CES ifo Working Papers

12447
2026

January 2026

The Hidden Cost of Minimum Energy Efficiency Standards

Sébastien Houde, Tobias Wekhof

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

The Hidden Cost of Minimum Energy Efficiency Standards

Sébastien Houde[a] and Tobias Wekhof[b]

January 31, 2026

Abstract

Minimum standard regulations for energy-using durables have long been suspected of having hidden costs: quality improvements in the regulated dimension reduce quality in other dimensions. We substantiate this claim for the U.S. clothes washer market, which has become a notorious example of the hidden cost phenomenon. We find that overall quality increased from 2001 to 2011, and these gains were primarily driven by improvements in energy efficiency. Quality in the non-energy dimensions declined or remained constant after the major standard change. These hidden costs, however, were quickly offset by energy-efficiency improvements in the new models.

JEL: Q48, Q55, L51, L68, D12

Keywords: Minimum Quality Standard, Hidden Cost, Energy Efficiency Regulations, Appliance Market, Ex Post Analysis.

[a]HEC Lausanne, Department of Economics, CESifo, sebastien.houde@unil.ch.

[b]Center for Energy Policy and Economics at ETH Zurich, twekhof@ethz.ch.

We would like to thank Anna Spurlock, Hung-Chia Yang, Larry Dale, and several seminar and conference attendees. The initial version of this work was developed with support from the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

1. Introduction

Minimum quality standards are a widely used policy tool to protect consumers, but they can generate unintended consequences. One concern is that they distort product design, a phenomenon referred to as the “hidden cost” of minimum standards: improvements in the regulated dimension lead to a reduction in quality in other dimensions. For example, a regulation aimed at lowering the energy use of durables, a particular type of quality standard, may unintentionally diminish performance.

This phenomenon may arise for two related reasons. First, most minimum standards are attribute-based regulations targeting a single performance outcome (e.g., energy consumption) but link compliance to multiple product attributes (Ito and Sallee 2018). Second, regulators rarely have perfect knowledge of manufacturers' abatement costs. Therefore, standards unintentionally influence product design along several dimensions. They reward some attributes while penalizing others, creating unforeseen trade-offs.

These trade-offs are especially prominent in the energy context, where minimum energy performance standards are among the main tools for regulating end-use energy demand. A rich literature argues that consumers under-invest in energy-efficient technologies, both from private and social perspectives (Jaffe and Stavins 1994; Allcott and Greenstone 2012), where hidden costs may be one reason for this under-investment (Gillingham, Newell, and Palmer 2009; Gerarden, Newell, and Stavins 2017; Gillingham and Myers 2025).

Recently, the hidden cost phenomenon has been appropriated by politicians who opposed federal energy regulations and has been part of the broader narrative on regulatory overreach. The Trump administration explicitly cited hidden costs when it ordered the Department of Energy to “eliminate restrictive water pressure and efficiency rules”, claiming they made appliances “less useful, more breakable, and more expensive to repair” (Trump 2025). At a time when thousands of federal regulations are under review (Davenport 2025), evidence-based ex post analysis is urgently needed.

This paper contributes to the debate by quantifying the hidden costs of minimum energy performance standards in the U.S. clothes washer market. Focusing on a highly debated 2004 standard change, we substantiate the various claims that this change led to significant hidden costs (Fraas and Miller 2020). Our goal is to provide a precise quantification of the phenomenon using an approach and data that can readily be expanded to other contexts and thus inform the debate about ex post valuation of such regulations. We develop a revealed-preference quality index and apply a decomposition method to show the evolution of quality dynamics. We find that energy efficiency gains were prominent, and vertical quality related to non-energy attributes declined, suggesting high hidden costs. However, these hidden costs were short-lived. We show that the vintage of models that reached the market a few years after the regulatory change was of the same quality, or better. Our findings contradict the narrative that, prior to the 2004 regulatory change, the standard was too stringent and would hurt consumers in the long term (Mer 2000). Manufacturers were able to adapt far more quickly to such regulations than anticipated.

The remainder of the paper proceeds as follows. Section 2 provides background on minimum energy performance standards. Section 3 outlines our empirical strategy, Section 4 describes the data, Section 5 presents results, and Section 6 concludes.

2. **Ex Post Analysis of Minimum Energy Performance Standards**

Since the energy crisis of the 1970s, minimum energy performance standards (MEPS) have been one of the main policy instruments for regulating energy use in durable goods. Their adoption and effectiveness, however, have remained the subject of intense debate. MEPS regulate the maximum energy consumption of a product based on a small set of observable characteristics. A central rationale of the U.S. Department of Energy (DOE) for employing attribute-based standards is to avoid narrowing the choice set or distorting non-energy dimensions of product quality (U.S. Department of Energy 2012). In other words, these standards were intended, at least in part, to preempt the hidden cost phenomenon.

Over the past five decades, however, MEPS have generated sustained debate. Early critics by Hausman and Joskow (1982) raised both economic and practical concerns about their design and effectiveness. When it comes to ex post analysis of these regulations, a rich body of research focuses on the automotive sector, particularly Corporate Average Fuel Economy (CAFE) standards. This literature consistently documents unintended product design responses and points to clear instances of hidden costs. For example, Whitefoot, Fowlie, and Skerlos (2011) and Knittel (2011) show that shifting to a carbon-footprint-based CAFE formula, which sets less stringent targets for larger vehicles, incentivized automakers to increase vehicle size. In Europe, Lin and Linn (2023) find that carbon emission standards reduced overall vehicle quality, offsetting welfare gains by roughly 26%. Similarly, Klier and Linn (2016) document reductions in horsepower and torque in response to U.S. and EU standards, while Ito and Sallee (2018) demonstrate how notches in attribute-based standards distorted vehicle design in Japan.

Similar to those studies, we document the evolution of specific attributes and the trade-offs between performance, energy-efficiency-related features, and other non-energy-related features at the time of standard changes. Relative to this previous work, our contribution is to tie attribute changes to a revealed-preference measure of quality that is welfare relevant. Our analysis shows that additional product features might have led to a higher overall perception of quality at the time of the purchase decision. Our proposed approach to measure quality and subsequently quality decomposition method to uncover the source of quality changes is readily applicable to other contexts. Overall, our approach provides an additional way to conduct ex post analyses of the effects of important standard changes in the energy context and beyond.

Evidence of hidden costs in the appliance market also remains limited. Although studies find that appliance standards affect prices and product variety (Spurlock 2013; Brucal and Roberts 2019), few analyses link these regulations to changes in non-energy product quality. To our knowledge, Taylor, Spurlock, and Yang (2015) provide the most comprehensive U.S. study to date. Using historical Consumer Reports data, they constructed reliability measures based on repair rates and found long-term declines across appliance categories, although without sharp breaks coinciding with new standards.[1]

The debate about the hidden costs of MEPS was especially salient in the early 2000s, when several major standard revisions targeted clothes washers. The 2004 standard was particularly controversial. It imposed stringent energy-efficiency requirements that disproportionately affected the incumbent high-energy-consuming top-loading design while favoring the already efficient then-emerging front-loaders. Some analysts predicted that top-loaders would be unable to comply and might disappear entirely from the market, significantly reducing consumer choice (Vaughn 2000). While the market share of top-loaders fell in the years surrounding the change, it eventually recovered (see Panel A of Figure 2).

Within the regulatory analysis community, the 2004 clothes washer standard is often cited as a textbook case of hidden costs (Fraas and Miller 2020). Regulators, lacking perfect information about firms' cost and quality trade-offs, unintentionally tilted the industry toward designs that reduced energy consumption but compromised performance along other valued dimensions. Despite widespread discussion, however, empirical evidence quantifying these hidden costs remains scarce. This paper seeks to address that gap.

[1]Their evidence suggests broad quality trends, but not definitive proof of hidden costs induced directly by regulatory changes.

3. Empirical Strategy

Our estimation approach builds on the empirical industrial organization and trade literature, which infers product quality from observed market shares using product fixed effects (e.g., Khandelwal 2010; Fajgelbaum, Grossman, and Helpman 2011; Jaimovich, Madzharova, and Merella 2023). The underlying micro-foundation consists of a demand model in which consumers choose among $\mathcal{J}$ differentiated products. We define a product as a bundle of attributes, where vertical quality captures the time-invariant characteristics valued by consumers. Formally, consumer i's utility from product j is:

$$U_{ij} = \gamma_j - \eta p_j + \epsilon_{ij},$$

where γ_j denotes the vertical quality of product j net of price, η is the marginal utility of income (price coefficient), and ϵ_{ij} captures idiosyncratic preferences.

Assuming ϵ_{ij} follows an i.i.d. extreme value Type I distribution, we apply Berry (1994)'s transformation to express market shares as a function of observable utility components:

$$\ln(\sigma_{jt}) = \gamma_j - \eta p_{jt} + \delta_t + \nu_{q(t-t_{0j})} + \lambda_1 N_{t,FL} + \lambda_2 N_{t,TL} + \xi_{jt},$$

where σ_{jt} is the market share of model j in period t, δ_t are month-of-sample fixed effects capturing the outside option, seasonality, and common shocks, and $\nu_{q(t-t_{0j})}$ controls for a product's time since market introduction (with year–quarter fixed effects to avoid multicollinearity). Following Ackerberg and Rysman (2005), we also control for crowding in the product space using the number of models available each period in the two main categories: front-loaders ($N_{t,FL}$) and top-loaders ($N_{t,TL}$).

Control Function Approach to Address Price Endogeneity

To obtain consistent estimates of model-specific quality γ_j, we address the potential endogeneity of price using the control function method proposed by Terza, Basu, and Rathouz (2008). This two-step approach proceeds as follows.

First stage. We model prices using a Gaussian Generalized Linear Model (GLM) with a log link, which implies log-normality of conditional price distributions. We base this distributional assumption on the non-negativity of prices and the presence of a right-skewed tail. The price equation is specified as:

$$p_{jt} = \alpha + \beta IV_{jt} + \lambda_1 N_{t,FL} + \lambda_2 N_{t,TL} + \nu_{q(t-t_{0j})} + \delta_t + u_{jt},$$

where IV_{jt} represents instrumental variables. Our instruments are in the spirit of Berry, Levinsohn, and Pakes (1995a) and use variation in product attributes as a supply-driven exogenous shock induced by price discrimination.[2] Specifically, we use the energy consumption and Energy Star status (in 2004) of other models, both within and across brands, as instruments. This approach builds on Spurlock (2014), who shows that clothes washer prices responded strongly and heterogeneously to the 2004 standard change.[3]

Second stage. We estimate the market share equation, controlling for the first-stage residuals:

$$\ln(\sigma_{jt}) = \gamma_j - \eta p_{jt} + \delta_t + \nu_{q(t-t_{0j})} + \lambda_1 N_{t,FL} + \lambda_2 N_{t,TL} + \phi u_{jt} + \xi_{jt}.$$

Constructing the Quality Index. Our parameter of interest is the product-specific quality γ_j. To express quality in money-metric terms, we divide by the marginal utility of income, η, yielding γ_j/η. We then aggregate to a sales-weighted, price-adjusted quality index:

$$Q_t = \sum_j s_{jt} \frac{\gamma_j}{\eta}, \tag{1}$$

where s_{jt} is the market share of product j at time t (equal to zero if not offered). This index tracks quality dynamics over time and provides the basis for decomposing the effects of energy standards. The quality index has a precise economic interpretation, which is the so-called mean utility term in the model of Berry, Levinsohn, and Pakes (1995b, 2004). It is a measure of quality entirely based on revealed preferences and captures consumers' willingness to pay for different vintages of models. Crucially, the product fixed effects are comparable as long as the different vintage models always overlap at least once in time. For instance, suppose that only three models

[2]The underlying idea of Berry, Levinsohn, and Pakes (1995a) is that product line decisions are made before pricing decisions and depend on manufacturers' costs. Hence, if competing products are located close/far in the product space, this will induce cost-driven price variation.

[3]The GLM with a log link implies:

$$\ln p_{jt} = X_{jt}\beta + u_{jt}, \quad u_{jt} \sim \mathcal{N}(0, \sigma^2),$$
$$\Rightarrow \quad p_{jt} = \exp(X_{jt}\beta) \cdot \exp(u_{jt}).$$

The residual $\hat{u}_{jt} = p_{jt} - \exp(X_{jt}\hat{\beta})$ captures unobserved price variation, which we include in the second-stage regression.

were offered, say Model A, Model B, and Model C, and the data hadfour time periods. If Model A was offered in periods 1 and 2, Model B was offered in periods 2 and 3, and Model C was offered in periods 3 and 4, the quality index is identified. That is, Model A can be compared to Model C, and the product fixed for each model can be estimated. Given that our product fixed effects are estimated with point-of-sale data, the quality we measure is the one perceived by consumers at the time of making a purchase. How consumers perceive the product when making their purchase decision may differ from the actual experience they have with it. If bad experiences with some models became public signals, through reviews or consumer reports, for instance, this would be accounted for in the estimates of quality.

The quality index can also be decomposed to identify which product characteristics contribute to the overall change in quality. In the present application, we are particularly interested in determining whether changes in overall quality are driven primarily by energy use and how standard revisions have affected quality in the non-energy dimension. Next, we show how we can modify this index to include price (yielding a price-inclusive index), isolate the role of energy efficiency, and decompose the dynamics of quality.

3.1. **Energy-Adjusted Quality**

To isolate the hidden costs associated with energy efficiency, we construct an energy-price-adjusted quality index. We regress the price-adjusted quality index on lifetime energy costs, assuming consumers value appliances based on discounted operating costs. Lifetime energy costs for product j are:

$$LC_{r,j} = \sum_{t=1}^{L} \rho^t C_{r,j} = \frac{1-\rho^L}{1-\rho} \cdot p_e \cdot e_j = \omega \cdot p_e \cdot e_j, \tag{2}$$

where L is product lifetime, $\rho = 1/(1+r)$ is the discount factor, ω denotes $\frac{1-\rho^L}{1-\rho}$, p_e is the electricity price, and e_j is reported annual energy use. The residual from regressing $\hat{\gamma}_j$ on $\omega \cdot p_e \cdot e_j$ yields a revealed-preference measure of non-energy quality:

$$\hat{\xi}_j = \hat{\gamma}_j - \hat{\theta} \cdot \omega \cdot p_e \cdot e_j, \tag{3}$$

3.2. **Decomposition of Quality**

To understand which margins, namely the introduction of new models versus retirement of old models, or simply changes in market shares for existing models, drive the quality dynamics and the speed at which manufacturers were able to adapt to the new regulation, we apply a decomposition method inspired by productivity studies (e.g., Foster, Haltiwanger, and Krizan 2001). We separate changes in the quality index into the following components:

$$\begin{aligned}\Delta Q_t = & \underbrace{\sum_{j \in C} \sigma_{jt-1} \Delta q_{jt}}_{\text{within}} + \underbrace{\sum_{j \in C} \Delta\sigma_{jt}(q_{jt-1} - Q_{t-1})}_{\text{between}} + \underbrace{\sum_{j \in C} \Delta\sigma_{jt} \Delta q_{jt}}_{\text{cross}} \\ & + \underbrace{\sum_{j \in N} \sigma_{jt}(q_{jt} - Q_{t-1})}_{\text{entries}} - \underbrace{\sum_{e \in X} \sigma_{jt-1}(q_{jt-1} - Q_{t-1})}_{\text{exits}} . \end{aligned} \tag{4}$$

In this formula, Q_t is our index of overall quality, σ_{jt} is the share of model offered j in period t, q_{jt} is an index of model-level quality, Δq_{jt} represents the change in quality for continuing models, $\Delta\sigma_{jt}$ represents the change in share for continuing models, C denotes continuing models, N denotes entering models, and X denotes exiting models.

The "within" variation can only be driven by a change in price in our context, where the "between" and "cross" variations are driven by changes in market shares, and thus demand. Finally, the "entries" and "exits" variations result from the entry and exit of new and old models, holding market shares and price constant.

For our standard errors, we implement our estimation with 500 bootstrap replications. In each iteration, 5% of unique models are randomly removed, subject to the condition that at least one model is present in consecutive years to preserve continuity in the quality index. The bootstrap distribution is then used to construct mean parameter estimates and standard errors.

4. **Data**

Our analysis relies on point-of-sale data provided by the NPD Group, a U.S.-based market research company. Each observation corresponds to the monthly national sales and revenues of a specific appliance model, identified by a unique manufacturer model number. The dataset spans 2001–2011 and is aggregated at the national level.[4] The data are highly disaggregated: the manufacturer model number directly maps to products offered in stores.

Matching with Energy Data

To measure energy performance, we match the NPD data with several publicly available sources. The initial linkage was constructed by Spurlock (2014), consisting of three data sources: first, the Federal Trade Commission (FTC) provides annual model-level energy consumption data displayed on EnergyGuide labels. Second, the ENERGY STAR program adds information on program certification. Third, the California Energy Commission (CEC) provides the energy-use metric used by DOE when setting minimum efficiency standards.

Prices are adjusted to 2011 dollars. To compute lifetime operating costs, we assume an average electricity price of $0.11 per kWh, a 15-year product lifetime, and a 3% discount rate, an assumption consistent with DOE regulatory analyses.[5] As shown in Appendix Table D, lifetime energy costs are a substantial share of ownership costs, particularly for top-loading washers: while front loaders had an average price of 757 USD, the lifetime energy costs were estimated at around 234 USD. In contrast, top-loaders had an average price of 443 USD but lifetime energy costs of around 581 USD.

Sample Construction

The initial dataset contains 20,722 model-month observations. After merging with energy-efficiency data, the sample is reduced to 14,147 observations. To ensure representativeness, we restrict attention to models that account for at least 95% of total sales in each year, yielding 494 unique washer models.[6]

Appendix Figure A.1 shows that excluded models are evenly distributed across years, while Appendix Figure A.2 demonstrates that market shares of top- and front-loaders in the restricted sample mirror those in the full dataset. Thus, our sample restriction does not distort aggregate market trends.

[4] The number of retailers sampled by NPD varies across years. Market coverage ranged from roughly 25% in the early 2000s to 80% by 2011, with steady improvements in coverage over time.

[5] See DOE technical support documents for appliance standards (U.S. Department of Energy 2001, 2012).

[6] Our results are robust to alternative thresholds. Due to administrative issues, data for December 2008 are missing. In this period, 14 models exited, 9 new models entered, and 150 models were continuously available.

5. Results

We proceed in three steps. First, we present demand estimates. Second, we document the evolution of quality and key attributes over time. Third, we decompose changes in aggregate quality into within-, between-, entry-, and exit-margins.

5.1. Demand Estimation

Table 1 reports the Two-Stage Residual Inclusion (2SRI) results. In the first stage, all instruments are statistically significant: own- and rival-model ENERGY STAR certifications in 2004 are associated with lower prices, whereas the instruments based on energy consumption have smaller effects.

In the second stage, we regress log market shares on price and the first-stage residual, controlling for the number of top- and front-loaders offered, model fixed effects, month-of-sample fixed effects, and flexible age controls.[7] The price coefficient is negative and statistically significant, implying an own-price elasticity of about -1.5 at the average price, consistent with estimates for other U.S. appliance markets (Houde and Myers 2021). The positive and significant coefficient on the first-stage residual confirms the importance of correcting for price endogeneity. We take the estimated product fixed effects $\hat{\gamma}_j$ from this regression as our measure of (price-adjusted) vertical quality.

In Appendix Table B.1, we regress $\hat{\gamma}_j$ on lifetime energy costs. The estimated coefficient is negative and economically large, roughly twice the magnitude of the price coefficient, indicating that quality responds strongly to energy efficiency improvements. We also show the average marginal effects of stage one in the Appendix in Table C.1.

[7]Given the log-linear specification with time fixed effects, using market shares or quantities is equivalent.

Table 1. Demand estimation results for clothes washers (2SRI)

	Second Stage OLS	First Stage log-normal
Price	-0.370^{***}	
	(0.051)	
Number TL	-0.010^{***}	-2.6×10^{-4}
	(0.003)	(4.2×10^{-4})
Number FL	-0.020^{***}	-0.005^{***}
	(0.004)	(2.9×10^{-4})
Own kWh		-1.2×10^{-4}
		(8.4×10^{-5})
Own ENERGY STAR (2004)		-0.161^{***}
		(0.024)
Rival kWh		-1.2×10^{-4}
		(8.4×10^{-5})
Rival ENERGY STAR (2004)		-0.156^{***}
		(0.024)
Residual (stage 1)	0.893^{***}	
	(0.057)	
Num. Obs.	14 147	14 147
R^2	0.705	0.922
	(0.002)	(0.001)

Note: Second-stage regression of log quantities on price and controls, with residuals from a first-stage log-normal GLM. Instruments include rival and own ENERGY STAR certification (2004) and energy consumption. All specifications include model fixed effects, month-of-sample fixed effects, and flexible age controls. Standard errors from 500 bootstrap replications in parentheses.

5.2. **Evolution of Quality**

Figure 1 plots energy use and prices (top panels) and three quality indices (bottom panels) separately for top-loaders (left) and front-loaders (right). The 2004 standard change coincides with a large drop in energy consumption for both technologies; the 2007 change has a smaller effect. Prices show no pronounced nonlinearity around 2004 except for a visible adjustment among front-loaders.

In the bottom panels, we display (i) the price-inclusive index, (ii) the price-adjusted index, and (iii) the energy-price-adjusted index. All indices are sales-weighted; the price-inclusive index is normalized to zero for the first month of the data and serves as the reference for the other indices. Price-adjusted quality rises steadily over the time period. Notably, its trend shifts in January 2004, indicating that standards are associated with increases in overall (vertical) quality once energy-efficiency gains are accounted for.

By contrast, the energy-price-adjusted index reveals a sharp decline in non-energy quality for top-loaders at the 2004 change, suggesting a clear manifestation of hidden costs. New top-loading models ultimately met the standard, but likely at the expense of distortions in the provision of other valued attributes. For front-loaders, energy-adjusted quality is comparatively flat around the standard changes, suggesting more limited trade-offs in non-energy dimensions.

FIGURE 1. Evolution of energy use, prices, and quality indices

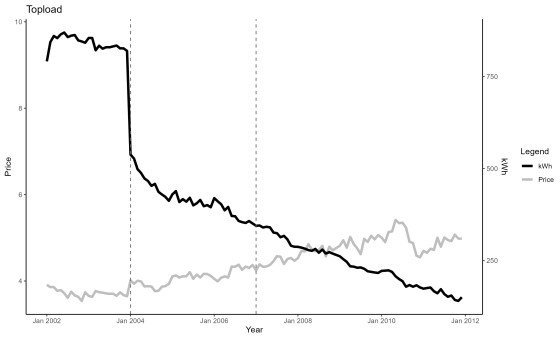

TL: Price and kWh/y

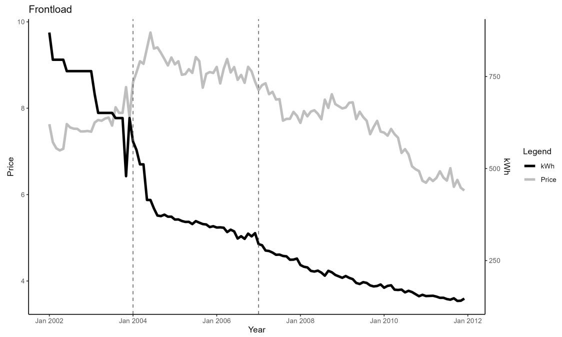

FL: Price and kWh/y

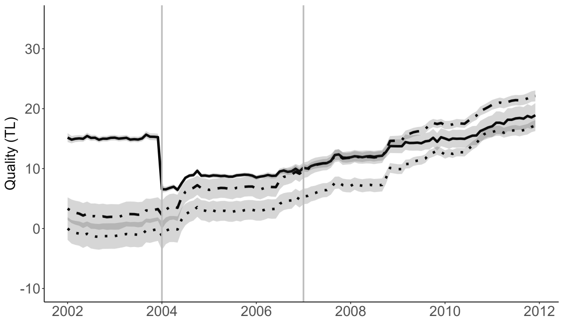

TL: Quality indices

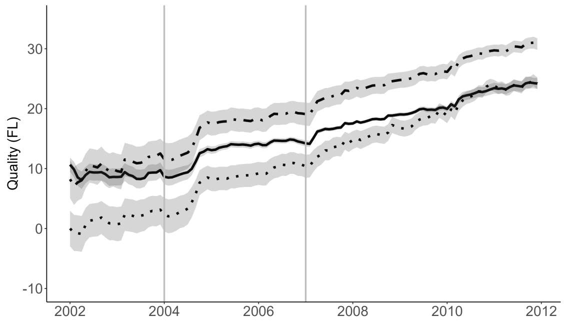

FL: Quality indices

Note: Top panels plot annual energy consumption (black) and average price (grey) for top- and front-loaders. Bottom panels show sales-weighted indices: price-inclusive (grey dotted), price-adjusted (grey dashed), and energy-price-adjusted (black). All estimates are from 500 bootstrap replications. Indices are normalized to price-inclusive quality in January 2002 = 0.

Table 2 quantifies these changes relative to December 2003 (the month prior to the 2004 change), evaluated at three time horizons: one month (in January 2004), six months (in July 2004), and twelve months (in January 2005) after the change. The first two rows (price-inclusive index) imply an initial reduction of roughly $98 for both top- and front-loaders in January 2004, followed by a rapid recovery: by six months, top-loaders are $245 higher than December 2003 and $282 higher after twelve months.[8] Front-loaders return to baseline by six months and surpass it by about $479 after one year. The price-adjusted index shows slightly smaller but qualitatively similar movements.

[8]Price is measured in $100 units; we therefore multiply the model coefficients by factor 100 to obtain price changes in dollars.

The energy-price-adjusted index tells a different story for top-loaders: non-energy quality drops by about $869 in January 2004 and remains depressed ($-678 at six months and $-651 at twelve months). For front-loaders, the initial drop is modest (about $109) and reverses by twelve months (+$333), aligning with much smaller trade-offs in non-energy attributes.

Table 2. Changes in quality relative to Dec 2003

	+ 1 month	+ 6 months	+ 12 months
TL: price-inclusive	-1.11 (0.65)	2.44 (0.63)	2.83 (0.68)
FL: price-inclusive	-0.82 (0.88)	0.19 (0.84)	5.11 (1.05)
TL: price-adjusted	-0.96 (0.66)	2.87 (0.62)	3.20 (0.66)
FL: price-adjusted	-0.95 (0.89)	0.14 (0.87)	4.97 (1.09)
TL: energy-price-adjusted	-8.70 (1.02)	-6.78 (0.68)	-6.51 (0.66)
FL: energy-price-adjusted	-1.09 (0.90)	-0.36 (0.88)	3.33 (1.02)

Note: Coefficients in units of $100 (standard errors in parentheses). Each cell reports the difference in the respective index relative to December 2003. Standard errors from 500 bootstrap replications.

5.3. **Additional Evidence: Quality and Features**

We use an additional data source to corroborate our revealed-preference measure of quality. We complemented the attribute information provided by NPD with information from the manufacturer's user manuals.[9] From these users' manuals, we observe detailed model-specific attribute information. In collecting the data, we carefully identified the various nomenclature used by manufacturers to describe a technology and tracked the incidence of each particular technology over the sample period. Finally, we consulted with appliance experts to distinguish energy efficiency-related characteristics from others. We now use this information to investigate how standards impact different product features and how the evolution of specific features correlates with the quality indexes.

Panel A of Figure 2 first shows the important market transformation at the time of the 2004' standard change when the market share of front-load models took over front-load models. In Panel B, we observe that size, measured by tube capacity, has also been steadily increasing, and the trend became more pronounced midway between the first and second revisions of the minimum standard.

[9]LBNL researchers, chiefly Anna Spurlock, collected the user manuals from various online sources and extracted the content from the PDF documents with the help of valuable research assistants.

Note that the minimum standard for clothes washers is set as a function of size. The expansion in tube capacity could then be partly the result of the new, stringent standards, which would have implicitly incentivized manufacturers to meet the energy efficiency standard with larger models.

Panel C shows the motor average spin speed—a proxy for motor performance. Spin speed has been steadily increasing, especially after the 2007 revision, particularly in top-load models. This suggests that manufacturers may have been able to deliver better motor performance but only after improving energy efficiency following the 2004 standard change. For top-load models, the trade-off between spin speed and energy efficiency was then the most important.

Panels D through F show the evolution of the number of features during the sample period. In all panels, we again distinguish between front-load and top-load models. To count the number of features in a given appliance model, we use a dummy variable that takes the value 1 if a particular feature is listed in the user manual and 0 otherwise. In these data, a feature is defined as a particular technology we track across the entire sample period. Panel D shows the average number of features per appliance model. We observe an increase in the number of features for both types of clothes washers, but a much larger increase for front-loaders, and again with a notable steady increase after the 2007 change. During this period, the front-load design was an important innovation. It is thus interesting to observe that innovation in the overall design was also accompanied by the addition of new features.

Panels E and F distinguish whether these new features were directed toward achieving energy efficiency or other dimensions. Panel E shows the evolution of energy-efficiency-related features identified by discussions with appliance experts as enabling higher energy-efficiency performance (see Table F.1, Appendix F). We find that for both front-load and top-load models, these features increased throughout the sample period, with a sharp increase at the 2004 revision for front-load models. This pattern is consistent with our revealed-preference quality index, which shows that the most important change in quality occurred at this moment and was driven by improvements in energy efficiency. Panel F shows non-energy efficiency-related features. For front-load models, we still observe an increase in the average number of features, but more gradual. For top-load models, we also observe an increase in non-energy-efficiency features, but of much smaller magnitude than the energy-efficiency-related features, and much later after the revision dates. Altogether, this suggests that for front-load models, innovation at the time of the revision in standards was directed toward both energy efficiency and other dimensions of the product space. On the other hand, for top-load models, innovation focused primarily on energy efficiency.

Figure 2. Evolution of Attributes and Features: Clothes Washers

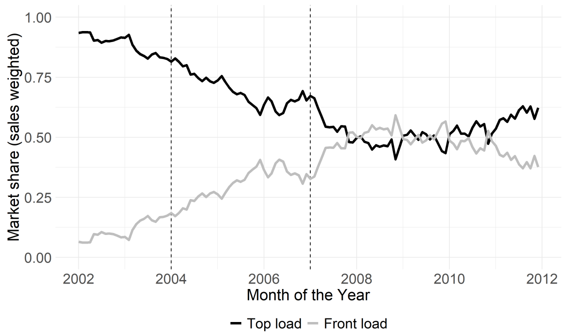

A) Market Share

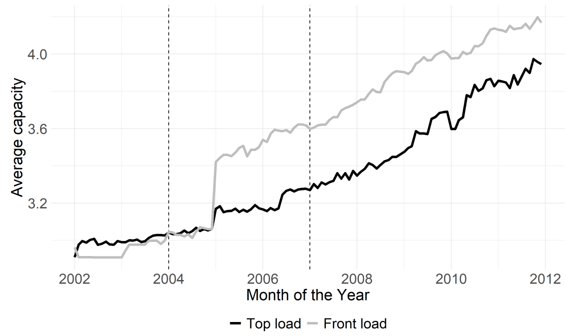

B) Capacity

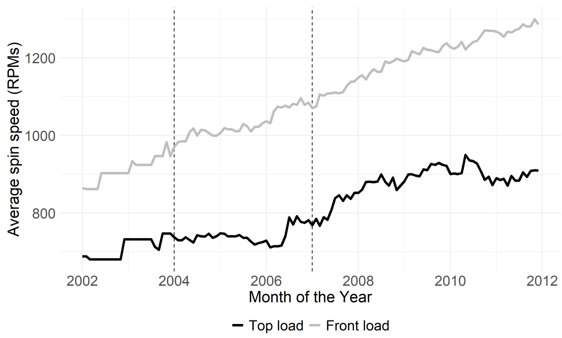

C) Motor Speed

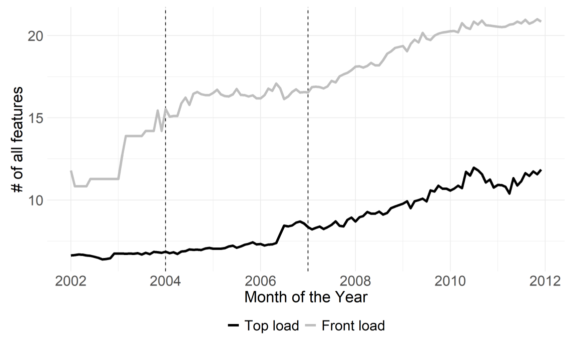

D) All Features

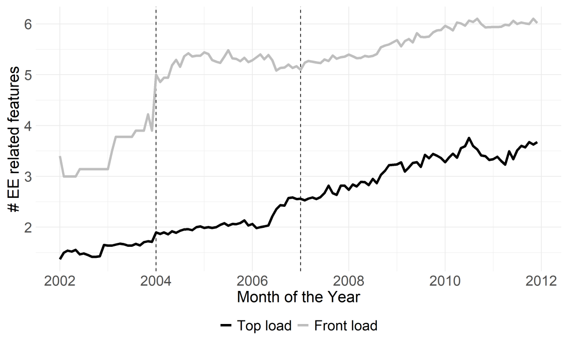

E) EE-Related Features

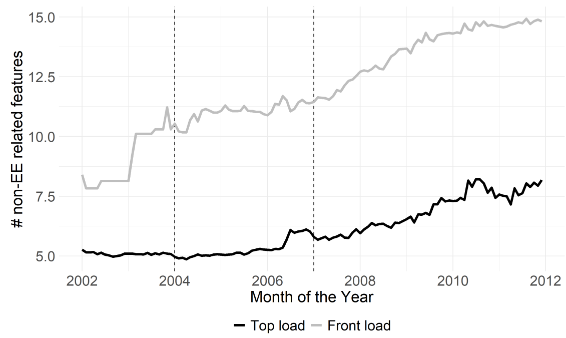

F) Non EE-Related Features

Note: This figure shows the monthly evolution for top and front-loaders for different variables: market share, capacity, motor speed, total number of features, number of energy-efficiency related features, and number of non-energy-efficiency related features.

5.4. Decomposition of Quality Dynamics

We next decompose changes in the sales-weighted, price-inclusive quality index around the 2004 standard change (Figure E.1). The goal is to attribute aggregate monthly movements to (i) within-model changes, (ii) shifts in market shares across continuing models (between), (iii) entry of new models, and (iv) exit of old models, plus a cross-model term.

Three findings emerge. First and most importantly, most of the movements in the quality dynamics occurred exactly at the time of the standard changes, especially in 2004. Second, quality changes around the 2004 standard are dominated by between-model dynamics. Third, new compliant models induce a sharp, temporary decrease in energy-adjusted quality; within/cross effects are negligible. The substitution effects and entry of new models were thus the main drivers of the hidden cost phenomenon.

Taken together, our results reveal a nuanced picture of the hidden cost phenomenon. The 2004 standard change induced a sharp decline in non-energy quality for top-loaders, even as overall vertical quality (inclusive of energy efficiency) improved. For front-loaders, the trade-offs were smaller and short-lived, with quality gains emerging within one year. Decomposition analysis reveals that most of the adjustment occurred through the reallocation of market shares across existing models and the introduction of new compliant designs, rather than through incremental improvements within models. Overall, while hidden costs were real and substantial for certain technologies, they were offset by welfare gains from energy efficiency in the medium run. Ultimately, despite concerns that the revised standards would change the market and harm consumers in the long term, manufacturers were very quick to adapt to the new standards, and consumer substitution played a role. The early generations of models, which had to meet more stringent regulations, appear to have suffered from lower quality, but this was corrected in the vintage of models that came to market a few years later.

6. Discussion and Conclusions

Our analysis of the U.S. clothes washer market demonstrates that minimum energy performance standards (MEPS) had heterogeneous effects across technologies. Between 2001 and 2011, overall product quality rose, mainly driven by gains in energy efficiency. Yet, once energy use is accounted for, non-energy quality either stagnated or declined, which provides evidence of the hidden cost phenomenon. These effects were particularly pronounced for top-loaders, the incumbent design most constrained by the 2004 standard, whereas front-loaders, already more energy-efficient, saw little change in non-energy quality. Decomposition analysis further reveals that the introduction of new compliant models was the primary channel through which these quality shifts occurred.

This paper's findings highlight a key trade-off in designing attribute-based regulation. MEPS can significantly cut energy use, but they might also alter other dimensions of product qualities that consumers value. The revealed preference approach allows for quantifying these hidden costs and is readily applicable to conduct ex post analysis of other regulatory changes to minimum standards. Our results point to the importance of incorporating potential trade-offs between energy efficiency and other attributes into ex ante regulatory design and evaluation. Better accounting of manufacturers' abatement cost structures and consumer preferences could help mitigate hidden costs, allowing standards to capture efficiency gains without compromising other dimensions of product quality. For this particular episode of standard changes, manufacturers' ability to quickly adjust appears to have been underestimated.

References

(2000): "Addendum to Public Interest Comment on the Department of Energy's Proposed Clothes-Washer Efficiency Standards," Regulatory Studies Program Addendum Docket EE-RM-94-403, Mercatus Center, George Mason University, Includes Rasmussen Research survey of 1,997 consumers.

Ackerberg, D. A., and M. Rysman (2005): "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects," *The RAND Journal of Economics*, 36(4), 771–788.

Allcott, H., and M. Greenstone (2012): "Is there an energy efficiency gap?," *Journal of Economic perspectives*, 26(1), 3–28.

Berry, S., J. Levinsohn, and A. Pakes (1995a): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, pp. 841–890.

——— (1995b): "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, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25(2), 242–262.

Brucal, A., and M. J. Roberts (2019): "Do energy efficiency standards hurt consumers? Evidence from household appliance sales," *Journal of Environmental Economics and Management*, 96, 88–107.

Davenport, C. (2025): "Inside Trump's Plan to Halt Hundreds of Regulations," *The New York Times date = 2025-04-15, url = https://www.nytimes.com/2025/04/15/us/politics/trump-doge-regulations.html, urldate = 2025-08-18*.

Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011): "Income distribution, product quality, and international trade," *Journal of political Economy*, 119(4), 721–765.

Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001): "Aggregate productivity growth: Lessons from microeconomic evidence," in *New developments in productivity analysis*, pp. 303–372. University of Chicago Press.

Fraas, A., and S. E. Miller (2020): "Measuring Energy Efficiency," *Economics of Energy & Environmental Policy*, 9(2), 181–198.

Gerarden, T. D., R. G. Newell, and R. N. Stavins (2017): "Assessing the Energy-Efficiency Gap," *Journal of Economic Literature*, 55(4), 1486–1525.

Gillingham, K., and E. Myers (2025): "Economics of energy efficiency," in *Handbook on Electricity Regulation*, pp. 254–267. Edward Elgar Publishing.

Gillingham, K., R. G. Newell, and K. Palmer (2009): "Energy Efficiency Economics and Policy," *Annual Review of Resource Economics*, 1(1), 597–620.

Hausman, J. A., and P. L. Joskow (1982): "Evaluating the Costs and Benefits of Appliance Efficiency Standards," *The American Economic Review*, 72(2), 220–225.

Houde, S., and E. Myers (2021): "Are consumers attentive to local energy costs? Evidence from the appliance market," *Journal of Public Economics*, 201, 104480.

Ito, K., and J. M. Sallee (2018): "The economics of attribute-based regulation: Theory and evidence from fuel economy standards," *Review of Economics and Statistics*, 100(2), 319–336.

Jaffe, A. B., and R. N. Stavins (1994): "The energy-efficiency gap What does it mean?," *Energy policy*, 22(10), 804–810.

Jaimovich, E., B. Madzharova, and V. Merella (2023): "Inside the white box: Unpacking the determinants of quality and vertical specialization," *European Economic Review*, 152, 104374.

Khandelwal, A. (2010): "The long and short (of) quality ladders," *The Review of Economic Studies*, 77(4), 1450–1476.

Klier, T., and J. Linn (2016): "The effect of vehicle fuel economy standards on technology adoption," *Journal of Public Economics*, 133, 41–63.

Knittel, C. R. (2011): "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector," *American Economic Review*, 101(7), 3368–3399.

Lin, Y., and J. Linn (2023): "Environmental regulation and product attributes: the case of European passenger vehicle greenhouse gas emissions standards," *Journal of the Association of Environmental and Resource Economists*, 10(1), 1–32.

Small, K. A., and H. S. Rosen (1981): "Applied welfare economics with discrete choice models," *Econometrica: Journal of the Econometric Society*, pp. 105–130.

Spurlock, C. A. (2013): "Appliance Efficiency Standards and Price Discrimination," Discussion Paper LBNL-6283E, Lawrence Berkeley National Laboratory, Berkeley, CA.

Spurlock, C. A. (2014): "Appliance Efficiency Standards and Price Discrimination," Discussion paper, Lawrence Berkeley National Laboratory.

Taylor, M., A. Spurlock, C. and H. Yang (2015): "Confronting Regulatory Cost and Quality Expectations: An Exploration of Technical Change in Minimum Efficiency Performance Standards," Discussion Paper 15-50, Resources for the Future, Washington, DC, Provides a detailed case study of the 2001–2007 clothes washer standards and the contemporaneous concern over consumer "hidden costs" and performance trade-offs; also issued as LBNL Report #LBNL 1000576.

TERZA, J. V., A. BASU, AND P. J. RATHOUZ (2008): "Two-stage residual inclusion estimation: addressing endogeneity in health econometric modeling," *Journal of Health Economics*, 27(3), 531–543.

TRUMP, D. J. (2025): "Rescission of Useless Water Pressure Standards," Presidential Memorandum.

U.S. DEPARTMENT OF ENERGY (2001): "Energy Conservation Program for Consumer Products: Clothes Washer Energy Conservation Standards," *Federal Register*, 66(9), 3314–3333, Table 1 lists "Life expectancy: 14.1 years.".

——— (2012): "Energy Conservation Program: Energy Conservation Standards for Residential Clothes Washers — Direct Final Rule," *Federal Register*, 77(105), 32308–32380, RIN 1904-AB90.

VAUGHN, G. (2000): "Public Interest Comment on DOE's Proposed Clothes-Washer Efficiency Standards," Regulatory Studies Program Comment Docket EE-RM-94-403, No. 224, Mercatus Center, George Mason University.

WHITEFOOT, K., M. FOWLIE, AND S. J. SKERLOS (2011): "Product Design Response to Industrial Policy: Evaluating Fuel Economy Standards Using an Engineering Model of Endogenous Product Design," *Energy Institute at Haas Working Paper WP-214*.

Appendix A. Full sample statistics

FIGURE A.1. Models with and without capacity and kwh

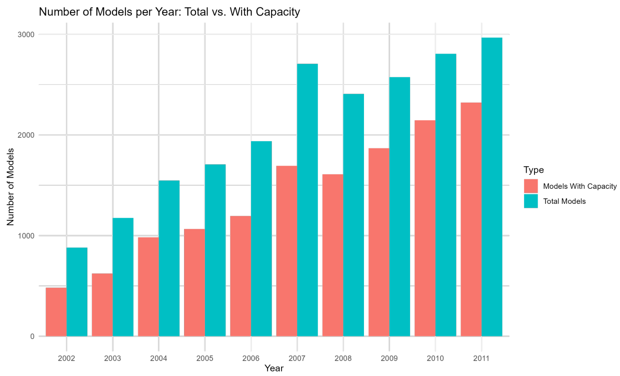

Note: This figure shows all models in a year and those with full capacity and kWh data.

FIGURE A.2. Market shares of top versus front loaders under different sample definitions.

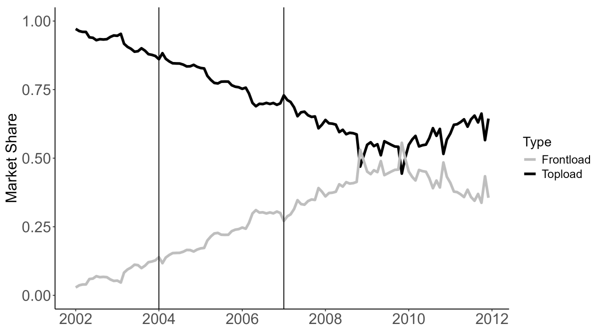

(a) Unrestricted sample

(b) Restricted sample

Note: Comparison of the evolution of market shares for unrestricted sample and restricted sample.

Appendix B. Energy adjusted quality estimation

Below, we present the regression results we used to construct the energy-price-adjusted quality index, where we regress the quality index, $gam\hat{m}a$, on the constructed measure of lifetime energy costs.

Table B.1. OLS: Quality on energy-cost

	Energy-Price-Adjusted Quality
Discounted energy cost	−0.714***
	(0.070)
Intercept	3.680***
	(0.838)
Num.Obs.	14 147
R2	0.509
	(0.032)

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

Note: This table presents the estimation results for an OLS regression with Quality as the dependent variable and discounted energy costs per appliance as the independent variable. Coefficients and standard errors were estimated with 500 bootstrap samples.

Appendix C. Demand estimation stage 1 average marginal effects

Table C.1. Stage 1 Average Marginal Effects

	First Stage AME
Number TL	−0.001
	(0.002)
Number FL	−0.028***
	(0.001)
Own kwh	−0.001
	(0.000)
Own estar 2004	−0.951***
	(0.137)
Rival kwh	−0.001
	(0.000)
Rival estar 2004	−0.924***
	(0.137)
Num.Obs.	14 147
R2	0.923

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

Note: This table presents the average marginal effects for stage 1 of the estimation (the estimation was performed on the whole sample without bootstrapping.

Appendix D. Summary Statistics: Main Sample

This table summarizes key differences between top-loaders (TL) and front-loaders (FL) washers between 2001–2011. The product mix is broad for both technologies, with TLs offering about one-third more unique models than FLs. Product shelf-life is similar: models in both groups remain available for roughly a year and a half on the market.

Market outcomes show a district segmentation between the two washer types. TLs capture a larger per-model market share; by contrast, FLs show higher prices—both in means and medians.

Energy performance diverges as well. FLs use less than half the annual electricity compared to TLs on average, translating into lifetime operating costs that are also less than half. The dispersion in energy use is wider for TLs, which is likely due to technology evolution following energy standards.

TABLE D.1. Summary statistics for clothes washers (restricted sample, 2001–2011)

	Top-loaders (TL)	Front-loaders (FL)
Sample & identification		
Unique models (N)	283	211
Avg model age on market (months)	18.32 (14.73)	19.14 (14.08)
Market outcomes		
Market share per model (%)	1.05 (1.71)	0.61 (1.18)
Price (2011 $)	443.13 (192.36)	757.86 (253.81)
Median	386.62	720.94
Energy & operating costs		
Annual energy use (kWh/yr)	442.74 (210.85)	178.18 (70.32)
Lifetime energy cost (2011 $)	581.40 (276.89)	233.98 (92.34)
Competitive environment		
# models offered per month	62.45 (15.35)	56.44 (39.41)

Note: This table presents the summary statistics of the sample between 2002 and 2011. Means (SD) unless noted. Prices in 2011 dollars. Lifetime energy cost uses $0.11/kWh, 15-year life, 3% discount rate. Monthly model counts ($N_{t,TL}, N_{t,FL}$) are averaged over months. "Avg model age on market" is the average number of months a model j was already on the market in month t.

Appendix E. Decomposition of Quality Dynamics

The upper left panel shows the decomposition of the price-inclusive measure of quality. Most variation over the entire period stems from between model dynamics, with yearly upward peaks. The positive effect of the between-model dimension is especially pronounced one year after the standard introduction. At the time of the introduction, the impact of new model entries into the market led to a slight downward peak in this dimension. While the models leaving the market immediately after the standard introduction also created an even smaller downward peak, this effect reversed one year after the standard introduction, with exiting models increasing the quality evolution. The other two dimensions, within and cross-models, do not significantly affect the evolution of quality.[10]

In the upper right panel, we show the decomposition for the energy-price-adjusted quality. In contrast to before, the drop due to new models at the time of the standard is significantly larger. Similarly, the existing models negatively impact quality evolution, albeit with a smaller magnitude than the new model's effect. The between-dimension follows a similar pattern for the overall quality but has a lower magnitude than the impact of new models.

The lower two panels show the decomposition for the energy-price-adjusted quality for top-loaders on the left and front-loaders on the right. Top loaders follow the same patterns as the joint graph in the upper right panel. Front-loaders follow a similar pattern, but the magnitude is smaller by a factor of 100; moreover, the effect of new models following the standard introduction is smaller in relative terms, and the impact of exiting models due to the standard is positive.

Overall, the decomposition that quality dynamics show that hidden cost phenomenon can thus be attributable to new entrants that had to meet the new standards, and incumbent models that had high quality, but did not the new energy efficiency requirement. Manufacturers, however, were able to recover relatively quickly and offer models that met previous quality levels in the non-energy dimension. The negative effects were nonetheless persistent given the nature of the durable good purchasing decision.

[10]In the appendix, we also compute a quality metric computed with Small and Rosen (1981)'s measure for logit-based discrete choice models. We report the evolution of quality, which follows a similar path to the price-inclusive quality index.

Figure E.1. Decomposition of quality dynamics around the 2004 standard change

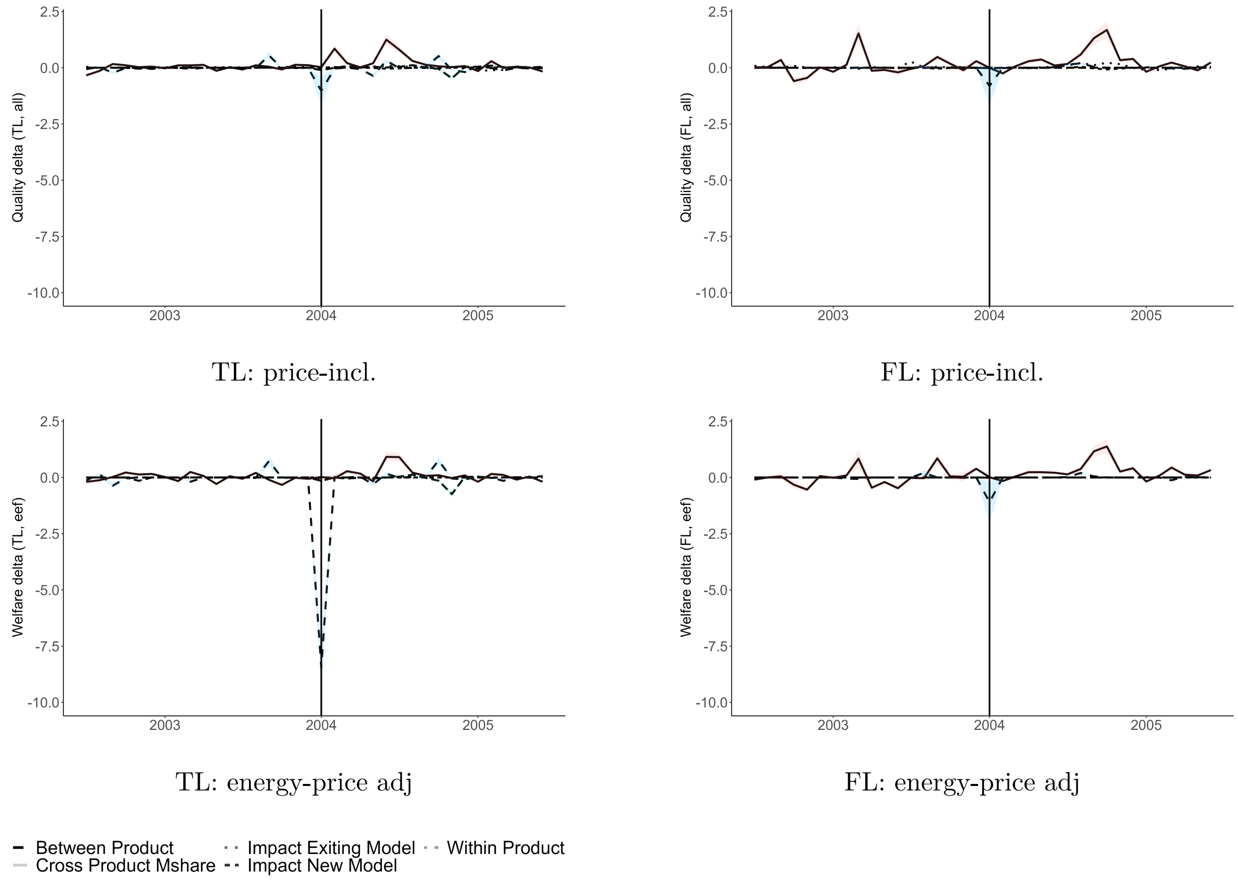

Note: This figure shows the decomposition of quality dynamics following equation 4 for different subsets of clothes washers. The upper left panel shows the entire sample, and the upper right panel shows the price and energy-adjusted quality. The lower left panel shows the price and energy-adjusted quality for top loaders, and the lower right panels show the same plot for front loaders. All values and standard errors were estimated with 500 bootstrap iterations.

The intuition from the four graphs in Figure E.1 is reflected in the quality evolution six months after the standard introduction, shown in Table E.1.

For the top-loader price-inclusive measure of quality in column (1), quality increased by 246$ from January 2004 to July 2004. This increase could be almost entirely attributed to the between-product component of quality. In comparison, column (2) shows the price-inclusive index for front loaders with nearly no change in the same period; market share shifts between models offset the moderate quality decrease from new models.

In column (3), we show the energy-price-adjusted quality for top-loaders, which shows a decrease of 392$ in the same period. Here, the impact between models is about 30% smaller. The impact of new models is six times larger than before and dominates the other effects. Column (4) shows the energy-price-adjusted quality decomposition for front loaders, which shows the same pattern as top loaders but with a smaller magnitude, making the quality evolution nearly flat.

Table E.1. Quality decomposition: 6-month difference Jan to July 2004

Quality	TL price-incl.	FL price-incl.	TL price-energy-adj	FL price-energy-adj
Delta quality	2.44 (0.63)	0.19 (0.84)	-6.78 (0.68)	-0.36 (0.88)
New models	-1.17 (0.65)	-0.65 (0.86)	-8.54 (0.99)	-1.07 (0.89)
Exiting model	0.07 (0.03)	-0.01 (0.01)	-0.21 (0.04)	-0.01 (0.01)
Between models	3.35 (0.64)	0.73 (0.26)	2.09 (0.57)	0.72 (0.27)
Within model	0.23 (0.01)	0.14 (0.02)	-0.12 (0.02)	-3.0e-04 (1.2e-04)
Cross models	-0.04 (0.01)	-0.02 (0.01)	-2.1e-03 (0.01)	1.7e-04 (9.9e-05)

Note: This table presents the quality decomposition from Dec 2003 to July 2004 (hence, 6 months after the standard introduction). The first column shows the price-inclusive index for top-loaders and column 2 for front-loaders. Columns 3 and 4 present the energy-adjusted quality index for top and front-loaders. All values and standard errors were estimated with 500 bootstrap iterations.

Appendix F. Detailed Appliance Features

Using the manufacturer's owner manual, we extracted features for a large number of models. The table below lists the features and their classification, where we distinguish energy and non-energy-related features.

Table F.1. Attributes of Clothes Washers)

Attribute	Description	Coding	EE-Related	Subject to Trademark
kWh/y	Yearly Electricity Consumption Reported by Manufacturers to FTC	Continuous	Yes	
Size (Cu. Ft.)	Overall Capacity	Continuous		
# Cycles	Number of washing cycles	Continuous		
ES-certified	ENERGY STAR certified	0-1 Dummy		
NSF Certified	National Sanitation Foundation (NSF) certified	0-1 Dummy		
Brand	Brand dummies for Maytag, Roper, Samsung, Whirlpool, Bosch, Estate, Frigidaire, GE, and LG	0-1 Dummy		
Remote Laundry Monitoring		0-1 Dummy		
Add a Garnment		0-1 Dummy		Yes
Clean Action		0-1 Dummy		
Electronic Control		0-1 Dummy		
Programmable Control		0-1 Dummy		Yes
Cycle Status End Signal		0-1 Dummy		Yes
Cycle Status Remaining Time		0-1 Dummy		Yes
Cycle Status Lights		0-1 Dummy		Yes
Delay Start		0-1 Dummy		
Bleach Dispenser		0-1 Dummy		
Detergent Dispenser		0-1 Dummy		
Fabric Softener Dispenser		0-1 Dummy		
Injection Dispenser		0-1 Dummy	Yes	Yes
Other Dispenser		0-1 Dummy	Yes	
Special Door Access		0-1 Dummy		
Dryer Ready		0-1 Dummy		
Heater		0-1 Dummy	Yes	Yes
Water Level Selector		0-1 Dummy	Yes	Yes
Water Level Sensor		0-1 Dummy	Yes	
Water Saving Technology		0-1 Dummy	Yes	
Advanced Motor Features		0-1 Dummy	Yes	
Extra Rinse		0-1 Dummy		
Sanitize Heat		0-1 Dummy		Yes
Sanitize Silver Ion		0-1 Dummy		
Sanitize Steam Technology		0-1 Dummy		
Sanitize Cycle		0-1 Dummy		
Smooth Balance		0-1 Dummy		
Smooth Suspension		0-1 Dummy		
Smooth Suspension		0-1 Dummy		
Soil Level Selector		0-1 Dummy		Yes
Soil Level Sensor		0-1 Dummy	Yes	
Maximum Spin Speed		0-1 Dummy	Yes	Yes
Spin Speed Option		0-1 Dummy		
Spin Timer Option		0-1 Dummy		
Cold Temperature Default		0-1 Dummy	Yes	
Temperature Selection		0-1 Dummy	Yes	Yes
Temperature Sensor		0-1 Dummy	Yes	
Automatic Timer		0-1 Dummy	Yes	
Quickwash		0-1 Dummy	Yes	Yes
Other Features Tub		0-1 Dummy		
Stainless Tub		0-1 Dummy	Yes	

The table lists the clothes washer attributes used in the analysis. Product characteristics include energy use, capacity, and number of cycles (continuous measures), as well as a set of binary indicators capturing certifications, brands, and specific washer features.