

The Effect of Product Misperception on Economic Outcomes: Evidence from the Extended Warranty Market *

Jose Miguel Abito Yuval Salant

January 25, 2017

Abstract

We use panel data on extended warranty purchases to investigate three potential drivers of the economic outcomes in this market: (1) pricing power due to consumers' search costs and the add-on feature of the warranty, (2) standard risk aversion, and (3) consumers' misperception of the insured product's failure probability. We estimate a model that incorporates these drivers, and perform counterfactual experiments to quantify their significance. We find that "shutting down" probability distortions leads to a decrease of more than 80% in market volume and profit, and to a dramatic increase in consumer welfare. We also find that probability distortions are reduced with information and experience.

*Abito: University of Pennsylvania, Wharton School, Business Economics and Public Policy, abito@wharton.upenn.edu. Salant: Northwestern University, Kellogg School of Management, Department of Managerial Economics and Decision Sciences, y-salant@kellogg.northwestern.edu. We thank Heski Bar-Isaac, JF Houde, Katja Seim, Eduardo Azevedo, Joe Harrington, Jeremy Tobacman, Tom Baker, Chris Conlon, and Nicola Persico for helpful comments.

1 Introduction

An extended warranty is an insurance contract that protects against the failure of a durable good such as a consumer electronic. The market for these warranties is highly profitable,¹ and has caught the attention of consumer protection and competition authorities. In the UK, the Office of Fair Trading observed that “there is insufficient competition and information to ensure that consumers get good value (in the extended warranty for consumer electronics market), and that many electrical retailers may make considerable profits on the sale of extended warranties” (UK Competition Commission, 2003). The UK Competition Commission has consequently conducted a comprehensive investigation of this market, and expressed in its final report serious concerns about the lack of market competition due to the way warranties are sold, and the lack of available information at the point of sale about the reliability of the insured good and the cost of repair (UK Competition Commission, 2003). The Federal Trade Commission in the US has also looked into this market, and is advising consumers to obtain information about the likelihood of product failure and the potential cost of repair before buying an extended warranty.²

There are potentially two forces that drive the economic outcomes in the extended warranty market: pricing power on the supply side and consumer behavior on the demand side. On the supply side, retailers benefit from significant pricing power because the warranty is an add-on product, whose price is not advertised and is not easily accessible. The warranty is usually offered to consumers immediately after they finalize their decision to buy the insured product at a stage in which it is costly for consumers to search for another product or switch to another retailer.³ This search and switching cost gives retailers significant pricing power when selling warranties a-la Ellison’s (2005) add-on pricing game.⁴

On the demand side, consumers’ choice behavior reveals very high willingness to pay for warranties. For example, in our panel data of extended warranty purchases described below, one in four TV buyers purchases an extended warranty. On average, these buyers are willing to pay \$90 or more to insure themselves against a loss of at most \$400 with 7% probability. Put differently, one in four TV buyers is willing to purchase insurance against a relatively small risk of a relatively small loss at a price that is three times the actuarially fair value.

A very high degree of risk aversion in the form of diminishing marginal utility for wealth is

¹The UK Competition Commission (2003) estimates that the top five consumer electronics retailers in the UK earned 116 to 152 million pounds annually on the sale of extended warranties in the early 2000s. Similarly, analysts in the US estimate that extended warranties accounted for almost half of BestBuy’s operating income in 2003, and that profit margins on warranties ranged from 50% to 60% (See “The Warranty Windfall”, *Business Week* (Dec 19, 2004)).

²See <https://www.consumer.ftc.gov/articles/0240-extended-warranties-and-service-contracts>

³For example, BestBuy trains its sales people to offer the warranty to a consumer only after the consumer finalizes the TV purchase decision. For BestBuy’s presentation of selling skills, see https://www.extendingthereach.com/wps/PA_VCorationFramework/resource?argumentRef=static&resourceRef=/files/Best_Buy_Vendor_Selling_Skills.pdf

⁴According to the UK Competition Commission (2003), “Most customers shop around for electronic goods; the retail price is a major factor in their choice. However, extended warranty buyers do not often plan to buy an extended warranty (less than half of consumers who bought an extended warranty said that they had planned to do so before they went into the store), and many are unaware of the existence of alternatives to taking the EW offered at POS.”

required to explain this behavior. On the other hand, the recent empirical work of Barseghyan et al (2013) highlights the importance of probability distortions in insurance choices. Barseghyan et al (2013) find that upward distortion of small claim probabilities and mild insensitivity to probability changes play an important role in explaining deductible choices in home and auto insurance. One possible reason for this distortion is that consumers know failure probabilities but overweight them as predicted by Prospect Theory (Kahneman and Tversky (1979)). Another reason — which is closely related to the concerns of competition authorities mentioned above — is that consumers in the extended warranty market do not have information about failure probabilities and overestimate them. Either way, the upward distortion of failure probabilities may cause consumers to perceive the warranty as more attractive than the objective failure probabilities imply, and thus influence their willingness to pay.

The goal of this paper is to study the relative importance of pricing power on the supply side, and standard risk aversion and probability distortions on the demand side in determining the economic outcomes in the extended warranty market.

Our first main finding is that probability distortions are an important driver of the high profits in the extended warranty market. Specifically, by “shutting down” probability distortions, market volume and profit would drop by more than 80%. Our second main finding is that probability distortions are reduced with information and experience suggesting that the source of the distortion is overestimation of unknown probabilities. Our third main finding is that consumer welfare in this market is negative, and policy interventions that focus on providing information to consumers are more effective in enhancing consumer welfare than interventions that focus on improving market competition.

Our model of consumer choice and market competition is based on Ellison’s (2005) add-on pricing game.⁵ Retailers decide on product prices and warranty prices. Consumers then decide which retailer to visit. When making this decision, consumers observe product prices but not warranty prices. Consumers visit the retailer of their choice, and learn the price of the warranty. They then decide whether to buy the product, the product and the warranty, or visit another retailer at a cost. Following Barseghyan et al’s (2013) analysis and findings, consumers are risk-averse expected utility maximizers who may distort failure probabilities.

To estimate consumers’ risk aversion and probability distortion and retailers’ cost of supplying the warranty, we use panel data on household-level product and extended warranty purchases from a large US electronics retailer. The dataset documents about 45000 transactions made by almost 20000 households between 1998 and 2004. Almost 30% of the transactions involved the purchase of an extended warranty. Our structural analysis focuses on TVs, which constitute about 11% of the data, due to the availability of TV failure rates from *Consumer Reports*.

On the consumer side, we construct the demand for warranties from consumers’ choice behavior,

⁵See also Gabaix and Laibson’s (2006) shrouded attributes model in which retailers can shroud information about the add-on, and myopic consumers do not anticipate the possible purchase of the add-on prior to visiting the retailer.

and separately identify and estimate consumers' degree of risk aversion and failure rate distortions. The identification strategy relies on how a household's willingness to pay for a warranty varies across products that have the *same* failure rate but *different* repair costs. We show that a single-crossing property of the willingness to pay function is sufficient for separate identification, and provide examples of utility functions that satisfy the single-crossing property, including the second-order Taylor approximation (Cohen and Einav (2007), Barseghyan et al (2011), and Barseghyan et al (2013)) that we use in the estimation.

The estimates indicate that there is a substantial upward distortion of failure probabilities that decreases as the failure probability increases. For example, a 5% objective failure probability is perceived as a 13% failure probability whereas a 10% failure probability is perceived as a 17% failure probability. These estimates are similar to the estimates of Barseghyan et al (2013) in the context of home and auto insurance. Standard risk aversion, on the other hand, plays an insignificant role in consumers' decision making: the average willingness to pay of consumers with our estimated risk aversion parameter who evaluate probabilities correctly is close to actuarially fair rates, consistent with the behavior of a risk neutral consumer. This estimate diverges from Barseghyan et al (2013) probably because the monetary losses associated with a TV failure are much smaller than home- or auto-associated monetary losses.

We proceed to demonstrate that probability distortions decrease with information and experience. First, we partition the households in the sample to two equally sized groups based on a measure of information that rewards households who visit the store frequently, and re-estimate the model separately for each group. The group of uninformed households, i.e. those who visit the store less frequently, distorts failure probabilities substantially more than the group of informed households that evaluate failure probabilities accurately. Second, regression analysis indicates that the likelihood of purchasing an extended warranty drops by 3 percentage points for every warranty a household bought in the past when controlling for household characteristics. This implies a decrease of more than 10% in the likelihood of buying a warranty based on past experience. Third, there is an even larger drop of 8 percentage points in the likelihood of buying a warranty for any warranty that a household returned in the past. Finally, about 33% of the returns of warranties *do not* involve the return of the insured product. Put differently, one in three consumers that returns an extended warranty within the retailer's 30-day return window does so without returning the product that the warranty insures.

These findings are relevant for identifying the mechanism behind probability distortions, and hence their welfare and policy implications. Kahneman and Tversky (1979) distinguish probability overweighting, which refers to assigning high decision weights to known low probability events, from probability overestimation, which refers to estimating incorrectly unknown low probability events. Overweighting is a preference parameter and thus welfare analysis is silent about it, whereas overestimation is a mistake in decision making that leads to welfare-inferior choices. To the extent that overestimation improves with information and experience but overweighting does not, the

above findings suggest that overestimation drives consumers' purchase behavior in the data, and that there is room for welfare-improving policies that attenuate the distortion.

On the retailer's side, the search and switching costs of consumers imply that retailers benefit from monopoly pricing power when selling warranties (Ellison (2005)). We use the first-order condition of the monopolist profit maximization problem to estimate the retailer's cost and profit margin. We estimate that the retailer's margin is about 46% of the price of the warranty, and that the remaining 54% include payment to third-party warranty providers, commissions to sales people and other costs.

After estimating the risk aversion, probability distortion, and cost parameters, we perform counterfactual experiments in order to quantify the effect of probability distortions on market outcomes. In the experiments, we compare outcomes in the existing market to outcomes in a counterfactual market in which retailers have the same pricing power but the distortion is "shut down" and consumers estimate failure probabilities correctly. Shutting down the distortion implies a reduction in consumers' willingness to pay, and hence an inward shift in the demand for warranties. This shift is expected to lead to a profit and price decrease but the effect on volume and welfare is ex-ante ambiguous.

A key assumption in the experiments is that TVs are priced close to cost whether consumers distort failure probabilities or not. To motivate this assumption, we first observe that there is little differentiation between electronics retailers, and that TV prices are easily observable. Thus, in the absence of extended warranties, TVs would have been priced close to marginal cost. The introduction of extended warranties complicates things because it implies that retailers may have incentives to price TVs below cost in order to stimulate demand for warranties. There are several reasons to believe that such price cutting, if it takes place, is insignificant. First, TV manufacturers have tight controls on the pricing and marketing practices of retailers to prevent retailers from pricing below cost. Second, as Ellison's (2005) analysis shows, cutting the prices of TV in order to stimulate the demand for warranties has an "adverse selection" effect: it attracts a disproportionate volume of consumers with low willingness to pay who will not buy warranties. This reduces retailers' incentives to cut prices on TVs. Third, when comparing the 2003 TV prices at BestBuy, which sold extended warranties then, and Target, which did not, prices are very similar: TV prices at BestBuy are on average 0.4 percentage points higher than in Target, and the median price difference is 0.

The counterfactual experiments demonstrate that probability distortions have a very large effect on market volume and profit. When shutting down probability distortions, volume and profit decrease by about 80%. Shutting down the distortion also leads to a reduction in prices that translates to a reduction of about 25% in the retailer's margin.

Shutting down the distortion also has a very large effect on consumer welfare. Consumer welfare in the existing market is negative because the positive welfare generated by consumers with true willingness to pay above price is dominated by the decrease in welfare due to consumers with true willingness to pay below the price who mistakenly buy warranties as a result of overestimation.

Shutting down the distortion has two positive effects on consumer welfare. First, fixing prices at the existing levels, consumers make better choices because they estimate failure probabilities correctly. Second, retailers optimally reduce prices in response to the shift in demand, which further increases consumer welfare. The counterfactual experiments demonstrate that the magnitude of the first effect is more than ten-fold larger than the magnitude of the second effect.

We use our estimates to evaluate possible policy interventions in the extended warranty market. Armstrong (2008) discusses two types of policy tools that competition authorities use to enhance market performance and consumer welfare. The more prominent tool is competition policies, which aim to enhance competition and drive down prices. But, as Armstrong (2008) argues, there are many reasons that prevent markets from delivering good outcomes to consumers even when competition is intense, which raises the need for consumer policies that aim to enhance consumer decision making.

In the context of extended warranties, a leading example of a competition policy is the proposal of the UK Competition Committee (2003) that retailers advertise and post the price of the warranty alongside the price of the product it insures. This policy has the potential to intensify competition and drive prices down to cost because it enables consumers to observe warranty prices prior to visiting a retailer or making a purchase decision, and thus eliminates consumers' search costs. An example of a consumer policy is to disclose the failure rate of the product to consumers prior to their decision to purchase the warranty. This policy has the potential, if fully effective, to eliminate probability distortions.

Counterfactual experiments illustrate that the above competition policy has a mild effect on consumer welfare, which remains negative even if warranties are priced at cost. This is because any gain in consumer welfare due to lower prices is dominated by the negative welfare effect generated by more consumers buying warranties even though they should not. On the other hand, the above consumer policy has a dramatic effect on consumer welfare, which becomes positive. This is because consumers make better choices when they estimate failure probabilities correctly. Thus, consumer policies may be more effective in improving consumer welfare than competition policies in markets with uninformed or biased consumers.

We proceed as follows. Section 1.1 reviews the related literature. Section 2 presents the data and reports empirical regularities on household purchase behavior over time and across outlets. Section 3 presents the model and the identification strategy. Section 4 describes the estimation procedure. Section 5 reports the results of the estimation. Section 6 describes the counterfactual experiments. Section 7 discusses the implications of our findings to consumer policies.

1.1 Related literature

This paper is related to the empirical literature on estimating risk preferences (see Barseghyan et al (forthcoming) for a recent survey.) Cohen and Einav (2007) and Sydnor (2010) use auto and home insurance deductible choices to estimate standard risk aversion. They observe that a very

high degree of risk aversion is required to explain consumer behavior in their data. Several other empirical studies incorporate probability distortions to the estimation of risk preferences, and find evidence of their relevance in various contexts including financial markets (Kliger and Levy (2009)), betting markets (Jullien and Salanié (2000)), Snowberg and Wolfers (2010), Chiappori et al (2012), Gandhi and Serrano-Padial (2012)), and auto and home insurance (Barseghyan et al (2013)).

Our paper is closest to Barseghyan et al (2013), who develop a structural model of risky choice with standard risk aversion and probability distortions, and estimate it using data on auto and home insurance deductible choices. Barseghyan et al's (2013) main finding is that probability distortions have an important role in deductible choices. We incorporate their utility specification to the add-on pricing game of Ellison (2005), and focus on how probability distortions affect market performance and consumer welfare.

We make three contributions to the literature on risk preferences and probability distortions. First, we quantify the effect of probability distortions on market outcomes including prices, volume, and profit. We are able to make progress on this question because electronics retailers have (1) monopolistic pricing power on the sale of warranties enabling us to estimate their cost, but (2) little flexibility in cutting TV prices below cost enabling us to conduct counterfactual experiments without adjusting TV prices. Second, by taking advantage of the fact that households in our data make multiple purchases, we provide evidence that the source of the distortion is overestimation of unknown probabilities rather than probability overweighting. This contribution is essential for assessing consumer welfare and potential policy interventions. Third, we estimate that consumer welfare in the extended warranty market is negative, and demonstrate using counterfactual experiments that consumer policies are more effective than competition policies in improving consumer welfare.

Another related literature is the marketing and experimental literature on why consumers buy extended warranties (Chen et al, 2009; Huysentruyt and Read, 2010; Jindal, 2014). Chen et al (2009) use purchase data of about 600 households from an unspecified US electronics retailer over the period November 2003 to October 2004 to study how the insured product characteristics (hedonic vs utilitarian) and marketing actions by retailers affect the likelihood of purchasing an extended warranty. Huysentruyt and Read (2010) use survey data of hypothetical choices to demonstrate that consumers overestimate the likelihood of washing machine breakdowns and the cost of repair. Jindal (2014) uses different survey data to highlight the role of loss aversion in the context of extended warranties for washing machines. Washing machines have failure rates of about 20% to 30% in the first four years of service, so it is expected that probability distortions have a less significant role in this context.⁶

⁶For example, in the Prelec (1998) specification, the weighting function gets flatter and flatter, and it crosses the 45 degree line at $1/e \approx 37\%$.

2 Data Characteristics

We use the INFORMS Society of Marketing Science Durables Dataset 1, which is a panel data of household durable goods transactions from a major U.S. electronics retailer. The full sample contains about 140,000 product-level transactions made by almost 20,000 households across the retailer’s 1,176 outlets and its online store. Prices across outlets and the online store are essentially identical. Transactions took place between December 1998 and November 2004.

There are four main types of transactions in the data. About 117,000 transactions involve the purchase of a specific product other than an extended warranty. About 15,000 transactions involve the purchase of an extended warranty. About 5000 transactions involve the return of a product other than an extended warranty and about 1000 transactions involve the return of an extended warranty. For each transaction, we observe the product ID, the price of the product, the brand, and the category and subcategory of the product.

A shopping trip is a collection of transactions made by a given household at a given store in a given date and time. For each household and shopping trip, we observe the buyer’s gender, the age and gender of the head of the household, income group⁷, and whether there are children in the household.

There are two data issues that we have to deal with. First, the data only tells us the product subcategory (e.g. 9-16 inch TVs) for which the warranty is for. We restrict our sample to shopping trips in which there is a clear one-to-one mapping between the extended warranty and the corresponding product. For example, we drop shopping trips involving a purchase of two 9-16 inch TVs but only one extended warranty purchased for this subcategory. We lose about 2,000 observations for this reason.

Second, if a household did not purchase an extended warranty for a given product, we do not observe the warranty’s price. To identify the warranty price in such cases, we match the non-warranty transaction with a corresponding warranty transaction for the same product ID from the closest transaction date. After dropping transactions for which we cannot find a corresponding warranty transaction, we end up with a sample of about 45,000 observations.⁸

2.1 Attachment rates, prices, and approximate profit margins

Table 1 shows the fraction of consumers who bought extended warranties (henceforth, the attachment rate) and the extended warranty-to-product price ratio for each product category. Attachment rates range from about 20% for items such as VCRs (VIDEO HDWR), music CDs and video games (MUSIC), to as high as 40% for items like car stereos and speakers (MOBILE). Warranties are priced on average at 24% of the price of the insured product, and the standard deviation of the

⁷Income group is a number from 1 to 9 where 9 is the highest income group. We do not have additional information on income within each group.

⁸We also drop the less than 1000 observations, in which the price of the good is significantly less than the price of the warranty.

Table 1: Attachment rate and price ratio by product category

	Attachment rate	EW-Product price ratio	Obs
AUDIO	0.281	0.232	6450
DVS	0.295	0.207	1439
IMAGING	0.377	0.199	3001
MAJORS	0.356	0.197	864
MOBILE	0.398	0.249	5176
MUSIC	0.208	0.169	1189
P*S*T	0.245	0.237	3765
PC HDWR	0.258	0.274	8773
TELEVISION	0.311	0.217	6307
VIDEO HDWR	0.206	0.240	5828
WIRELESS	0.245	0.317	1485
OVERALL	0.287	0.239	44277

Notes: Overall attachment rates and EW-Product price ratios are sales-weighted averages.

Table 2: EW information for TVs

	Attach rate	TV price	EW-TV price ratio	Fail rate	Margin	Obs
9-16in	0.149	122.99	0.284	0.072	0.729	422
19-20in	0.176	173.97	0.240	0.065	0.710	1067
25in	0.270	245.33	0.220	0.069	0.643	522
27in	0.299	354.06	0.197	0.059	0.681	1477
>30in	0.348	812.53	0.219	0.076	0.619	1229
OVERALL (TV)	0.268	400.07	0.223	0.067	0.672	4717

Notes: Fail rates are from *Consumer Reports*. Overall numbers are sales-weighted averages. Margin is computed as

$(\text{EW price} - \text{fail rate} \times \text{TV price}) / \text{EW price}$.

warranty-to-product price ratio is 11% (see Figure 1 for the distribution of ratios.) There is no significant correlation at the product level between variations in the product price and variations in the warranty price.⁹

Our structural analysis focuses on TVs due to the availability of TV published failure rates from *Consumer Reports*. The failure rate of a product is the likelihood that the product needs repair within 3 to 4 years of purchase. Table 2 reports attachment rates, prices, extended warranty-to-product price ratios, and failure rates broken down by TV subcategory. Attachment rates range from 15% to 35%, with higher attachment rates for more expensive categories. The average price ratio for TVs is about 22% with a standard deviation of 8% (see Figure 1 for the distribution of ratios.)

Using the price of a TV multiplied by its failure rate as the expected cost of servicing a TV warranty, Table 2 also reports a “back of the envelope” profit margin on TV warranties. This margin ranges from 62% to 73% for different TV subcategories, which is close to what is cited in the popular press. We expect the seller in our dataset to have lower margins due to revenue sharing with warranty providers and commissions to sales people.

⁹We regress the log of the product price on the log of the warranty price for each product, and estimate an average coefficient equal to 0.046 with an average p-value of 0.26.

Table 3: Attachment rates by buyer and household characteristics

Characteristic	All		TV	
	Attach rate	Obs	Attach rate	Obs
Female	0.305	13976	0.286	1534
Male	0.280	26228	0.261	2768
Female (head of hh)	0.308	12412	0.285	1380
Male (head of hh)	0.280	24760	0.260	2620
Below median income (category < 5)	0.321	10404	0.310	1170
Above median income (category \geq 5)	0.276	33900	0.253	3547
Lowest income category (category = 1)	0.343	2656	0.340	300
Highest income category (category = 9)	0.253	6452	0.233	660
Over 50 (head of hh)	0.293	23259	0.282	2717
Under 50 (head of hh)	0.280	20882	0.250	1975
Has child in hh	0.282	13940	0.248	1279
No child in hh	0.296	6234	0.323	779

2.2 Buyers' characteristics

Tables 3 and 4 examine the relationship between attachment rates and buyers' characteristics for all product categories and for TVs. In Table 3, attachment rates are broken down by buyer gender, gender and age of the head of the household, income, and whether there is a child in the household. Income is the only characteristic that is strongly correlated with attachment rates for TVs and all other product categories. For example, when moving from the highest to the lowest income category, there is an increase of almost 11 percentage points in TV attachment rates. Having a child seems to decrease TV attachment rates by 7 percentage points, but this effect goes away once we introduce controls in a regression analysis.

Table 4 presents the results of regressing an extended warranty purchase dummy on buyers' and households' characteristics and their interactions with gender. The regressions include brand and subcategory fixed effects to account for average differences in purchasing behavior across these dimensions. Consistent with most of the raw means in Table 3, the only characteristic that is statistically and economically significant is income when including all product categories. Adjusted R^2 's are very small despite including subcategory and brand fixed effects. All in all, the two tables indicate that the above characteristics (except perhaps income) are not strongly correlated with warranty purchases.

2.3 Warranty purchases and returns over time

The data tracks households purchase and return decisions over time, so we can examine how past purchases and returns of warranties influence future purchases of warranties. Table 5 contains the results of regressing a warranty purchase dummy for a *given* product on whether an extended warranty was purchased and whether an extended warranty was returned in the past for *any other* product. We also include as a regressor a dummy, which is equal to 1 if the transaction was made in

Table 4: Regression of EW purchase dummy on buyer and household characteristics

Dependent variable: EW purchase dummy				
	All		TV	
	Coeff	SE	Coeff	SE
Male	-0.064	(0.039)	-0.069	(0.101)
Age (head)	0.001*	(0.0004)	0.001	(0.001)
Income	-0.014***	(0.003)	-0.013*	(0.007)
Has child in hh	$< 10^{-5}$	(0.014)	-0.017	(0.038)
Male \times Age	0.001	(0.001)	0.001	(0.001)
Male \times Income	0.002	(0.003)	$< 10^{-4}$	(0.009)
Male \times Child	0.004	(0.017)	-0.006	(0.045)
Subcategory FE	Y		Y	
Brand FE	Y		Y	
Adjusted R ²	0.08		0.06	
No. obs (good-hh-trip)	19375		1973	

Notes: Standard errors in parentheses are clustered at shopping trip level.

Significance level: ***1%, **5%, *10%

store as opposed to online. Our preferred specification is specification IV, which includes household, subcategory, brand, month and year fixed effects, and uses the number of past extended warranty purchases and returns as regressors.

The results provide evidence of learning from experience.¹⁰ Buying an extended warranty in the past is associated with a 15 percentage points decrease in the likelihood of buying a warranty today. This decrease is more than half of the average attachment rate across products (28.7%). When experience is measured in terms of the number of extended warranties bought in the past, buying an additional extended warranty in the past is associated with a decrease of 3 percentage points in the likelihood of buying a warranty today.

The effect of past returns on the likelihood of buying an extended warranty is even more profound. Returning an extended warranty in the past is associated with a 20 percentage points decrease in the likelihood of purchasing a warranty today, and each returned warranty is associated with a decrease of 8 percentage points in the likelihood of buying another warranty.

There are 1239 warranty return transactions in the data. About 67% of them are returns that accompany the insured product return. These returns are probably made due to the add-on feature of the warranty – it has no value if the insured product is returned. But 33% of the warranty returns are made without returning the main product.¹¹ These returns are another indicator that consumers learn ex-post that insuring against product failure is less attractive than they initially thought.

¹⁰When household fixed effects are not included, a past warranty purchase has a positive effect on the likelihood of buying a warranty today, contrary to learning. This reflects the classic problem of disentangling unobserved persistent heterogeneity and state dependence.

¹¹We run regressions similar to Table 4 and find that none of the buyer or household characteristics in the data is strongly correlated with this return behavior.

Table 5: Regression of EW purchase dummy on past EW purchases and returns

Dependent variable: Buy EW today?				
	I	II	III	IV
Bought EW before?	0.257*** (0.009)	. (.)	-0.152*** (0.024)	. (.)
Returned EW before?	0.041** (0.017)	. (.)	-0.197*** (0.031)	. (.)
No. of EW bought before	. (.)	0.071*** (0.004)	. (.)	-0.033*** (0.007)
No. of EW returned before	. (.)	-0.015 (0.012)	. (.)	-0.082*** (0.023)
In-store?	0.154*** (0.016)	0.169*** (0.015)	0.181*** (0.039)	0.173*** (0.038)
Household FE	N	N	Y	Y
Subcat & brand FE	Y	Y	Y	Y
Month & Yr FE	Y	Y	Y	Y
No. obs	14878	14878	14878	14878
(good-hh-trip)				
No. HHs	6321	6321	6321	6321

Notes: Standard errors in parentheses are clustered at shopping trip level. Dependent variable refers to a given product while “Bought” and “Returned” dummy regressors refer to buying and returning an EW for any product at some point in the past.

Models II and IV use the number of EWs bought and returned as regressors.

Significance level: ***1%, **5%, *10%

Table 6: Regression of extended warranty purchase dummy on shopping mode

Dependent variable: EW purchase dummy					
	I	II	III	IV	V
In-store?	0.247*** (0.022)	0.200*** (0.027)	0.180*** (0.027)	0.175*** (0.027)	0.166*** (0.027)
Household FE	N	Y	Y	Y	Y
Subcategory FE	N	N	Y	Y	Y
Brand FE	N	N	N	Y	Y
Month FE	N	N	N	N	Y
Year FE	N	N	N	N	Y
No. obs	44304	44304	44304	44304	44304
(good-hh-trip)					
No. HHs	17158	17158	17158	17158	17158

Notes: Standard errors in parentheses are clustered at shopping trip level.
Significance level: ***1%, **5%, *10%

2.4 In-store versus online transactions

About 1% of the transactions in the data were made online. The attachment rate for these transactions is about 4%, which is one-seventh of the in-store attachment rate. To examine what drives this sevenfold difference and its robustness, we explore various regressions in Table 6. The first model does not include any controls so it gives numbers that are very similar to the raw attachment rates. The other models turn on various fixed effects. Subcategory and brand fixed effects allow us to soak up any differences in mean purchasing behavior induced by the nature of the product. We also include household, month and year fixed effects as further controls.

We see a drop in the effect of in-store purchases as we add fixed effects. Including a household fixed effect reduces the effect by about 5 percentage points. Including product-related fixed effects reduces the effect by additional 2.5 percentage points. Including all the fixed effects leads to a reduction in the effect from 25 to 17 percentage points. That is, the likelihood of purchasing an extended warranty jumps from 12% to 29% when being in the store.

As an additional robustness check, Table 7 contains the results of regressing an extended warranty purchase dummy on shopping mode broken down by product category. The left specification is a simple OLS regression, the middle one adds household characteristics as controls, and the right one adds household fixed effects while removing the household characteristics. There is significant variation in the effect of in-store purchases across the product categories but overall the effect remains large. The effect survives even when we include household characteristics. Although we lose statistical significance once we include household fixed effects due to a small number of observations, the magnitudes are similar across the different regression specifications.

Table 7: Regression of EW purchase dummy on shopping mode broken down by product category

Dependent variable: EW purchase dummy									
	OLS	se	Obs	OLS with char	se	Obs	FE (HH)	se	Obs
AUDIO	0.16***	(0.05)	6450	0.12*	(0.07)	2517	0.07	(0.10)	6450
DVS	0.05	(0.23)	1439	0.34	(0.46)	586	0.50*	(0.28)	1439
IMAGING	0.36***	(0.07)	3001	0.39***	(0.12)	1197	.	(.)	3001
MOBILE	0.40**	(0.19)	5176	0.37	(0.28)	1883	.	(.)	5176
MUSIC	0.16*	(0.09)	1189	0.20	(0.15)	469	0.30	(0.24)	1189
P*S*T	0.23***	(0.06)	3765	0.23***	(0.08)	1519	0.19*	(0.11)	3765
PC HDWR	0.18***	(0.05)	8773	0.09	(0.08)	3471	0.03	(0.11)	8773
TV	0.31***	(0.08)	6307	0.32**	(0.14)	2360	0.22	(0.20)	6307
VIDEO HDWR	0.21***	(0.04)	5828	0.22***	(0.07)	2151	0.17*	(0.09)	5828
WIRELESS	0.25	(0.18)	1485	0.27	(0.31)	483	0.57**	(0.25)	1485

Notes: The left specification is simple OLS. The middle one adds household characteristics as controls. The right specification adds household fixed effects while removing household characteristics.

Significance level: ***1%, **5%, *10%

3 Model and Identification

We consider an add-on pricing model a-la Ellison (2005) and Ellison and Ellison (2009).

There are several sellers of a main product M and an extended warranty EW for the product M . Each seller sets a price p for the product that is observable to buyers, and a price t for the warranty that is not.

The assumption that the product price is observable and the warranty price is not seems to be the case in practice. For example, BestBuy advertises product prices but not warranty prices. In fact, BestBuy trains its sales people to offer warranties and other add-ons to buyers only after they finalize their decision to purchase the product.¹²

Buyers decide which seller to visit after observing the price of the product across sellers, and forming rational expectations about warranty prices.¹³ Buyers visit the seller of their choice at a cost s and learn the price of the warranty. The cost s corresponds to the hassle or time involved in visiting a store and going through the purchase process. Buyers then decide whether to buy the main product, the main product and the warranty, or visit another store at a cost of s , where they will face the same decision.

Relevant equilibrium properties. As Ellison (2005) shows, there are two properties of any pure strategy sequential equilibrium of the above game. The first is that the price of the warranty set by any seller is the monopoly price relative to the demand for warranties among buyers of the

¹²See footnote 3.

¹³The assumption that buyers form rational expectations about warranty prices is not necessary for our empirical analysis. One could alternatively assume, as in Gabaix and Laibson (2006), that buyers do not plan to purchase a warranty prior to visiting the seller, and decide which seller to visit based solely on the price of the main product. See footnote 4 for evidence that many buyers indeed do not plan to buy a warranty. In this alternative specification, buyers form rational expectations about warranty prices of other sellers after visiting the first seller and being offered the warranty.

product. Otherwise, as in Diamond (1971), the seller can raise the price of the warranty by some $\epsilon < s$ and buyers will not switch to another seller. We use the first order condition of this monopoly pricing problem to estimate the seller’s cost of providing warranties.¹⁴

The second property is that buyers visit only one seller and always buy the product in equilibrium. This is because buyers incur a cost of visiting a seller. Thus, if they anticipate they will not buy the main product, they will not visit the store. We therefore focus below on buyers’ decision to buy the warranty conditional on already purchasing the product.

Warranty purchase decision. Following Barseghyan et al (2013), we model buyers as risk-averse expected utility maximizers who may distort failure probabilities.

Let W denote the buyer’s wealth after buying the main product, t the price of the warranty, and $u(\cdot, r)$ the buyer’s concave utility over wealth levels that is parameterized by r , the buyer’s degree of risk aversion around W .

A buyer’s utility if he purchases the warranty is $V_{EW} = u(W - t; r)$.¹⁵ A buyer’s utility if he does not purchase the warranty is $V_{NW} = \omega(\phi)E(u(W - X; r)) + (1 - \omega(\phi))u(W; r)$ where $\omega(\phi)$ is the probability distortion function, which increases in ϕ , and X is the random cost of repair. We assume that X is smaller than the price of the main product p because the buyer can always buy a new product instead of fixing the existing one. Thus, the buyer’s utility if he does not purchase the warranty is bounded below by $\omega(\phi)u(W - p; r) + (1 - \omega(\phi))u(W; r)$. We will identify V_{NW} with this lower bound in our estimation, i.e., we will have $V_{NW} = \omega(\phi)u(W - p; r) + (1 - \omega(\phi))u(W; r)$.¹⁶

The non-standard component in the buyer’s utility is the probability distortion function $\omega(\cdot)$. There are at least two potential sources of probability distortions in the context of warranties. The first is overestimation of unknown small probabilities. Estimating failure probabilities is not straightforward; buyers usually have limited personal experience with failures of durable goods, and credible information on failure rates is not readily available at the point of sale. The common view is that lack of experience or information leads to overestimation of failure probabilities.¹⁷ This is probably why the FTC encourages consumers to research the likelihood of product failure before buying a warranty.¹⁸

The second potential source of probability distortions is overweighting of known small probabilities. Prospect Theory (Kahneman and Tversky (1979)) proposes that individuals incorporate probabilities in decision making by using decision weights. In particular, individuals tend to put too much weight on low probability events, such as the failure probability of a durable good.

Demand for warranties. Observationally equivalent households may make different purchase

¹⁴We do not use the first order condition to estimate the demand side parameters.

¹⁵This assumes that there is no deductible associated with using the warranty as is often the case in practice.

¹⁶This implies that we likely underestimate ω and r because using a higher repair cost makes the purchase of the warranty more attractive for any ω and r . We discuss the robustness of our estimates to this assumption in Section 5.

¹⁷For example, the New York Times on August 28, 2014 writes: “The company selling the warranty has the information on failure rates. You don’t....That’s not easy to find out. Companies aren’t in the habit of telling you that their products fail 4 percent or 12 percent of the time. Failure rates are usually low. Warranty companies know that. And they know, too, that consumers tend to think the failure rate is higher.”

¹⁸See footnote 2.

decisions due to various unobserved factors such as a preference for flexibility in choosing a repair facility, doing-it-yourself, or buying a new product instead of repairing. We account for this unobserved heterogeneity by incorporating additively separable individual choice shocks, ϵ_{EW} and ϵ_{NW} , to V_{EW} and V_{NW} . Assuming these shocks are iid Type I Extreme Value with scale parameter σ and normalizing the buyer population to 1, we can derive the demand for warranties:

$$\begin{aligned} D(t; r, \omega(\phi), p, \sigma) &= \Pr(\epsilon_{NW} - \epsilon_{EW} \leq \Omega(t; r, \omega(\phi), p, \sigma)) \\ &= \frac{\exp \Omega(t; r, \omega(\phi), p, \sigma)}{1 + \exp \Omega(t; r, \omega(\phi), p, \sigma)} \end{aligned} \quad (1)$$

where

$$\Omega(t; r, \omega(\phi), p, \sigma) \equiv \frac{V_{EW} - V_{NW}}{\sigma}.^{19} \quad (2)$$

3.1 Identification of risk aversion and probability distortion

We focus on the identification of risk aversion and probability distortions from willingness to pay (WTP).²⁰

Fix a product M with price p_M and failure rate ϕ , and let $\omega = \omega(\phi)$. The willingness to pay $WTP(p_M, r, \omega)$ of buyers with risk aversion r and the distorted probability ω for a warranty to product M is the price t that solves $V_{EW}(t, r) = V_{NW}(p_M, r, \omega)$. The identification problem is that the same WTP can be explained by a continuum of pairs $(r, \omega(r))$ where r is the degree of risk aversion and $\omega(r)$ is the distorted probability as a function of r . This is because an increase in r can be undone by a decrease in ω .

We exploit variation in prices for two products with the *same* failure rate ϕ in order to uniquely identify the pair (r, ω) . Because the failure rate is the same, the same pair (r, ω) should explain the different WTP for warranties to these two products. The pair (r, ω) can then be identified uniquely if the two iso-WTP “curves”, i.e., the two continuums of pairs $(r, \omega(r))$ that explain the different WTPs, cross each other exactly once.

Figure 2 provides graphical intuition. The solid curve is the iso-WTP curve for product M with price p_M . Without additional variation, we cannot uniquely identify the pair (r, ω) against another pair such as (r', ω') because the pairs lie on the same iso-WTP curve, and so can rationalize the same willingness to pay. If, however, we also have data on the WTP for another product M' with the same failure probability but with a different price, then the pair (r, ω) can be uniquely identified as long as the iso-WTP curve for product M' (the dashed curve in Figure 2) intersects the iso-WTP curve for product M exactly once.

Proposition 1 presents a sufficient condition on a family of instantaneous utility functions that guarantees this “single-crossing” property and hence unique identification.

¹⁹The utility specification we will use in estimation imposes a specific normalization so we can identify the scale parameter σ . This scale parameter is the inverse of the marginal utility of income.

²⁰WTP can be uniquely obtained from choice probabilities.

Proposition 1 *Let $\{u(\cdot, r)\}_r$ be a family of utility functions parametrized by the degree of risk aversion r such that a larger r is associated with more aversion to risk. The pair (r, ω) is uniquely identified if the marginal utility of wealth $u_x(x; r)$ does not increase in r .*

Proof. Following the discussion in the main text, it suffices to prove that as we change the price p_M of the main product, the slope of the iso-WTP curves $\frac{d\omega}{dr}$ changes monotonically. We will do so by showing that

$$-\frac{\frac{\partial WTP(p_M, r, \omega)}{\partial r}}{\frac{\partial WTP(p_M, r, \omega)}{\partial \omega}} = \frac{d\omega}{dr}$$

is strictly monotone in p_M .

Let $\Omega = u(W - t; r) - \omega u(W - p_M; r) - (1 - \omega) u(W; r)$. The WTP is defined as the price t for which $\Omega = 0$. Thus, by the implicit function theorem,

$$\begin{aligned} \frac{\partial WTP}{\partial r} &= \frac{\partial \Omega / \partial r}{\partial \Omega / \partial t} = \frac{u_r(W - t; r) - \omega u_r(W - p_M; r) - (1 - \omega) u_r(W; r)}{u'(W - t; r)}, \text{ and} \\ \frac{\partial WTP}{\partial \omega} &= \frac{u(W; r) - u(W - p_M; r)}{u'(W - t; r)} > 0. \end{aligned}$$

The numerator in the first expression $\partial \Omega / \partial r$ is positive because larger r implies more risk aversion and hence a lower certainty equivalent for the lottery $(\omega, -p_M)$. Thus, Ω goes up when r goes up.

Observe that

$$\begin{aligned} \frac{\partial}{\partial p_M} \left[\frac{\partial WTP / \partial r}{\partial WTP / \partial \omega} \right] &= \frac{\partial}{\partial p_M} \left[\frac{u_r(W - t; r) - \omega u_r(W - p_M; r) - (1 - \omega) u_r(W; r)}{u(W; r) - u(W - p_M; r)} \right] \\ &= \frac{\omega u'_r(W - p_M; r) [u(W; r) - u(W - p_M; r)]}{(u(W; r) - U(W - p_M; r))^2} \\ &\quad - \frac{[u_r(W - t; r) - \omega u_r(W - p_M; r) - (1 - \omega) u_r(W; r)] u'(W - p_M; r)}{(u(W; r) - U(W - p_M; r))^2}. \end{aligned}$$

Let us examine the second term first. As indicated above the term inside the brackets is positive. The derivative of u with respect to W is positive, and hence the entire expression with the minus before is negative. Monotonicity is thus guaranteed if the cross derivative of u with respect to (W, r) is non-positive. ■

The family of CARA utility functions $\{-e^{-rx}\}_r$ satisfies the condition of Proposition 1 for sufficiently large wealth levels. This is because the cross derivative with respect to x and r , $e^{-rx} - rxe^{-rx}$, is negative for $x > \frac{1}{r}$. It is also straightforward to verify that the utility specification we use in the estimation satisfies the condition of the proposition.

4 Estimation

We first describe how we estimate the risk aversion parameter and the probability distortion function, and then how we estimate the seller's cost.

Following Cohen and Einav (2007), Barseghyan et al (2011), and Barseghyan et al (2013), we use a second order Taylor approximation of buyers' utility function $u(\cdot)$ in estimating the model. The main benefit of using this specification is that it does not require data on wealth.

The second order Taylor approximation of $u(\cdot)$ around W for some wealth deviation Δ is given by

$$u(W + \Delta) \approx u(W) + u'(W)\Delta + \frac{u''(W)}{2}\Delta^2.$$

Dividing by $u'(W)$ and letting $r = -u''(W)/u'(W)$ denote the Arrow-Pratt coefficient of absolute risk aversion,²¹ we obtain that

$$\frac{u(W + \Delta)}{u'(W)} \approx \frac{u(W)}{u'(W)} + \Delta - \frac{r}{2}\Delta^2.$$

Using this specification to evaluate the utility difference Ω_j between purchasing and not purchasing a warranty for product j (equation 2), we obtain that:

$$\Omega_j = \frac{\omega_j p_j - t_j + \frac{r}{2}(\omega_j p_j^2 - t_j^2)}{\sigma}.$$

Let D_j be the observed attachment rate for product j . Our choice model implies that²²

$$\log \frac{D_j}{1 - D_j} = \Omega_j = \frac{\omega_j p_j - t_j + \frac{r}{2}(\omega_j p_j^2 - t_j^2)}{\sigma}. \quad (3)$$

The decision weight ω_j acts like a (non-additive) product effect. We decompose this effect to

$$\omega_j = \omega(\phi_j) + \xi_{k(j)} + \eta_j \quad (4)$$

where $\omega(\cdot)$ is some unknown function of ϕ , $\xi_{k(j)}$ is a subcategory-level effect, and η_j is a random shock. The parameters $\xi_{k(j)}$ allow decision weights to vary between subcategories, thus capturing the possibility that consumers may apply different decision weights for TVs of different sizes, different projection technology, etc.

Using equation 3, we can express the ω_j 's as a function of the unknown parameters (r, σ) and

²¹Strictly speaking, the Arrow-Pratt coefficient of absolute risk aversion can vary with income if the utility specification is not CARA.

²²A complication in linking Ω_j to product-level attachment rates arises because p and t vary at the product level. With infinite data, we could compute $D_{j|p,t} = \Pr(d_i = 1 | j(i) = j, p, t)$ and then invert this equation to get Ω_j for each pair (p, t) . But with our finite data, we need to calculate the product-level attachment rate by aggregating over all price pairs (p, t) for the particular product. As for prices, we use the highest p and the lowest t in the estimation because these make the purchase of the warranty most attractive even without appealing to risk aversion and probability distortions.

the data:

$$\omega_j = \frac{\sigma\Omega_j + t_j + \frac{r}{2}t_j^2}{p_j + \frac{r}{2}p_j^2}. \quad (5)$$

We construct moment conditions involving ω_j to estimate r and σ .²³ Once we have these parameters, we calculate ω_j using equation 5. Our assumption regarding the error structure in equation 4 implies the following moment condition

$$E[\omega_j - \omega_{j'} | \phi_j = \phi_{j'}, k(j) = k(j'), p_j, p_{j'}, t_j, t_{j'}] = 0$$

since $\omega_j - \omega_{j'} = \eta_j - \eta_{j'}$ for j, j' such that $\phi_j = \phi_{j'}$ and $k(j) = k(j')$. Thus, as long as failure rates for *any* two products j and j' belonging in the same subcategory are the same, decision weights associated with extended warranties for these products will be equal, on average.

For the purpose of the counterfactual experiments, we fit the non-parametric product effects ω_j 's with the one-parameter Prelec (1998) function $\omega(\phi) = \exp[-(-\log(\phi))^\alpha]$, which is used in the literature to model probability distortions.

We also estimate the marginal cost of the seller c_j of selling and serving a warranty for product j from the first order condition of the seller's profit maximization problem:

$$\frac{t_j - c_j}{t_j} = \frac{1}{|\mathcal{E}(t_j; r, \omega(\phi), p_j, \sigma)|} \quad (6)$$

where $\mathcal{E}(t_j; r, \omega(\phi), p_j, \sigma)$ is the price elasticity of demand for warranties to product j .

For the purpose of the counterfactual experiments, we fit the non-parametric cost estimates c_j 's using the cost function $c(p_j, \phi) = p_j f(\phi)$. That is, the cost of selling and serving the warranty increases linearly with the price of the main product p_j , where the exact slope $f(\phi)$ is determined by the failure rate. The cost function increases in p_j because p_j is likely positively correlated with the production cost of product j , which in turn is likely positively correlated with the repair cost of product j conditional on failure. Having a constant in the formulation of f allows capturing cost components that are incurred regardless of product failure, e.g. commissions to sales people.

5 Results

We first report the estimate of the probability distortion function. We then report the estimate of the risk aversion parameter. Finally, we report the estimated cost of the seller.

5.1 The relevance of probability distortions

Figure 3 includes a scatter plot of the estimated non-parametric product effects ω_j 's, and a local linear fit based on the non-parametric values with a bootstrapped 95% pointwise confi-

²³Because of the large variation in prices across categories, we estimate separate scale parameters for each subcategory, and we allow for heteroskedasticity. Specifically we let $\sigma_k(p) = \sigma_k \sqrt{p}$ for subcategory k .

Table 8: Estimate of Probability Distortion

Fail rate intervals	Average Fail rate	Average Distorted Prob	Convexity measure
4–6	5.0	13.4	—
6–8	6.8	16.5	1.68
8–10	9.3	17.5	0.40
≥ 10	13.3	17.9	0.11

Notes: Columns 1,2, and 3 are in %. In column 1, the upper bound is not included in the bin.

dence band. The figure also includes a fit according to the one-parameter Prelec (1998) function $\omega(\phi) = \exp[-(-\log(\phi))^\alpha]$ in which $\alpha = 0.652$. The Prelec function falls within the 95% confidence band, it is below the local linear fit for smaller probabilities, and above it for larger probabilities.

Table 8 presents the estimated product effects from a different perspective. It partitions failure probabilities to bins of 2%. For each bin, it reports the average failure rate and the average estimated distorted probability in the bin. It also reports the convexity measure $(\omega(\phi') - \omega(\phi)) / (\phi' - \phi)$, where $\phi' > \phi$ are two adjacent average failure rates in the Table. This convexity measure is the increment in the average distortion divided by the increment in the average failure rate.

Two patterns emerge in the Figure and the Table. First, there is substantial upward distortion of failure probabilities. This is illustrated in Figure 3 by the vast majority of the estimated w_j 's and the 95% confidence band around the local linear fit lying above the 45 degree line. This is also illustrated in Table 8 where, for example, products with an average failure rate of 5% are perceived as products with a 13.4% failure probability. Second, the degree of the upward distortion decreases as the failure probability increases. This is illustrated in Figure 3 by the concavity of the local linear fit, and in Table 8 by the decrease in the convexity measure. These two patterns are in line with the shape of the probability weighting function in Prospect Theory, and with the recent empirical findings of Barseghyan et al (2013) on probability distortions.

Based on the reduced form evidence regarding the effect of past warranty purchases and returns on the likelihood of future warranty purchases, and based on the concerns of competition authorities in the UK and the US, we conjecture that the source of the distortion is households' lack of information on failure rates. In particular, uninformed households distort failure probabilities more than informed households. To examine this conjecture, we construct a crude measure of informativeness, divide our sample to two equal groups of households based on this measure, and re-estimate the model for each group separately.

We measure how informed a household is by the number of shopping trips made by members of the household divided by the time that has passed since the household's first transaction date in the dataset. We interpret a larger value of the measure as reflecting a more informed household for three reasons. First, holding a household's first transaction date constant, additional visits to the store make buyers more informed about product characteristics and sales practices. Of course, frequent visits to the store also indicate that the buyer is electronics-savvy, which is likely also positively correlated with being more informed. Second, fixing the number of visits to the store,

visits are likely more informative when the time gap between them is smaller. This is because inputs to decision making decay over time according to leading cognitive theories.²⁴ Third, the extended warranty market received increasing attention in the popular press and in consumers' forums in later years relative to earlier years in the dataset. All else equal, this likely makes households that entered the dataset later more informed on average than households that entered the dataset earlier.

We first regress a household's extended warranty purchase dummy on the informativeness measure. We find that an increase of one standard deviation in the measure decreases the likelihood of purchasing an extended warranty by 2 percentage points.

We then split the sample into two groups. A household is labeled as informed if its informativeness measure is above the median value of the measure, and as uninformed if its measure is below the median value. When comparing the two groups, informed households are 3 percentage points less likely to buy an extended warranty than uninformed households.

We re-estimate the structural model separately for the two groups of informed and uninformed households. The results appear in Figure 4, which plots the local linear fits of the probability distortion function for the two groups. The figure also includes a 95% bootstrapped confidence band around each fit. Clearly, uninformed households distort failure probabilities substantially more than informed households, which evaluate failure probabilities quite accurately.

5.2 The irrelevance of risk aversion

We estimate the risk aversion parameter r to be practically zero. Specifically, r is approximately $\approx 10^{-6}$ with a 95% confidence interval that has width of less than 10^{-6} . We obtain similar estimates when the sample is partitioned according to the informativeness measure, or when we allow r to depend on gender, age, income and having a child. The parameter r becomes economically relevant only when we turn off the probability distortion by imposing $\omega(\phi) = \phi$ in the estimation. Its value is 0.036 with 95% confidence interval of (0.026, 0.046).

To interpret the effect of risk aversion on buyers' behavior, Table 9 presents the willingness to pay (WTP)²⁵ for an extended warranty for a product worth \$400 under various failure rates. Columns 2 and 3 present the WTP using the estimated risk aversion parameter from the full model (using the Prelec distortion function). In column 2, we compute WTP with our estimated distortion function, and in column 3, we impose $\omega(\phi) = \phi$. Column 4 uses the estimated risk aversion parameter from the standard model without probability distortions.

Columns 2 and 3 illustrate that in the full model, the contribution of probability distortions to WTP is significantly larger than that of standard risk aversion. In fact, when we turn off the probability distortion in column 3, WTP is essentially equal to the actuarially fair rate. On the

²⁴An example is the class of Adaptive Character of Thought (ACT) theories, developed by Anderson and colleagues (see e.g. Anderson (1993)).

²⁵The average WTPs are computed by setting the choice shocks, ϵ_{EW} and ϵ_{NW} , and the shock that enters the distortion function, η_j , to zero.

Table 9: WTP for EW on a good with value \$400

Failure rate	Full model with ω	Full model with $\omega(\phi) = \phi$	Standard model with $\omega(\phi) = \phi$
0.04	46.97	16.00	78.88
0.06	56.23	24.00	96.92
0.08	64.24	32.01	112.13
0.10	71.50	40.01	125.53
0.12	78.24	48.01	137.64
0.14	84.61	56.01	148.78

other hand, Column 4 illustrates that when we do not account for probability distortions in the estimation, the estimated risk aversion parameter implies a high degree of aversion to risk that is more sensitive to variations in the failure probability than in the full model.

In the estimation, we use the main product's price p as buyers' perceived cost of repair conditional on failure. But buyers may view the repair cost as smaller since they can always buy a new product instead of repairing the existing one. It is possible that as buyer's perceived repair cost decreases, the relative importance of risk aversion increases and that of the probability distortion decreases. As a robustness check, we re-estimate the model allowing the repair cost to vary as a function of p . Similarly to above, we use the Prelec function to fit the non-parametrically estimated product effects.

Table 10 reports the estimates of the risk aversion parameter r and the Prelec parameter α as we vary the repair cost. The risk aversion parameter continues to be economically non-significant while the probability distortion becomes more important in the sense that α decreases. Put differently, as we reduce the repair cost and thus make purchasing a warranty less attractive, the estimated degree of probability distortion increases until its explanatory power is exhausted, but the risk aversion parameter continues to be insignificant.

Table 10: Parameter estimates with varying expected cost

% of product price	r	Prelec α	$\omega(0.05)$
100%	$\approx 10^{-6}$	0.652	0.129
90%	$\approx 10^{-6}$	0.597	0.146
80%	$\approx 10^{-6}$	0.532	0.166
70%	$\approx 10^{-6}$	0.455	0.192
60%	$\approx 10^{-6}$	0.367	0.224
50%	$\approx 10^{-6}$	0.268	0.261
40%	$\approx 10^{-6}$	0.178	0.297
30%	$\approx 10^{-6}$	0.119	0.320
20%	$\approx 10^{-6}$	0.079	0.336
10%	$\approx 10^{-6}$	0.049	0.348

Table 11: Retailer’s cost and margin

	TV price	Fail rate	EW-TV price ratio	Overhead cost	Repair cost	Implied margin
9-16in	122.99	0.072	0.28	0.15	0.29	0.56
19-20in	173.97	0.065	0.24	0.18	0.31	0.51
25in	245.33	0.069	0.22	0.20	0.36	0.44
27in	354.06	0.059	0.20	0.22	0.35	0.43
>30in	812.53	0.076	0.22	0.20	0.39	0.41
OVERALL	400.07	0.067	0.22	0.19	0.35	0.46

Notes: Marginal cost of EW is computed as $c(p_j, \phi) = p_j f(\phi)$ where $f(\phi) = 0.04 + 1.32\phi - 2.48\phi^2$. Overhead costs represent costs that are not a function of failure rates, e.g. sales commissions. Costs are represented as fractions of the EW price.

5.3 Retailer’s cost

Figure 5 includes a scatter plot of $f(\phi) = \frac{c(p_j, \phi)}{p_j}$, and a local linear fit with a bootstrapped 95% confidence band. We also include a fit for the specification $f(\phi) = \mu_0 + \mu_1\phi + \mu_2\phi^2$, which is very similar to the local linear fit. In this specification, we treat μ_0 as reflecting costs that are involved in selling the warranty such as sales commission to sales people and other overhead, and we estimate it to be 0.04. We treat $\mu_1\phi + \mu_2\phi^2$ as reflecting costs that are involved in repairing the product, and we estimate it to be $1.32\phi - 2.48\phi^2$.

Using the parametric cost specification, we can estimate the seller’s profit margin, overhead cost μ_0 , and repair cost $\mu_1\phi + \mu_2\phi^2$. Table 11 reports these numbers where cost is expressed as a fraction of the warranty price.

5.4 Model fit

To examine how well the model fits the data, we compare the attachment rates and the warranty prices predicted by the model to those in the data. The prices that we use in the comparison are those that we used in estimation, i.e., the maximum observed TV price and the minimum observed warranty price.²⁶

To compute predicted attachment rates at the product level, we plug into equation (1) the non-parametric estimate of the probability distortion function, the estimated risk aversion and scale parameters, and the observed TV and warranty prices from the data. The model predicts a mean attachment rate of 0.276, which is very similar to the 0.268 mean attachment rate in the data. Figure 6 depicts that the distributions of the predicted and observed attachment rates are also very similar.

To compute the predicted warranty prices, we construct the demand for warranties at the product level based on the non-parametric estimated distortion function, estimated risk aversion and scale parameters, and the observed TV price from the data. Using this demand and the non-parametric cost estimate, we derive the retailer’s profit maximizing warranty price.

²⁶These prices increase the attractiveness of purchasing a warranty even without appealing to risk aversion and probability distortions. See footnote 22 for a discussion.

The model predicts a mean warranty-to-product price ratio of 0.175, which is very similar to the 0.171 price ratio in the data.²⁷ Figure 7 depicts that the distributions of the predicted and observed price ratios are also very similar.

6 Counterfactual Analysis

Two forces contribute to the economic outcomes in the extended warranty market. The first is retailers’ pricing power due to the add-on feature of the warranty and buyers’ search costs. The second is the overestimation of failure probabilities, which increases consumers’ willingness to pay for warranties. The goal of the counterfactual analysis is to quantify the importance of these two forces.

To this end, we consider three scenarios. In the first “Pricing power with Biased consumers” scenario (henceforth, scenario PB), the retailer has monopolist pricing power and buyers distort failure probabilities. In the second “Pricing power with Unbiased consumers” scenario (scenario PU), the retailer continues to benefit from monopolist pricing power but buyers do not distort failure probabilities. Comparing the second scenario to the first enables us to study the importance of probability distortions while fixing the competitive environment. In the third “Competitive market with Unbiased consumers” scenario (scenario CU), unbiased consumers purchase warranties in a competitive retail market. Because of competition, market prices are equal to the retailers’ marginal cost and retailers make zero profit. Comparing this scenario to the second scenario enables us to study the importance of pricing power in this market in the absence of probability distortions.

We assume that any change in the main product’s price across the three scenarios is negligible. To motivate this assumption, we first observe that there is little differentiation between electronics retailers, and that TV prices are easily observable. Thus, in the absence of extended warranties (or other add-ons), TVs would have been priced close to marginal cost. The introduction of extended warranties complicates things because it implies that retailers may have incentives to cut TV prices below cost in order to increase the demand for warranties.

There are several reasons to believe that such price cutting, if it happens, is insignificant. First, on the theoretical side, Ellison’s (2005) analysis indicates that the incentive to cut prices below marginal cost is mitigated by an “adverse selection” effect, i.e., cutting the prices of TVs will attract a disproportionate volume of consumers with low willingness to pay who will not buy warranties. Second, on the institutional side, TV manufacturers have tight controls on the pricing and marketing practices of retailers in order to prevent such price cutting. Two examples of controls are the Minimum Advertised Price (MAP) policy, whereby retailers cannot advertise below the manufacturer’s suggested retail price, and the Unilateral Manufacturer’s Retail Price (UMRP) policy, which imposes a penalty when retailers set prices below the UMRP. Third, on the empirical

²⁷Note that these price ratios are smaller than the 0.223 attachment rate in Table 2 because of our choice of observed prices.

side, we hand collected and compared forty TV prices from BestBuy and Target in 2003. BestBuy offered extended warranties in 2003, while Target started offering them in October 2006. To the extent that Target faces TV wholesale prices that are similar or higher than those BestBuy faces, we would expect TV prices at BestBuy to be lower than in Target if BestBuy’s warranty business affected their TV pricing. We find, however, that TV prices are very similar across retailers. On average, TV prices at BestBuy are 0.4 percentage points higher than in Target (standard error of 0.9 percentage points), and the median price difference is 0.

In the three counterfactual scenarios, cost and demand depend on the TV price p and the failure rate ϕ .²⁸ We use the average TV price of \$400 and consider the failure rates $\{0.04, 0.05, \dots, 0.15\}$. For the cost specification, we use our parametric estimate of $c(p_j, \phi)$ with $p_j = 400$. To construct demand in scenario PB, we use the Prelec distortion function with $\alpha = 0.652$ and the risk aversion parameter from the full model. In the two other scenarios we turn off the bias by setting $\omega(\phi) = \phi$ but keep the same risk aversion parameter. Given cost and demand, we calculate optimal prices, attachment rates, and profit margins in the different scenarios. To derive a yearly dollar equivalent measure of profit, consumer welfare, and total welfare, we assume that there are 30 million buyers of TVs every year.²⁹

6.1 Prices, attachment rates, and profits

Figure 8 plots the extended warranty to TV price ratio as a function of failure rates for the three scenarios. The price ratios are increasing and concave in the failure rate. They are highest in scenario PB, lower in scenario PU, and lowest in scenario CU. Moreover, the gap between scenarios PB and PU is smaller than the gap between scenarios PU and CU.

Thus, probability distortions lead to a price increase, but this price increase is smaller than the price increase due to search costs and the add-on feature of the product. For example, at the average failure rate of 7%, the price ratio in scenario PB is about 40% higher than in scenario CU. Almost three-quarters of this gap are attributed to pricing power (gap between PU and CU), and the rest is attributed to probability distortions. Since prices in scenario CU are equal to marginal cost, the same conclusions apply to the effect of probability distortion and pricing power on profit margins: pricing power is responsible for almost three-quarters of retailers’ profit margin.

Figure 9 plots attachment rates as a function of failure rates. Attachment rates in scenario PB are four- to five-times larger than in scenario PU. That is, fixing the competitive environment, probability distortions have a dramatic effect on quantity. For example, at a 7% failure rate, the attachment rate in scenario PB is more than five times higher than in scenario PU.

Attachment rates in the competitive market scenario CU lie between those in scenarios PB and

²⁸Cost and demand are also functions of the scale parameter $\sigma = \sigma_k \sqrt{p}$ where k is the TV category and p is the TV price. To have a single σ in the counterfactual analysis, we set $p = 400$ and calculate a sales-based weighted average of the estimated σ_k ’s. Alternative specifications yield qualitatively similar results.

²⁹This number is somewhat smaller than the yearly TV shipments in the US between 2010 to 2013, which ranged from 37 million to 40 million. (*CNN*, “With new TVs, size matters”, June 26, 2013).

PU. For most failure rates, the gap between CU and PB is larger than the gap between CU and PU. That is, taking scenario CU as a benchmark, probability distortions lead to a larger quantity increase than the quantity decrease due to retailers pricing power for most failure rates.

Finally, Figure 10 plots profits as a function of failure rates. An overwhelming 80 to 85 percent of retailers' profit is due to probability distortions.

To sum up, probability distortions have a dramatic effect on the volume of warranties sold in the marketplace and on retailers' profit. Probability distortions also affect warranty prices, but to a much smaller extent than retailers' pricing power.

6.2 Consumer and total welfare

There are at least two approaches to evaluating consumer welfare in the current context. The first is that probability distortions are a preference parameter. Consumers overweight failure probabilities because it improves their well-being relative to when they evaluate failure probabilities correctly. It may not be straightforward to reconcile this approach with the empirical evidence on warranty returns, warranty purchase behavior over time, and the difference in estimating failure probabilities between informed and uninformed consumers.

The second approach is that probability distortions are a mistake in decision making. Due to lack of information or experience, consumers overestimate failure rates at the point of sale. They thus make choices that do not reflect their underlying preferences, which are measured with respect to the objective failure rates. We will evaluate consumer welfare according to this approach, which seems consistent with the empirical evidence presented above and with the concerns of competition and consumer protection authorities.

Figure 11 plots consumer welfare as a function of failure rates. Consumer welfare is highest in scenario CU. This is because consumers evaluate failure probabilities correctly and prices are lowest. Consumer welfare is lower in scenario PU. This is because prices are higher due to retailers' pricing power leading to lower attachment rates. In both scenarios, consumer welfare is positive because only consumers with willingness to pay above price buy warranties.

Consumer welfare is negative in scenario PB, which is identical to the existing market. Put differently, consumer welfare would increase if the market for extended warranties did not exist. This is because probability distortions are sufficiently large, so that any positive surplus generated by consumers with true willingness to pay above price who buy warranties is dominated by the decrease in consumer welfare due to consumers with true willingness to pay below price who buy warranties.

Figure 11 illustrates that the reduction in consumer welfare due to probability distortions (the gap between PU and PB) is much larger than the reduction in consumer welfare due to pricing power (the gap between CU and PU). For example, at a 7% failure rate, there is a 235 million dollar decrease in consumer welfare when moving from CU to PB. Of this 235 million dollar gap, about 191 million dollars (or more than 80%) is attributed to probability distortions.

There are two positive welfare effects when moving from scenario PB to scenario PU (see Figure 12 for a graphic illustration). First, holding the warranty price constant, there is a shift inward of the demand curve when shutting down probability distortions. Consumers who now forgo buying the warranty are exactly those who overpaid for warranties based on their true willingness to pay for it, thus increasing consumer welfare. We refer to this effect as the *ripoff effect*. Second, warranty prices go down when shutting down probability distortions. Additional consumers, who estimate failure probabilities correctly, would now buy warranties. This implies an increase in attachment rates and in consumer welfare. We refer to this effect as the *price effect*. Figure 13 plots the dollar equivalent value of the two effects for various failure rates. The ripoff effect clearly dominates the price effect for all failure rates.

The effect of probability distortions on total welfare is ex-ante ambiguous. The first-best level of insurance is the one in scenario CU. This is because all consumers with true willingness to pay above cost are insured. In scenario PU, there is under-insurance relative to the first-best. This is because retailers price warranties above cost, and thus consumers with true willingness to pay above cost but below market prices are not insured. In scenario PB, there is over-insurance relative to the first-best. This is because consumers with true willingness to pay below cost purchase insurance due to probability distortions. Figure 14 depicts qualitatively the welfare losses in scenarios PU and PB relative to the first-best.

Figure 15 plots total welfare as a function of failure rates. For failure rates below 12%, total welfare in scenario PU is larger than in scenario PB. Put differently, shutting down probability distortions leads to an increase in total welfare. Only for failure rates above 12%, the increase in attachment rates due to probability distortions compensates for the negative effect on welfare of pricing power. This is because as failure rates increase, probability distortions become less significant, and thus lead to a smaller increase in attachment rates.

To sum up, probability distortions result in negative consumer welfare for all failure rates. This implies that consumers would have been better off if extended warranties were not sold at all in the market place. For the majority of failure rates, total welfare is smaller when consumers distort probabilities than when they do not. This implies that for the majority of failure rates, the gain in consumer welfare due to the shutting down of probability distortions dominates the reduction in retailers' profit.

7 Concluding Comment on Policy Implications

This paper studied the effect of probability distortions on economic outcomes in the extended warranty market. Our first main finding is that probability distortions are responsible for more than 80% of retailers' profit in this market. Our second main finding is that probability distortions are reduced with information and experience suggesting that the source of probability distortions is overestimation of unknown probabilities.

These findings are relevant for evaluating consumer policies in this market. An example of such a policy is the proposal of the UK Competition Commission (2003) that electronics retailers “display the price of an applicable EW (extended warranty) alongside the DEG (domestic electrical good) in store and in press advertisement and other publicity”. This policy has the potential to eliminate consumers’ search costs, and is thus expected to intensify competition. But the intense competition will not necessarily lead to significant increase in consumer welfare. This is because even though more consumers will buy warranties at lower prices, many of them will do so because they overestimate failure probabilities and not because their true willingness to pay for warranties is above price. An alternative policy that requires retailers to post the failure rate of the relevant product in their advertisements or in store may be more effective in increasing consumer welfare.

To illustrate this point, we compare the existing market (scenario PB) to two alternative scenarios using the same parameters of the counterfactual analysis. In the first “Competitive market with Biased consumers” scenario (scenario CB), retailers price at marginal cost and consumers overestimate failure probabilities. This scenario corresponds to the optimal market outcome following the implementation of the UK Competition Commission (2003) proposal to post warranty prices alongside electronics prices. In the second “Pricing power with Unbiased consumers” scenario (scenario PU), which was discussed in Section 6, retailers benefit from pricing power and consumers evaluate failure probabilities correctly. To the extent that probability distortions result from overestimating unknown probabilities, this scenario corresponds to the optimal market outcome following the alternative proposal to post failure rates.

Figure 16 illustrates that market volume will double following the implementation of the UK Competition Commission proposal. Figure 17 illustrates that consumer welfare will also improve relative to the existing market. However, consumer welfare will still be negative because many extended warranty buyers should not purchase them. Total welfare will also be negative because firms make zero profit when the Competition Commission proposal is implemented. On the other hand, in the alternative policy, both consumer welfare and total surplus will be positive. Thus, policies that aim to intensify price competition may lead to suboptimal results in markets with uninformed or biased consumers relative to policies that aim to improve consumers’ decision making.

References

- Anderson, J. R. (1993), “Rules of the Mind,” Psychology Press, New York, NY.
- Andrikogiannopoulou A. and F. Papakonstantinou (2016), “Heterogeneity in Risk Preferences: Evidence from a Real-World Betting Market”, *Working paper*.
- Barseghyan, L., F. Molinari, T. O’Donoghue, and J. C. Teitelbaum (2013), “The Nature of Risk Preferences: Evidence from Insurance Choices,” *American Economic Review*, 103(6), 2499–2529.
- Barseghyan, L., F. Molinari, T. O’Donoghue, and J. C. Teitelbaum (Forthcoming), “Estimating Risk Preferences in the Field”, *Journal of Economic Literature*.
- Barseghyan, L., J. Prince, and J. C. Teitelbaum (2011), “Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data,” *American Economic Review*, 101(2), 591–631.
- Chen, T.C., A. Kalra and B. Sun (2009), “Why Do Consumers Buy Extended Service Contracts,” *Journal of Consumer Research*, 36(4), 611–623.
- Chiappori, P. A., B. Salanié, F. Salanié, and A. Gandhi (2012), “From Aggregate Betting Data to Individual Risk Preferences,” *Working paper*.
- Cohen, A. and L. Einav (2007), “Estimating Risk Preferences from Deductible Choice,” *American Economic Review*, 97(3), 745–788.
- Diamond, P. A. (1971), “A Model of Price Adjustment,” *Journal of Economic Theory*, 3, 156–168.
- Ellison, G. (2005), “A Model of Add-On Pricing,” *Quarterly Journal of Economics*, 120(2), 585–637.
- Ellison, G. and S. F. Ellison (2009), “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, 77(2), 427–452.
- Gandhi, A. and R. Serrano-Padial (Forthcoming), “Does Belief Heterogeneity Explain Asset Prices: The Case of the Longshot Bias,” *Review of Economic Studies*.
- Huysentruyt, M. and D. Read (2010), “How do people value extended warranties? Evidence from two field surveys,” *Journal of Risk and Uncertainty*, 40(3), 197–218.
- Jindal, P. (2014), “Risk Preferences and Demand Drivers of Extended Warranties,” *Marketing Science*, 34(1), 39–58.
- Kahneman, D. and A. Tversky (1979), “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47(2), 263–291.

- Jullien, B. and B. Salanié (2000), “Estimating Preferences under Risk: The Case of Racetrack Bettors”, *Journal of Political Economy*, 108(3), 503–530.
- Kliger, D. and O. Levy (2009), “Theories of Choice under Risk: Insights from Financial Markets,” *Journal of Economic Behavior and Organization*, 71(2), 330–346.
- Prelec, D. (1998), “The Probability Weighting Function,” *Econometrica*, 66(3), 497–527.
- Snowberg, E. and J. Wolfers (2010), “Explaining the Favorite-Long Shot Bias: Is It Risk-Love or Misperceptions?”, *Journal of Political Economy*, 118(4), 723–746.
- Sydnor, J. (2010), “(Over)insuring Modest Risks,” *American Economic Journal: Applied Economics*, 2(4), 177–199
- UK Competition Commission (2003), “Extended warranties on domestic electrical goods: A report on the supply of extended warranties on domestic electrical goods within the UK – Volumes 1, 2 and 3.”

Figures

Figure 1: Histogram of extended warranty to product price ratio




Figure 2: Identification: Single-crossing of willingness to pay





Figure 3: Estimated distortion function: All consumers



Notes: For presentation purposes, the graph is truncated at 0.25. The largest value of ω is about 0.34.

Figure 4: Estimated distortion function: Informed vs uninformed consumers

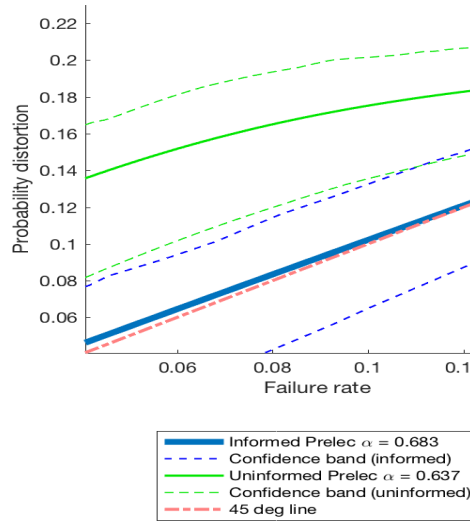



Figure 5: Retailer's cost



Notes: For presentation purposes, the graph is truncated at 0.25. The largest value of the normalized cost is about 0.35.

Figure 6: Model fit: Attachment rates

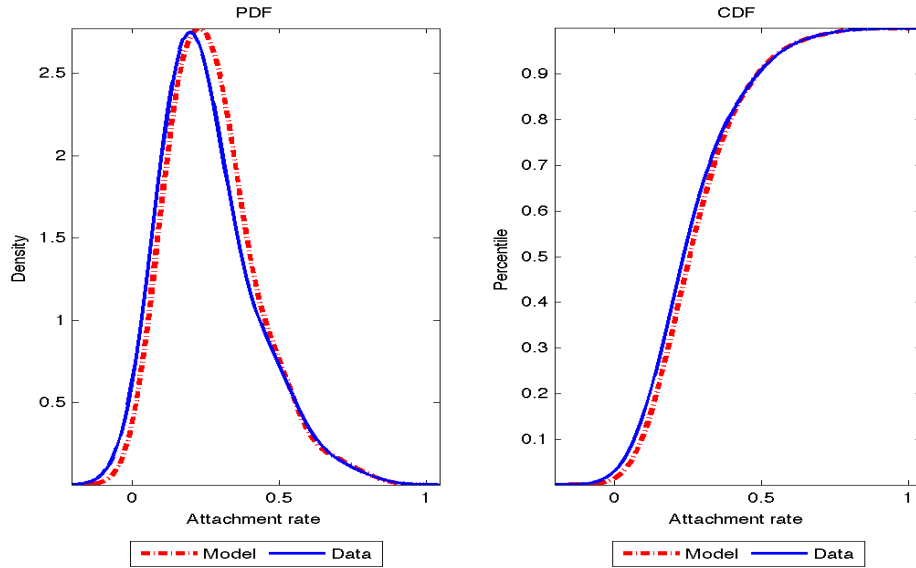


Figure 7: Model fit: Warranty to TV price ratios

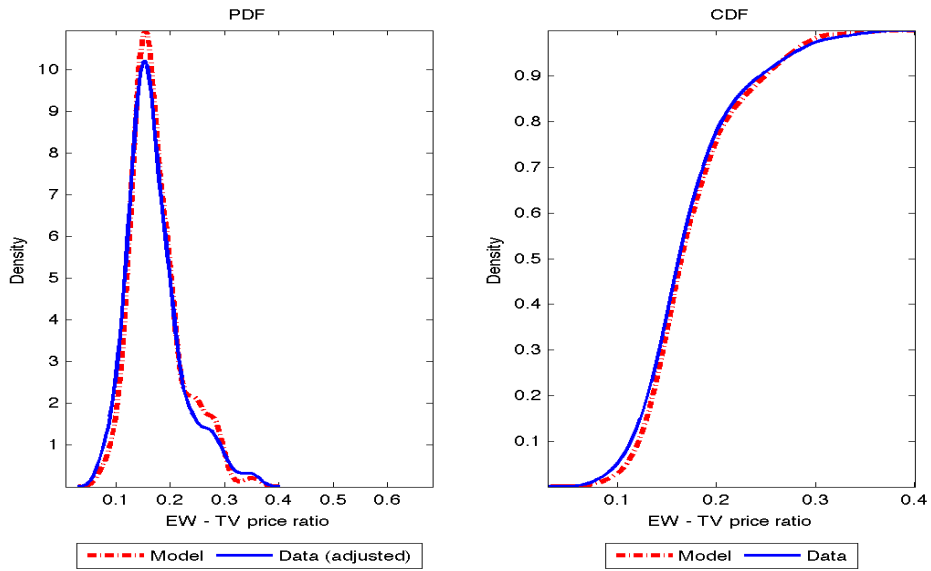


Figure 8: Counterfactuals: Extended warranty to TV price ratios




Figure 9: Counterfactuals: Attachment rates




Figure 10: Counterfactuals: Profits




Figure 11: Counterfactuals: Consumer welfare




Figure 12: Ripoff and price effects of shutting down distortion

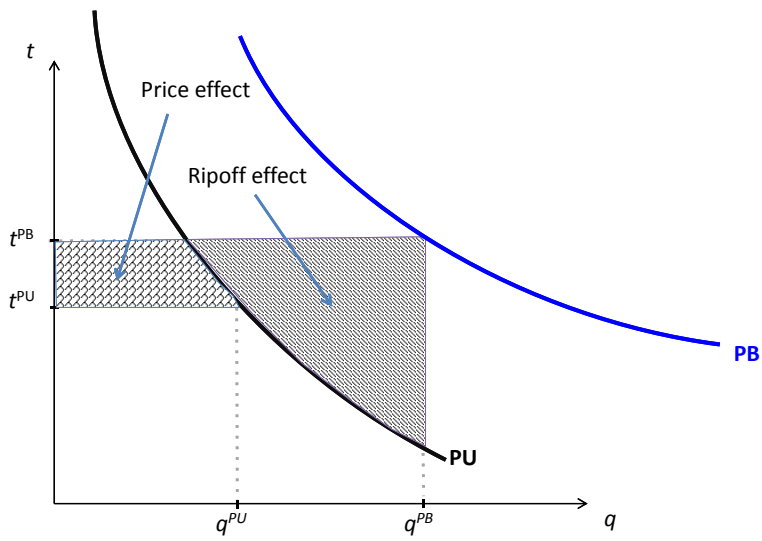


Figure 13: Counterfactuals: Magnitude of ripoff and price effects




Figure 14: Sources of welfare loss

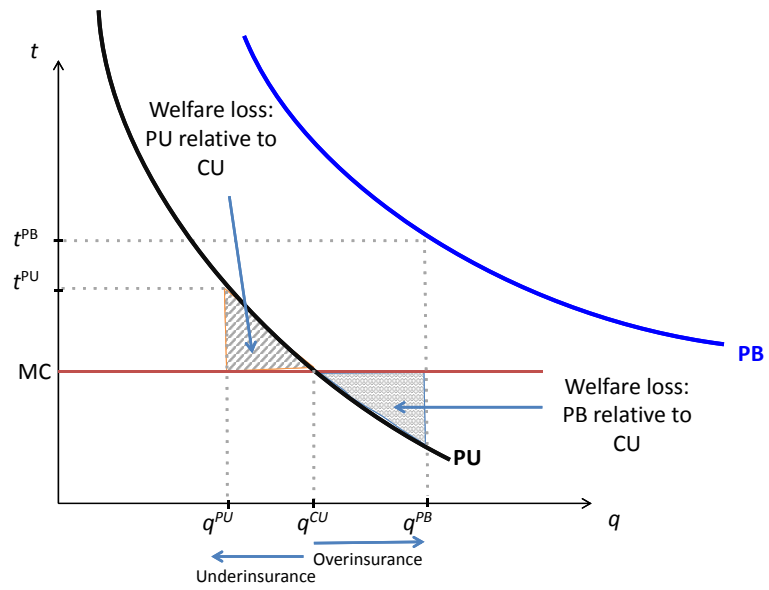


Figure 15: Counterfactuals: Total welfare




Figure 16: Comparing policy interventions: Attachment rates




Figure 17: Comparing policy interventions: Consumer welfare

