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

Felix Montag
NYU Stern School of Business*

This version: November 14, 2024

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 Whirlpool's acquisition of Maytag in the household appliance industry. I compare the observed acquisition to a counterfactual one by a foreign buyer. While Whirlpool's acquisition decreased consumer welfare by $230 million annually, it preserved 900 domestic jobs. Therefore, these jobs need to be worth above $260,000 annually for the domestic employment effects to offset the consumer harm.

JEL classification: F61, L13, L40

*Email: felix.montag@nyu.edu. I am deeply grateful to my advisors Monika Schnitzer and Ali Yurukoglu for their patient guidance, constant encouragement and invaluable advice. I thank Alessandra Allocca, Antonin Bergeaud, Laura Grigolon, Anna Gumpert, Ying Fan, Teresa Fort, Mitsuru Igami, Torsten Jaccard, Claudia Steinwender, Fabian Waldinger, Lulu Wang, Martin Watzinger, and participants at many seminars and conferences for helpful comments and discussions. Financial support from 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, 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 a merger's consumer welfare and employment effects. I incorporate two key margins of firm decision-making in response to a merger: price adjustments and product portfolio changes. 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. My key methodological innovation is to construct a model that captures the interactions between the product market and labor demand. Using this model, I estimate the average value of a domestic job such that gains to workers offset losses to consumers.

To illustrate this worker–consumer trade-off, I use my 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 and a hypothetical purchase of Maytag by the alternative potential buyer at the time, the Chinese company Haier. Production locations are exogenous. 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.[4] While there was no threat to consumers from a Haier acquisition because the firm had no prior presence in

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

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

[3] Many jurisdictions, such as Australia, China, France, Germany, Italy, Korea, the United Kingdom or South Africa incorporate public interest considerations into merger control (see OECD, 2016). Although this is currently not the case for the European Union or in the United States, workers have become increasingly important in merger control in these jurisdictions.

[4] This assumption is based on anecdotal evidence on the Haier bid, described in Section 2.

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

Three main findings emerge: first, the Whirlpool acquisition is always worse for consumers. The merging parties increase prices of washing machines without inducing substantial rival product entry, leading to an annual consumer welfare loss of $230 million. Second, the Whirlpool purchase preserves approximately 900 more U.S. jobs than the Haier purchase. Finally, I find that each additional job maintained by Whirlpool (relative to employment under the Haier acquisition) must be worth approximately $260,000 per year to offset the consumer harm.

While I find that the worker–consumer trade-off in this application strongly depends on the presence of an alternative buyer with different offshoring plans, this feature of a proposed merger need not be present for the trade-off to arise.[5] 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.

The counterfactual simulations show that incorporating the pricing and product assortment decisions substantially changes the employment effects. This matters for many studies that analyze the production decisions of multinational enterprises (MNEs).[6] A single industry study allows focusing on the key trade-offs manufacturing MNEs face in their global production decisions. In the appliance industry, MNEs produce different product groups in different locations. A production cost change for a product group (such as a local wage change) leads MNEs (and their rivals) to adjust prices and product portfolios, which affects production and labor demand. This margin of adjustment complements work by Muendler and Becker (2010), who study how employment at MNEs responds to wages at the extensive margin (opening a new foreign location) and intensive margin (changing employment at an existing location) when different locations can complement or substitute each other's production. Whereas I focus on intensive-margin adjustments only, I account for how the product market equilibrium adjusts after an MNE experiences a cost change at one of its locations. The findings also complement the results of Igami (2018), who studies the interplay between offshoring and market structure without analyzing employment.

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

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

[6]See, for example, Tintelnot (2017) or Arkolakis, Eckert, and Shi (2023).

difference-in-differences identification strategy, similarly to Ashenfelter, Hosken, and Weinberg (2013). I show that there was a significant increase in concentration within the U.S. market for washing machines. While prices remained stable in the U.S. after the merger, they fell for washers across European countries affected by similar cost trends but where Maytag did not have a presence. New product introductions by foreign rivals occurred, but this was also the case in the control market. Finally, I show a rise in local unemployment and a decrease in employment for counties that experienced the closing of a Maytag plant.[7]

The descriptive evidence raises several questions: Was the rival product entry induced by the merger? If consumers were harmed by the Whirlpool acquisition, was this harm offset by benefits to U.S. workers? Answering these questions requires a model.

The model I elaborate here features manufacturers and consumers. Manufacturers choose their product portfolios and prices. Consumers make purchase decisions. The model is set up as a three-stage game, with the final two stages entering the estimation. At the beginning of the game, each manufacturer chooses its production locations. In the second stage, it 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. Each firm chooses which potential product to introduce into the market, at a per product fixed and sunk entry cost.[8] The firm has expectations about the quantities it will sell of each product, based on which it hires workers to manufacture these quantities. Whether a job is created domestically or abroad depends on the production location for each product. In the third stage, firms set prices, and consumers make purchases. I model consumer demand using a static random coefficients discrete choice model, where consumers' price sensitivity depends on income and some consumers have an unobserved taste for front-loading washers.

The demand estimation is in the spirit of Berry, Levinsohn, and Pakes (2004). Informally, the nonlinear 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. I use this as an instrumental variable for price, which is exogenous to product-level demand conditions (see Goldberg and Verboven, 2001 and 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. A growing literature is concerned with

[7]This is in line with recent evidence that the presence of multinational firms affects wages at other firms (see Card et al., 2018, Alfaro-Ureña, Manelici, and Vasquez, 2021, and Setzler and Tintelnot, 2021).

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

estimating bounds on the fixed costs of introducing a new product 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 the approach of Eizenberg (2014).

I collect plant-level data to estimate the relationship between employment and output using a Leontief production function. I find that there is a fixed number of workers needed to run a plant with no output after which output grows linearly in employment. Combining this estimate with the product-level plant locations, I can estimate how different product market equilibria affect the demand for domestic manufacturing workers. Finally, I decompose marginal costs into different components, including location-specific labor and transportation costs, and estimate how the marginal costs react to a change in production location.

I face several empirical challenges. A first challenge is to identify the set of products that multiproduct firms can introduce.[9] Studying an unconditionally cleared merger allows me to overcome this issue. 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.[10]

A second empirical challenge is the multiplicity of equilibria in 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 and 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 simulation results show that the average saved job would need to be worth approximately $260,000 per year to offset the associated consumer harm ($123,000 if we consider the sum of consumer welfare and industry profits). In comparison, Setzler and Tintelnot (2021) find that the total wage bill in a local labor market increases by approximately $113,000 per year for each job created by a foreign MNE. This estimate does not account for other employment benefits, which further increase the value of a job. The employment effects are

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

[10] Eizenberg (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 for studying the introduction of new products.

sizable, and the welfare effects are unequally distributed. Whereas the losses to consumers are spread across the country, the employment effects are concentrated in a few local labor markets. A policy implication of this finding is that procompetitive competition enforcement could be paired with targeted labor market policies like the wage insurance under the Trade Adjustment Assistance Program (see Hyman, Kovak, and Leive, 2023) to compensate losers. 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 and Flaaen, Hortaçsu, and Tintelnot, 2020) on workers and consumers. Among these estimates, I identify the lowest job values necessary to offset consumer harm.

Finally, I contribute novel evidence on how endogenous product portfolio choices change the consumer welfare effects of mergers.[11] Despite the centrality of entry arguments to the clearing decision for the Whirlpool–Maytag merger, I find that endogenous portfolio adjustments increase the harm to consumers. This is because rival product entry is mostly independent of the merger whereas the merging parties offer fewer products post-merger. Existing studies mostly consider hypothetical changes in concentration and find mixed results.[12] Caradonna, Miller, and Sheu (2021) show that, under certain conditions, 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, approximately 90 percent of washing machines 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.

[11] A related literature (e.g., 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 multiproduct 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).

[12] Fan 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.

2.1 Whirlpool's acquisition of Maytag

Prior to its acquisition by Whirlpool, Maytag had been struggling financially. Although the company had already cut costs by reducing its workforce by 20 percent, in 2004, it continued to struggle with cost pressure and 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 $1.1 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 receiving an unconditional merger clearance from the Department of Justice (DoJ).

Haier's bid came at a time when the Chinese government was pushing its large companies to make foreign acquisitions to obtain access to foreign markets, particularly in the European Union and the United States, for Chinese manufactured goods.[13] Since Chinese acquirers were 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 consumer resistance to Chinese brands in the product market.[14] With its weak financial performance and its strong brand portfolio, Maytag fit the bill. Haier, which previously had negligible sales in the U.S. appliance market, planned to use Maytag's brands, repair network and distribution channels while 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. In its investigation of the acquisition, the DoJ focused on residential washing machines and dryers.[15] This merger would shrink the number of firms manufacturing laundry products for the U.S. market from four to three firms, with Whirlpool and Maytag being the largest and second largest manufacturers. Sears, with its Kenmore brand, was another large brand owner in the U.S. 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) call this decision "[...] a highly visible instance of underenforcement", and Baker

[13] This was part of China's "Go Out Policy" promoting investments abroad (Goodman and White, 2005).
[14] A famous example is the 2005 acquisition of I.B.M.'s personal computer division by Lenovo.
[15] See Armington, Emch, and Heyer (2006) for a case description from the perspective of the DoJ.

and Shapiro (2008b) describe 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 despite the fact that entrants had thus far achieved only relatively low market shares.

2.2 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 among the data sources that I assemble is *TraQline*. This data set is well known across the appliance industry and is used by retailers and brands in the industry as a source for market insights.[16] In every quarter, a representative sample of approximately $150,000$ U.S. households is asked about appliance purchases. The survey is a repeated cross-section, and in total, approximately $600,000$ households are surveyed every year. The data span 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 the product is Energy Star certified), other brands that the household considered buying, the retailer at which the appliance was bought, and a detailed set of household demographics. The data include information for washers and dryers and for freestanding ranges. I aggregate products to the national level because product entry is determined for each major retailer at this level. I also aggregate responses to 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 preferences of 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 washers, this includes whether the 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

[16] The 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 is the basis of the analysis by Ashenfelter, Hosken, and Weinberg (2013).

[17] In 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 stands for quality and craftsmanship, Whirlpool for innovation and Maytag for reliability.

definition by using information on the retailer at which the product is sold. 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 of variation in observed and unobserved characteristics, 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. The households selected to answer the more detailed characteristics questions do not have the possibility to opt out, ruling out any selection problems. For washers, these more detailed characteristics include whether the washer has a child lockout, the number of special programs, whether it is a stacked pair and 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. A report on washing machines is released at least once a year and includes an overview of brand-level repair rates. These data are based on responses to the Annual Product Reliability Survey conducted by the Consumer Reports National Research Center for more than $100,000$ washers. I digitize this information to create a measure of the perceived product reliability of a brand over time.

Annual information on advertising expenditures between 2005 and 2015 comes from Kantar AdSpender. 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 includes household demographics only for respondents who 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.

In the descriptive analysis, I compare washing machine prices and sales in the U.S. to those in several European markets unaffected by the merger. To do this, I use monthly data on washing machine sales at the model level for Germany, France, Great Britain, and the Netherlands between 2005 and 2008, which were compiled by the *Gesellschaft für Konsumforschung*. In addition to volume and value sales, the data set includes detailed information

[18]For retailers, I distinguish Best Buy, H.H. Gregg, Home Depot, Lowe's, Sears, and all others. The latter group predominantly includes smaller, regional retailers.

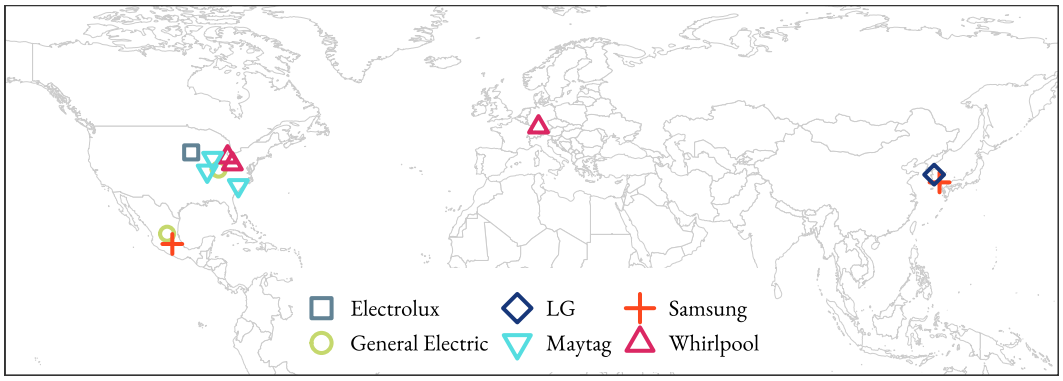
about product characteristics, such as brand, spin speed, and water consumption.

2.3 Production locations, cost shifters and employment

The core of the supply-side data consists of hand-collected locations of plants manufacturing washing machines for the U.S. market at the product level. This data set serves two purposes. First, it allows me to construct a product-level instrumental variable for prices based on differences in production costs. Second, it allows me to simulate how the number of U.S. washing machine manufacturing jobs changes across counterfactual scenarios.

Figure 1 shows the plant locations of major washing machine 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.

Figure 1: Washing machine plants manufacturing for the U.S. market, 2005



Notes: The map shows all plants manufacturing washing machines 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.

Plant locations are based on the investigation by the U.S. International Trade Commission (USITC) into imports of large residential washing machines from Mexico and South Korea, firm-level washing machine imports based on the PIERS data set, which is reported in Flaaen, Hortaçsu, and Tintelnot (2020), annual reports, and news reports. Details on the construction of this data set and the weights on the share of the products produced in each production location are described in Appendix I.B.

I construct a product-level cost shifter using the product-level weighted average real exchange rate between the U.S. and the countries in which the product is produced. 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. I calculate this 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). This measure therefore captures differences in relative price levels and serves as a proxy for the local wage level and fluctuations in the nominal exchange rate. Figure A.7 in Appendix I.C shows the evolution of the average RER over time and illustrates the source of the variation.

Estimating the relationship between plant-level output and employment requires plant-level data on both. Since these data are only sparsely available for plants producing for the U.S., I also consider BSH and Whirlpool plants that produce for the European market. In practice, the technology at a Whirlpool plant producing washing machines for the U.S. market in Germany appears to be similar to the technology at a Whirlpool plant producing washing machines for the European market. I collect data on annual plant output quantities and manufacturing workers employed between 1995 and 2017. In total, there are 25 plant–year combinations. Four of these belong to Whirlpool plants producing for the U.S. market, whereas 21 belong to BSH and Whirlpool plants producing for the European market. Details on the construction of this data set are described in Appendix I.B.

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 county-level employment and wage data as reported by employers. I use the wages per employee and number of persons employed, disaggregated by county and industry. The wages include total compensation, bonuses, stock options, severance payments, the cash value of meals and lodging, tips, and other gratuities. I annualize these wages to ease interpretation. The Local Area Unemployment Statistics (LAUS) aggregate data from state-level workforce agencies. They include monthly information on the number of unemployed individuals for every U.S. county.

To identify counties similar to the counties affected by plant closures, I construct a matched control group based on relevant county characteristics that are not outcome variables. I use annual QCEW data for 2003–2006 on the county-level industry composition of employment, decomposed into 21 industries. I also use information on county-level age structure, racial composition, and educational attainment. These data are not available at annual frequency, which is why I use the 2005–2009 average from the 2009 ACS.

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
ΔHHI	1,149			1,124		

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

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 U.S. horizontal merger guidelines, the acquisition therefore potentially posed 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 share and the HHI gradually declined (relative to the

[19]The HHI is calculated as the sum of squared market shares in whole percentages.

[20]U.S. horizontal merger guidelines identify mergers with a pre-merger HHI between $1,500$ and $2,500$ and an HHI increase of more than 100 as potentially posing significant competitive concerns.

post-merger HHI based on pre-merger market shares), the increase in concentration due to the merger remains substantial and persistent.

Hereafter, I focus exclusively on washers. The main reason for this is data limitations. Whereas I have access to detailed sales data for washing machine markets unaffected by the merger, I do not have comparable data for dryers. Similarly, I collect detailed information on the production locations of washers, whereas this is more difficult for dryers. On the product market side, the results can be expected to be similar. As shown in Table 1, market shares for washers and dryers evolve similarly. In Appendix II, I show that some of the descriptive evidence for price effects and entry by LG and Samsung is similar for washers and dryers.

3.2 Evolution of prices

Next, I turn to the descriptive evolution of prices around the time of the acquisition. There are two obvious strategies for analyzing the price effects of the Maytag–Whirlpool merger descriptively. The first strategy is used by Ashenfelter, Hosken, and Weinberg (2013), who compare the evolution of Maytag and Whirlpool product prices for appliance categories with a large increase in concentration to that of categories with low increases in concentration. I repeat this exercise in Appendix II using freestanding ranges as a comparison appliance and also estimate price effects for dryers. The results are mostly in line with those of Ashenfelter, Hosken, and Weinberg (2013).[21] The second strategy is to compare prices in affected markets to those in geographic markets unaffected by the merger. Here, I follow the latter strategy.

To estimate the price effects of the Maytag–Whirlpool merger, I compare the logarithm of prices for Maytag and Whirlpool washers in the U.S. after controlling for key characteristics to the logarithm of prices for washers from other manufacturers in several key European markets. As Maytag had no significant presence in Europe, the latter serve as an appropriate control group subject to similar industry trends but unaffected by the merger.

I use quarterly data for the years 2005–2008 to estimate a separate event study for every country. Specifically, I estimate the parameters of the following model:

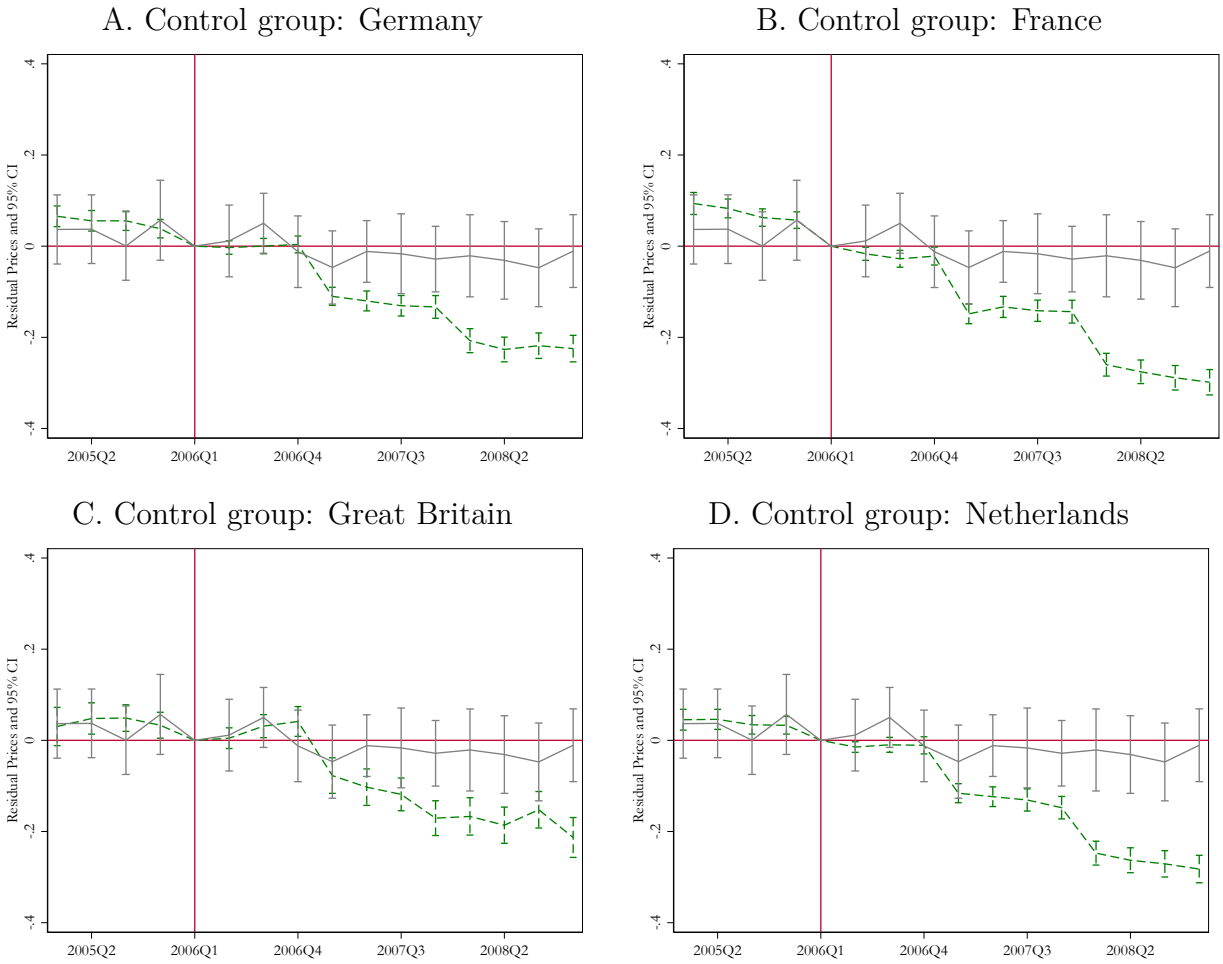
$$ln(p_{it}) = \beta x_{it} + \delta_b + \gamma_t + \epsilon_{it}\,, \tag{1}$$

where i denotes a model of brand b at time t, x_{it} is product characteristics, δ_b is brand fixed effects, γ_t is quarterly time fixed effects, and ϵ_{it} is idiosyncratic product-level shocks. Different characteristic variables are available in the different data sets. For Germany, France,

[21] Since Ashenfelter, Hosken, and Weinberg (2013) use data from *NPD*, to which I do not have access, whereas I use data from *TraQline*, I am not able to conclusively identify the source of the remaining differences in results. The main results and economic conclusions between the two analyses are the same, however.

Great Britain and the Netherlands, x_{it} includes whether the washer is a front-loader, the load capacity, the spin speed and the water consumption. For the U.S., it includes whether the washer is a front-loader, whether a top-loader has an agitator, the brand repair rate, and the total advertising expenditure.

Figure 2: Price effects of the merger for washing machines



Notes: The figure shows the residualized logarithm of prices for Maytag and Whirlpool washing machines in the United States relative to prices of washing machines from other manufacturers in control countries unaffected by the merger. The solid line shows the evolution of residualized prices in the United States and the dashed line that for the control group. Confidence bounds are at the 95 percent level. Standard errors are clustered at the model level.

Figure 2 plots the quarterly time fixed effects δ_t, where Q1 of 2006, the quarter in which Whirlpool acquired Maytag, is normalized to zero. Standard errors are clustered at the model level.

There are two key takeaways from the descriptive evidence: First, the washing machine prices of the merging parties' products in the U.S. and of other manufacturers in key European markets remained stable and evolved in parallel throughout 2005 and 2006. Whereas prices after the merger fell in 2007 and 2008 in control countries, prices in the United States

did not. Thus, if the washing machine markets in the control countries are subject to similar industry trends on all aspects other than the merger, this suggests that the Maytag acquisition by Whirlpool significantly increased prices relative to their counterfactual evolution.[22]

By simply studying the price evolution in the U.S., one might incorrectly conclude that the merger did not lead to price increases. Two industry trends are plausible candidates for the observed price decreases in European markets: The first is the general trend in the industry during the 2000s to relocate production to low-cost countries. As shown in Figure A.7, most major manufacturers relocated increasing shares of their production to lower-cost countries between 2007 and 2010.[23] Second, entry by new competitors from Asia, such as LG and Samsung, occurred around the time of the merger.

3.3 Product entry

To assess the effects of the merger on consumers and workers, it is crucial to understand whether rival product entry would have occurred even in the absence of the merger. Rival product entry could affect the estimated price effects of the merger in two distinct ways: First, if the merger led to merger-specific product entry, this could increase competition and decrease prices. Second, if there was merger-independent product entry by rivals around the time of the merger, this would also increase competition and reduce prices.

To assess whether the product entry by LG and Samsung occurred because of the merger, I document the evolution of the two firms' market shares in the U.S. washing machine market versus that for Germany, France, Great Britain, and the Netherlands.

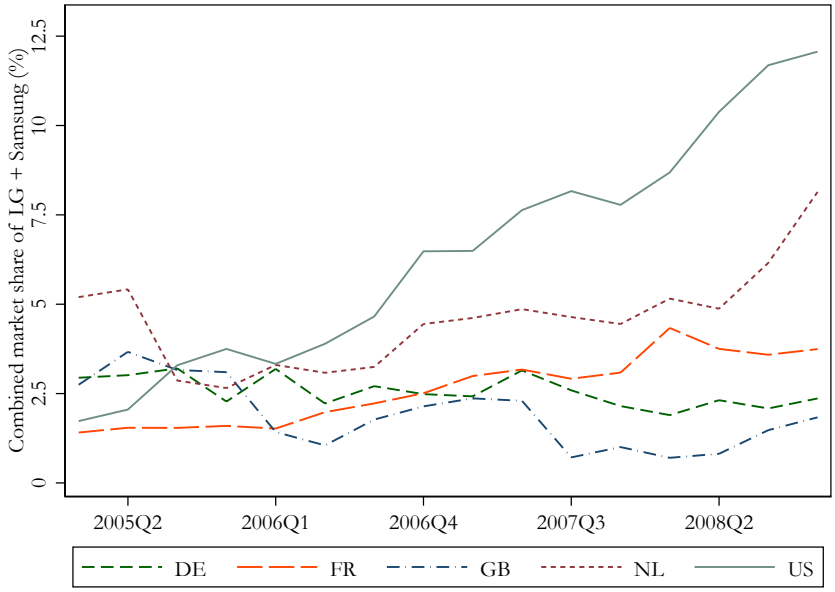
Figure 3 shows that LG and Samsung were present across all markets in 2005, before the merger. Whereas their combined share stagnated in Germany and Great Britain between 2005 and 2008, it increased in France, the Netherlands, and the United States. This suggests that expansion by LG and Samsung into the U.S. market was, to the very least, not exclusively driven by the Maytag–Whirlpool merger and that, in the absence of the merger, residual prices in the U.S. should have fallen instead of remained stable. At the same time, among these countries, the increase in market share was strongest in the U.S. This raises the question of whether, at least in part, the merger induced stronger product entry and expansion by LG and Samsung than would otherwise have happened.

In Appendix II, I show that LG and Samsung strongly increased their market share for washers and dryers in the U.S. but achieved a market share for freestanding ranges of

[22] In Appendix Figure A.9, I plot the evolution of deflated consumer price indices in the U.S. and EU between 2001 and 2015. The results corroborate the findings of the descriptive price analysis using GfK data.

[23] Appendix I.C documents quotes from the annual reports of, e.g., Electrolux and Maytag, to show that this was a key part of their business strategy at that time.

Figure 3: Combined market share of LG and Samsung in the U.S. and Europe



Notes: The figure shows the evolution of the combined market share in terms of sales volumes in the U.S. and key European control markets between 2005 and 2008.

only 2 percent at the end of 2008. This suggests either that the merger induced entry by LG and Samsung or that freestanding ranges may be a less appropriate control group for the evolution of prices in the U.S. washer market than the washing machine markets in Europe.

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 the Maytag plant closures by Whirlpool post-acquisition affected employment, unemployment, and wages of the employed in the affected counties.

Although Whirlpool maintained some of Maytag's manufacturing plants (e.g., in Amana, Iowa, and 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), and manufacturing and headquarters operations in Newton, Iowa (1,000 manufacturing and 1,800 corporate jobs). At the same time, Whirlpool announced the addition of 1,500 jobs at two existing plants in Ohio. There were thus three mass layoffs and two new job events. All of these were announced in 2006 and occurred in 2007.

To quantitatively assess the local labor market effects of changes in Maytag employ-

ment, I build on the methodology set forth by Gathmann, Helm, and Schönberg (2020), who study local employment effects of mass layoffs. I compare labor outcomes in treated counties to those in a set of matched control counties using an event study methodology.

First, I follow Gathmann, Helm, and Schönberg (2020) and match a control group for every treated county that is similar on pre-treatment industry composition (21 different industries), as well as age structure (age groups 15–24, 25–49, and 50–66), education (low, medium, and high), and race (white, Black, and any other race). I do not match on any of the outcome variables of the employment analysis. Unlike Gathmann, Helm, and Schönberg (2020), since I observe only a small number of events, I use the ten closest counties as a control group for every treated county instead of only the closest.

Specifically, I calculate the normalized squared distance in the population share by characteristic of the treated counties to any other county in the United States.[24] For the industry composition, I use annual data for the four years prior to treatment (i.e., 2003–2006). The county-level demographic characteristics are not available at this frequency, which is why I use the 2005–2009 average from the 2009 ACS. In practice, these variables should not vary much year on year.

Second, I estimate the following event study:

$$\mathrm{Y}_{ct} = \sum_{\tau=2003}^{2005} \alpha_\tau \mathrm{Event}_{ct}^{\tau} + \sum_{\tau=2006}^{2008} \beta_\tau \mathrm{Event}_{ct}^{\tau} + \delta_c + \gamma_t + \epsilon_{ct}\,, \tag{2}$$

where Y_{ct} is the logarithm of unemployment, manufacturing employment or annualized manufacturing wages in county c at time t, $\mathrm{Event}_{ct}^{\tau}$ is an indicator equal to one if time t is in year τ and county c is a treatment county, δ_c is county fixed effects and γ_t is time fixed effects. Whereas data on unemployment are available at monthly level, the other outcome variables are available only at quarterly level. In each case, I use the most disaggregated level.

In the baseline specification, I use a simple difference-in-differences strategy, where I use data for 2003–2006 as the pre-treatment period and data for 2007 and 2008 as the post-treatment period.[25] Appendix II includes graphical results for the full event study analysis, where there are separate treatment effects for the years prior to treatment (2003–2005) and the treatment years 2007 and 2008.

I group the counties into two different treatment groups and estimate separate regres-

[24]The distance equation for all treatment counties i and control counties j is given by $distance_{ij} = \sum_n \omega_n \sum_t \frac{(x_{int}-x_{jnt})^2}{s_{nj}^2}$, where s_{nj} is the standard deviation of characteristic n in the control counties j and ω_n is the weights of the different characteristics. Industry composition together has a weight of 50 percent, and each of the demographic characteristics accounts for a third of the rest. See Gathmann, Helm, and Schönberg (2020) for more details on the methodology.

[25]To avoid adding noise from the great recession, I do not use data for 2009.

Table 2: Labor market effects of mass layoffs and new jobs

	Unemployment		Employment		Wages	
	(1)	(2)	(3)	(4)	(5)	(6)
Mass layoff	0.154***		-0.336***		-0.015	
	(0.056)		(0.107)		(0.049)	
New jobs		0.018		0.085***		-0.005
		(0.032)		(0.026)		(0.018)
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,160	1,368	720	456	720	456

Notes: The dependent variable is the logarithm of county-level unemployment in Columns (1) and (2), employment in manufacturing in Columns (3) and (4), and wages in manufacturing in Columns (5) and (6). The analysis is at monthly level for Columns (1) and (2) and at quarterly level for Columns (3) to (6). Standard errors are clustered at county level and reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

sions. The first treatment group consists of Jasper County, in which there was a shutdown of both manufacturing and corporate operations, as well as White and Williamson Counties, in which manufacturing plants were shut down. The second group consists of Marion and Sandusky Counties, where Whirlpool created new jobs.

Columns (1) and (3) in Table 2 show that the mass layoffs led to a statistically significant increase in local unemployment and a decrease in local manufacturing employment. Column (5) shows no significant effect on local average wages in manufacturing. Columns (2) and (4) show that although the newly created Whirlpool jobs did not affect local unemployment, they led to a statistically significant increase in local manufacturing employment. The results in Column (6) suggest that this did not affect average local manufacturing wages.

The event study plots in Figure A.11 confirm that the effects on unemployment and manufacturing employment coincided with the timing of the employment changes at Maytag and Whirlpool and are not the consequence of pre-existing trends. They also show that the treatment effects were still present in 2008, the second year after treatment.

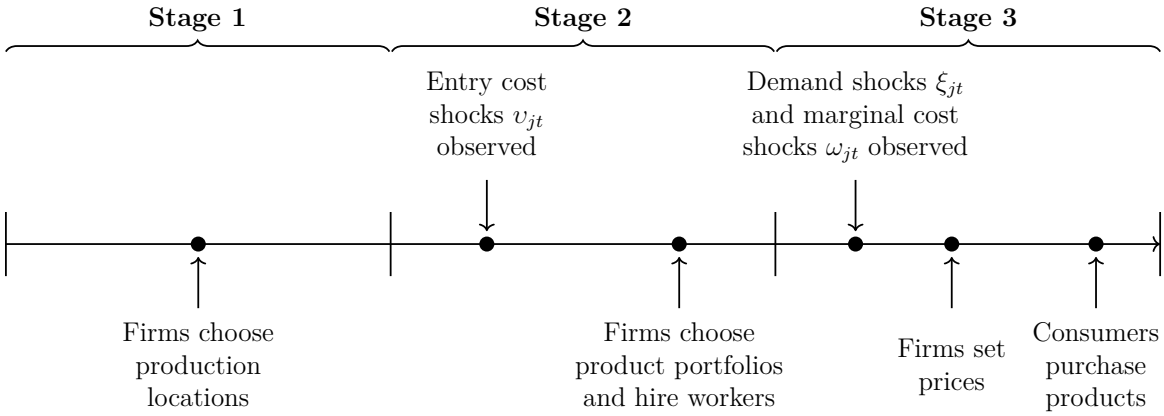
Overall, these results suggest that appliance manufacturing jobs matter for local labor markets and can have persistent effects on employment and unemployment. This conclusion is further supported by Gathmann, Helm, and Schönberg (2020), who show that mass layoffs also lead to sizable and persistent negative spillover effects to the local economy, and Jacobson, LaLonde, and Sullivan (1993), who show that high-tenure workers suffer long-term

earnings losses after separating from distressed firms.

4 Model

Three observations emerge from the preceding analysis. First, entry played a crucial role in the product market, so we need to understand entry to assess the effects of the merger. Second, although there were many changes in the product portfolios of the 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.

These observations inform my model, which features manufacturers and consumers. The manufacturers choose their product portfolios and prices, and the consumers make purchase decisions. The model proceeds in three stages. In the first stage, firms choose their production locations. This stage serves the sole purpose of demonstrating the trade-offs that firms consider in this decision and is not part of the empirical estimation. In the second 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, decide which products to offer, and hire workers. At this stage, firms do not observe transitory demand or marginal cost shocks and only form expectations about them. In the final 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.



I solve this game backward by searching for the subgame perfect Nash equilibria (SPNE) of the game.[26] To estimate the parameters of the game, I require that an SPNE

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

exists but not that it is unique.

4.1 First stage: Production locations

At the start of every period, firm f chooses the set of locations C_{flt} where it produces washers of load type l, that is, front loaders or top loaders. The set of feasible production locations is denoted by $\mathcal{C}$. Let us consider the firm's production location choice problem in period t, where it maximizes expected discounted profits:

$$\Pi(C_{flt-1}) = \max_{C_{flt} \subseteq \mathcal{C}} \left\{ \mathbb{E}\left[\pi(\boldsymbol{p}, \boldsymbol{J_{ft}}) - \sum_{c \in C_{flt}} PC_{cflt} \right] - \sum_{c \in C_{flt}} RC_{cflt} \mathbb{1}\{c \notin C_{flt-1}\} + \tau \Pi(C_{flt}) \right\},$$

where $\pi(\boldsymbol{p}, \boldsymbol{J_{ft}})$ is the period t profits of firm f conditional on its product portfolio $\boldsymbol{J_{ft}}$ and a price vector $\boldsymbol{p}$, PC is the fixed production cost of producing washers of load type l in plant c, RC is the fixed cost of relocating production for washers of load type l to plant c, and τ is an intertemporal discount factor. The relocation costs are multiplied by an indicator function that is one if load type l was not produced in plant c in period $t-1$.

The production location choice is not part of the estimation. Instead, formalizing it reveals two observations: First, as production costs increase with every additional plant in which a firm produces washers of a particular load type, firms have an incentive to reduce the number of production locations. Empirically, producers rarely have more than two plants making washers of a particular load type for the same geographic market. Second, since relocation costs are incurred only if the firm does not yet produce washers of that particular load type in plant location c, the plant location network in $t-1$ matters for the relocation decisions in t. Since Haier and Whirlpool had different plant location networks prior to their proposed acquisition of Maytag, they would also have different relocation costs post-acquisition. Haier, with its existing production locations in China, therefore had a larger incentive to relocate Maytag's production to China than Whirlpool.

4.2 Second stage: Product portfolios and production

In the second stage, firms decide which products to offer and how to produce. At the outset, each firm is endowed with a set of potential products that it can offer in market t, $\mathcal{J}_{ft}$.[27] This can be thought of as the set of products that the firm is technologically capable of producing. The set includes products that it sells already at a different retailer or in a

[27] As I do not have cross-sectional variation in markets, only variation over time, I use the terms "market" and "time" interchangeably.

different market and new products resulting from minor adjustments to existing products that it could make 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) and marketing or retailer investments. Empirically, I analyze markets at the yearly level. At the same time, Ashenfelter, Hosken, and Weinberg (2013) show that the sales volumes of a particular model rapidly decline by the twelfth month after its introduction. 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-specific component F_b and a mean-zero idiosyncratic product- and market-specific fixed cost shock υ_{jt}. Thus, $F_{jt} = F_b + \upsilon_{jt}$ and $E[\upsilon_{jt}|j \in \mathcal{J}_{ft}] = 0$. Before deciding on its product portfolio, a firm observes the fixed cost shocks related to all products that 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. Specifically, it solves the following maximization problem:

$$\max_{J_{ft} \subseteq \mathcal{J}_{ft}} \left\{ \pi_{ft} = E\left[VP(p)|J_{ft}\right] - \sum_{j \in J_{ft}} F_{jt} \right\}. \tag{3}$$

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

Once firms choose their product portfolios, they decide how to produce the products. This decision can be thought of as simultaneous to the portfolio choice, as no additional information is revealed. At this stage, firms have expectations about e_{jt}, though these shocks are not yet realized and observed.

I model the production function for washing machines as Leontief:

$$q_{jt} = \min\left\{\lambda_f K_{jt}, \psi'_f\left(L_{cft}\right) L_{cjt}, \rho_{jt} M_{jt}\right\}. \tag{4}$$

To produce a particular quantity q_{jt} of washer model j at time t, a firm needs a certain amount of capital K_{jt}, labor L_{jt}, and material M_{jt}. Firms cannot substitute among capital, labor, and materials conditional on the production parameters. Firms can, however, use different mixes of capital, labor, and materials to produce the same quantity of washers because of differences in the parameters λ_f, $\psi'\left(L_{cft}\right)$, and ρ_{jt}.

Firms can differ in the amount of capital used to produce per quantity produced, which is captured by the firm-specific parameter λ_f. One reason for this could be differences in the capital intensity of production (e.g., due to differences in the degree of automation). Another could be differences in the efficiency of capital use.

$\psi'_f(L_{cft})$ is the marginal product of labor. In the empirical application, I estimate the relationship between plant-level employment and output.

Since firms cannot substitute among capital, labor, and materials, marginal costs are additive in the factors of production. Including transportation costs, the marginal costs of product j at time t are

$$mc_{jt} = \frac{r}{\lambda_f} + \frac{w_{ct}}{\psi'(L_{cft})} + sd_{jt} + \frac{z}{\rho_{jt}} + \omega_{jt}\,, \tag{5}$$

where r is the cost of capital, w_{ct} is the cost of labor, z is the cost of material, d_{jt} is the distance between the production location of product j and the destination market, s is the per-unit transportation cost for every unit of distance, and ω_{jt} is a transitory product-level unobserved marginal cost shock realized after the choice of product portfolios and input factors is made. The latter can be thought of as an unanticipated shock to per-unit material costs. Firms are price takers in input markets and cannot affect the local cost of labor, capital, or material.

Finally, the production function for washers allows us to derive the demand of firm f for domestic workers. This can be written as

$$LD_{ft} = \sum_{j \in J_{ft}} \frac{E[q_{jt}]}{\psi'(L_{cft})} \times \mathbb{1}\{c \in \text{domestic}\}\,, \tag{6}$$

where $\mathbb{1}\{c \in \text{domestic}\}$ is an indicator function equal to one if plant c in which product j is produced is a domestic plant. As firms hire workers before the realization of the pair of transitory demand and supply shocks e_{jt}, they do so based on the expected equilibrium quantities in the product market. Wages affect equilibrium quantities through marginal costs and are thereby part of labor demand.

4.3 Third stage: Pricing and demand

In the third 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

$$VP_{ft} = \sum_{j \in J_{ft}} (p_{jt} - mc_{jt}) s_{jt} S_t\,, \tag{7}$$

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

$$p_{jt} = mc_{jt} - [(\nabla_p s \bullet \Lambda)^{-1} s]_{jt}\,, \tag{8}$$

where Λ is the ownership matrix and $\nabla_p s$ is the matrix of partial derivatives of market shares with respect to prices.[28]

Having analyzed the behavior of firms, let us turn to the demand for washers by consumers. Demand is a household-level discrete choice between different 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 chooses to purchase one of the different washing machines on offer in the market or to not purchase a washer at all—i.e., it chooses the outside good. This can be thought of as the household's keeping the washer it already owns or not owning a washer at all.

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

$$u_{ijt} = x_{jt}\beta + \sigma^{FL}\nu_{it}^{FL}x_{jt}^{FL} - (\alpha + \kappa_\alpha \min(\$400\text{k}, z_{it}))p_{jt} + \xi_{jt} + \epsilon_{ijt}\,. \tag{9}$$

The vector x_{jt} includes nonprice product characteristics, such as whether the washer can be loaded from the front or the number of special programs it includes. It also includes indicator variables for the brand and retailer at which the washer was purchased, year fixed effects and brand time trends.[29] p_{jt} is the price of a 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 nonprice characteristics are captured by α and β, respectively. x_{jt}^{FL} is an indicator variable for whether a particular 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 very high-income households that 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 across households. κ_α captures how the

[28] The 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.

[29] The full list of product characteristics includes the price, brand repair rate, total advertising expenditure at brand level, and indicator variables for whether the 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, and has a child lockout. Finally, the vector includes retailer, brand and year fixed effects and linear brand time trends.

sensitivity to prices varies with household income.

The remaining part of the utility function consists of an unobservable component constant across households, ξ_{jt}, and 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 across products but are common across households. Finally, ϵ_{ijt} is an i.i.d. draw from a type I extreme value (Gumbel) distribution.

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 $\theta = (\theta_1, \theta_2)$ that 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 identify utilities only 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 me to derive the familiar logit choice probabilities from this specification. When I integrate 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

$$s_{jt} = \int \frac{exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_{k \in J_t} exp(\delta_{kt} + \mu_{ikt})} P_D(D) P_\nu(\nu) \, . \tag{10}$$

4.4 Connection between labor demand and product market

After seeing the model, one might think that the decision by manufacturers to hire more workers is only distantly related to the product assortment and pricing decisions. To understand why these are fundamentally related, let us think about how a manufacturer and its competitors would react to an exogenous increase in wages in one of its production locations.

Let us assume that a manufacturer produces top-loaders in the U.S. and front-loaders in China, that wages in China increase, and that the manufacturer is the only one producing in China. If wages in China rise sufficiently, the manufacturer may decide to relocate its front-loader production to a different location, provided that the change in expected profits is sufficiently large to justify the relocation costs.

Next, let us assume that this is not the case and analyze the product assortment decision. With an increase in wages, marginal costs of front-loaders increase, and the expected variable profits of front-loader models decrease. It becomes less attractive for the manufacturer to launch front-loaders and relatively more attractive to launch top-loaders produced in the U.S. if these are substitutes. It also becomes more attractive for rivals to launch

front-loaders, as the Chinese front-loaders become less competitive.

Finally, manufacturers set prices. This affects the number of workers hired through the price and quantity expectations in the previous stage. The change in wages affects quantities and workers directly through marginal costs since an increase in prices of Chinese front-loaders leads to substitution away from these. It also affects quantities and workers indirectly through its effect on the product portfolio decisions. Understanding the hiring responses of multinational manufacturers after a labor cost shock in one of its production locations or a merger therefore requires modeling the product market implications.

5 Estimation

In this section, I describe how to estimate the model parameters. To reflect how the estimates build on each other, I proceed with the description in reverse order, beginning with the demand parameters.

5.1 Demand

My estimation of the demand parameters is similar to the approach of Berry, Levinsohn, and Pakes (2004). In a first step, I estimate the nonlinear 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 purchased washer being a front-loader and the share of front-loaders among the second-choice brand. Respondents to the *TraQline* survey are asked only which other brands they considered buying, not which exact model. Some brands have both front-loaders and top-loaders. However, the share of front-loading washers differs greatly among brands. 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 an 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 purchased washer. The correlation coefficient between the two is 0.5.[30] On average, the higher the income of a household, the higher is the price of a purchased washer. This suggests that high-income households are less sensitive to prices.

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

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 nonprice 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 since 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}. To obtain an unbiased estimate of the reaction of quantities to price changes, I need an instrument for price that is unrelated to unobserved demand shocks (satisfying the exogeneity condition) and has a sufficiently strong effect on prices (satisfying the relevance condition). The ideal instrument is a variable that captures differences in product-level marginal costs while being unrelated to demand. The product-level RER, described in Section 2, does just that. It captures product-level marginal cost shocks that depend on the production location of a particular product and fluctuations in the RER between the U.S. and the country of production. These should affect marginal costs but should be unrelated to transitory demand shocks in the U.S. that have different effects on different products. Since I deflate prices to 2012 dollars, I also deflate the RER in the same way.

The estimated market size is an essential input in the demand estimation. This is the fraction of the population for which I observe purchases that was considering buying a washer. The *TraQline* survey includes information on the total number of survey respondents (including those who did not buy a washer). Next, I assume that every seventh household is a potential purchaser of a washer in a particular year. The life expectancy of a washer is approximately ten years. Many consumers consider purchasing a washer before this date, whereas others completely forgo the purchase of a washer. The market size then simply is one-seventh of the number of responding households in a particular year.

For the linear and nonlinear demand parameters, I cluster standard errors at the brand level using the residual bootstrap.

5.2 Marginal costs and production

I compute the 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.

Before attributing product-level marginal costs to different input factors, we need to estimate how the marginal product of labor changes with plant size. I use self-collected data on annual plant output quantities and manufacturing workers employed between 1995 and 2017. Previewing the results reported in Section 6, I observe that the relationship between output and employment is approximately linear and becomes steeper over time. Since I am interested in the employment effects around the time of the merger, I approximate ψ by a linear function and estimate its parameters using data from between 2000 and 2009 only:

$$q_{cft} = \psi(L_{cft}) = \kappa_0 + \kappa_1 L_{cft} + \eta_{cft}\,, \tag{11}$$

where η_{cft} is an error term that occurs because firms hire workers before observing transitory shocks, thus knowing only $E\left[q_{jt}\right]$, not q_{jt} itself. q_{cft} is the total output quantity of firm f at plant c and time t.

The decomposition of marginal costs closely follows the marginal cost function derived in Equation 5. I estimate the parameters of the following equation:

$$mc_{jt} = FE_f + \gamma_1 \frac{RER_{ct}}{\hat{\kappa}_1} + \gamma_2 d_{jt} + \zeta x_{jt} + \omega_{jt}\,. \tag{12}$$

The firm fixed effects FE_f capture firm-level differences in the capital intensity of production. RER_{ct} is the real exchange rate between the U.S. and the location of the production plant, which is also used as an instrumental variable for price and is a proxy for the local wage level. $\hat{\kappa}_1$ is the estimated marginal product of labor for products for which plant-level employment data are available. d_{jt} captures the distance between the plant location and the center of the United States. Finally, the nonprice characteristics vector x_{jt} captures that material costs vary with product characteristics, and ω_{jt} captures transitory material cost shocks. Finally, we can compute the difference between $\widehat{mc}_{jt}$ and mc_{jt} to estimate ω_{jt}.

5.3 Fixed cost bounds

The entry model in Section 4.2 provides inequality conditions for profitable entry, making it impossible 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 determine the set of potential products of each firm. I refer to all products that a firm could have added as potential products, to the potential products that it actually added as active products, and to the potential products that it chose not to add as 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 $\tilde{J}_{ft}$.

The set of active products includes products I observe in the data. Before I determine the set of inactive products, it is worth remembering that the goal is to estimate the fixed costs of adding 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-loaders among its active products, I do not consider that it could have added 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, additional retail floor space or customizing the product for the clientele of a particular retailer but not those of developing new technologies.

The estimation of the fixed cost bounds closely follows the procedure described by Eizenberg (2014). If the observed product entry is a pure strategy SPNE, 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.[31]

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 from offering that product, holding other products fixed. That is,

$$F_{jt} \leq E[VP_{ft}(J^*_{ft}) - VP_{ft}(J^*_{ft} - \mathbf{1}^j_{ft})] \equiv \bar{F}_{jt}\,, \tag{13}$$

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

For each inactive product, a lower bound on the fixed cost is the expected incremental

[31] In principle, I could add further restrictions on fixed cost bounds due to the lack of profitable multistep deviations. In practice, restrictions based on multistep deviations may be difficult to use since the additional inequalities would include idiosyncratic fixed cost shocks υ_{jt} for each product.

profit from offering that product, holding other products fixed. That is,

$$F_{jt} \geq E[VP_{ft}(J^*_{ft} + \mathbf{1}^j_{ft}) - VP_{ft}(J^*_{ft})] \equiv \underline{F}_{jt}\,, \tag{14}$$

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

Constructing the upper bound on F_b based on active products and the lower bound based on inactive products is inadmissible since product portfolio decisions are not independent of υ_{jt}, i.e., $E[\upsilon_{jt}|j \in J_{ft}] \neq 0$. Recall, however, that $E[\upsilon_{jt}|j \in \mathcal{J}_{ft}] = 0$, which means that the product-level fixed cost shock has mean zero conditional on the 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 obtain an unbiased estimate of the bounds on the set of brand-level average fixed costs F_b. I fill in the missing bounds using two further assumptions.

Assumption 2. $\sup_{j\in\mathcal{J}_{bt}} F_{jt} = F^U_{bt} < \infty$ *and* $\inf_{j\in\mathcal{J}_{bt}} F_{jt} = F^L_{bt} > -\infty$ *(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 met in all contexts. If F_{jt} is the cost of developing a new breakthrough technology, it could be that no amount of money in the world would make 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, an upper bound on the necessary fixed costs seems plausible. 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 on the fixed cost support, F^L_b, is an innocuous assumption.

Assumption 3. $[F^L_b, F^U_b] \subset$ *supp(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 the 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 the fixed costs of introducing different potential products of a particular brand come from the same distribution and there exists a blockbuster product that increases the 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 from 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 that

it would never be introduced, then I always observe this product as an inactive product, and the expected change in variable profit from adding this product must be lower than the fixed cost of introducing any potential product.

I fill in the missing lower bound for active products by using the minimum change in firm-level expected variable profits among inactive products of the same brand. I fill in the missing upper bound 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 the fixed costs for active and inactive products are defined as

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

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

$$E[L_{jt}(\theta)] \leq F_{bt} \leq E[U_{jt}(\theta)] \quad \forall j \in \mathcal{J}_{bt}\,. \tag{15}$$

Let us turn to estimating the different components of Equation 15. In Section 7, I am interested in the incentive of brand owners to add products in different acquisition scenarios around the time of the merger. Since I restrict the set of potential products for this exercise to potential products in 2005 (pre-merger) and 2007 (post-merger), I use the same two years to estimate bounds on the fixed cost of adding products. Using data only for these two years for the fixed cost bounds ensures that I do not bias the product portfolio choices just after the merger by basing them on fixed cost estimates for later years.

To estimate the lower bound on fixed costs for inactive products, I replace the population vector θ by the first-stage estimator $\hat{\theta}$, such that $L_{jt}(\theta) = \underline{F}_{jt}(\hat{\theta})$ if $j \in \tilde{J}_{bt}$. Similarly, for active products, $U_{jt}(\theta) = \bar{F}_{jt}(\hat{\theta})$ if $j \in J_{bt}$. I replace the expectation by the sample average across 500 draws from the joint empirical distribution of the demand and marginal cost shocks, e_{jt}. Next, I use $\min_{j\in\tilde{J}_{bt}}\{\underline{F}_{jt}(\hat{\theta})\}$ as an estimator for $V_{bt}^{L}(\theta)$ and $\max_{j\in J_{bt}}\{\bar{F}_{jt}(\hat{\theta})\}$ as an estimator for $V_{bt}^{U}(\theta)$.

The estimated set for every brand is $\left[\bar{l}_n^b(\hat{\theta}), \bar{u}_n^b(\hat{\theta})\right]$, where

$$\bar{l}_n^b(\hat{\theta}) = (1/n^b)\sum_{j=1}^{n^b} L_j(\hat{\theta}), \quad \bar{u}_n^b(\hat{\theta}) = (1/n^b)\sum_{j=1}^{n^b} U_j(\hat{\theta}) \tag{16}$$

and n^b is the number of potential products that belong to brand b. This estimation procedure produces unbiased estimates and overall leads to wide and conservative fixed cost bounds.

Finally, I follow the arguments of Imbens and Manski (2004) and Eizenberg (2014) to construct a $(1-\alpha)\times 100\%$ confidence set for F_b. This is based on the one-sided intervals around the point-identified upper and lower bounds of the fixed cost interval:

$$\left[\bar{l}_n^b(\hat{\theta}) - \frac{S_l(\hat{\theta})}{\sqrt{n^b}} z_{1-\alpha}, \bar{u}_n^b(\hat{\theta}) + \frac{S_u(\hat{\theta})}{\sqrt{n^b}} z_{1-\alpha}\right] \tag{17}$$

where $S_l(\hat{\theta})$ and $S_u(\hat{\theta})$ are estimators of the standard deviation of L_j and U_j. I obtain these by repeating the estimation for the estimated set using the same 200 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.

6 Estimation Results

6.1 Demand

Table 3 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 of a full unit leads to an increase in washer prices of \$196. The F-statistic is approximately 24, suggesting that the IV is relevant.

Column (2) includes the reduced-form estimates after I regress the outcome variable (the average utility that consumers obtain from purchasing washer j at time t, δ_{jt}) on the instrument. The higher the RER, the lower is the purchasing utility for a consumer. In Columns (3) and (4), I report the price coefficient for the simple logit demand model using ordinary least squares (OLS) and the IV, respectively. When I account for the endogeneity of prices, the average product-level own-price elasticity of residual demand changes from -0.96 to -2.19. Finally, I report the price effects for the full mixed logit model using the 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 have an above-average unobserved preference for other front-loaders. Accounting for these effects, I estimate that the average product-level own-price elasticity of demand reduces to -2.89.[32]

[32]These elasticity estimates are comparable in magnitude to the results of 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, based on weekly data.

Table 3: Demand estimates

	(1)	(2)	(3)	(4)	(5)
	First stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	
Linear parameters					
Real exchange rate	1.956***	-0.729**			
	(0.399)	(0.354)			
Price ('00 2012 $)			-0.164**	-0.373*	-0.588***
			(0.062)	(0.192)	(0.023)
Nonlinear parameters					
Income effect κ_α					0.089***
					(0.011)
Unobserved taste σ^{FL}					2.495***
					(0.017)
Characteristics	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Retailer FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand time trends	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,590	1,590	1,586	1,590	1,590
Kleibergen–Paap F-statistic	24.082				
Avg. own-price elasticity			-0.964	-2.186	-2.887

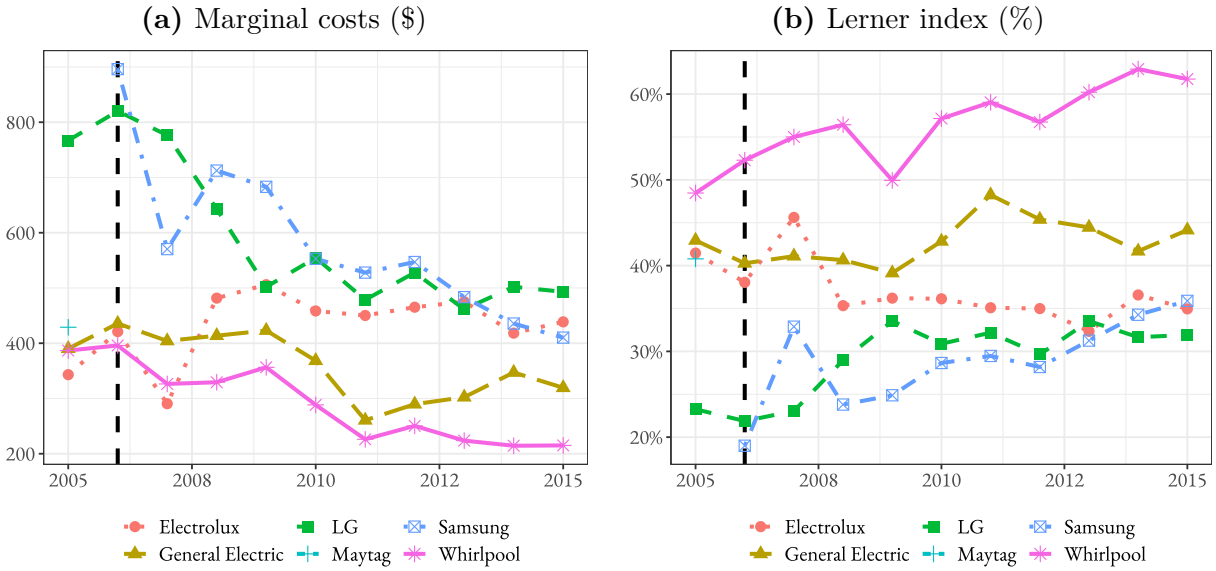
Notes: Column (1) presents results for the first-stage regression of prices on the real exchange rate. Column (2) includes reduced-form estimates for the simple logit model. Column (3) reports demand estimates for the simple logit without a price instrument. Column (4) presents demand estimates for the 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 the simulated method of moments and the linear parameters using linear IV regression. Standard errors are clustered at brand level. The own-price elasticity of residual demand is computed at product level, and the average is calculated by weighting products according to their sales volume. Estimates for nonprice characteristics are reported in Table A.3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6.2 Marginal cost and production

Across brands and time, I find that average marginal costs are approximately \$384 and range between close to zero and approximately \$1,500. The average Lerner index is 45 percent.

Figure 4 shows the evolution of marginal costs and the Lerner index by brand owner. The merger coincides with an industry-wide fall in marginal costs. As we do not expect cost efficiencies from 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 in the absence of the merger. The fall in marginal costs is consistent with the fall in prices in Europe documented in Section 3. At the same time, profit margins increased.[33]

Figure 4: Evolution of marginal cost and Lerner index by brand owner



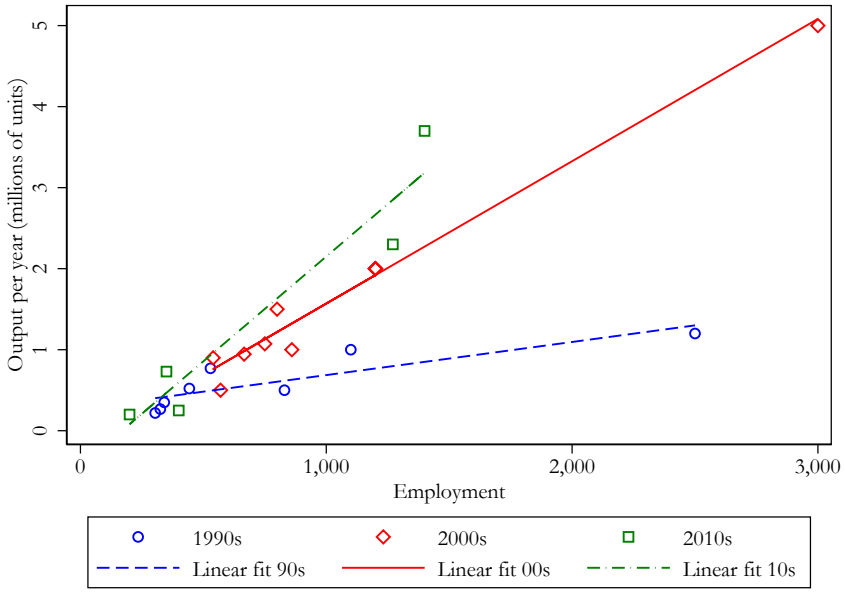
Notes: The plots show the evolution of marginal costs (left) and the Lerner index (markup over price; right) by brand owner over time. The vertical line shows the time of the Maytag acquisition by Whirlpool. For years from 2006 onward, Whirlpool also includes former Maytag products.

Figure 5 shows the relationship between plant-level output and employment. Each data point represents a plant in a particular year for which I have data on employment and output. Two observations stand out: First, within a decade, the relationship between employment and output is approximately linear. A plant needs a certain number of manufacturing workers to operate, after which output increases linearly in the number of workers. Though marginal costs are constant, variable costs decrease with scale. Second, this relationship becomes steeper over time. Thus, technological change increases output per worker over time.

[33]The large movements in the Lerner index for Samsung between 2006 and 2007 should be interpreted with caution since these are based on the relatively few Samsung products available at the time.

To estimate the employment effects of the potential Maytag acquisitions, I use the estimates for the relationship between plant-level output and employment using plant–year observations from the 2000s. The x-intercept is at 107 workers, and each additional worker increases annual output by 1757 units (with a standard error of 86).

Figure 5: Plant-level output and employment



Notes: The figure shows the relationship between plant-level output and employment.

Table 4 shows how marginal costs vary with labor costs, distance between the production location and the U.S., product characteristics, and firm fixed effects.

Although the effect of distance on transportation cost is very noisy, the point estimate lines up with the estimated transportation cost from other sources. Based on the point estimate, shipping a washer from Germany to the U.S. would increase the shipping cost by around $50 (in 2012 prices). This closely aligns with the $50 reported as Whirlpool's cost of shipping front-loaders from its plant in Germany to the U.S. in 2005.[34]

The deflated RER captures country differences in labor costs. For 2006, the deflated RER of China is 0.28, of Mexico 0.64, and of the U.S. 1 (by virtue of its being deflated with U.S. prices). This means that moving production from the US to Mexico would decrease labor costs by $63 per washer and from the U.S. to China by $126. At the same time, relocating production from the U.S. to China would increase transportation costs by $67.

Moving production from the US to China vs Mexico in 2006 would therefore have had very similar effects on marginal costs. Since I find that average cost decreases in plant size,

[34]https://www.nytimes.com/2005/06/17/business/worldbusiness/globalization-its-not-just-wages.html

Table 4: Marginal cost decomposition

	Marginal costs (2012 $)
Real exchange rate	175.152***
	(33.849)
Distance ('000 km)	6.274
	(7.379)
Front-loader	17.484
	(20.139)
Agitator	-243.226***
	(26.731)
Characteristics	*Yes*
Retailer FE	*Yes*
Brand FE	*Yes*
Brand time trends	*Yes*
Year FE	*Yes*
N	1,586

Notes: The table presents results for a regression of product-level marginal costs on proxies for labor and shipping costs as well as product characteristics, brand time trends, and fixed effects.

a potential acquirer's decision to relocate Maytag's production to China or Mexico would depend on where it had large plants that could achieve the necessary scale.

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, recall that a product is defined as the combination of a brand, a retailer and major washing machine 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 5 describes the 95 percent confidence sets on the fixed costs of adding new products. I find that the range of plausible fixed costs of adding products involves higher values for brands with large market shares (e.g., Maytag and Whirlpool) than for brands with lower market shares (e.g., KitchenAid, Hotpoint and Westinghouse). This could be because the former are offered at a retailer only if this involves a full range of washing machines within that product category, requiring more floor space and higher marketing expenditures.

Table 5: Brand-level fixed costs of adding a product ($M)

Brand owner	Brand	95% confidence sets	n^b
Maytag	Admiral	[4.7, 11.2]	10
	Amana	[1.5, 4.6]	15
	Maytag	[4.8, 22.4]	30
Whirlpool	KitchenAid	[0.8, 3.1]	10
	Roper	[1.3, 10.0]	20
	Whirlpool	[12.6, 26.5]	30
General Electric	General Electric	[2.2, 16.9]	30
	Hotpoint	[0.6, 2.2]	15
Electrolux	Frigidaire	[2.3, 10.4]	20
	Westinghouse	[0.5, 2.0]	20
LG	LG	[4.8, 14.3]	20
Samsung	Samsung	[0.7, 5.7]	10

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

Having estimated the demand- and supply-side parameters, 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 changes, an acquisition by Haier is equivalent for consumers to a standalone Maytag remaining in the product market. Differences arise only on the employment side. With marginal cost changes, there are differences on the consumer and employment sides.

7.1 Counterfactual comparisons

For the following analysis, I compare the consumer and employment effects of a Whirlpool vs. Haier acquisition of Maytag under four scenarios. First, I perform a traditional merger simulation, keeping the product portfolios of all firms fixed and changing only ownership and the location of jobs. Second, I repeat this analysis but allow labor and shipping costs to change with the production location, using the cost estimates reported in Table 4. Third, I perform a merger simulation in which I let the firms reoptimize their product portfolios but assume marginal costs do not vary by production location. Fourth, I repeat this analysis and allow costs to change when production is relocated.

In all simulations, I use the parameter estimates for the full 2005–2015 demand and supply estimation. This assumes that preferences for characteristics do not vary over time. A brand time trend lets preferences for different brands vary over time. Year fixed effects on

the demand side allow for differences in the value of the outside good. On the supply side, they allow for differences in cost factors, such as commodity prices, over time.

The traditional merger simulation is akin to the ex ante merger simulations often carried out in merger cases. I use the 2005 (pre-merger) product portfolios of all firms and fix the brand time trends, deflated RER and year fixed effects to their 2005 levels.

When I allow firms to adjust their product portfolios, I fix the market environment (i.e., time trends and fixed effects) to 2007 and define the set of potential products as all the firms' potential products in 2005 and 2007.

The set of potential products of each player consist of its potential products in 2005 (pre-merger) and 2007 (post-merger).[35] Since I observe the acquisition scenario with the highest increase in market power, I observe the state of the world in which the incentive of rivals such as LG or Samsung to add more products at more retailers was greatest. The higher the increase in market power, the lower the intensity of competition becomes, and so the higher are the incentives for rivals to add new products. Thus, any product not added by rivals after the merger is also unlikely to have been added without the merger.

To estimate how U.S. employment differs between the acquisition scenarios, I assume that Haier offshores Maytag's production to China and Whirlpool implements the post-merger production plan for Maytag that I observe in the data, offshoring front-loader production to Mexico and keeping top-loader production in the U.S. Since most manufacturers have a single production location for front-loaders or top-loaders post-merger, a potential front-loader/top-loader not yet introduced to the market would be produced in the same production location.[36]

7.2 Portfolio choice algorithm

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.[37] I also fix the products of players at retailers not among the five major retailers. This results in 135 potential products and 69 exogenously active products (products of nonplayers and products of players at smaller retailers).

A well-known feature of product entry games is that there can be many potential equilibrium product portfolios. One way to identify the set of potential equilibria is to

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

[36] When I see two production locations (e.g., Samsung's production of front-loaders in Mexico and South Korea in 2007), I assume that the new front-loader is produced in both locations. This does not matter for the domestic employment effects as I do not observe any product group made in the U.S. and abroad.

[37] This includes Electrolux, General Electric, LG, Maytag, Samsung, and Whirlpool.

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 an SPNE of the entry game.[38] 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 closely related to the heuristic algorithm of 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. Specifically, if there are multiple equilibria of the post-merger entry game, it is plausible to assume that equilibria closer in the 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 brief, 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 makes 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.[39]

7.3 Effects on consumers, workers, and firms

Table 6 summarizes the effects of Maytag's acquisition by the different acquirers on consumers, workers, and firms. All results are differences in the outcome between an acquisition by Whirlpool versus one by Haier. The first two columns show the results of a traditional merger simulation, in which I fix product portfolios to their pre-merger (2005) state. The final two columns present the results when firms can choose their product portfolios in the two acquisition scenarios. Columns (2) and (4) incorporate cost changes from offshoring due to changes in labor and shipping costs, whereas Columns (1) and (3) do not.

In the traditional merger simulation where product portfolios are fixed and there are no offshoring cost changes, I find that, after a Whirlpool acquisition, average prices increase by 3.3 percent and consumer welfare decreases by $147 million annually. There are 573 more jobs in the U.S. after a Whirlpool acquisition. For these to outweigh the consumer harm of

[38] This 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.

[39] Further technical details and a discussion of these details can be found in Appendix V.A. I report welfare results for alternative fixed cost draws in Appendix Tables A.7 and A.8.

Table 6: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	-1.6 [-3.5, 0.3]	-1.4 [-3.5, 0.8]
Whirlpool	– –	– –	-2.1 [-4.6, 0.4]	-2.7 [-5.2, -0.2]
LG + Samsung	– –	– –	0.3 [-0.8, 1.4]	0.5 [-0.9, 1.8]
Electrolux + GE	– –	– –	0.9 [-1.3, 3.0]	1.0 [-1.0, 3.1]
Overall industry	– –	– –	-2.5 [-6.4, 1.3]	-2.6 [-5.8, 0.6]
Prices and consumer welfare				
Average price	3.3% [1.4%, 5.2%]	3.5% [-1.7%, 8.7%]	2.7% [1.4%, 4.1%]	3.0% [1.4%, 4.6%]
Consumer welfare	$-147M [$-191M, $-104M]	$-156M [$-227M, $-85M]	$-215M [$-260M, $-169M]	$-228M [$-264M, $-191M]
Employment				
Domestic jobs maintained	573 [422, 724]	708 [518, 898]	700 [485, 915]	897 [665, 1129]
Industry profits				
Variable profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$57M [$28M, $86M]	$76M [$52M, $100M]
Total profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$96M [$71M, $121M]	$120M [$93M, $146M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-1.8% [-7.0%, 3.6%]	– –	-1.6% [-2.1%, -1.1%]
Variable profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$-13M [$-54M, $27M]	$5M [$-35M, $44M]
Total profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$32M [$13M, $51M]	$57M [$38M, $76M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$257k [$176k, $338k]	$220k [$117k, $323k]	$315k [$187k, $444k]	$260k [$167k, $352k]
Consumer welfare + total industry profits	$127k [$33k, $221k]	$31k [$-114k, $175k]	$175k [$97k, $253k]	$123k [$75k, $171k]

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

an acquisition by Whirlpool, each job needs to be worth at least $257,000 per year if we consider consumer welfare and jobs only and $127,000 if we also consider industry profits.

Accounting for Maytag labor and shipping cost changes due to different offshoring plans by Whirlpool and Haier exacerbates the difference in the impact of the two potential acquisitions on consumer welfare.[40] This is because a Whirlpool acquisition in this case not only leads to a reduction in competition but also the forgoing of cost reductions from the offshoring of a greater fraction of Maytag's production abroad.[41] At the same time, the difference in the number of jobs maintained in the U.S. grows as well. Overall, if we consider consumer welfare only, each job needs to be worth $220,000 per year for the Whirlpool acquisition to be better for U.S. domestic welfare if we consider consumer welfare and jobs only and $31,000 if we also consider industry profits.

The scenarios in which firms can endogenously adjust their product portfolios allow us to examine the DoJ's claim that if Whirlpool were to try to increase prices, this would increase its rivals' incentive to add products to the U.S. market. Column (3) shows that without any cost effects, an acquisition of Maytag by Whirlpool leads to a reduction in the number of products by the merging parties as they limit cannibalization. At the same time, the merger induces no appreciably greater incentive for rivals, in particular LG and Samsung, to launch new products than they would have under a Haier acquisition or no acquisition. The results look very similar with offshoring cost changes.

Since product portfolio adjustments lead to a reduction in available products after a Whirlpool acquisition, these make a Whirlpool acquisition even worse for consumers. Differences in U.S. manufacturing jobs increase to 700 (without cost changes) and 897 (with cost changes). This shows that the product assortment decision has an economically significant impact on the employment effects. To offset consumer harm, each protected job needs to be worth at least $315,000 per year (without cost changes) or $260,000 (with cost changes).

Four conclusions can be drawn from these results. First, a Whirlpool acquisition is always worse for consumers than an Haier acquisition: it leads to a larger increase in market power, little additional entry by rivals, and fewer products overall. This does not change if we consider the sum of consumer welfare and total profits as the relevant outcome.

Second, offshoring cost changes exacerbate the welfare differences between a Maytag acquisition by Whirlpool and one by Haier. Accounting for labor and shipping cost changes

[40] Since the distance coefficient in Table 4 is imprecisely estimated because of the limited variation in the distance variable, I fix this coefficient to the point estimate in Table 4 for all bootstrap draws. In Appendix Table A.4, I repeat the analysis without this restriction. The only implication is wider confidence intervals on the estimates involving offshoring cost changes.

[41] Comparing the first two columns of Table 6 does not allow us to assess the total effect of offshoring cost changes as these are present in both merger scenarios. Appendix Tables A.5 and A.6 compare acquisitions of Maytag by Whirlpool and Haier relative to no acquisition with and without offshoring cost changes.

due to a change in production location leads to larger volume-weighted marginal cost reductions for Maytag products after a Whirlpool than after a Haier acquisition. However, pass-through is higher after an acquisition by Haier, which means that consumers benefit more from these cost savings than after a Whirlpool acquisition. U.S. workers prefer cost decreases abroad accruing to Whirlpool instead of Haier, since the latter use these to compete with U.S. production more intensly.

Third, accounting for endogenous product portfolio adjustments does not make the Whirlpool acquisition look more favorable to U.S. consumers than it does under the traditional merger simulation. In fact, since Maytag and Whirlpool remove products from the market that consumers value while the incentive for rivals to add new products does not change, this makes a Whirlpool acquisition even worse for U.S. consumers. My results therefore align with the findings of Fan and Yang (2020) and Fan and Yang (2021) and contrast with those of Wollmann (2018) in other settings.

Fourth, there is a significant difference between the two potential acquirers in the number of jobs maintained in the U.S. Based on the changes in consumer welfare, industry profits, and domestic employment, I can calculate the value per job required to offset the consumer welfare loss. Depending on the scenario, the point estimates range between $\$220,000$ and $\$315,000$ per year. On their own, these figures are not sufficient to assess whether the employment effects could have plausibly offset the product market effects. Ideally, we would like to know the value of these jobs to the local labor markets. The descriptive analysis in Section 3 indicates that these jobs are valuable, but the exact value remains unclear. Simply analyzing average wages is insufficient, too, as these jobs likely have a multiplier effect. Instead, I compare the estimates to prior results in the literature: first, estimates of the value of jobs created by foreign multinationals to local labor markets, and second, the job values necessary to offset the consumer effects 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. The literature on earnings losses of displaced workers after mass layoffs suggests that these effects are highly persistent. Jacobson, LaLonde, and Sullivan (1993) find short-term earnings reductions for displaced workers of approximately 40 percent and long-term losses of 25 percent using data from Pennsylvania in the 1970s and 1980s. Couch and Placzek (2010) reproduce this analysis using data from Connecticut for 1993–2004 and find short-term earnings losses of approximately 32 percent and long-term losses of approximately 15 percent. Schmieder, Wachter, and Heining (2023) find reductions in annual earnings of approximately 15 percent lasting a minimum of 15 years after a mass layoff during a recession. They document this

using a detailed linked employer–employee panel spanning 1975–2009 for Germany.

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 to offset the losses to U.S. consumers from the 2011 safeguard tariffs on tire imports from China, each job protected would have had to be worth at least $\$900,000$ per year. Even more relevant, Flaaen, Hortaçsu, and Tintelnot (2020) estimate that to offset the losses to consumers from the 2018 U.S. global safeguard tariffs on washing machine imports, each of the domestic jobs created 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 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. Although this is toward the lower end of the job values necessary to offset consumer harm that I estimate, the order of magnitude is similar.

7.4 Unequal distribution of welfare effects

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

For the U.S. economy as a whole, 1,000 additional washing machine 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 employment at the county level even two years after the plant closure. As illustrated in Figure 1, washer manufacturing plants are concentrated in a few counties in Illinois, Iowa, Michigan, Ohio, and South Carolina.

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 the employment effects are not concentrated among the very wealthy, then the acquisition by Whirlpool may be better for the sum of workers and consumers even if the increase in the total domestic wage bill relative to that under a Haier acquisition is lower than the relative consumer welfare loss. Furthermore, consideration of other nonwage outcomes related to job loss, such as mental or physical health, could improve the domestic welfare effects of a Whirlpool relative to a Haier acquisition.

8 Conclusion

I propose a framework to quantify the trade-off that policy choices can have for workers and consumers. To consider how policies 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. I find that capturing the interactions between the product market and labor demand is key to estimating the employment effects of the policy choice.

Empirically, I focus on how foreign competition changes the consumer welfare and domestic employment effects of a merger. I find that Whirlpool's acquisition of Maytag led to a large decrease in consumer welfare. The incentive for rivals to launch new products in response to the merger was small and did not mitigate the associated consumer harm. Whereas Whirlpool's acquisition of Maytag led to more consumer harm than a potential acquisition of Maytag by Haier would have produced, the latter would have led to a larger decrease in domestic employment. Finally, I calculate the average value that a job would have to be worth such that the benefits to domestic workers offset the consumer harm. Although these estimates vary between scenarios the harm to U.S. consumers and the benefits to U.S. workers are of a similar order of magnitude.

Blocking acquisitions that can lead to offshoring or allowing anticompetitive mergers that may protect domestic jobs is unlikely to be optimal. Policymakers could, however, use the framework in this paper to quantify the trade-off between workers and consumers and identify instances in which employment effects are of first-order importance. A procompetitive merger decision could then be complemented by more targeted labor market policies such as the Trade Adjustment Assistance Program (see Hyman, Kovak, and Leive, 2023). Pairing these policies could help prevent opposition to procompetitive merger enforcement.

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

Arkolakis, Costas, Fabian Eckert, and Rowan Shi. 2023. "Combinatorial Discrete Choice: A Quantitative Model of Multinational Location Decisions." *NBER Working Paper 31877.*

Armington, Elizabeth, Eric Emch, and Ken Heyer. 2006. "The Year in Review: Economics at the Antitrust Division, 2005-2006." *Review of Industrial Organization* 29 (4): 305–326.

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.

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.

Couch, Kenneth A., and Dana W. Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *The American Economic Review* 100 (1): 572–589.

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.

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

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

Garrido, Francisco. 2020. "Mergers Between Multi-Product Firms With Endogenous Variety: Theory and an Application to the RTE-Cereal Industry." *mimeo.*

Gathmann, Christina, Ines Helm, and Uta Schönberg. 2020. "Spillover Effects of Mass Layoffs." *Journal of the European Economic Association* 18 (1): 427–468.

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.*

Hyman, Benjamin, Brian Kovak, and Adam Leive. 2023. "Wage Insurance for Displaced Workers." *Working Paper.*

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.

Muendler, Marc-Andreas, and Sascha O. Becker. 2010. "Margins of Multinational Labor Substitution." *The American Economic Review* 100 (5): 1999–2030.

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.*

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

Schmieder, Johannes, Till von Wachter, and Joerg Heining. 2023. "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany." *The American Economic Review* 113 (5): 1208–1254.

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.

Tintelnot, Felix. 2017. "Global Production With Export Platforms." *The Quarterly Journal of Economics* 132 (1): 157–209.

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

Appendix for Online Publication

I Appendix to Section 2: Details on Data Set Construction

I.A Product market data set

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

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

Figure A.1 illustrates the difference between a front-loader and a top-loader. Where 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 washer), is more water and energy efficient, cleans better, and is usually more expensive than a top-loader. 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.[42]

Within a market (here, national at the yearly level), I group responses that are the same along these three dimensions.[43] In doing so, I end up with $2,939$ products between 2005 and 2015. Under this product definition, the market share for many products is 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" and all products with a volume share of the 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 washing machine sales in the *TraQline* data. Dropping very rare products should therefore not bias the estimation results.

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

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

[43] I classify Maytag products as belonging to Whirlpool for the first quarter of 2006, when the acquisition had not yet been carried out. This is to avoid artificially inflating the number of washing machine products in that year. In addition, since the merger talks were public since mid-2005, it seems unlikely that Maytag and Whirlpool would still be competing heavily in the first quarter of 2006.

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

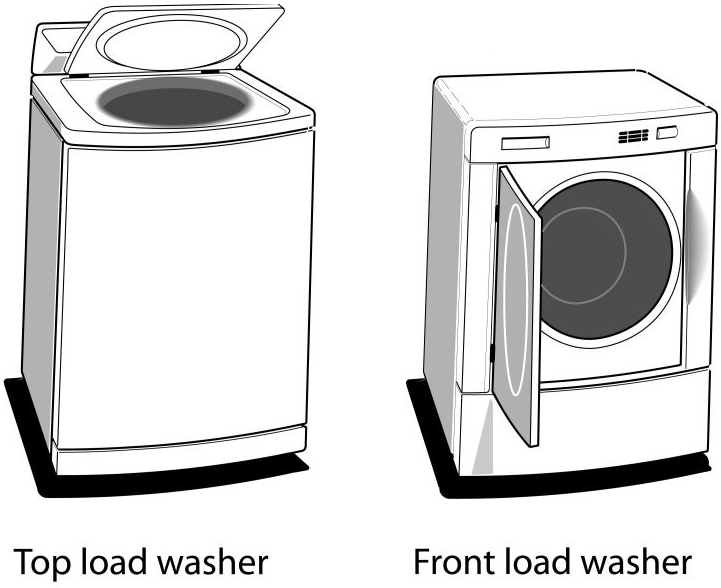
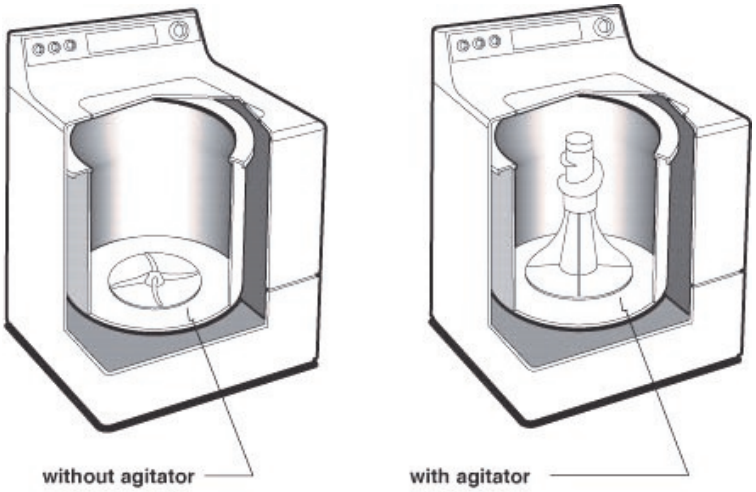


Figure A.2: Top-loaders with and without agitator



include whether a 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; and the number of special programs it has.

Household income. Whereas the CPS data include the exact income of the sampled households, the *TraQline* data include only 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 midpoint of the nonoverlapping 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 washing machines for the U.S. market involves three steps: First, I use information from various sources, such as annual reports, news articles and transcripts from the USITC's antidumping hearing on imports of large residential washing machines from Mexico and South Korea to identify the location of major manufacturers' washing machine plants.

For LG and Samsung, the production locations before 2012 are mostly based on the USITC investigation. For 2012–2015, the production locations for LG and Samsung are based on firm-level 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 its plant locations on a combination of documents from the USITC investigation and news reports. Finally, to make sure that the plants

produce washing machines for the U.S. market, I check the plant locations against import data split by top-loading and front-loading washers at the country level from the USITC.

In many cases, this information is insufficient for me to identify whether a plant produces washing machines for the U.S. or for another market. Second, I use information on the general imports of front-loader and top-loader washing machines to the U.S. split by source country over time. I use these data to eliminate any plants that cannot plausibly produce substantive volumes for the U.S. market. Finally, I use the 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 washing machines.

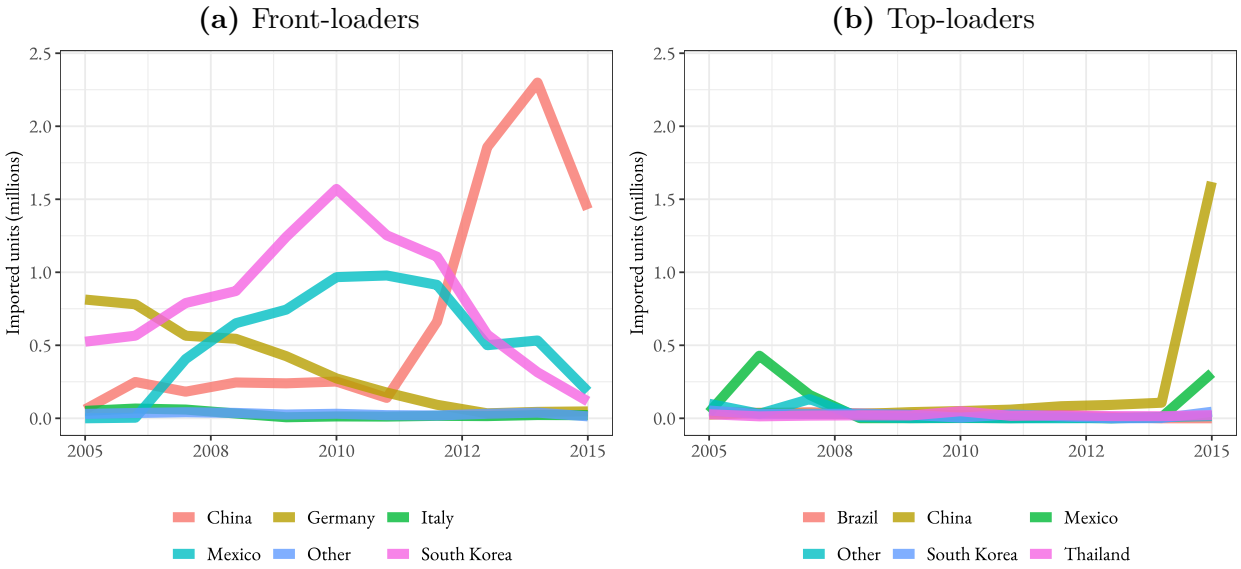
Figure A.3 shows the evolution of annual imports of front-loaders and top-loaders into the U.S., split by source country. Across the sample period, more than half of the front-loaders sold in the U.S. were imported. In 2005, Germany was the largest exporter of front-loaders into the United States. These were produced not 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 their front-loaders from South Korea and, like other manufacturers such as General Electric or Whirlpool, also from Mexico. After the imposition of antidumping duties on large residential washing machines 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 began increasing their sales of top-loaders in the U.S. and imported most of these from China.

For reference, according to Appliance Portrait (2006), 9.3 million washing machines were sold across the U.S. in 2005. Of those, according to the *TraQline* data, approximately one-third were 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 approximately 25 percent in 2015. This suggests that although substantial numbers of front-loaders were imported into the U.S. throughout the sample period, most top-loaders were produced domestically.

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

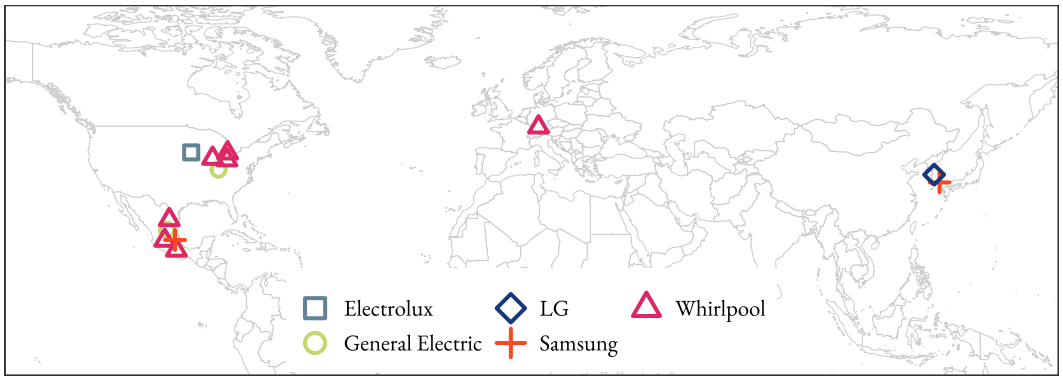
Plant location weights. Finally, Table A.1 summarizes the plant location weights used to calculate the average real exchange rate for each product. Based on the plant locations, the aggregate USITC import data shown above, and the firm-level washer imports for 2012–2015 based on PIERS bill of landing data and reported in Flaaen, Hortaçsu, and Tintelnot

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



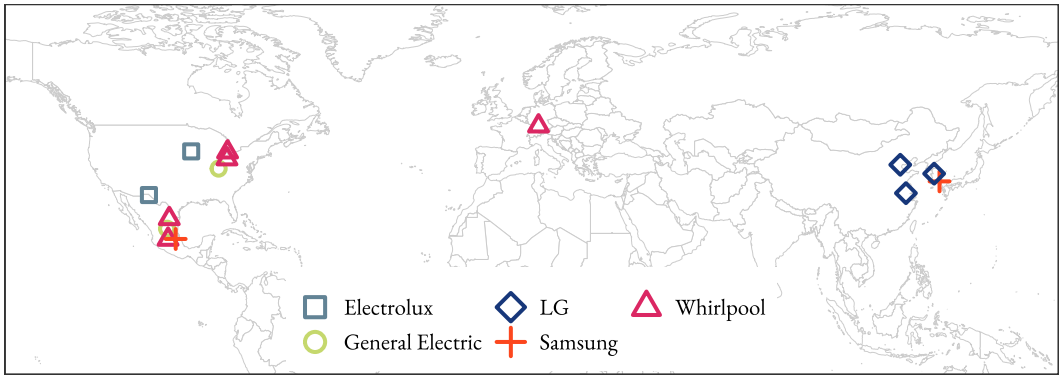
Notes: The left panel plots the annual general imports in terms of volume of front-loader washing machines (HS8450110080, HS8450200080, HS8450200090) imported into the U.S. by source country. The right panel plots the annual general imports in terms of volume of top-loader washing machines (HS8450110040, 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 come from the United States International Trade Commission.

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



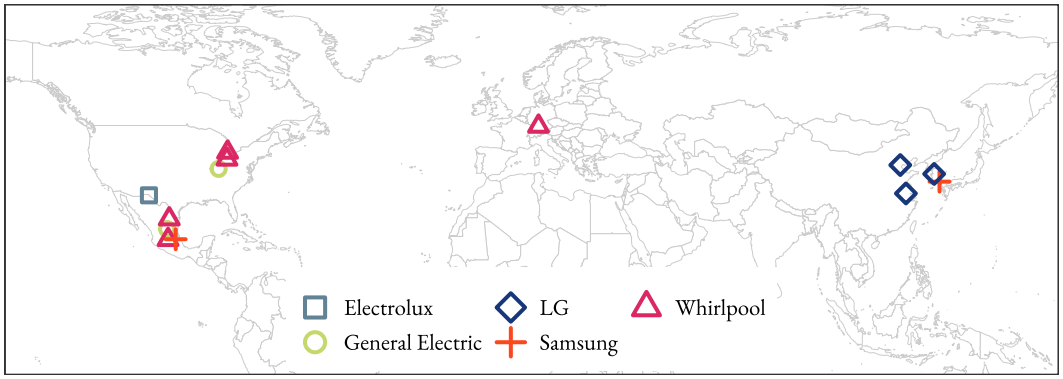
Notes: The map shows all plants manufacturing washing machines for the U.S. market in 2007 by manufacturers with a market share of more than 3 percent in any year in the sample.

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



Notes: The map shows all plants manufacturing washing machines 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: Washer plants manufacturing for the U.S. market, 2011



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

(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

Plant-level data on output and employment. Information on output and employment for the Whirlpool plants in Clyde, Ohio, and Schorndorf, Germany—both producing for the U.S. market—as well as Radomsko, Poland—producing for the European market—are based on news reports. Information on output and employment for the Whirlpool plants in Amiens, France, and Poprad, Slovakia, are based on Ferencikova (2002) and Rubens, Ferencikova, and Bardy (2019), which are case studies on the relocation of washing machine production by Whirlpool from France to Slovakia. The authors report that they cooperated with Whirlpool Slovakia to obtain data on employment and output.

Information on output and employment for the BSH plants in Berlin, Germany, and Alcalá, Spain, are based on historical internal documents from BSH, such as local employee newsletters, which I obtained from the BSH company archive.

I.C Details on the instrumental variable for price

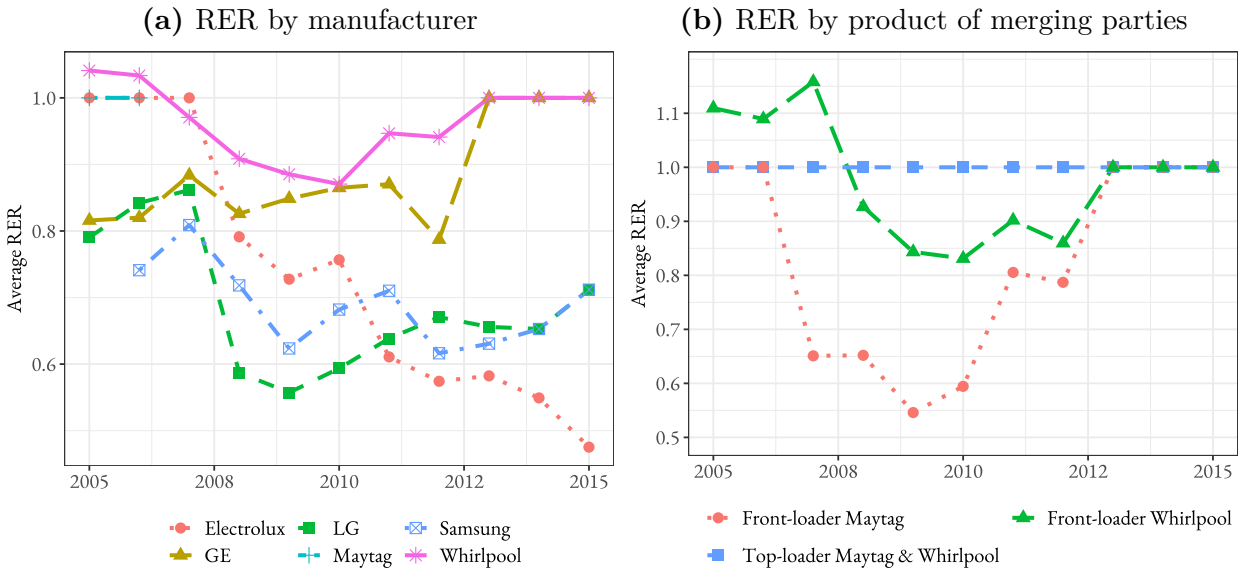
Figure A.7 shows the evolution of the deflated average RER over time and illustrates the source of the variation. The left panel plots the deflated average RER of all the production locations for a particular manufacturer. The deflated average RER is based on the country-level RER of the different plant 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. Since prices are converted to 2012 U.S. dollars, I also use the U.S. GDP deflator to deflate all RERs into 2012 terms. Although this measure masks within-manufacturer variation in the RER, there is still significant variation at this level. In the right panel, I disentangle the average RER for Whirlpool and Maytag products.[44] 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.

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 part because of its lack of international production.[45] In a similar

[44]Maytag products include all products marketed under the brands owned by Maytag pre-acquisition (i.e., Admiral, Amana, Magic Chef and Maytag), and Whirlpool products include all other brands owned by Whirlpool.

[45]This 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

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 products include all products marketed under the brands owned by Maytag pre-acquisition (i.e., Admiral, Amana, Magic Chef and Maytag), and Whirlpool products include all other brands owned by Whirlpool.

spirit, Electrolux launched its global cost-cutting program in 2004, with the aim of offshoring more than half of its production to low-cost countries by 2009 (Electrolux, 2007).[46] Both firms exclusively served the U.S. washing machine 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 and changes in the production locations.

II Appendix to Section 3: Additional Descriptive Results

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

To make my descriptive results comparable to those of the analysis by Ashenfelter, Hosken, and Weinberg (2013) and provide evidence on dryers, I repeat the analysis using freestanding ranges in the U.S. as a control market.[47] This is an appropriate control group if, in the absence of the merger, prices would have evolved similarly in the treatment and control groups.

I begin by estimating the event study plot in Equation 1 with freestanding ranges as the comparison group. The results are shown in Figure A.8. There are two key takeaways. First, the evolution of prices for washers and dryers is very similar. Second, there are no price increases for washers and dryers relative to freestanding ranges.

Next, I estimate differences in the evolution of prices 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):

$$log(p_{it}) = \alpha_1 \text{Maytag}_{it} \times \text{post}_t + \alpha_2 \text{Whirlpool}_{it} \times \text{post}_t + \beta x_{it} + \tau_i + \gamma_t + \epsilon_{it} \,. \tag{18}$$

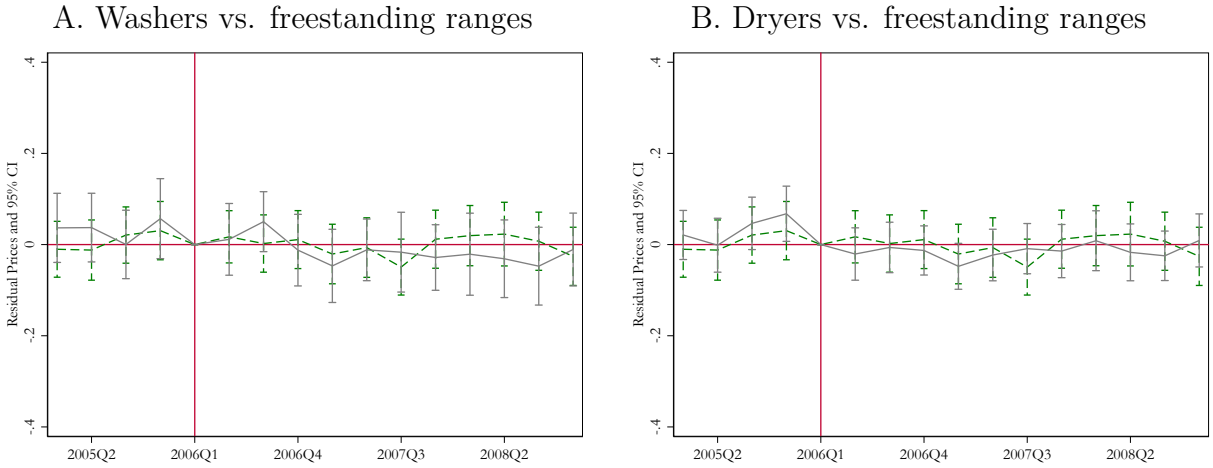
The parameters of interest are α_1, which captures the average price increase for Maytag products, and α_2, which captures the average price increase for Whirlpool products. Like Ashenfelter, Hosken, and Weinberg (2013), I restrict the observation period to Q2 of 2005

are responding" (Maytag, 2005, p. 3).

[46] By 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.

[47] Ashenfelter, Hosken, and Weinberg (2013) use ranges, cooktops, ovens and freezers as the comparison group.

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



Notes: The figure shows the residualized logarithm of prices for Maytag and Whirlpool washers and dryers in the United States relative to the prices of freestanding ranges from other manufacturers in the United States unaffected by the merger. The solid line shows the evolution of the residualized prices in the United States and the dashed line that for the control group. Confidence bounds are at the 95 percent level. Standard errors are clustered at model level.

and Q3 of 2008. In contrast to Ashenfelter, Hosken, and Weinberg (2013), I have only quarterly instead of monthly data.

Table A.2 shows the average difference in the evolution of prices for Maytag and Whirlpool washers and dryers relative to that of freestanding ranges from competitors. Columns (1) and (4) include estimates from a regression pooling Maytag and Whirlpool products. These results suggest that there was no price increase for washers or dryers. Columns (2) and (5) disaggregate the effects by Maytag and Whirlpool products. Columns (3) and (6) repeat the previous analysis but swap the brand fixed effects for more granular product fixed effects. The results suggest that Maytag and Whirlpool washing machines and dryers did not increase prices or even decreased prices relative to freestanding ranges.

When comparing washing machines to ranges, Ashenfelter, Hosken, and Weinberg (2013) find a decrease in prices for old and new Whirlpool washers and a zero effect for the price of old and new Maytag washers. Although it is not possible to fully reconcile the different results as the data sources are different and the *TraQline* data do not allow me to control for product age and differentiate old and new products, the results are broadly in line.

For dryers, Ashenfelter, Hosken, and Weinberg (2013) find a price increase for new Whirlpool dryers, a price decrease for old Whirlpool dryers, no price effect for old Maytag dryers and a small price effect for new Maytag dryers. Depending on the relative importance of new vs. old dryers, these results may be consistent with each other.

A causal interpretation of these results could lead to the conclusion that Whirlpool's

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

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

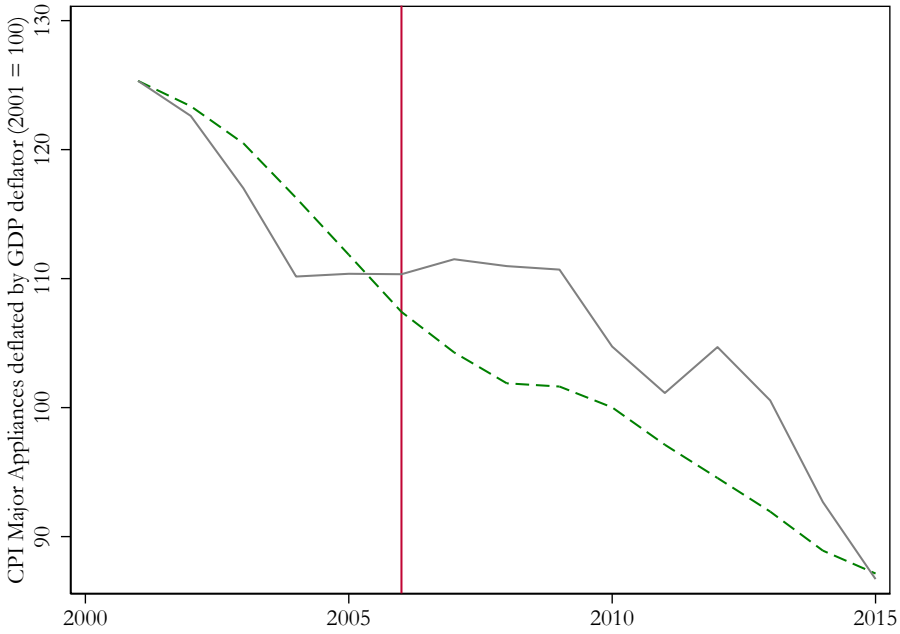
Notes: Columns (1) to (3) compare the logarithm of prices for washers and freestanding ranges. Columns (4) to (6) compare the logarithm of prices for dryers and freestanding ranges. The differences in observations in Columns (3) and (6) from the preceding columns are due to the iterative dropping of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

acquisition of Maytag did not increase prices for laundry products. This is unlikely to be true, however, as comparing washer prices in the U.S. to those in other geographic markets very similar to the U.S. except for the merger indicate that there should have been a strong decrease in prices in the U.S. in the absence of the merger. It thus seems more likely that the market for freestanding ranges is affected by different industry trends than the market for laundry products.

Since pre-treatment price data only starts in 2005, just before the merger, one might still wonder how good of a control group washer prices in the European Union are for washer prices in the U.S. Although there is no consumer price index (CPI) available for washing machines, I can compare the evolution of the CPI for major appliances in the U.S. to the CPI for household appliances in the European Union.

Figure A.9 shows the results of this comparison. In line with the descriptive evidence in Section 3 and the anecdotal evidence, deflated prices in the U.S. and EU are falling before Whirlpool's acquisition of Maytag. In 2005, when negotiations between Whirlpool and Maytag began, these price decreases stop for the U.S., whereas prices in the EU continue to fall. After 2010, deflated prices start falling again in the U.S.

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



Notes: The figure shows the consumer price indices for major appliances in the U.S. and EU between 2001 and 2015. The solid line shows the evolution in the United States and the dashed line that for the European Union. The CPIs are deflated by the U.S. GDP deflator. Both series are indexed such that 100 corresponds to the year 2001.

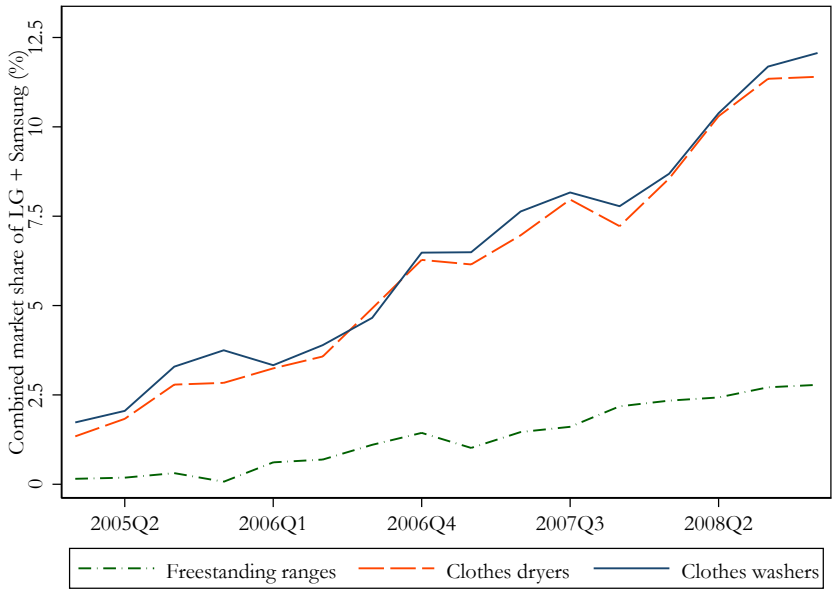
II.B Additional descriptive results on product entry

Figure A.10 shows the evolution of LG and Samsung's market shares for washers, dryers, and freestanding ranges in the United States. Whereas LG and Samsung increased their combined market share for washers and dryers from approximately 2 percent to approximately 12 percent between 2005 and 2008, they had a market share of only approximately 2 percent for freestanding ranges at the end of 2008.

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

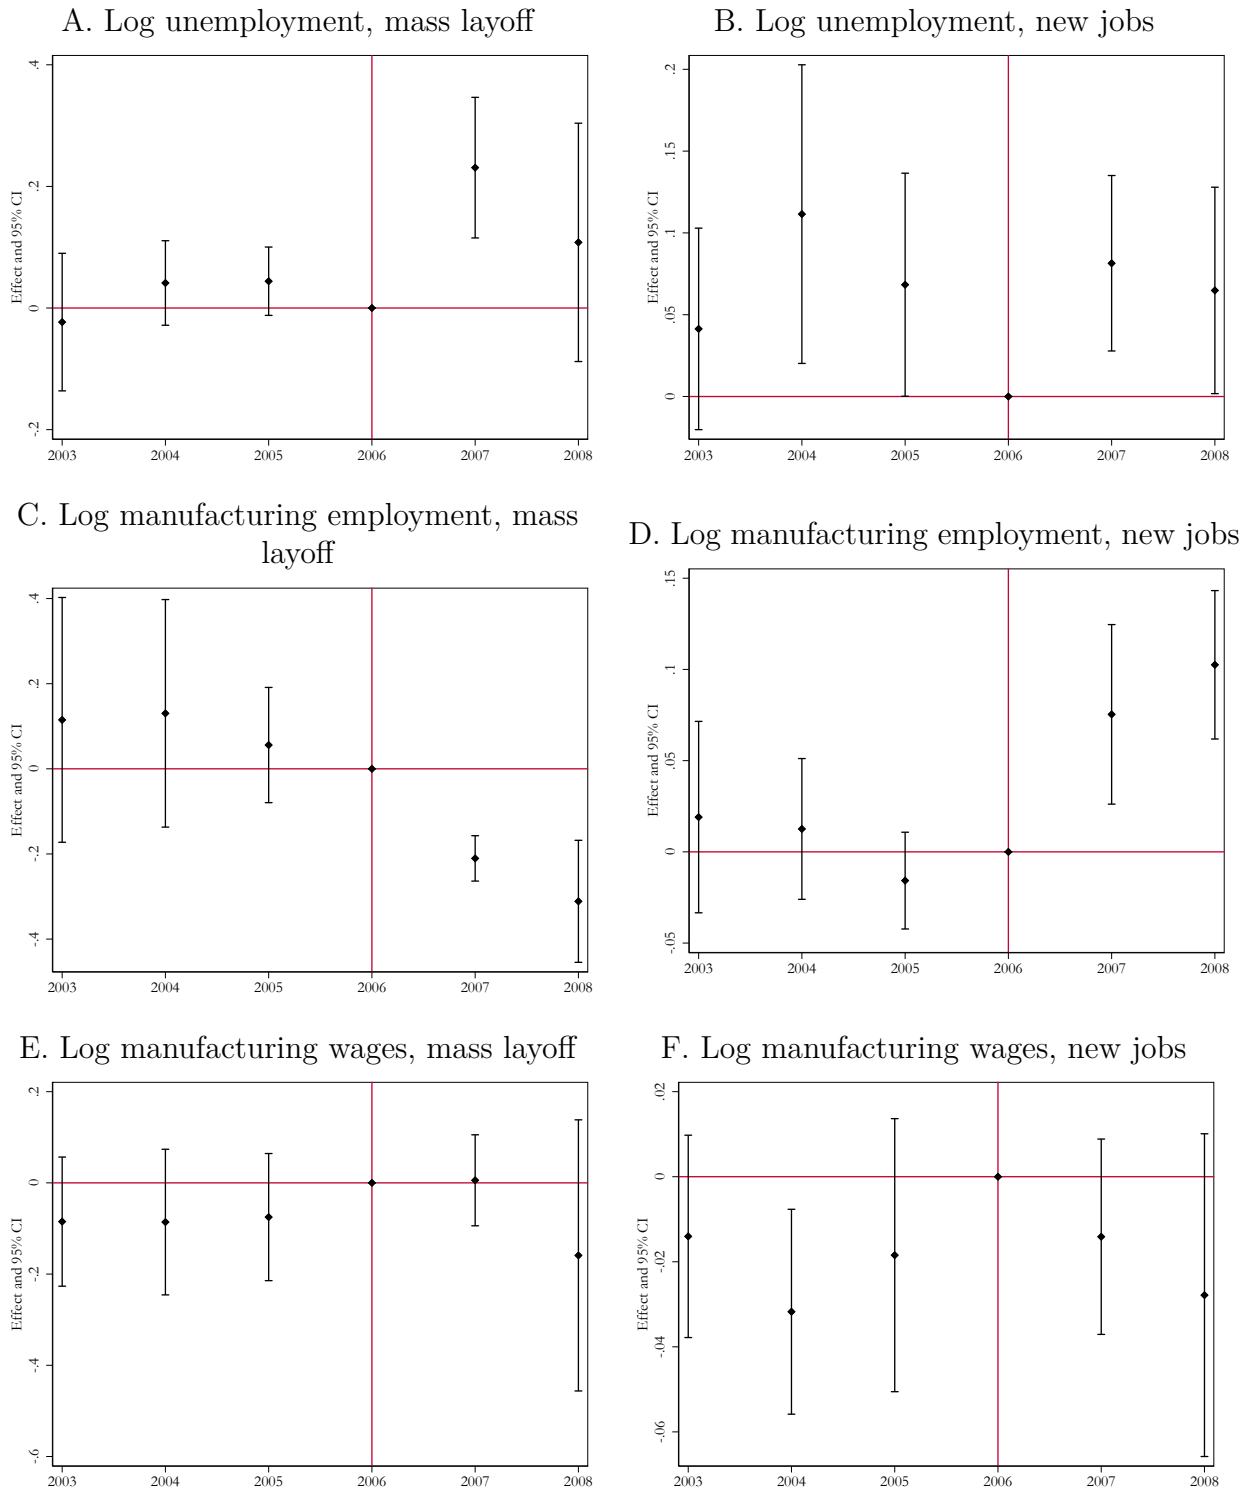
Figure A.11 shows event study plots for the effect of the mass layoffs and new jobs post-merger on unemployment, manufacturing employment, and manufacturing wages at county level. The key outcomes for which I find treatment effects in the main analysis, unemployment and manufacturing employment for the mass layoffs, and manufacturing employment for the new jobs, do not exhibit any pretrends. This indicates the robustness of the estimates from the main analysis.

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



Notes: The figure shows the evolution of the combined market share in terms of sales volumes in the U.S. and key European control markets between 2005 and 2008.

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



Notes: The figure shows the yearly change in labor market outcomes in counties that experienced a mass layoff or new jobs relocated to a Whirlpool plant in the county relative to the change in the matched control group. Panels A, C, and E depict the effect of the closure of Maytag plants and the Maytag HQ. Panels B, D, and F depict the effect of new jobs being relocated to existing Whirlpool plants. Standard errors are clustered at county level.

III Appendix to Section 5: Details on Estimation Procedures

III.A Details on estimating product characteristics for potential products

Potential products are all the products that brand owners added to the market (active products) and all the 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 firms' decision 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 are not among 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.[48]

Product characteristics can vary mildly between retailers. That is, Amana top-loaders sold at Best Buy might modestly differ in their characteristics from Amana top-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 those of that sold 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 the 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

[48] The major retailers are Best Buy, H.H. Gregg, Home Depot, Lowe's and Sears.

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

The first set of moments equates the observed market shares in the data with the simulated market shares from the demand model. To obtain 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.[49] Second, I use the fixed-point formulation from Berry, Levinsohn, and Pakes (1995) to estimate $\hat{\delta}_j$. I use the SQUAREM described in Reynaerts, Varadha, and Nash (2012) to accelerate the convergence of the fixed-point iterations. As convergence is not guaranteed, whenever it fails, I revert to the contraction mapping in Berry, Levinsohn, and Pakes (1995), which has guaranteed convergence. Finally, I speed up the inversion of market shares by using the reformulation of the contraction mapping in terms of consumer-specific choice probabilities for the outside option, described by Brunner et al. (2020).

To 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 10. 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 washer and the average income of households purchasing the product. I compute the moment

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

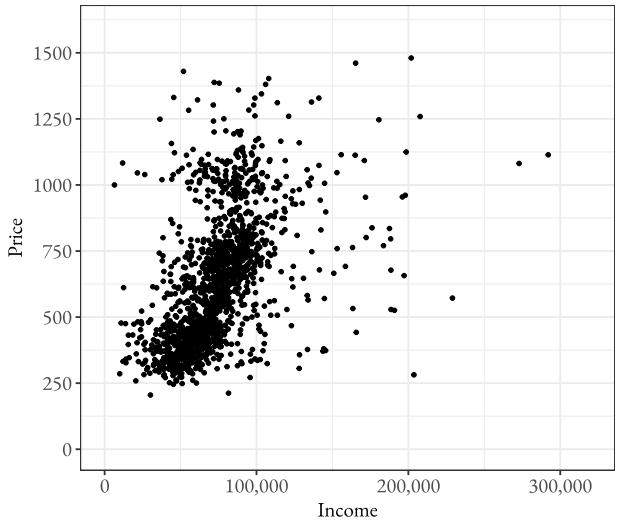
as follows:

$$\sum_j \frac{n_j}{n} p_j \left\{ \left(\frac{1}{n_j} \sum_{i s.t. y_i^1 = j} z_i \right) - E\left[z | y^1 = j, \theta \right] \right\}, \tag{19}$$

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 denotes the income of household i.

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

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



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

The third set of moments fits the covariance between whether the first-choice 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, only the second-choice brand. In particular, I use the following moment condition:

$$\sum_j \left(\frac{n_j}{n} x_j^{FL} \sum_{b' \neq b_j} x_{b'}^{FL} \left\{ \frac{n_{jb'}}{n_j} - E\left[\mathbb{1}\left(b^2 = b' | y^1 = j, \theta \right) \right] \right\} \right), \tag{20}$$

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 19 and 20. Since there are two nonlinear parameters and two

moment conditions, the parameters are just-identified, and I estimate $\hat{\theta}_2$ using the method of simulated moments. I therefore estimate

$$\hat{\theta}_{2,MSM} = \operatorname{argmin} \hat{m}\left(\theta_2\right)' \hat{m}\left(\theta_2\right) . \tag{21}$$

Solving the minimization problem above allows me to recover not only 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:

$$\hat{\delta}_j = x_j\beta - \alpha p_j + \xi_j . \tag{22}$$

As explained in Section 5, I assume that the nonprice 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.

III.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 washer in a particular year. According to Consumer Reports, in 2009, the average life expectancy of a washing machine was ten years. Many households will consider buying a washer before the end of the life expectancy of their washer, e.g., to obtain 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 in the market for washing machines in a particular year. There are two alternative estimation methods: I can take the total number of U.S. households in a particular year and assume that the market size is one-seventh the number of these households. Alternatively, I can use estimates of the annual total washing machines 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 approximately 15 million households.

III.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.[50]

Computing equilibrium prices. Each draw of the second-stage marginal cost and demand shocks e_{jt} requires re-estimation of 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 recompute equilibrium prices 500 times for each product portfolio. Speeding up this process is therefore crucial. Furthermore, not all methods to recompute equilibrium prices necessarily converge.

Morrow and Skerlos (2011) compare different numerical methods to recompute 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 compute equilibrium prices by using fixed-point iteration on the ζ-markup equation.

Drawing e_{jt}. The heuristic algorithm to choose product portfolios requires comparison of the expected profits of the current product portfolio to the expected profits of any product portfolio 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 in computing the expected profits of different product portfolios. In terms of economics, this is desirable because there is no good reason 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 need to compute the expected profits of all firms for a given set of product portfolios only once. Every time the algorithm revisits the particular set of product portfolios, I can reuse the memorized expected profits and do not need to recompute equilibrium prices and expected profits.

[50] As 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.

IV Appendix to Section 6: Further Results of the Structural Estimation

IV.A Demand estimation

Table A.3: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)	(5)
	First stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Real exchange rate	1.956***	-0.729**			
	(0.399)	(0.354)			
Price ('00 2012 $)			-0.164**	-0.373*	-0.588***
			(0.062)	(0.192)	(0.023)
Front-loader	0.183	0.273	0.358	0.342	-0.711***
	(0.207)	(0.266)	(0.244)	(0.218)	(0.021)
Korean	-0.575***	1.737***	1.569***	1.523***	1.495***
front-loader	(0.181)	(0.353)	(0.349)	(0.346)	(0.011)
Fisher & Paykel	-4.514***	-0.636	-1.455***	-2.319***	-2.963***
front-loader	(0.332)	(0.410)	(0.480)	(0.834)	(0.088)
European high-end	0.096	1.217***	1.192***	1.253*	1.242***
front-loader	(1.272)	(0.320)	(0.438)	(0.667)	(0.025)
Agitator	-2.511***	0.952***	0.540**	0.016	-0.334***
	(0.276)	(0.270)	(0.252)	(0.513)	(0.056)
Stacked pair	0.492*	-0.225	-0.147	-0.042	0.004
	(0.279)	(0.149)	(0.149)	(0.192)	(0.010)
Stainless steel	0.476	-0.053	0.009	0.124	0.153***
exterior	(0.602)	(0.248)	(0.270)	(0.342)	(0.010)
White exterior	-0.286	0.675***	0.624***	0.568***	0.523***
	(0.360)	(0.131)	(0.101)	(0.125)	(0.009)
Energy Star	0.023	0.089	0.092	0.098	0.113***
	(0.182)	(0.125)	(0.126)	(0.135)	(0.004)
Extra noise	0.399*	0.247**	0.312**	0.395**	0.451***
insulation	(0.207)	(0.125)	(0.120)	(0.156)	(0.009)

continued

Table A.3: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)	(5)
	First stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Number of special programs	0.008	0.051	0.052	0.054	0.051***
	(0.058)	(0.035)	(0.039)	(0.045)	(0.001)
Child lockout	-0.074	0.206	0.200	0.178	0.181***
	(0.163)	(0.172)	(0.167)	(0.170)	(0.004)
Repair rate	-2.433	2.057	1.627	1.150	0.839***
	(3.156)	(3.263)	(2.957)	(2.788)	(0.131)
Total advertising expenditure	-0.006	0.004	0.003	0.002	0.001***
	(0.005)	(0.002)	(0.002)	(0.002)	(0.0002)
Retailer Best Buy	-0.098	-1.045***	-1.062***	-1.082***	-1.093***
	(0.085)	(0.299)	(0.307)	(0.307)	(0.002)
Retailer H.H. Gregg	-0.369***	-1.903***	-1.963***	-2.040***	-2.087***
	(0.119)	(0.282)	(0.299)	(0.278)	(0.007)
Retailer Home Depot	-0.161	-0.738**	-0.765**	-0.798**	-0.820***
	(0.106)	(0.314)	(0.321)	(0.322)	(0.003)
Retailer Lowe's	-0.180**	-0.302	-0.334	-0.369*	-0.393***
	(0.091)	(0.230)	(0.231)	(0.222)	(0.004)
Retailer Sears	0.015	-0.436	-0.435	-0.431	-0.427***
	(0.114)	(0.442)	(0.445)	(0.443)	(0.001)
Brand FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand time trends	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,590	1,590	1,586	1,590	1,590
Kleibergen–Paap F-statistic	24.082				
Own-price elasticity			-0.964	-2.186	-2.887

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 brand level. The own-price elasticity of residual demand is computed at product level, and the average is calculated by weighting products according to their sales volume. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

V Appendix to Section 7: Details on Welfare Effects

V.A Technical details on the portfolio choice algorithm

After the algorithm is initialized, 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 from adding each inactive product separately to the brand's product portfolio and the expected change in firm-level profits from removing each active product separately.[51] 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 me to compute 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 be an important factor as firms segment their products by brand and so products within a brand are much closer substitutes than those between the same firm's brands.

Another way in which I reduce the computational burden is by considering only one-step deviations and disregarding multistep deviations. This is necessary because checking for multistep deviations is also computationally infeasible in this case.[52] It could thus be that although there is no profitable one-step deviation, there nevertheless exists a profitable multistep deviation. To assess whether this could be an important problem, it is helpful to consider when such a situation could arise. Since washing machines are substitutes in the marketplace, if it is not profitable to add a particular washer, it is also not profitable to add that and another potential washer. The same logic applies to the removal of active washing machines from the product portfolio. It is, however, possible that although adding a particular washer is not profitable, it would be profitable to add that washer and remove another washer from the product portfolio simultaneously. Similarly, it could be that it is profitable to add a washer and remove two washers simultaneously. Overall, however, it may

[51] Since 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.

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

not be desirable to consider multistep 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 I estimate 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 to the midpoint of the confidence bounds. I draw the idiosyncratic product- and market-specific fixed cost shock υ_{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.

V.B Additional results

Table A.4 repeats the analysis in Table 6 but without constraining the shipping cost parameter in the bootstrapped samples to the point estimate using the main sample. As expected, this leads to wider confidence intervals in Column (2). The results in Column (4) are unaffected, as I estimate the endogenous portfolio algorithm on the main sample only.

Table A.5 shows the results for the different merger simulations when comparing an acquisition of Maytag by Whirlpool to a standalone Maytag. The assumption here is that Maytag could have continued operating at the same costs and in the same locations as the pre-merger Maytag but would not have been able to offshore production and take advantage of offshoring cost changes. Two things stand out: First, the cost changes from offshoring Maytag front-loaders to Mexico is sizeable. Although top-loader production is kept in the U.S., this leads to a volume-weighted average marginal cost reduction of 15% when holding the product portfolio fixed. Second, with an average marginal cost reduction of 15% for Maytag products, a traditional merger simulation still finds that a Whirlpool acquisition leads to consumer harm.

Table A.6 shows the results for the different merger simulations when comparing an acquisition of Maytag by Haier to a standalone Maytag. As expected, without cost changes, the product market implications of these two scenarios are the same. With offshoring cost changes, U.S. consumers benefit from an acquisition of Maytag by Haier, whereas U.S. workers lose out. The results in Column (4) show that a Haier acquisition increases consumer welfare by $80 million per year, whereas there are 1750 fewer U.S. manufacturing workers. These results suggest that the harm to U.S. workers outweighs the benefits of a Haier ac-

Table A.4: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier (bootstrapping shipping cost parameter)

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	-1.6 [-3.5, 0.3]	-1.4 [-3.5, 0.8]
Whirlpool	– –	– –	-2.1 [-4.6, 0.4]	-2.7 [-5.2, -0.2]
LG + Samsung	– –	– –	0.3 [-0.8, 1.4]	0.5 [-0.9, 1.8]
Electrolux + GE	– –	– –	0.9 [-1.3, 3.0]	1.0 [-1.0, 3.1]
Overall industry	– –	– –	-2.5 [-6.4, 1.3]	-2.6 [-5.8, 0.6]
Prices and consumer welfare				
Average price	3.3% [1.4%, 5.2%]	3.5% [-4.3%, 11.2%]	2.7% [1.3%, 4.1%]	3.0% [1.4%, 4.6%]
Consumer welfare	$-147M [$-191M, $-104M]	$-156M [$-318M, $6M]	$-215M [$-260M, $-169M]	$-228M [$-264M, $-191M]
Employment				
Domestic jobs maintained	573 [422, 724]	708 [486, 931]	700 [485, 915]	897 [665, 1129]
Industry profits				
Variable profits	$75M [$39M, $110M]	$134M [$44M, $225M]	$57M [$28M, $86M]	$76M [$52M, $100M]
Total profits	$75M [$39M, $110M]	$134M [$44M, $225M]	$96M [$71M, $121M]	$120M [$93M, $146M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-1.8% [-22.1%, 18.6%]	– –	-1.6% [-2.1%, -1.1%]
Variable profits	$18M [$-16M, $51M]	$111M [$-24M, $246M]	$-13M [$-54M, $27M]	$5M [$-35M, $41M]
Total profits	$18M [$-16M, $51M]	$111M [$-24M, $246M]	$32M [$13M, $51M]	$56M [$38M, $76M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$257k [$176k, $338k]	$220k [$-20k, $460k]	$315k [$186k, $444k]	$260k [$167k, $352k]
Consumer welfare + total industry profits	$127k [$33k, $221k]	$31k [$-361k, $422k]	$175k [$97k, $253k]	$123k [$75k, $171k]

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

Table A.5: Simulated effects of Maytag acquisitions by Whirlpool vs. no acquirer

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	-1.6 [-3.5, 0.3]	0.2 [-1.9, 2.3]
Whirlpool	– –	– –	-2.1 [-4.6, 0.4]	-3.0 [-5.6, -0.5]
LG + Samsung	– –	– –	0.3 [-0.8, 1.4]	0.3 [-0.8, 1.4]
Electrolux + GE	– –	– –	0.9 [-1.3, 3.0]	0.7 [-1.2, 2.6]
Overall industry	– –	– –	-2.5 [-6.4, 1.3]	-1.8 [-5.3, 1.7]
Prices and consumer welfare				
Average price	3.3% [1.4%, 5.2%]	1.7% [-0.4%, 3.7%]	2.9% [1.5%, 4.3%]	1.9% [0.3%, 3.5%]
Consumer welfare	$-148M [$-192M, $-104M]	$-68M [$-119M, $-17M]	$-217M [$-262M, $-171M]	$-148M [$-193M, $-102M]
Employment				
Domestic jobs maintained	-629 [-759, 499]	-542 [-673, -411]	-992 [-1197, -787]	-853 [-1083, -624]
Industry profits				
Variable profits	$75M [$40M, $111M]	$148M [$90M, $207M]	$58M [$30M, $87M]	$134M [$104M, $164M]
Total profits	$75M [$40M, $111M]	$148M [$90M, $207M]	$97M [$72M, $122M]	$163M [$135M, $191M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-15.0% [-21.7%, -8.3%]	– –	-11.1% [-12.0%, -10.2%]
Variable profits	$18M [$-15M, $52M]	$125M [$62M, $189M]	$-13M [$-54M, $28M]	$81M [$37M, $125M]
Total profits	$18M [$-15M, $52M]	$125M [$62M, $189M]	$33M [$14M, $51M]	$117M [$95M, $140M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	– –	– –	– –	– –
Consumer welfare + total industry profits	– –	– –	– –	– –

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

quisition to U.S. consumers so long as a job is worth at least $\$45,000$ per year ($\$71,000$ if industry profits count as well).

Table A.7 repeats the analysis in Table 6 drawing fixed cost from a distribution around the 25^{th} percentile of the distance between the lower- and upper-bound brand-level fixed costs, instead of its midpoint. The results are similar to the baseline results. All of the point estimates are within the confidence sets of the main estimation results in Table 6.

Table A.8 repeats the analysis in Table 6 drawing fixed cost from a distribution around the 75^{th} percentile of the distance between the lower- and upper-bound brand-level fixed costs, instead of its midpoint. The results are similar to the baseline results. All of the point estimates are within the confidence sets of the main estimation results in Table 6.

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

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	0 [0, 0]	1.6 [-0.8, 4.0]
Whirlpool	– –	– –	0 [0, 0]	-0.3 [-2.0, 1.4]
LG + Samsung	– –	– –	0 [0, 0]	-0.1 [-1.3, 1.0]
Electrolux + GE	– –	– –	0 [0, 0]	-0.3 [-1.5, 1.0]
Overall industry	– –	– –	0 [0, 0]	0.8 [-1.9, 3.5]
Prices and consumer welfare				
Average price	0.0% [-0.0%, 0.0%]	-1.7% [-3.5%, 0.1%]	0.2% [0.1%, 0.2%]	-1.1% [-2.4%, 0.3%]
Consumer welfare	$-817k [$-3M, $866k]	$88M [$5M, $170M]	$-2.1M [$-2.3M, $-2.0M]	$80M [$48M, $112M]
Employment				
Domestic jobs maintained	-1202 [-1466, -938]	-1250 [-1520, -979]	-1692 [-1799, -1585]	-1750 [-1955, -1545]
Industry profits				
Variable profits	$854k [$-38k, $2M]	$14M [$-43M, $71M]	$1.5M [$1.3M, $1.6M]	$58M [$40M, $76M]
Total profits	$854k [$-38k, $2M]	$14M [$-43M, $71M]	$1.5M [$1.3M, $1.6M]	$44M [$26M, $61M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-13.3% [-25.0%, -1.5%]	– –	-9.5% [-10.2%, 8.7%]
Variable profits	$854k [$-44k, $2M]	$15M [$-79M, $109M]	$0.7M [$0.6M, $0.8M]	$97M [$67M, $127M]
Total profits	$854k [$-44k, $2M]	$15M [$-79M, $109M]	$0.7M [$0.6M, $0.8M]	$79M [$69M, $88M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	– –	– –	– –	– –
Consumer welfare + total industry profits	– –	– –	– –	– –

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

Table A.7: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier, fixed cost draws around 25^{th} percentile

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	-1.0 [-3.1, 1.1]	-0.9 [-2.7, 0.9]
Whirlpool	– –	– –	-2.1 [-4.8, 0.6]	-2.5 [-5.0, 0.0]
LG + Samsung	– –	– –	0.3 [-0.7, 1.3]	0.3 [-1.1, 1.7]
Electrolux + GE	– –	– –	0.5 [-2.0, 3.0]	0.8 [-1.3, 2.8]
Overall industry	– –	– –	-2.3 [-6.0, 1.4]	-2.3 [-6.2, 1.6]
Prices and consumer welfare				
Average price	3.3% [1.4%, 5.2%]	3.5% [-1.7%, 8.7%]	3.2% [2.1%, 4.3%]	3.4% [2.1%, 4.7%]
Consumer welfare	$-147M [$-191M, $-104M]	$-156M [$-227M, $-85M]	$-224M [$-247M, $-201M]	$-230M [$-248M, $-213M]
Employment				
Domestic jobs maintained	573 [422, 724]	708 [518, 898]	681 [532, 830]	880 [757, 1013]
Industry profits				
Variable profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$79M [$61M, $97M]	$103M [$93M, $114M]
Total profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$100M [$87M, $112M]	$119M [$107M, $131M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-1.8% [-7.0%, 3.6%]	– –	-1.7% [-2.0%, -1.5%]
Variable profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$14M [$-14M, $42M]	$39M [$24M, $54M]
Total profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$36M [$29M, $43M]	$56M [$49M, $63M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$257k [$176k, $338k]	$220k [$117k, $323k]	$334k [$232k, $436k]	$263k [$220k, $307k]
Consumer welfare + total industry profits	$127k [$33k, $221k]	$31k [$-114k, $175k]	$186k [$124k, $247k]	$128k [$103k, $152k]

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

Table A.8: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier, fixed cost draws around 75^{th} percentile

	Pre-acquisition product portfolios		**Endogenous product portfolios**	
Cost adjustments:	None	Offshoring	None	Offshoring
Number of products				
Maytag	– –	– –	-1.2 [-3.2, 0.7]	-1.1 [-2.8, 0.6]
Whirlpool	– –	– –	-2.3 [-4.6, -0.1]	-2.6 [-4.7, -0.4]
LG + Samsung	– –	– –	0.8 [-0.7, 2.2]	0.5 [-1.4, 2.3]
Electrolux + GE	– –	– –	1.2 [-0.8, 3.2]	1.3 [-1.4, 4.1]
Overall industry	– –	– –	-1.6 [-5.1, 1.9]	-1.9 [-5.9, 2.2]
Prices and consumer welfare				
Average price	3.3% [1.4%, 5.2%]	3.5% [-1.7%, 8.7%]	2.4% [0.9%, 3.8%]	2.2% [0.3%, 4.2%]
Consumer welfare	$-147M [$-191M, $-104M]	$-156M [$-227M, $-85M]	$-192M [$-253M, $-131M]	$-207M [$-265M, $-149M]
Employment				
Domestic jobs maintained	573 [422, 724]	708 [518, 898]	862 [594, 1130]	1142 [860, 1424]
Industry profits				
Variable profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$33M [$-7M, $72M]	$50M [$13M, $87M]
Total profits	$75M [$39M, $110M]	$134M [$88M, $181M]	$81M [$50M, $112M]	$105M [$74M, $137M]
Maytag + Whirlpool profits				
Maytag cost change	– –	-1.8% [-7.0%, 3.6%]	– –	-1.9% [-2.7%, -1.1%]
Variable profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$-41M [$-105M, $24M]	$-23M [$-84M, $38M]
Total profits	$18M [$-16M, $51M]	$111M [$58M, $164M]	$26M [$6M, $46M]	$46M [$27M, $66M]
Value per job required to offset consumer welfare loss				
Consumer welfare only	$257k [$176k, $338k]	$220k [$117k, $323k]	$228k [$123k, $334k]	$184k [$111k, $258k]
Consumer welfare + total industry profits	$127k [$33k, $221k]	$31k [$-114k, $175k]	$132k [$63k, $201k]	$91k [$42k, $140k]

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