and the value of energy information.” *The RAND Journal of Economics* 49 (2): 453–477.
- Hufbauer, Gary, and Sean Lowry. 2012. “US Tire Tariffs: Saving Few Jobs at High Cost.” *Peterson Institute for International Economics*.
- Igami, Mitsuru. 2018. “Industry Dynamics of Offshoring: The Case of Hard Disk Drives.” *American Economic Journal: Microeconomics* 10 (1): 67–101.
- Imbens, Guido, and Charles Manski. 2004. “Confidence Intervals for Partially Identified Parameters.” *Econometrica* 72 (6): 1845–1857.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. “Earnings Losses of Displaced Workers.” *The American Economic Review* 83 (4): 685–709.
- Jaravel, Xavier, and Erick Sager. 2020. “What are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models.” *mimeo*.
- Jeziorski, Przemislaw. 2015. “Empirical Model of Dynamic Merger Enforcement—Choosing Ownership Caps in US Radio.” *mimeo*.
- Kroger. 2022. “Kroger and Albertsons Companies Announce Definitive Merger Agreement.” *Company Release* (October). <https://ir.kroger.com/CorporateProfile/press-releases/press-release/2022/Kroger-and-Albertsons-Companies-Announce-Definitive-Merger-Agreement/default.aspx>.
- Lee, Robin S., and Ariel Pakes. 2009. “Multiple Equilibria and Selection by Learning in an Applied Setting.” *Economics Letters* 104 (1): 13–16.
- Legere, John. 2019. “Just the Facts on Jobs: The New T-Mobile Will Create Jobs From Day One.” *T-Mobile Press Release* (April). <https://www.t-mobile.com/news/un-carrier/new-t-mobile-creating-jobs>.
- Li, Sophia, Joe Mazur, Yongjoon Park, James Roberts, Andrew Sweeting, and Jun Zhang. Forthcoming. “Repositioning and Market Power After Airline Mergers.” *The RAND Journal of Economics*.
- Marinescu, Ioana, and Herbert Hovenkamp. 2019. “Anticompetitive Mergers in Labor Markets.” *Indiana Law Journal* 94 (3): 1031.
- Maytag. 2005. *2004 Annual Report*. Newton, IA: Author.
- Mazzeo, Michael J. 2002. “Product Choice and Oligopoly Market Structure.” *The RAND Journal of Economics* 33 (2): 221–242.
- 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.
- Nevo, Aviv. 2001. “Measuring Market Power in the Ready-to-Eat Cereal Industry.” *Econometrica* 69 (2): 307–342.
- Noble, Breana. 2021. “Stellantis CEO says PSA-FCA merger is a ‘shield’ for auto jobs, but doesn’t rule out cuts.” *The Detroit News* (January). <https://www.detroitnews.com/story/business/autos/chrysler/2021/01/19/stellantis-ceo-carlos-tavares-says-psa-fca-merger-shield-auto-jobs/4203304001/>.
- OECD. 2016. “Public Interest Considerations in Merger Control.” *DAF/COMP/WP3/WD(2016)30*.
- Owen, Art B. 2017. *A randomized Halton algorithm in R*.
- Pakes, A., J. Porter, K. Ho, and J. Ishii. 2015. “Moment Inequalities and Their Application.” *Econometrica* 83 (1): 315–334.
- Prager, Elena, and Matt Schmitt. 2021. “Employer Consolidation and Wages: Evidence from Hospitals.” *American Economic Review* 111 (2): 397–427.
- 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.
- Seim, Katja. 2006. “An Empirical Model of Firm Entry with Endogenous Product-Type Choices.” *The RAND Journal of Economics* 37 (3): 619–640.
- Setzler, Bradley, and Felix Tintelnot. 2021. “The Effects of Foreign Multinationals on Workers and Firms in the United States.” *The Quarterly Journal of Economics* 136 (3): 1943–1991.
- Shapiro, Carl. 2019. “Protecting Competition in the American Economy: Merger Control, Tech Titans, Labor Markets.” *Journal of Economic Perspectives* 33 (3): 69–93.
- Sweeting, Andrew. 2010. “The Effects of Mergers on Product Positioning: Evidence From the Music Radio Industry.” *RAND Journal of Economics* 41 (2): 372–397.
- Werden, Gregory J, and Luke M Froeb. 1998. “The Entry-Inducing Effects of Horizontal Mergers: An Exploratory Analysis.” *The Journal of Industrial Economics* 46 (4): 525–543.
- Wollmann, Thomas G. 2018. “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles.” *American Economic Review* 108 (6): 1364–1406.

Appendix

I Appendix to Section 2: Details on data set construction

I.A Product market data set

In this section, we add further details on the construction of the product market data set.

Product data. As described in Section 2, for clothes washers, a product is defined as the combination of a brand, a retailer and whether the clothes washer is a front-loader, a regular top-loader (with an agitator) or a high-efficiency top-loader (without an agitator). For clothes washers, these are the key differentiating characteristics between products.

Figure A.1 illustrates the difference between a front-loader and a top-loader. Whether the former can be loaded from the front, the latter is loaded from the top. The former can therefore be stacked (i.e. a front-loading dryer can be placed on top of a front-loading washing machine), is more water and energy efficient, cleans better, and is usually more expensive than top-loaders. The latter can never be stacked, however, for top-loaders, there is an important distinction related to whether they have an agitator, which is illustrated in Figure A.2. top-loaders without an agitator are also called high-efficiency top-loaders. In all respects but stacking, they are in between regular top-loaders and front-loaders.³⁷


Within a market (here, national at the yearly-level), I group responses that are the same along these three dimensions.³⁸ Doing so, I end up with 2,939 products between 2005 and 2015. Using this product definition, many products are often very small and based only on a single responding household. Some responses also do not contain information on the brand. I therefore drop all products whose brand response is “Other Brands” or “Store Brand/Generic”, as well as all products with a volume share of the clothes washer market of less than 0.01 percent. This results in a final product data set with 1,590 products. Throughout the years, the remaining products account for between 97.3 and 99.0 percent of the volume share of all clothes washer sales in the *TraQline* data. Dropping very rare products should therefore not bias the estimation results.

For other characteristics, which are only available for a random subset of *TraQline* respondents, I calculate the within-group average of responses for that characteristic. These


³⁷See, for example, McCabe (2016) for a detailed comparison of the different clothes washer types.

³⁸For 2006, I classify Maytag products as belonging to Whirlpool also for the first quarter, where the acquisition was not yet carried out. This is to avoid artificially inflating the number of clothes washer products in that year. Also, since merger talks were public since mid-2005, it seems unlikely that Maytag and Whirlpool would still compete heavily in the first quarter of 2006.

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




Top load washer




Front load washer

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



without agitator



with agitator

include whether a clothes washer is part of a stacked pair, whether its exterior is made of stainless steel, is white, or of a different color, whether it is Energy Star certified, has additional noise insulation or a child lockout, as well as the number of special programs it has.

Household income. Whereas the CPS data includes the exact income of the sampled households, the *TraQline* data only includes an income range for each household. To estimate how the price sensitivity of households depends on household income using a single parameter only, I need an exact income for each household. For this, I randomly draw a household income for each respondent based on the empirical distribution of household incomes and the income range that the household falls into.

This involves the following steps:

1. Compute the mid-point of the non-overlapping household income buckets for each response.
2. For each year, fit a log-normal income distribution to the observed household-level income range mid-points.
3. Draw 1,000,000 incomes from the fitted log-normal income distribution.
4. Allocate each income draw to a particular income bucket.
5. For each household, sample with replacement an income from the set of incomes that correspond to its income bucket.

I.B Plant locations and plant location weights

Plant locations. Constructing the data set on plants manufacturing clothes washers for the U.S. market involves three steps: First, I use information from various sources, such as annual reports, news articles or the United States International Trade Commission's (USITC) anti-dumping hearing transcripts into imports of large residential clothes washers from Mexico and South Korea to identify the location of clothes washer plants by the major manufacturers.

For LG and Samsung, the production locations before 2012 are mostly based on the investigation by the USITC . For 2012 until 2015, production locations for LG and Samsung are based on firm-level clothes washer imports based on the PIERS data set, which uses bill of landing documents and is reported in Flaaen, Hortaçsu, and Tintelnot (2020).

For Electrolux, Maytag and Whirlpool, the bulk of the information on manufacturing plant locations is based on information in their annual reports. Since General Electric is

not primarily an appliance manufacturer, its annual report does not contain information on appliance plant locations. I therefore base plant locations on a combination of documents from the USITC investigation and news reports. Finally, to make sure that plants produce clothes washers for the U.S. market, I check plant locations against import data split by top-loading and front-loading clothes washer at the country-level from the USITC.

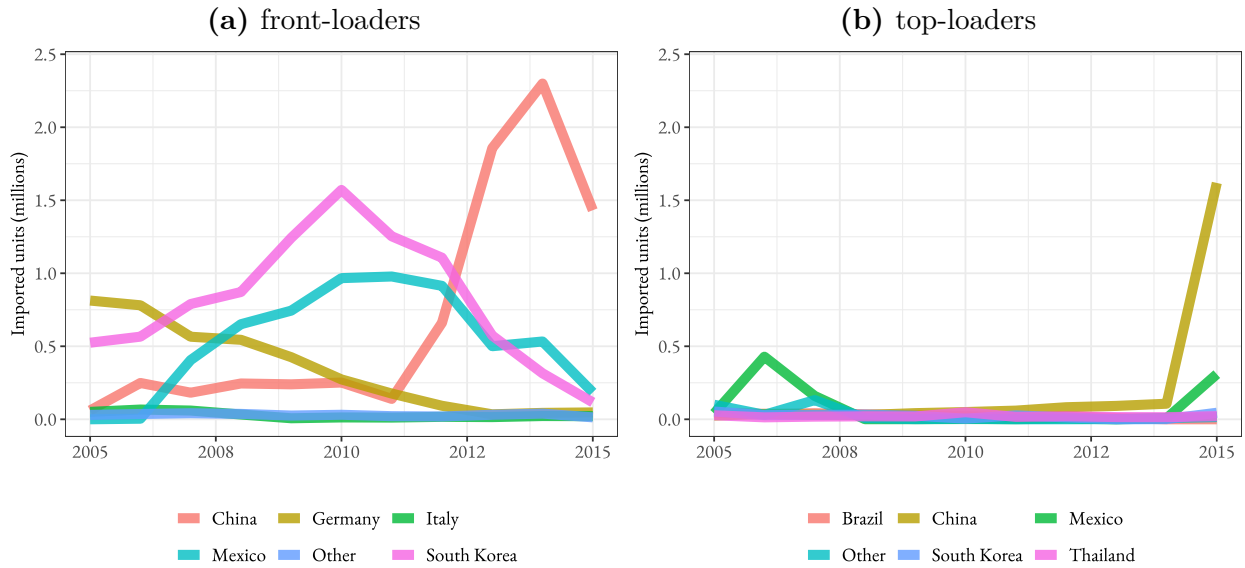
In many cases, this is insufficient to know whether a plant produces clothes washers for the U.S. or for another market. Second, I use information on the general imports of front-loader and top-loader clothes washers to the U.S. split by source country over time. I use this data to eliminate any plants that cannot plausibly produce substantive volumes for the U.S. market. Finally, I use this data to verify that there are production plants that can plausibly be responsible for the imported volumes for each country from which the U.S. imports substantial numbers of clothes washers.

Figure A.3 shows the evolution of annual imports of front-loaders and top-loaders into the U.S., split by source country. Across the sample period more than half of the front-loaders sold in the U.S. are imported. In 2005, Germany is the largest exporter of front-loaders into the United States. These are not produced by a German manufacturer, but by Whirlpool in its plant in Schorndorf, which was closed in 2012. Until 2012, LG and Samsung imported many of its front-loaders from South Korea and, like other manufacturers such as General Electric or Whirlpool, also from Mexico. After the imposition of anti-dumping duties on large residential clothes washers from Mexico and South Korea in 2012, imports from both countries declined and LG and Samsung moved their production to China (see Flaaen, Hortaçsu, and Tintelnot, 2020 for an in-depth discussion). In contrast, no country exported more than 50,000 top-loaders to the U.S. until 2011, aside from a temporary spike in top-loader imports from Mexico in 2006 and 2007. Thereafter, LG and Samsung begin increasing their sales of top-loaders in the U.S. and import most of these from China.

For reference, according to Appliance Portrait (2006), 9.3 million clothes washers were sold across the U.S. in 2005. Of those, according to the *TraQline* data, around one-third are front-loaders and the rest top-loaders. The share of front-loaders gradually increased to over 40 percent in 2010 and then decreased again to around 25 percent in 2015. This suggests that although substantial amounts of front-loaders were imported into the U.S. throughout the sample period, most top-loaders were produced domestically.

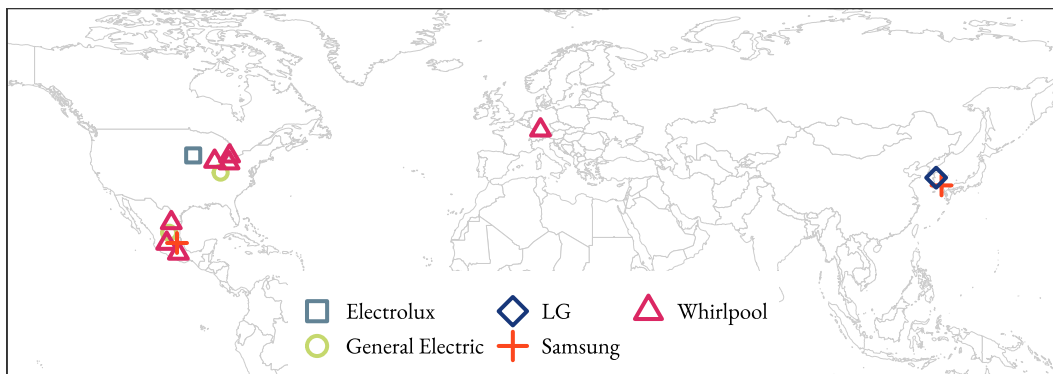
By combining the clothes washer plant locations of major manufacturers with the USITC import data, I can identify which plants manufacture clothes washers for the U.S. market. Figures A.4, A.5, and A.6 show the locations of clothes washer plants for all manufacturers that have a volume share of more than 3 percent of the U.S. clothes washer market in any year in the sample.

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



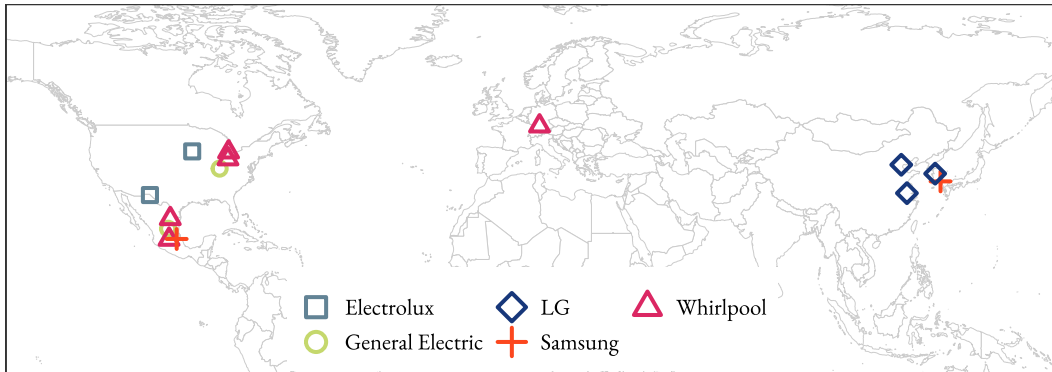
Notes: The left panel plots the annual general imports in terms of volume of front-loader washing machines (HS8450110080, HS8450200080, HS8450200090) imported into the U.S. by source country. The right panel plots the annual general imports in terms of volume of top-loader washing machines (HS8450110040, HS8450200040) imported into the U.S. by source country. The graphs include the top six importing countries for each product class and groups all other importing countries into “Other”. The data comes from the United States International Trade Commission.

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




Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2007 by manufacturers with a market share of more than 3 percent in any year in the sample.

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



Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2009 by manufacturers with a market share of more than 3 percent in any year in the sample.

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



Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2011 by manufacturers with a market share of more than 3 percent in any year in the sample.

Plant location weights. Finally, Table A.1 summarizes the plant location weights used to calculate the average real exchange rate for each product. Based on the plant locations, the aggregate USITC import data shown above, and the firm-level clothes washer imports for 2012 until 2015 based on PIERS bill of landing data and reported in Flaaen, Hortaçsu, and Tintelnot (2020), these are best estimates of which share of a product is sourced from which country in a particular year.

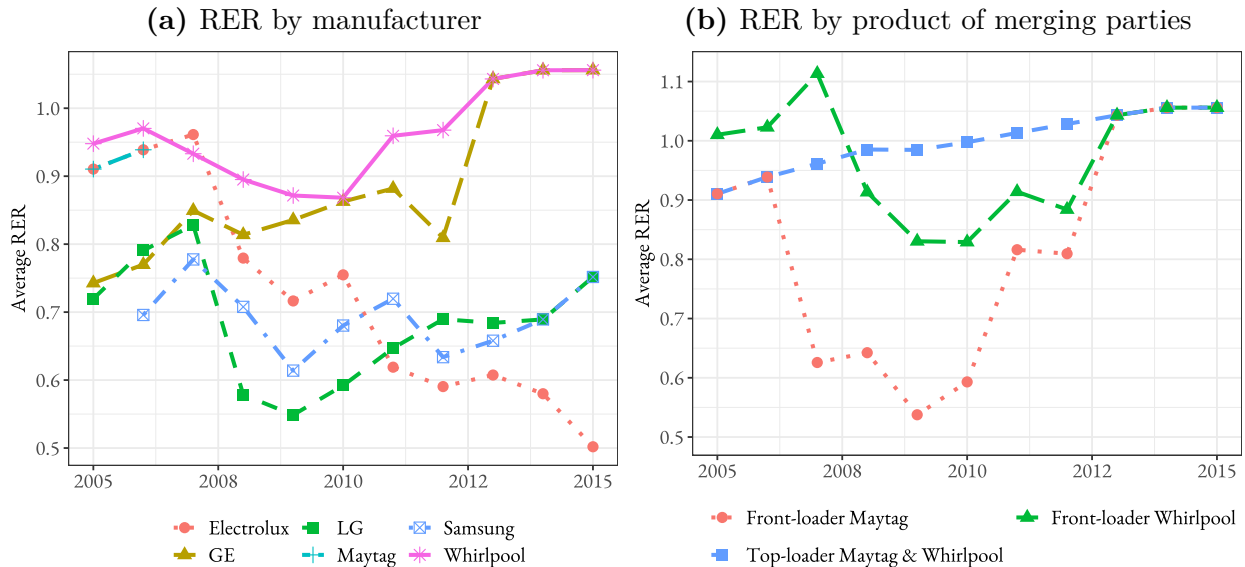
Table A.1: Plant location weights

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	All brands	Front Loader	2005		1			
Whirlpool	All other WP brands	Front Loader	2006-2007		1			
Whirlpool	All other WP brands	Front Loader	2008-2010		0.5	0.5		
Whirlpool	All brands	Front Loader	2011		0.33	0.33		0.33
Whirlpool	All brands	Front Loader	2012-2015					1
Whirlpool	Admiral, Amana, Maytag	Front Loader	2006-2010			1		
Whirlpool	Admiral, Amana, Maytag	Front Loader	2010			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				1	
Samsung	All brands	Front Loader	2012	0.5			0.5	
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

I.C Details on the instrumental variable for price

Figure A.7 shows the evolution of the average RER over time and illustrates the source of the variation. The left panel plots the average RER of all production locations for a particular manufacturer. The average RER is based on the country-level RER of different plant locations of a manufacturer for a product in a particular year, weights that capture which share of a product is produced by a particular plant, and weights based on the sales volume of different products sold by a manufacturer. Although this masks within-manufacturer variation in the RER, already at this level there is significant variation. In the right panel, I disentangle the average RER for Whirlpool and Maytag products.³⁹ This shows that there is additional variation in the RER below the manufacturer level, because the same manufacturer produces different products in different countries. For example, whereas all Maytag and Whirlpool top-loaders are produced in the U.S., over the sample period Maytag front-loaders were produced in the U.S. and Mexico and Whirlpool front-loaders in the U.S., Mexico and Germany.

Figure A.7: Average real exchange rate over time



Notes: The left panel plots the average real exchange rate of all production locations by manufacturer over time. It includes the RER for all manufacturers with a market share of at least 3 percent in any year in the sample. The right panel plots the average RER of all production locations by product of the merging parties. The average RER is based on the plant locations in a particular year, the plant weights and the country-level RER. In the right panel, Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, MagicChef and Maytag) and Whirlpool includes all other brands owned by Whirlpool.

³⁹Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, MagicChef and Maytag) and Whirlpool includes all other brands owned by Whirlpool.

The large variation in the RER over time is consistent with anecdotal evidence about the importance of the local cost of production for appliance manufacturers. One of the principal reasons why Maytag was struggling financially pre-merger was that its production costs were too high, in parts due to its lack of international production.⁴⁰ In a similar spirit, Electrolux launched its global cost-cutting program in 2004, with the aim to offshore more than half of its production to low-cost countries by 2009 (Electrolux, 2007).⁴¹ Both firms exclusively served the U.S. clothes washer market from the U.S. until 2007. This highlights the importance of production locations for costs and competitiveness in the appliance industry and also describes the source of variation in the cost measure: Changes in the RER between the U.S. and a particular production location over time, as well as changes in the production locations.

II Appendix to Section 5: Details on the estimation procedures

II.A Details on estimating product characteristics for potential products

Potential products are all products that brand owners added to the market (active products), as well as all products that they could have added but did not (inactive products). Estimating the former is easy, since we can simply observe these in the market. Estimating the latter is more complicated.

The focus of the analysis in this paper lies on the decision of firms to add or remove products that they are technologically already capable of making. For example, if a firm does not carry front-loading washing machines, these will also not be part of its potential products. If, for example, Maytag sells regular top-loading washing machines under its Amana brand at Best Buy and Lowe's, but not at other major retailers, Amana regular top-loaders at other major retailers are potential products.⁴²

Product characteristics can mildly vary between retailers. That is, Amana top-loaders sold at Best Buy might modestly differ in their characteristics compared to Amana top-

⁴⁰This was highlighted throughout Maytag's 2004 annual report, as for example in the following: "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).

⁴¹By the end of the sample period, Electrolux had lost most of its share of the U.S. laundry market and served its remaining customers from low-cost countries.

⁴²Major retailers are Best Buy, H. H. Gregg, Home Depot, Lowe's and Sears.

loaders sold at Lowe’s. In the example, Amana regular top-loaders at Sears are an inactive product. To determine the exact product characteristics of this inactive product, I need to decide whether to attribute it the characteristics of the Amana regular top-loader sold at Best Buy or at Lowe’s.

Whenever a particular combination of brand and key characteristic exists at two or more retailers, I use the following ordering of “closest” retailers to match other product characteristics:

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

II.B Details on the demand estimation

The estimation of the demand parameters follows Berry, Levinsohn, and Pakes (2004) and proceeds in two steps. First, I search for estimates $\hat{\kappa}_\alpha$ and $\hat{\sigma}_{FL}$ (jointly denoted by $\hat{\theta}_2$ of the non-linear parameters, as well as of the vector of mean utilities δ . Next, I estimate $\hat{\beta}$ for the vector of linear demand parameters. Wherever possible, I implement the best practices described by Conlon and Gortmaker (2020). For notational simplicity, I omit the time subscript t in this section. The details of the technical implementation should thus be seen as valid for a single market t and then repeated and averaged over markets.

The estimation of the non-linear 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 non-linear parameters. In the outer loop, I search for the non-linear parameters that minimize the objective function, solving the inner loop at each step.

The first set of moments equates the observed market shares in the data with the simulated market shares from the demand model. To get an estimate $\hat{\delta}$ of the mean utilities, I proceed as follows: First, as described by Berry (1994), I 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 .⁴³ Second, I use the fixed-point formulation due to Berry, Levinsohn,

⁴³Note, that $s_j(\delta_j; \theta)$ also depends on the product and household characteristics, which I omitted to simplify notation.

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 this is not guaranteed to converge, whenever convergence 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 estimate the market shares implied by the estimate $\hat{\theta}$ of the parameter vector, the model and the data, I need to solve the integral in Equation 4. As is standard in the literature, I approximate this integral using Monte Carlo simulations by drawing household demographics and unobserved taste shocks from the joint empirical distribution for 1000 households. Household demographics come from the CPS. I draw unobserved taste shocks from a standard normal distribution, using scrambled Halton draws (see Owen, 2017).

The second set of moments fits the covariance between the price of the first-choice clothes washer and the average income of households purchasing the product. I compute the moment as follows

$$\sum_j \frac{n_j}{n} p_j \left\{ \left(\frac{1}{n_j} \sum_{i.s.t. y_i^1=j} z_i \right) - E[z|y^1 = j, \theta] \right\}, \quad (12)$$

where J continues to denote a product, n denotes the total number of households, n_j denotes the number of households buying good j , y_i^1 denotes the first choice product of household i , p_j continues to denote the price of product j , and z_i the income of household i .

Figure A.8 shows a scatter plot of the relationship between the household income and price.


The third set of moments fits the covariance between whether the first-choice clothes washer is a front-loader and the share of front-loaders among products of the second choice brand. In contrast to Berry, Levinsohn, and Pakes (2004), I do not observe the exact second-choice product but only the second-choice brand. In particular, I use the following moment condition

$$\sum_j \left(\frac{n_j}{n} x_j^{FL} \sum_{b' \neq b_j} x_{b'}^{FL} \left\{ \frac{n_{jb'}}{n_j} - E[\mathbf{1}(b^2 = b'|y^1 = j, \theta)] \right\} \right), \quad (13)$$

where b denotes a brand, b_j denotes the brand of product j , b^2 denotes the brand of the second choice, x_j^{FL} indicates whether product j is a front-loader and $x_{b'}^{FL}$ denotes the volume-weighted share of front-loaders among products sold of brand b .

The objective function that I minimize in the outer loop to estimate $\hat{\theta}_2$ consists of the moments in Equations 12 and 13. Since there are two nonlinear parameters and two moment conditions, the parameters are just-identified and we estimate $\hat{\theta}_2$ using the method

Figure A.8: Correlation of average purchaser household income and price by product



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

of simulated moments. I therefore estimate

$$\hat{\theta}_{2,MSM} = \operatorname{argmin} \hat{m}(\theta_2)' \hat{m}(\theta_2) . \quad (14)$$

Solving the minimization problem above does not only allow recovering the nonlinear parameters of the demand model, but also the mean utilities $\hat{\delta}$. In the final step, I estimate the linear parameters of the demand model using the following specification:

$$\hat{\delta}_j = x_j \beta - \alpha p_j + \xi_j . \quad (15)$$

As explained in Section 5, I assume that the non-price product characteristics are independent of unobserved quality differences ξ_j , whereas the price can be correlated with these unobserved differences. To solve the endogeneity problem, I use an instrumental variables estimator, where the product-level real exchange rate serves as a cost shifting instrumental variable for price, as described in Section 2.

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

To compute the total market size, I assume that every seventh household is a potential purchaser of a clothes washer in a particular year. According to Consumer Reports, in 2009 the average life expectancy of a clothes washer was ten years. Many households will consider

buying a clothes washer already before the end of the life expectancy of their washer, e.g. to get a new washer with novel features. Some households will consider new washing machines for multiple years. Households that recently purchased a washer are unlikely to be on the lookout for a new one immediately. It therefore seems plausible that the true market size is somewhere between a fifth and a tenth of the number of households. The results are robust to alternative market size assumptions.

To compute firm profits, consumer welfare and estimate entry cost bounds in Dollar terms for the U.S. population, I need to scale the estimates by the number of households that are in the market for clothes washers in a particular year. There are two alternative estimation methods: We can take the total number of U.S. households in a particular year and assume that the market size is one seventh of these households. Alternatively, I can use estimates of the annual total clothes washers shipped as reported by Appliance Portrait and divide this by the share of the inside good. Both methods yield similar results for the years around the merger date and so I assume that the total market size in the U.S. is around 15 million households.

II.C Speeding up the computation of expected profits

Both, the estimation of fixed costs, and the heuristic entry algorithm require computing the expected profits of firms for many different product portfolios. This is computationally costly and since it has to be repeated many times, speeding up this process is crucial. In the following, I briefly describe the key elements that helped speed up the computations for this paper.⁴⁴

Computing equilibrium prices. Each draw of the second-stage marginal cost and demand shocks e_{jt} requires re-estimating the equilibrium price vector for all active products. Since I use 500 draws of e_{jt} to approximate the expected variable profits for a single product portfolio, it is also necessary to re-compute equilibrium prices 500 times for each product portfolio. Speeding up this process is therefore crucial. Furthermore, not all methods to re-compute equilibrium prices necessarily converge.

Morrow and Skerlos (2011) compare different numerical methods to re-compute equilibrium prices using the Nash-Bertrand first order conditions. They find that applying Newton methods to this problem is reliable but slow. On the other hand, they show that fixed point iteration on the BLP-markup equation need not converge and is slow. Instead, they propose a reformulated markup equation, the ζ -markup, which is fast and reliable. I therefore

⁴⁴As noted in the Online Appendix to Wollmann (2018), implementing the computations in Julia has significant speed advantages, as it can handle loop commands at comparable speed to “vectorized” code in Matlab.

compute equilibrium prices by using fixed point iteration on the ζ -markup equation.

Drawing e_{jt} . The heuristic algorithm to choose product portfolios requires comparing the expected profits of the current product portfolio to the expected profits of any product portfolio that is within a one-step change of the current product portfolio. This involves revisiting the same product portfolios many times.

An important feature of the heuristic portfolio choice algorithm is to use the same e_{jt} draws for the same product when computing the expected profits of different product portfolios. In terms of economics, this is desirable because there is no good reason for why a firm should form its expectation about demand and cost shocks for a product differently based on what other products are in the market. In terms of computations, this is desirable because it means that I only need to compute expected profits of all firms for a given set of product portfolios once. Every time that the algorithm re-visits the particular set of product portfolios, I can re-use the memorized expected profits and do not need to re-compute equilibrium prices and expected profits.

II.D Details on the fixed cost estimation

I follow the approach proposed by Eizenberg (2014) and fill the missing bounds by adding two further assumptions.

Assumption 2. $\sup_{j \in \mathcal{J}_{bt}} F_{jt} = F_{bt}^U < \infty$ and $\inf_{j \in \mathcal{J}_{bt}} F_{jt} = F_{bt}^L > -\infty$ (bounded support)

Assumption 2 states that the fixed costs associated with introducing a new product have a bounded support. This assumption does not need to be fulfilled in all contexts. If F_{jt} is the cost of developing a new breakthrough technology, it could be that no money in the world makes the necessary invention possible. Since I consider F_{jt} to be the cost of introducing a product at a new retailer and developing new products interior to a firm's technological capability frontier, it seems plausible that there exists an upper-bound to the necessary fixed costs. At the same time, the cost of developing and introducing a new product in this context should never be negative and so the existence of a lower-bound of the fixed cost support, F_b^L , is an innocuous assumption.

Assumption 3. $[F_b^L, F_b^U] \subset \text{supp}(\text{expected change in variable profit due to the elimination or addition of a single product of brand } b)$.

Assumption 3 adds further restrictions on the support of F_{jt} . For each brand b , the support of the fixed costs of introducing any potential product is contained within the support of expected changes in variable profits of firm f if any potential product of brand b is introduced. The intuition behind this assumption is quite simple. If fixed costs of introducing

different potential products of a particular brand come from the same distribution and there exists a blockbuster product that increases expected variable products of the firm so much, that it would always be introduced, then I observe this product as an active product in the data and the expected change in variable profit of adding this product must be higher than the fixed cost of introducing any potential product. Similarly, if there exists a product that has such a small impact on the expected change in variable profit, such that it would never be introduced, then I will always observe this product as an inactive product and the expected change in variable profit of adding this product must be lower than the fixed cost of introducing any potential product.

With these additional assumptions, I can fill the missing upper- and lower-bounds on the fixed costs of potential products. I fill the missing lower-bound on fixed costs for active products by using the minimum change in firm-level expected variable profits among inactive products of the same brand. I fill the missing upper-bound on fixed costs for inactive products by using the maximum change in firm-level expected variable profits among active products of the same brand. The product-level bounds on fixed costs for active and inactive products are defined as

$$L_{jt}(\theta) = \begin{cases} VP_{bt}^L(\theta) & j \in J_{bt} \\ \underline{E}_{jt}(\theta) & j \in \tilde{J}_{bt} \end{cases} \quad U_{jt}(\theta) = \begin{cases} \bar{F}_{jt}(\theta) & j \in J_{bt} \\ VP_{bt}^U(\theta) & j \in \tilde{J}_{bt} \end{cases} .$$

Since $E[v_{jt}|j \in \mathcal{J}_{ft}] = 0$, and with estimates on the upper- and lower-bound on fixed costs for all $j \in \mathcal{J}_{ft}$, I can now apply an unconditional expectation, such that

$$E[L_{jt}(\theta)] \leq F_{bt} \leq E[U_{jt}(\theta)] \quad \forall j \in \mathcal{J}_{bt} . \quad (16)$$

To estimate the set in 16, I replace the true parameter vector θ by the first stage estimator $\hat{\theta}$ and estimate the change in firm-level variable profits of removing any active product and adding any inactive product in the data. I use $\min_{j \in \tilde{J}_{bt}} \{\underline{E}_{jt}(\hat{\theta})\}$ as an estimator for $VP_{bt}^L(\theta)$ and $\max_{j \in J_{bt}} \{\bar{F}_{jt}(\hat{\theta})\}$ as an estimator for $VP_{bt}^U(\theta)$.

Finally, I compute the within brand and market sample average across $L_{jt}(\hat{\theta})$ and $U_{jt}(\hat{\theta})$, to estimate bounds on the set of brand- and market-level fixed costs. This estimation procedure produces unbiased estimates and overall leads to wide and conservative fixed cost bounds.

II.E Details on the employment calibration

To simulate the employment effects of the different hypothetical acquisitions, I need an estimate of how many clothes washers a manufacturing worker produces on average per year. Since I do not have systematic data on employment by manufacturer and appliance category, I calibrate the number of clothes washers produced by manufacturing workers based on different sources.

In 2005, Maytag produced clothes washers and dryers in Newton, Iowa (1,000 manufacturing jobs) and Herrin, Illinois (1,000 manufacturing jobs) and dryers in Searcy, Arkansas (700 manufacturing jobs).⁴⁵ In addition, there was a small plant manufacturing clothes washers and dryers in Florence, South Carolina (60 manufacturing jobs).⁴⁶ According to Appliance Portrait (2006), Maytag shipped 1.75 million clothes washers and 1.6 million dryers in 2005. On average, these are around 1,200 clothes washers and dryers per manufacturing worker per year.

In 2011, the Whirlpool plant manufacturing front-loading clothes washers in Schorndorf, Germany, had 500 employees and produced 200,000 clothes washers.⁴⁷ This amounts to 400 clothes washers per manufacturing worker per year.

To simplify matters, I assume that the number of employees necessary to produce clothes washers linearly increases in the number of clothes washers and that this technology is constant over time and across manufacturers, products, and production locations. With richer data and depending on the institutional context, all of these assumptions can be relaxed.

Based on the evidence described above, I calibrate that a manufacturing worker produces on average around 1,000 clothes washers per year. Among clothes dryers, top-loading washers, and front-loading washers, the first are the simplest products to produce and the last the most complex. It therefore seems plausible that the estimate for Whirlpool front-loaders is an overall underestimate of the number of clothes washers produced by worker and the estimate based on Maytag washers and dryers an overestimate. Either way, choosing a relatively high number of clothes washers per manufacturing worker is a conservative approach, since it likely underestimates the employment effects of either acquisition.

⁴⁵See <https://www.nbcnews.com/id/wbna12718867>.

⁴⁶See <https://www.twice.com/news/maytag-close-florence-laundry-facility-27876>.

⁴⁷See <https://www.stuttgarter-zeitung.de/inhalt.bauknecht-in-schorndorf-konzern-gibt-den-standort-auf.2559fd28-6719-48b9-a055-5956c7f61c03.html>.

III Appendix to Section 6: Further results of the structural estimation

III.A Demand estimation

Table A.2: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)	(5)
	First-stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
<i>Dependent variable:</i>	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Real exchange rate	1.909*** (0.398)	-0.787** (0.358)			
Price ('00 2012 \$)			-0.164** (0.062)	-0.412** (0.202)	-0.637*** (0.024)
Front-loader	0.174 (0.205)	0.267 (0.267)	0.358 (0.244)	0.339 (0.215)	-0.715*** (0.021)
Korean front loader	-0.563*** (0.179)	1.746*** (0.353)	1.569*** (0.349)	1.514*** (0.348)	1.484*** (0.012)
Fisher & Paykel front-loader	-4.506*** (0.331)	-0.624 (0.412)	-1.455*** (0.480)	-2.481*** (0.859)	-3.165*** (0.093)
European high-end front-loader	0.071 (1.311)	1.235*** (0.314)	1.192*** (0.438)	1.264* (0.715)	1.256*** (0.025)
Agitator	-2.510*** (0.276)	0.952*** (0.270)	0.540** (0.252)	-0.083 (0.532)	-0.457*** (0.060)
Stacked pair	0.493* (0.280)	-0.225 (0.149)	-0.147 (0.149)	-0.022 (0.202)	0.028** (0.010)
Stainless steel exterior	0.481 (0.603)	-0.052 (0.247)	0.009 (0.270)	0.146 (0.362)	0.180*** (0.010)
White exterior	-0.289 (0.360)	0.677*** (0.130)	0.624*** (0.101)	0.558*** (0.131)	0.510*** (0.009)
Energy Star	0.023 (0.182)	0.089 (0.126)	0.092 (0.126)	0.099 (0.138)	0.114*** (0.004)
Extra noise insulation	0.395* (0.207)	0.248** (0.125)	0.312** (0.120)	0.411** (0.162)	0.470*** (0.010)

continued

Table A.2: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)	(5)
	First-stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
<i>Dependent variable:</i>	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Number of special programs	0.009 (0.058)	0.050 (0.035)	0.052 (0.039)	0.054 (0.047)	0.051*** (0.001)
Child lockout	-0.073 (0.164)	0.204 (0.172)	0.200 (0.167)	0.174 (0.171)	0.176*** (0.005)
Repair rate	-2.397 (3.156)	2.048 (3.272)	1.627 (2.957)	1.060 (2.793)	0.727*** (0.133)
Total advertising expenditure	-0.006 (0.005)	0.004 (0.002)	0.003 (0.002)	0.001 (0.002)	0.001*** (0.0002)
Retailer Best Buy	-0.097 (0.086)	-1.045*** (0.299)	-1.062*** (0.307)	-1.085*** (0.309)	-1.098*** (0.002)
Retailer H. H. Gregg	-0.369*** (0.119)	-1.903*** (0.282)	-1.963*** (0.299)	-2.054*** (0.278)	-2.105*** (0.008)
Retailer Home Depot	-0.162 (0.106)	-0.738** (0.314)	-0.765** (0.321)	-0.804** (0.324)	-0.828*** (0.004)
Retailer Lowe's	-0.180** (0.091)	-0.301 (0.230)	-0.334 (0.231)	-0.375* (0.224)	-0.401*** (0.004)
Retailer Sears	0.015 (0.113)	-0.436 (0.442)	-0.435 (0.445)	-0.430 (0.445)	-0.426*** (0.001)
Brand FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Brand time trends	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	1,590	1,590	1,586	1,590	1,590
F-statistic	22.979				
Own-price elasticity			-0.964	-2.416	-3.301

Notes: Column (1) presents results for the first stage regression of prices on the real exchange rate. Column (2) includes reduced form estimates for the simple logit model. Column (3) reports demand estimates for the simple logit without a price instrument. Column (4) presents demand estimates for the simple logit model using the RER as an IV for price. Column (5) shows demand estimates for the mixed logit model presented in Section 4 and using the RER as an IV. Standard errors are clustered at the brand level. The own-price elasticity of residual demand is computed at the product level and the average is calculated by weighting products according to their sales volume. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

IV Appendix to Section 7: Details on the welfare effects

IV.A Technical details on the portfolio choice algorithm

After initializing the algorithm, there is an inner and an outer optimization loop to find a one-step equilibrium in portfolio choices. In the inner loop, a particular player computes both the expected change in firm-level profits of adding each inactive product separately to the brand's product portfolio, as well as the expected change in firm-level profits of removing each active product separately.⁴⁸ If there is at least one profitable one-step deviation, the player implements this deviation and changes her product portfolio accordingly. I repeat this process until the player has no profitable one-step deviation left. In the outer loop, I repeat this process for each player.

In practice, I can considerably reduce the computational burden by optimizing product portfolios brand-by-brand instead of firm-by-firm. This requires computing fewer potential one-step deviations for every portfolio adjustment. Although I fully take into account how the introduction or removal of a product impacts the firm's expected profit (and not just that of the brand), the downside to this approach is that if products of two brands of the same firm are very close substitutes, the order of play could matter for which product enters. This is unlikely to play an important factor, as firms segment their products by brands and so products within a brand are much closer substitutes than between brands of the same firm.

Another way in which I reduce the computational burden is by only considering one-step deviations and disregarding multi-step deviations. This is necessary because checking for any multi-step deviations is also computationally infeasible in this case.⁴⁹ It could thus be that although there is no profitable one-step deviation, there nevertheless exists a profitable multi-step deviation. To assess whether this could be an important problem, it is helpful to consider when such a situation could arise. Since clothes washers are substitutes in the marketplace, if it is not profitable to add a particular clothes washer, it is also not profitable to add that and another potential clothes washer. The same logic applies to the removal of active clothes washers from the product portfolio. It is, however, possible that although adding a particular clothes washer is not profitable, it would be profitable to add the clothes washer and remove another washer from the product portfolio simultaneously.

⁴⁸Since I do not observe realized demand and supply shocks for potential products, I estimate the expected welfare effects based on 500 demand and supply residual draws for each product.

⁴⁹To illustrate this point, brands have up to 15 potential products. Checking for all multi-step deviations would thus require checking up to $2^{15} = 32,768$ candidate deviations at each brand iteration.

Similarly, it could be that it is profitable to add a clothes washer and remove two washers simultaneously. Overall however it may not be desirable to consider multi-step deviations with many different portfolio adjustments simultaneously, since it is more difficult to make many portfolio adjustments at the same time.

Finally, as I only set identify fixed costs, I repeatedly apply the portfolio choice algorithm for 50 different fixed cost draws for each product. Although there are no restrictions on the within-brand distribution of fixed costs when estimating the fixed cost bounds, I need to make a distributional assumption for the estimation of counterfactuals. In the spirit of Wollmann (2018), I assume that F_{bt} is equal the midpoint of the confidence bounds. I draw the idiosyncratic product- and market-specific fixed cost shock v_{jt} from a mean-zero normal distribution with a standard deviation equal to 25 percent of the difference between the upper- and lower-bounds of the 95 percent confidence sets of brand-level fixed costs. In all scenarios, I report 95 percent confidence sets for the welfare effects across fixed cost draws.