Market Regulation of Voluntary Add-on Contracts

Christian Michel*
Universitat Pompeu Fabra and Barcelona GSE

April 5, 2017

Abstract

This paper analyzes contract choices and the effectiveness of consumer protection policies when firms can offer voluntary add-on insurance for their products. We develop a model in which a base product can be sold together with a voluntary extended warranty contract that insures consumers against the risk of product breakdown. Some consumers do not pay attention to extended warranties before making base product choices, but overestimate the value of such warranties at the point of sale. Under retail competition, the consumers' option to buy multiple base products can endogenously create a base price floor that may prevent firms from redistributing the full warranty profits via loss-leadership. Inducing competition in the warranty market weakly increases consumer welfare and weakly outperforms a minimum warranty standard, which can even reduce consumer surplus. The results are consistent with the effects of recent changes regarding extended warranty regulation by UK legislators.

JEL Codes: D18, D21, L13

*Email: christian.michel@upf.edu. I would like to thank Larbi Alaoui, Susanna Esteban, Jana Friedrichsen, Michael Grubb, Paul Heidhues, David Laibson, Volker Nocke, Marco Ottaviani, Andre Stenzel, and seminar participants and conference audiences for their helpful comments. This paper is based on my 2013 thesis chapter "Contractual Structures and Consumer Misperceptions – Warranties as an Exploitation Device." I gratefully acknowledge financial support from the Marie Curie Intra-European Fellowship FP7-PEOPLE-2013-IEF-623912 and from the Spanish Ministry of Economy and Competitiveness Grant ECO2013-43011.

1 Introduction

In many retail industries, it is common for firms to offer consumers the option to buy an extended warranty for a product that they have just agreed to purchase. This is especially prevalent for many consumer electronics products, such as televisions, notebooks, and mobile phones. Arnum (2013) estimates a total of $37.9 billion paid for extended warranty premiums in the US in 2013 across all industries.

Extended warranty practices have come under scrutiny from policy makers. In its 2002 consumer electronics extended warranty investigation, the British Office of Fair Trading concluded that "consumer protection in this market is inadequate" (OFT, 2002). OFT (2012) reports claims ratios, defined as the incurred claim costs relative to the insurance payments paid, for extended warranties in the UK consumer electronics industry. For 2002, the ratio is approximately 20 percent, compared with a 74 percent claims ratio in the UK's general insurance industry. Extended warranty contracts are among the largest drivers of retail profits in many consumer electronics segments. Such contracts are often recommended by salespeople in stores but are rarely advertised in advance. According to Warranty Week (2005), extended warranty sales accounted for 3 to 4 percent of the revenues and more than 50 percent of the profits of the US retailer Best Buy in 2003.[1]

These facts lead us to two main questions: Is a special form of consumer protection needed for voluntary add-on contracts? And if so, what is the optimal form of protection?

Section 2 introduces our baseline model in which retail firms compete in selling a homogeneous base product to consumers. The product has an exogenous probability of breakdown that is known to consumers. At the point of sale, retail firms can offer voluntary extended warranty contracts that insure consumers against the risk of product breakdown. A fraction of the consumers is naive and underestimates the costs of returning a faulty product. Evidence indicates that very few consumers in the consumer electronics industry shop around at different retailers with respect to extended warranty contracts.[2] Furthermore, extended warranties are of very low priority with respect to consumers' base product choices. We therefore assume that consumers are locked in at the point of sale with respect to buying

[1] Competition Commission (2003) reports an estimated value of £900 million for the British consumer electronics extended warranty market, with 18.5 million extended warranties sold in a year. Businessweek (2004) reports that extended warranty contracts account for nearly 100 percent of the retailer Circuit City's profits, with an estimated average profit margin between 50 and 60 percent; see also Chen *et al.* (2009).

[2] According to Rogers (2003b), only 4 percent of consumers said that they considered obtaining an extended warranty from a different provider than the one they had chosen. Using survey data, Westra (2002) reports that across all consumer electronics categories, roughly half of the consumers who purchased a consumer electronics product did not anticipate buying an extended warranty before purchasing the electronic product.

extended warranty contracts.[3] This gives firms the opportunity to profitably sell extended warranty contracts to naive consumers. As a consequence, firms compete harder to attract these consumers by further decreasing the price of the base product below its marginal costs.

Section 3 derives our main model predictions and analyzes the effects of different consumer protection policies. Because each product has a risk of breaking down, there is also a positive value for a consumer of buying multiple products. When selling multiple base products as a loss-leader while not increasing the extended warranty profits, the overall markup per consumer for a firm decreases. We show that in equilibrium, the threat of consumers buying multiple base products can endogenously generate a lower bound for the base product price. Such a *price floor* can keep firms from re-transferring the whole extended warranty profits via loss-leadership upfront. This can occur if it is less profitable for a firm to capture the whole market while selling multiple base products per consumer compared to capturing part of the market when selling a single quantity per consumer at the price floor. Thus, an equilibrium in which firms make positive profits while consumers only buy a single product can arise because consumers have the possibility to buy multiple products.

The model setting relates to models with shrouded costly add-on prices and base product price floors; see, for example, Heidhues *et al.* (2016), building on the seminal work of Gabaix and Laibson (2006). There are several differences from these models to our model. In our model, naive consumers make a voluntary choice about whether to buy an add-on extended warranty contract and that–apart from the return costs–all contract details are known. Furthermore, because of the possibility of product breakdown, owning multiple base products can increase the expected consumption utility which in turn can lead to an endogenous price floor. We show that our results also hold qualitatively for alternative explanations for why consumers overpredict the values of extended warranties, such as the overprediction of failure rates and loss aversion.

In our model, the base products are not deceptive. However, the voluntary add-on contracts are overpriced because of consumer naivete and the non-competitive market structure for these contracts. The key question that we assess is how different protection policies for voluntary add-on contracts change firm behavior and consumer choices. A policy that induces extended warranty competition at each retail store through the presence of independent warranty providers always leads to zero profits for all firms and weakly increases consumer surplus. This occurs because of the presence of independent insurance firms that prevents

[3]Rogers (2003a) asks consumers to rank the most important attributes for their product choices, and finds that consumers pay the most attention to product reliability, durability, ease of use, and price. Moreover, consumers consider size, color, and brand. Both the manufacturer's guarantee and extended warranties always rank behind all of these attributes in the range of the least important attributes.

a retailer monopoly for insurance contracts at the point of sale. Therefore, retail firms no longer have an incentive to sell the base product as a loss leader. Overall, naive consumers weakly benefit from such a policy, while sophisticated consumers are weakly hurt by it.

A minimum warranty level has ambiguous effects on consumer surplus. It decreases naive consumers' predicted value of a higher extended warranty contract. However, when sophisticated consumers return a defective product with positive probability and incur return costs, these frictions have a negative effect on consumer surplus. If these frictions are sufficiently high relative to the reduction in the prices of extended warranty contracts, the overall effect of this policy on consumer welfare is negative. A policy that allows canceling an extended warranty contract for a fixed period of time after signing the contract often induces a change in firms' contract offerings only if all consumers are willing to change the contract. If this does not occur, the policy is inferior to inducing competition at the point of sale.

We provide several conditions under which a "shrouding" equilibrium persists in which firms do not advertise their extended warranty quotes to consumers. This is the case, for example, when consumers become aware of a competitive insurance market or when they do not expect the purchase of an extended warranty contract prior to visiting a store to buy a base product. Overall, our results clearly indicate that inducing competition at the point of sale is the most favorable protection policy when possible, as the policy always maximizes consumer welfare while avoiding the risk of negatively affecting the market.

Our model predictions are broadly consistent with the effects of the UK 2005 Extended Warranty Order. This law was introduced after the UK's Competition Commission concluded in an in-depth market investigation that there is insufficient competition for extended warranties in the consumer electronics industry and that a "complex monopoly exists." The Competition Commission argued that prices for extended warranties have typically not been displayed, cancellation periods for such contracts have been short, and consumers have lacked information about different extended warranty providers at the point of sale.

The law includes a cancellation right for consumers, the right to shop around for other extended warranty quotations, and a mandatory information leaflet distributed before signing. After the introduction of the policy, claims ratios in the market increased, although they were still significantly lower than the ratios in most other industries.[4] This indicates an

[4]OFT (2012) finds an increase in the claims ratio from 20 percent to approximately 50 percent six years after the law was introduced. Although it is increasing, this number is still relatively low compared with other markets: in the UK general insurance industry, the OFT finds on average a 74 percent claims ratio, including a 62 percent claims ratio for property insurance. The fraction of consumers who shop around also increased from 4 to 25 percent in the same period. Unsurprisingly, those consumers who researched products online were particularly likely to receive multiple offers and more likely to buy extended warranty contracts online. However, only 18 percent of retail extended warranty customers stated that they had shopped around. In comparison, the OFT reports that 69 percent of general

increased level of competitiveness in the market for extended warranties with some remaining frictions.[5]

Section 4 provides an overview of the related theoretical literature. Section 5 concludes.

2 Baseline model

There are $M \geq 2$ retail firms in the market, all selling an identical base product with quality q and marginal cost c. In this context, the marginal cost of production c can also be interpreted as the exogenous upstream price that each firm must pay to a manufacturer. The quality q reflects the probability of the product working properly and is known to consumers. Each firm i sets a base product price p_i and offers an individual menu of extended warranty contracts $\Gamma_i = \{\gamma_0, \gamma_i\} = \{(0,0), (x_i, w_i)\}$. An extended warranty contract $\gamma_i \equiv (x_i, w_i) \in \Gamma_i$ grants a consumer a payment w_i in case of product breakdown at an extra price x_i that he must pay upfront in addition to the product price p_i. For notational simplicity, a base product is associated with a zero manufacturer warranty, such that all positive extended warranties are offered by the retailer. This is consistent with many consumer electronics retailers offering their own extended warranty contracts; see, for example, OFT (2012) for recent evidence. Thus, we require a costless zero default extended warranty contract $\gamma_0 \equiv (0,0)$ to be included in the menu of warranty contracts for each product i. We assume that there is a finite maximum warranty level $\overline{w}$ that firms can offer.[6]

Consumers are risk neutral.[7] They value the consumption of a properly working product with utility $I > 0$, and derive zero utility from consuming a malfunctioning product. Further, they derive disutility $-p$ from paying a product price p per unit they buy, and $-x$ per extended warranty contract priced at x they purchase. While consumers have the possibility to purchase multiple base products, they only derive positive utility from a single functioning product. We assume that consumers are not able to resell base products that they do not use.[8] Thus, the expected utility when buying n products without any extended warranty is

insurance customers shopped around the last time they renewed their contract.

[5] Our setting can furthermore be used to explain the profitability of money-return mail-in rebate policies under industry competition. Also see Bar-Gill (2007) for a discussion and an example of a rebate misperception problem. Jolson *et al.* (1987) provide further evidence that firms use redeemable rebates as a promotional tool as well because many consumers do not use them and end up paying the full price.

[6] In equilibrium, no firm can increase its profits by offering more options in its menu of extended warranty contracts. The upper bound on warranties may not be needed, but it eliminates outcomes in which firms offer infinite warranty contracts. One natural boundary is $\overline{w} = I$, i.e., the maximum willingness to pay for a working product.

[7] In Appendix B.2, we discuss microfoundations of return cost misperceptions, such as hyperbolic discounting, loss aversion, and misprediction of failure rates. We show that these specifications qualitatively lead to similar outcomes.

[8] Otherwise if the consumer could check whether a product is functioning, he could resell the remaining products.

$[1-(1-q)^n]I - np$, which is equal to $qI - p$ when buying a single base product. In case consumers buy multiple base products, an extended warranty contract has to be registered with a specific base product. This implies that when consumers want to insure multiple base products, they have to buy multiple extended warranty contracts. Firms can only offer linear warranty contracts, i.e. cannot give quantity discounts in case consumers would like to buy multiple extended warranty contracts. While this is a restrictive assumption from a theoretical viewpoint, it matches evidence from the consumer electronics industry.[9]

A key feature of our model concerns the payment of a warranty in case of product breakdown. To receive the warranty payment, a consumer must send the product back to the firm, which incurs socially wasteful costs r to the consumer. Return costs are distributed according to the differentiable cumulative distribution function (henceforth, cdf) $F(r)$. For simplicity, we assume that the distribution of return costs does not depend on the specific warranty contract and thus cannot be influenced by the firm. The draw of r is unknown to each consumer prior to the purchase of both the product and warranty contract. This eliminates any selection effects from privately known consumer return costs.

The two consumer types in the market differ in how they anticipate return costs. A fraction $1-\theta$ of consumers is "sophisticated" in the sense that they correctly predict the distribution of return costs. The remaining fraction θ of consumers is "naive" in that they erroneously underestimate the costs of returning the product, i.e., the anticipated distribution of return costs. Despite having the same return costs cdf $F(r)$ as the sophisticated consumers, the naive consumers believe that the distribution of return costs r in the whole population can be described by the differentiable cdf $\tilde{F}(r)$. We assume that F first-order stochastically dominates $\tilde{F}(r)$ for the full support of the functions, i.e., $\tilde{F}(r) \geq F(r)\ \forall r \in (0,\infty)$, with $\tilde{F}(r) > F(r)$ for $r \in (0,\overline{w}]$. This implies that naive consumers underestimate their return costs for all positive warranty levels and that there is always a positive probability of facing return costs $r > 0$. This formalization allows for relatively flexible consumer return patterns.[10] At the time of product breakdown, consumers learn their true cost draw r. They

However, our setting includes only a single period for buying a base product.

[9]Prices for extended warranty contracts in the consumer electronics industry are usually not only linear, but they are very often pooled across many different product classes, see for example OFT (2012). Such pricing ignores different breakdown probabilities for products, and quantity discounts. There is also a very low volatility in the prices of extended warranty contracts. OFT (2012), for example, notes that between 2003 and 2011, the UK retailer Dixons changed its prices for extended warranty contracts only five times. This assumption allows to easily characterize the pricing of extended warranty contracts in case multiple warranty contracts are purchased. Furthermore, it implies that the value of buying a single extended warranty contract is not increasing in the number of base products purchased.

[10]Using our specification for the cost distributions allows both for equilibrium warranty levels that are below the maximum warranty level $\overline{w}$ and for consumers to return a product in equilibrium with positive probability. When we allow for only one cost outcome r that naive consumers underestimate, in equilibrium, naive consumers never return

claim the warranty payment when it is at least as high as the return cost draw, i.e., when $w \geq r$. In case consumers purchase multiple extended warranty contracts and experience multiple product breakdowns, they experience an independent cost draw for each product.

Throughout the main text we make two additional parametric assumptions. First, we assume that the product has a probability of working without a defect $q \geq \frac{1}{2}$. This is sufficient for firms not being able to increase their extended warranty profits when selling multiple base product quantities per consumer.[11]

Second, we assume that the marginal costs are such that each consumer chooses to buy a single quantity of a base product when it is priced at marginal costs and derives a positive utility from it, i.e. $qI > c > \frac{I}{4}$. If the first inequality does not hold, consumers do not derive a positive utility from the base product when it is priced at marginal costs. If the second inequality does not hold, there are cases in which consumers purchase multiple units of a base product even if the product is not sold as a loss-leader. We characterize all possible equilibra without these assumptions in Appendix B.1.

Consumer utility A sophisticated consumer's expected utility of consuming a single base product of quality q at price p associated with warranty contract γ, $V(q, p, \gamma)$, can be written as

$$V(q, p, \gamma) = qI - (p + x) + (1 - q) \int_0^w [w - r] f(r) dr. \tag{1}$$

The first part on the right-hand side reflects the expected utility from consuming a product. The second part is the disutility that a consumer derives from the price of both the base product and warranty. The third part is the expected utility from the warranty payment conditional on returning a defective product, where $f(r)$ is the return cost probability density function. With $\tilde{f}(r)$ denoting a naive consumer's predicted return cost probability density function, the analogous expected consumption utility of a naive consumer can be written as

$$\tilde{V}(q, p, \gamma,) = qI - (p + x) + (1 - q) \int_0^w [w - r] \tilde{f}(r) dr. \tag{2}$$

Similarly a naive consumer's utility from buying k base products and $m \leq k$ extended contracts products becomes $[1 - (1 - q)^k]I - kp - mx + \sum_{i=1}^{m} (1 - q)^m \int_0^w [w - r] \tilde{f}(r) dr$.

Timing of the game Firms simultaneously set their base product prices and menus of extended warranty contracts. Consumers initially observe only the base product prices and

a faulty product if this cost is sufficiently high. This is not a desirable property given the small but yet positive claim ratios in related industries.

[11] While not needed, this assumption highly facilitates the analysis, see Appendix B.1 for the general case.

either choose a retailer to buy at least one base product from or choose not to buy any product. At this stage, no consumer observes the firms' menus of extended warranty contracts. Sophisticated consumers can correctly infer the equilibrium menus of extended warranty contracts, while naive consumers are myopic with respect to the existence of these contracts at this stage. The sequential modeling assumption is consistent with evidence that very few consumers shop around for extended warranties and that many of them do not anticipate their purchase of extended warranties. After making a base product purchase from retailer i at the point of sale, consumers observe their retailer's menu of extended warranty contracts Γ_i. At this point, consumers decide whether to buy an additional extended warranty contract.

Equilibrium Definition We seek a combination of base product prices and menus of extended warranty contracts that maximizes firms' profits given their competitive constraints. We assume that the two consumer types are indistinguishable from one another before making purchase decisions. The consumer types can differ in their utility predictions for each contract because naive consumers underestimate their return costs.

Definition (Industry Equilibrium) An equilibrium for $M \geq 2$ firms consists of a set of base product prices $\{p_1, .., p_M\}$ and a set of associated menus of extended warranty contracts $\{\Gamma_1, .., \Gamma_M\}$ with the following properties:

1. *(Firm profit maximization)* For any firm i, given its rivals' base product prices and menus of extended warranty contracts, there is no combination of base product price p_i and menu of extended warranties Γ_i that leads to a higher expected profit.

2. *(Optimal decision for sophisticated consumers)* Sophisticated consumers make base product and extended warranty contract choices that yield the highest expected consumption utility. These consumers do not observe the menu of extended warranty contracts before making base product choices, and they claim a warranty payment when the payment is at least as high as their return costs.

3. *(Hypothetical optimality for naive consumers)* Naive consumers make base product and extended warranty contract choices that yield the highest expected consumption utility in the virtual case in which their return costs are represented by the cdf $\tilde{F}$. These consumers do not anticipate the existence of extended warranties before making base product choices, and they claim a warranty payment when the payment is at least as high as their return costs.

Under this definition, both consumer types maximize their initial predicted utilities. Firms and sophisticated consumers are aware of naive consumers' false beliefs, while naive consumers believe that all consumers and the firm share the same beliefs about the return cost distribution. In section 3.3, we also analyze the case in which firms can choose to "unshroud" the extended warranty in the industry contracts before consumers make a base product purchase. In that section, we show the conditions for which a shrouding equilibrium persists. One advantage of our main modeling assumption compared to an exogenous overprediction of failure rates is that consumers in our model correctly compute their expected consumption utility of the base product, such that the willingness to overpay can only be attributed to the extended warranty contracts. In particular, exogenous failure rate misperceptions can further change the number of purchased base products in equilibrium.[12]

3 Industry equilibrium and evaluation of protection policies

3.1 Derivation of industry equilibrium

In this subsection we derive the industry equilibrium in absence of any consumer protection policy. We first focus on the firms' optimal choice of extended warranty contracts. This is followed by the derivation of our main channel for endogenous price floors via the consumers' possibility to buy multiple base products.

Optimal extended warranty contracts Because consumers make an extended warranty choice only after making the decision to purchase the base product, this implies that retailers have a quasi-monopoly for extended warranty contracts at the point of sale. Due to the underestimation of return costs, there are cases in which retailers can sell extended warranty contracts to naive consumers at positive markups. Intuitively, this can occur when a naive consumer's predicted net utility from an extended warranty is higher than a retail firm's expected cost of paying the warranty.[13] Formally, this is the case if

$$\tilde{\pi}(q,w) \equiv (1-q)\left[\tilde{F}(w)\tilde{E}[w-r|w>r] - F(w)w\right] > 0,\ w \in (0,\overline{w}].$$

[12] See Michel (2016) for a model in which firms can use warranties upfront to endogenously induce consumer quality misperceptions because of consumers underestimating their return costs. The paper provides supporting evidence for consumers underestimating return costs in case of product breakdown from customer reviews for knife sets that are sold via infomercials on TV.

[13] While "peace of mind" is an often stated reason for consumers to buy an extended warranty contract, this does not eliminate the possibility of underestimating the return costs. Rogers (2003b) reports the most important reasons cited by sales personal while advising consumers to buy an extended warranty. The four most frequently used reasons remembered by consumers were "no repair costs" (17 percent), "peace of mind" (17 percent), "good value for money" (11 percent), and "free/no cost" (6 percent).

When there exists at least one positive warranty level for which this inequality holds, each firm sells positive extended warranties to naive consumers in equilibrium. The firm sets a warranty level $\tilde{w}$ to maximize the value of the left-hand side of the above inequality. As a price for the extended warranty, the firm charges a naive consumer's predicted willingness to pay, $(1-q)\tilde{F}(\tilde{w})\tilde{E}[\tilde{w}-r|\tilde{w}>r]$.

We next show that a firm can never increase its extended warranty profits by selling multiple extended warranty contracts to naive consumers even if they buy multiple base products. Given the linearity of extended warranty contracts, if a firm wants to sell two extended warranty contracts to consumers that buy multiple base products, it has to reduce the extended warranty price for any warranty level w to $\tilde{x}_2 = (1-q)^2\tilde{F}(w)\tilde{E}[w-r|w>r]$. Its profit turns to $2(1-q)^2\tilde{F}(w)\tilde{E}[w-r|w>r] - [(1-q)+(1-q)^2]F(w)w$. This can be rewritten as $2(1-q)\tilde{\pi}(q,w) - q(1-q)F(w)w$, which leads to an upper bound for the extended warranty profits when selling two contracts of $2(1-q)\tilde{\pi}(q,\tilde{w})$. It follows that because $q \geq \frac{1}{2}$, firms are at least as well off selling a single extended warranty contract as when selling two contracts to naive consumers. It is easy to show that the higher profitability of selling a single extended warranty contract also holds compared to any higher number of extended warranty contracts.[14] Throughout our analysis we make the convention that a firm always participates in the market if this leads to zero equilibrium profits. This excludes cases in which a firm sets a base product price such that no consumer buys from it while all firms in the industry make zero equilibrium profits. Under this convention, all equilibria have a symmetric base product price.

Endogenous price floor To attract naive consumers who buy extended warranties, retailers compete harder on the base product price, which leads to a decrease in the base product price below marginal costs. Because firms cannot distinguish between naive and sophisticated consumers ex ante, sophisticated consumers benefit from this loss leadership by having to pay a lower price for the base product. Therefore, consumer naivete causes a transfer in utility from naive to sophisticated consumers. However, when the base product price is sufficiently low, consumers demand multiple quantities. The expected consumption utility when buying two base products at price p without any additional warranty is $[1-(1-q)^2]I - 2p$, while the expected utility when buying a single product is $Iq - p$. It follows that consumers prefer buying two base products rather than one whenever $q(1-q)I > p$. The expression $q(1-q)I$ represents the highest base product price such that no consumer hase an incentive to buy

[14] The expected utility of buying the m^{th} extended warranty contract is $(1-q)^m F(w)E[w-r|w>r] - x$, which is exponentially decreasing in m. I would like to thank an anonymous referee for his suggestions regarding this section.

multiple quantities, which we define as the potential base product price floor

$$\underline{b} \equiv q(1-q)I.$$

Because $c > \frac{I}{4}$, consumers are only willing to buy multiple quantities of the base product if it is offered as a loss-leader.[15] If consumers buy multiple products instead of a single one, it thus follows that a firm's markup per consumer decreases. We next show that the threat of consumers buying multiple quantities of the base product can result in an endogenous base product price floor from which under some circumstances no firm has an incentive to deviate by lowering the base price to capture the whole market and selling two base products per consumer instead of only obtaining a share of $\frac{1}{M}$ of consumers. This occurs whenever $\frac{1}{M}(\underline{b} - c + \tilde{\pi}(q, \tilde{w})) \geq 2(\underline{b} - c) + \tilde{\pi}(q, \tilde{w})$. Rewriting this weak inequality leads to our first proposition.

Proposition 1. *(Price floor equilibrium) If* $\underline{b} \leq c - \frac{M-1}{2M-1}\theta\tilde{\pi}(q, w)$*, there is an equilibrium in which each firm* i *sets a base product price* $p_i = \max[\underline{b}, c - \theta\tilde{\pi}(q, \tilde{w})]$*, together with a menu of extended warranty contracts* $\Gamma_i = \{(0,0), (\tilde{x}, \tilde{w})\}$*, where*

$$\tilde{w} = \inf\left[\arg\max_{w \in [0, \overline{w}]} \tilde{\pi}(q, w)\right], \text{ and } \tilde{x} = (1-q)\tilde{F}(\tilde{w})\tilde{E}[\tilde{w} - r | \tilde{w} > r]. \tag{3}$$

If $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$*, the equilibrium base product price is unique.*

Note that if multiple profit-maximizing warranty levels exist, there can be asymmetric equilibria in which firms set the same base product price but different warranty levels in their extended warranty contracts. In this case the symmetric equilibrium in which all firms offer a warranty level $\tilde{w}$ in their extended warranty contracts is still the consumer surplus maximizing equilibrium, because it induces the lowest socially wasteful return costs.

Chen *et al.* (2009) find that promotions of a base product increase the probability of selling extended warranties. The authors attribute this effect to an increase in spending resulting from an emotional gain because of the decrease in the base price. In our model, the reasoning is the opposite: because firms earn higher profits from products that are sold with extended warranties, they compete harder for them by giving base product discounts.

If the price floor $\underline{b}$ does not bind, i.e. if the marginal costs of production are higher than the profits that can be extracted because of the extended warranties, then firms never obtain positive profits. This is because all extended warranty profits are competed away via loss leadership of the base product. In this latter case, sophisticated consumers always benefit from each additional naive consumer.

[15] This follows from $q(1-q)I$ being maximized for $q = \frac{1}{2}$.

Corollary 1. *If* $\frac{c-\underline{b}}{\tilde{\pi}(q,\tilde{w})} > \theta$*, sophisticated consumers' utility is strictly increasing in* θ*.*

The higher the number of firms M in the market, the lower extended warranty profits can be for a single base product quantity equilibrium to exist. The intuition for this is that the equilibrium profits decrease in the number of firms while the optimal deviation profits are invariant to the number of firms in the market. If $c - \theta\frac{\tilde{\pi}(q,\tilde{w})}{2} < \underline{b} \leq c - \frac{M-1}{2M-1}\theta\tilde{\pi}(q,w)$, there also exists at least one more equilibrium in which consumers buy multiple products. This is because firms can possibly decrease the base product price to sell multiple products. This equilibrium leads to lower profits compared to the equilibrium in which consumers only buy a single base product each. If $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$, however, there is no equilibrium in which consumers purchase multiple base product quantities: Selling multiple base product quantities per consumer at a price below $\underline{b}$ in this case does not allow firms to fully recoup their losses via selling extended warranties. This leads to the following sufficient condition for firms always making positive profits in equilibrium.

Corollary 2. *If* $\frac{c-\underline{b}}{\tilde{\pi}(q,\tilde{w})} < \theta < \frac{2(c-\underline{b})}{\tilde{\pi}(q,\tilde{w})}$*, firms' equilibrium profits are positive.*

As M goes to infinity, the loss per base product has to be at least one half of the warranty profits per consumer for no firm having an incentive to deviate from a price floor equilibrium. The fact that the industry equilibrium characteristics depend on the number of retail firms in the industry also has important implications regarding merger policy. Consider the case in which there are M firms in the market, and they price their products such that consumers buy multiple products. After a merger, with only $M-1$ firms in the market, the multi-product case may not be feasible anymore. As a consequence, firms could increase prices such that consumers only buy a single product. Formally, this always holds if $\frac{M-2}{2M-3} > \frac{c-\underline{b}}{\theta\tilde{\pi}(q,w)} > \frac{M-1}{2M-1}$.[16]

One might argue that when insuring themselves already against product breakdown by buying multiple base products, consumers are not willing to buy any positive extended warranty contract anymore. When making this simplifying assumption, the base product price floor becomes more rigid: In this case it is never possible for firms to sell multiple products to consumers without making losses, such that the product will always be priced at least at the price floor.

Besides the option to buy multiple products, another motivation for such a price floor

[16]Note that while in Gabaix and Laibson (2006) the number of firms in the market determines whether firms can make positive profits in equilibrium, this occurs for a different reason. In their model, a firm's decision to "unshroud" add-on costs and educate naive consumers depends on the number of firms. If the number of firms is sufficiently high in the industry, a firm has an incentive to unshroud the add-on prices, which results in a zero profit equilibrium. In contrast, in our model, the number of firms in the industry determines whether or not firms have an incentive to deviate from setting a the price equal to the potential price floor such that each consumer only buys a single product.

could of course be contracts between retailers and a manufacturer regarding the minimum retail price such that the manufacturer can preserve a high brand value.[17]

In case product quality is not observable to consumers, a warranty can be used to signal product quality upfront. When extended warranties are salient factors in the base product decision, this can naturally prevent the retailers' monopoly positions at the point of sale.

3.2 Evaluation of consumer protection policies

In this subsection, we analyze the effectiveness of three policies, namely, inducing competition at the point of sale, having minimum required warranty levels, and implementing a cancellation policy for extended warranty contracts.

Competition for extended warranties at the point of sale Suppose that every retail firm is required to show consumers warranty quotes from independent extended warranty providers before it can sell an extended warranty contract. With competition among extended warranty providers at the point of sale next to retail competition for base products, the prices for extended warranty contracts decrease to the expected claim costs, $(1-q)F(w)w$. Base product prices increase to marginal costs, such that consumers always buy a single product in equilibrium. This eliminates all industry profits. Furthermore, this outcome stops the transfer from naive consumers to sophisticated consumers.

Proposition 2. *(Effects of extended warranty competition at point of sale) Suppose that there is competition for extended warranties at the point of sale. In equilibrium, each retail firm i sets a base product price $p_i = c$. Naive consumers buy an extended warranty contract $\gamma_c = (x_c, \tilde{w})$, where $\tilde{w}$ is defined by equation (3), and $x_c = (1-q)F(\tilde{w})\tilde{w}$. All firms make zero profits in equilibrium. If $\tilde{w} > 0$, the policy increases consumer surplus and does not change overall welfare.*

Minimum required warranty level In cases in which firms can profitably sell positive extended warranty contracts to naive consumers while consumers only buy a single base product, the adoption of a positive minimum warranty level $w^{min} \leq \tilde{w}$ has ambiguous effects. Such a policy lowers the appeal of a positive extended warranty contract purchase for naive consumers and thus also reduces the prices of such contracts. This strictly lowers

[17] Furthermore, a sufficiently low base product price might cause naive consumers to become suspicious about an offer being "too good to be true," thus preventing them from buying an extended warranty contract. Heidhues *et al.* (2016) present a microfoundation of consumer suspicion in which consumers are uncertain about firms' marginal costs. In their setting, consumer suspicion itself can be a reason for a base product price floor in equilibrium.

the degree of loss leadership of the base products if the warranty level is high enough that naive consumers believe they derive a positive utility from claiming the minimum warranty with positive probability. As a consequence, in this case cross-subsidization of sophisticated consumers through the extended warranty profits made by naive consumers is reduced. However, a positive minimum warranty level strictly increases the return costs that sophisticated consumers face if with positive probability they claim the warranty payment while at the same time facing positive return costs, which decreases overall welfare. In cases in which firms make zero profits in absence of the policy, the extra return costs paid by sophisticated consumers lead to a decrease in the overall consumer surplus.

In case a consumer buys a single product when both single- and multi-product quantity equilibria exist in absence of the minimum warranty level, a second positive effect of this policy can arise: there can be a switch from a single-product to a multi-product equilibrium, which causes a base price decrease from which even sophisticated consumers can benefit. Note, however, that when a minimum warranty level is introduced, the single-product equilibrium always remains feasible, while the possibility of a multi-quantity equilibrium can vanish because of the extended warranty profits weakly decreasing due to the policy.

Proposition 3. *(Effects of minimum warranty level) Suppose consumers only buy a single product in absence of any policy. Any positive minimum warranty level* w^{min}*,* $0 < w^{min} \leq \tilde{w}$*, has a weakly positive effect on naive consumers, and an ambiguous effect on sophisticated consumers and on overall consumer surplus. The policy weakly decreases welfare.*

If $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$, consumers only buy a single base product in equilibrium both with and without intervention, such that a minimum default warranty weakly decreases sophisticated consumers' surplus, and strictly if the base product price goes up because of lower redistributed extended warranty profits. Only if there are multiple equilibria both with and without the policy, and the policy induces a switch to more quantities being purchased in equilibrium, the policy can increase sophisticated consumers' surplus.[18]

In case consumers buy multiple base product quantities in equilibrium in absence of any policy, a positive minimum default warranty can even decrease the consumer surplus of both types. Consider for example the case in which because of the minimum warranty level there is a switch to a price floor equilibrium which increases firm profits. If the increase in the base product price is higher than the savings in the extended warranty prices for naive consumers, both consumer types' surplus decreases.

[18] In case $w^{min} > \tilde{w}$, without further assumptions, there are cases in which even naive consumers can become excluded from buying the product, such that the policy can decrease naive consumers' surplus. The other results of the proposition under this assumption remain qualitatively.

Free extended warranty contract cancellation policy Another form of consumer protection is a free cancellation policy. While such a policy usually expires after the first use of a product for physical products such as TVs and notebooks, such a strong expiration policy is not needed for add-on insurance contracts such as extended warranties and for many financial products. The UK's 2005 Extended Warranty Order grants consumers the right to cancel an extended warranty contract within the first 40 days after signing the contract. We assume that a fraction $\beta \in (0, 1]$ of consumers becomes aware of all other extended warranty contracts in the market after purchasing a base product and potential extended warranty contract at a retailer. In such a case, consumers can cancel their old contract without costs and switch to the new contract.[19] A potential product breakdown occurs only after this opportunity to switch or cancel the initial contract. Furthermore, we assume that there are at least two independent warranty providers in the market that do not offer any base product and that set their warranty contracts after observing the base product prices. This relates to a relatively competitive independent insurance market.

A free cancellation period does not change naive consumers' beliefs about their return costs. Therefore, these consumers are still willing to pay the same amount for an extended warranty contract as before and would cancel it only if they receive a contract with a better offer. For brevity, we restrict our analysis to cases in which there is a unique equilibrium in which consumers buy a single base product quality in absence of any policy, i.e. $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$. We later discuss the effects when this does not hold informally.

If $\beta < 1$, i.e., if some consumers do not understand the possibility of returning their products, then our results show that retail firms do not have any incentives to change their contract structures.[20] This lack of incentive is evident because retail firms know that they will never be able to derive a positive markup from naive consumers who become aware of rival contracts, and therefore, these firms focus on maximizing the profits from consumers who will never switch. Therefore, all naive consumers that become aware of the option to switch contracts immediately do so to an extended warranty contract that the independent insurers offer at zero markup. When some consumers do not become aware of rival offers, this policy is thus inferior to inducing competition at the point of sale. If $\beta = 1$, the policy is

[19] β can account either for consumers' unawareness of rival offers or for additional switching costs that prevent some consumers from changing their contracts.

[20] Inderst and Ottaviani (2013) analyze the effect of a money return policy in a cheap talk game in which a seller tries to convince a buyer of the suitability of a product. In their model, when consumers are "credulous", i.e., when they believe the advice of a self-interested salesperson, a minimum refund level enhances the overall surplus and efficiency. The results differ from ours because of the absence of externalities between different consumer types in their model and because returns for their products are not costly. Also see Petrikaite (2017) for a model with fixed anticipated return costs and sequential consumer shopping.

as effective as inducing competition at the point of sale. In this case, the policy may induce retail firms to offer competitively priced extended warranty contracts.

Proposition 4. *(Effects of extended warranty cancellation policy) Let* $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$. *Suppose there is a money return policy,* $M \geq 2$ *retail firms, and at least two independent extended warranty providers.*

1. If $\beta < 1$*, in equilibrium, each firm sets the same menu of extended warranty contracts as in the case without intervention. A fraction* β *of naive consumers cancels their initially purchased extended warranty contracts and switches to a competitively priced contract from an independent retailer.*

2. If $\beta = 1$*, in equilibrium, all firms earn zero profits, and all consumers end up purchasing a competitively priced extended warranty contract.*

The proposition shows that in a single quantity equilibrium, a free cancellation policy alone can affect the retail firms' incentives to change their warranty contract structures only if all naive consumers are aware of the possibility of return. In these situations, the policy leads to the same consumer surplus outcome as inducing competition at the point of sale. If some naive consumers are not aware of the cancellation option, those who are aware of the policy switch to an independent warranty provider. Those who are unaware of this option do not switch options and have the same expected utility as before. One positive aspect compared to a minimum warranty level is that this policy does not induce an additional welfare loss, because it does not lead to an increase in warranty levels purchased. Note, however, that we did not consider other costs of policy cancellation arising from negative emotions and opportunity costs of time. Such considerations could further reduce the policy's efficiency.

If there is no unique equilibrium both with and without the policy, i.e. if $\underline{b} \geq c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$ and $\beta = 1$, the results are as in the proposition. If $\beta < 1$, similar as in case of a minimum default warranty level, because of a multiplicity of equilibria, the policy can lead to ambiguous consumer surplus effects.

Overall, if competition at the point of sale can be effectively introduced, this is always at least as efficient from an overall consumer surplus perspective as the introduction of a minimum warranty standard or a free cancellation policy. Furthermore, induced competition does not have the risk of distorting the market negatively with respect to consumer surplus that a minimum default warranty level might have.

3.3 Unshrouding of extended warranty contracts

We next consider the case in which each firm has the opportunity to "unshroud" the menus of extended warranty contracts Γ of all firms and thus to make naive consumers aware of the

possibility of buying an extended warranty contract before they purchase a base product. This follows the work of Gabaix and Laibson (2006) and, in particular, Heidhues *et al.* (2016). The following corollary provides sufficient conditions for a shrouding equilibrium to exist, which will subsequently be discussed.

Corollary 3. *Assume at least one of the following three assumptions holds:*
1. When a retail firm unshrouds the extended warranty contracts, naive consumers become aware of a competitive insurance market for extended warranties.
2. Unshrouding the extended warranty contracts makes naive consumers fully sophisticated.
3. Naive consumers do not anticipate that an extended warranty can be of use to them before making a base product purchase, even if the extended warranty contracts are unshrouded. Then, the equilibrium in Proposition 1 is a shrouding equilibrium.

There are several reasons why disclosing extended warranty quotes may not be profitable for a retail firm. The existence of a competitive insurance market on the internet can lead to retailers shrouding their extended warranty contracts in equilibrium if unshrouding those contracts also induces consumer search for the lowest extended warranty price available. Assume that if a retailer discloses the extended warranty contracts in the industry before consumers enter a store, then consumers start searching for extended warranty quotes from an independent competitive insurance market, such as from providers on the internet. This leads to Bertrand competition for extended warranties and to zero firm profits for retailers. Thus, disclosing warranty quotes is not profitable for any retail firm in this case.

A second case in which unshrouding is not profitable for firms arises if disclosing the extended warranty quotes makes naive consumers fully sophisticated, i.e., causes them to have correct beliefs about their return costs. In this case, unshrouding the extended warranty quotes reduces the willingness to pay of naive consumers, such that a retail firm can never earn positive profits from selling a positive extended warranty.

Third, survey evidence suggests that extended warranties are of the lowest priority for consumers when it comes to making base product decisions among all available product attributes; see, for example, Rogers (2003b). This evidence suggests that even if firms advertised these quotes, many consumers would not pay attention to them before making a product purchase, such that lowering the prices of extended warranty contracts and disclosing them would not significantly affect the product's demand. Furthermore, in many cases, lowering extended warranty quotes for firms would result in the firm needing to increase the base product price. In cases in which consumers initially pay attention only to the base product price, this would result in zero demand for a firm that decreases the prices of its extended warranty contracts.

In practice, the unshrouding of rivals' warranty prices seems to be both difficult and costly for firms. Retailers often do not freely post their extended warranty quotes; therefore, gathering these quotes is difficult. Moreover, advertising these quotes publicly through advertising is costly in practice and might also confuse consumers because of the increased complexity of the quoted contracts, which is something we abstract from in our model. If the overall cost of unshrouding is greater than the market profits from extended warranties, the firm will never educate naive consumers about the extended warranty contracts in the market.

In all of the above cases, unshrouding the extended warranty quotes leads to a decrease in the profits of an unshrouding firm, such that each retail firm prefers the shrouding equilibrium to a non-shrouding equilibrium.

4 Related theoretical literature

Our paper is primarily related to the literature on hidden add-on fees and consumer myopia in behavioral industrial organization and to the literature on warranties in both law and industrial organization. Previous work in the behavioral industrial organization literature establishes conditions under which firms choose not to disclose add-on fees upfront and potentially obtain positive profits from naive consumers in equilibrium. Our key contribution related to this literature is to show how different consumer protection policies affect equilibrium outcomes and efficiency when consumers make voluntary choices to purchase add-on products, i.e., when warranty contracts are disclosed to them as options at the point of sale.

Gabaix and Laibson (2006) give a bounded rationality explanation for why firms shroud the prices of add-on products in equilibrium. They develop a competitive model in which the purchase of a base product implies subsequent add-on costs. Sophisticated consumers foresee the add-on costs even without advertising and can substitute them. In contrast, myopes are unaware of these costs if information on add-on prices is shrouded. When the number of firms is sufficiently low, shrouding can persist in equilibrium.

Heidhues *et al.* (2016) build on the Gabaix and Laibson framework and show that when there exists a price floor for base products, firms can earn profits by selling deceptive products to naive consumers, which can persist in equilibrium for any number of firms. When shrouding is costly, the deceptive shrouding equilibrium is unique. The authors provide several microfoundations for the base product price floors, such as adverse selection and consumer suspicion. From a modeling perspective, our model has the feature of Gabaix and Laibson (2006) with both sophisticated and naive consumers buying the same base products in equilibrium, and it can contain base price floors as in Heidhues *et al.* (2016). Unlike both of these models, an add-on in our model is reflected not only by an additional potentially shrouded

cost but also by a payment in case of product breakdown whose value is incorrectly predicted by naive consumers. This often makes it beneficial for firms to offer a menu of add-on products. In our model, the consumers' incentive to buy multiple products in case of a low base product price can endogenously create a base product price floor. Furthermore, unlike in Gabaix and Laibson (2006), there is no extra effort cost in not choosing an add-on. This is reflected in consumers' ability to simply choose no extended warranty contract. Unlike in Heidhues *et al.* (2016), both consumer types choose the same base product in equilibrium. This underlines the non-"deceptive" character of the base product.

Miao (2010) considers a dynamic duopoly model in which new consumers arrive every period and each consumer lives for two periods. In his second period, a then "old" consumer needs to buy an additional printer cartridge for a printer that is initially bundled with a single cartridge, which myopic consumers fail to anticipate ex-ante. In equilibrium, firms set the bundle price sufficiently high to keep old consumers from buying a new bundle in the second period to use the included cartridge. In contrast, from an ex-ante perspective in our static model, purchasing multiple base products for a sufficiently low product price can yield an increased inherent expected consumption utility.

In our companion paper Michel (2016), we show that when product quality is unobservable to consumers, a firm can offer lifetime warranties upfront together with strong exclusion fine print to profitably establish and exploit quality misperceptions among naive consumers. In contrast to the present paper, this relates to manufacturing firms that set a low product quality to save production costs and use an associated warranty to induce a high quality perception in naive consumers. Sophisticated consumers never gain from the presence of naive consumers. This is because the initially offered contracts already suffice to screen between the different consumer types. Furthermore, the main role of consumer protection policies also differs from the one in the present paper: instead of keeping consumers from overpaying for add-on contracts, the main goal of a protection policy in this case is to prevent consumers from buying "scam" products. The paper offers supporting evidence from infomercials for consumers not anticipating positive return costs in case of making a warranty claim, and for firms including several exception clauses in the contract fine print.

Armstrong and Vickers (2012) study forms of add-on price regulation in a similar setting as Heidhues *et al.* (2016). With regard to banking regulation, the authors argue that adopting specific regulations such as overdraft warnings and hard budget constraints intended to help naive consumers may result in overall market inefficiencies. Ko and Williams (2013) further find that the combination of add-on price disclosure and sufficient add-on price caps weakly increases consumer welfare, while in isolation, the effects are ambiguous. Grubb (2015) shows that a bill-shock regulation can decrease consumer surplus under sophistication, whereas it

increases consumer surplus if inattention with respect to past usage is not anticipated.

Inderst and Ottaviani (2013) study a relationship between a financial adviser and a private investor, and they show that granting a cancellation right to rational consumers who foresee the adviser's self-interest can make his cheap talk credible. If all consumers believe the advice the seller gives, then the seller always claims that a product is the most suitable for a consumer. A sufficiently high minimum refund level is both a consumer surplus and efficiency-enhancing tool, as it lowers the expected rents that the firm receives from false suitability claims. By contrast, in our model, such a policy can decrease efficiency because of the return costs and the interaction between both consumer types in a competitive environment.

Furthermore, the initial consumer neglect with respect to extended warranty prices when making a base product choice relates to the literature on context-dependent decision making and focusing, see for example Bordalo *et al.* (2012) and Kőszegi and Szeidl (2013). Bordalo *et al.* (2014) analyze a model of market competition in which strategic product positioning can affect the salience of different product features.

More generally, there is a growing body of literature on industrial organization models with a behavioral economics foundation; see Ellison (2006) and DellaVigna (2009) for broad literature reviews and the review of Armstrong (2008) for a summary of the consumer policy literature. Armstrong and Chen (2009) develop a model in which some consumers do not pay attention to the quality component of products when making their purchasing decision under firm competition. In a symmetric mixed-strategy equilibrium, the existence of inattentative consumers is responsible for positive firm profits. Sandroni and Squintani (2007) show in an insurance model of asymmetric information that if there is a high fraction of overconfident consumers, a policy intervention in terms of minimum insurance is socially wasteful.

The paper also relates to the literature on extended warranties in both law and industrial organization. Baker and Siegelman (2013) summarize the related law and economics literature with respect to extended warranties. Mann and Wissink (1990) assess the effectiveness of money-back and replacement warranties both when the product quality is observable and when it is not. The authors find that a money-back warranty is always better, except for an intermediate range of replacement costs for a firm. Lutz and Padmanabhan (1998) develop a model in which independent non-manufacturing firms can enter the market to sell extended warranties in the presence of a manufacturing monopoly. The effect on the manufacturer's profits is ambiguous.

5 Conclusion

In this paper, we analyze how different consumer protection policies affect market outcomes when retail firms can offer add-on insurance for their products. In our model, naive consumers' return cost misperceptions and the non-salience of warranties before a base product choice results in the profitability of extended warranty contracts at the point of sale. Under retail competition, this forces firms to sell a base product as a loss leader, from which sophisticated consumers benefit. We show that the threat of consumers buying multiple quantities of a base product can keep firms from fully redistributing the extended warranty profits via a lower base product price. In this case they set the highest possible base product price such that each consumer only buys a single base product. Such an endogenous price floor can thus lead to positive profits from selling extended warranty profits under retail competition.

The effectiveness of different consumer protection policies crucially depends on how they affect both the prices of extended warranties and the overall return costs incurred in the population. Our model yields the clear prediction that when possible, inducing competition for extended warranty contracts at the point of sale is the most favorable consumer policy because it leads retail firms to lower extended warranty prices down to the expected claim costs while avoiding increasing return costs in the population. While a minimum warranty level also decreases warranty contract prices, it results in increased social costs when leading sophisticated consumers to return faulty products in equilibrium. A policy that allows consumers to avoid extended warranty contracts increases consumer welfare but leads to a change in pricing only if there are no warranty return frictions in the population.

Our predictions are consistent with the effects of the introduction of policies that increase the competitiveness of the extended warranty market for consumer electronic products by UK legislators. One important empirical question for the future is whether the easy accessibility of independent insurers due to the internet, e.g., via smartphones, will translate into higher market shares. Recent evidence does not support this hypothesis: OFT (2012) reports that the largest UK independent insurance provider has only a 2 percent market share in the consumer electronics extended warranty industry, while almost 70 percent of the extended warranties are still sold at the point of sale. Furthermore, in many countries, amazon.com, as a leading retailer, only offers a single extended warranty provider for its electronic products, and consumers can buy such contracts only on the same day they buy the base product.

A particular aspect of selling products over the internet is that the point of sale is virtual. If needed, introducing competition for add-on insurance at the point of sale could, for example, be induced via a common mandatory insurance comparison website that must be shown by each provider. However, problems of policy legislation arise for products that are

sold by a single firm in different countries.

Finally, the industries in which add-on insurance is the most prevalent, such as retail markets for consumer electronics and low-cost airlines, typically share a pattern of relatively low profit margins. This is consistent with our model in which the base products are sold as a loss leader. An open question in this case is whether the introduction of add-on insurance by these firms increases the competitiveness of the associated base product markets or is rather a response to an increased industry competitiveness itself.

References

Armstrong, M. (2008). Interactions between competition and consumer policy. *Competition Policy International*, **4** (1), 97–147.

— and Chen, Y. (2009). Inattentive consumers and product quality. *Journal of the European Economic Association*, **7** (2-3), 411–422.

— and Vickers, J. (2012). Consumer protection and contingent charges. *Journal of Economic Literature*, **50** (2), 477–493.

Arnum, E. (2013). State of the industry & market sizing. *Extended Warranty and Service Contracts Conference*, `http://warrantyinnovations.com/pdf/2013-Speaker-Presentations/ericarnum.pdf/`.

Baker, T. and Siegelman, P. (2013). Protecting consumers from add-on insurance products: New lessons for insrance regulation from behavioral economics. *University of Pennsylvania Research Paper*.

Bar-Gill, O. (2007). The behavioral economics of consumer contracts. *Minnesota Law Review*, **92**, 749.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, **1243**, 1285.

—, — and — (2014). Competition for attention. *Review of Economic Studies, forthcoming*.

Businessweek (2004). The warranty windfall. *http://www.businessweek.com, December 20*.

Chen, T., Kalra, A. and Sun, B. (2009). Why do consumers buy extended service contracts? *Journal of Consumer Research*, **36** (4), 611–623.

Competition Commission (2003). Competition Commission report on extended warranties for domestic electrical goods. *Policy report*.

DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*, **47** (2), 315–372.

— and Malmendier, U. (2004). Contract design and self-control: theory and evidence. *Quarterly Journal of Economics*, **119** (2), 352–402.

Ellison, G. (2006). Bounded Rationality in Industrial Organization. In R. Blundell, W. Newey and T. Persson (eds.), *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Cambridge University Press.

Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, **121** (2), 505–540.

Grubb, M. D. (2015). Consumer inattention and bill-shock regulation. *The Review of Economic Studies*, **82** (1), 219–257.

Heidhues, P., Kőszegi, B. and Murooka, T. (2016). Inferior products and profitable deception. *The Review of Economic Studies, forthcoming*.

Huysentruyt, M. and Read, D. (2010). How do people value extended warranties? evidence from two field surveys. *Journal of risk and uncertainty*, **40** (3), 197–218.

Inderst, R. and Ottaviani, M. (2013). Sales talk, cancellation terms and the role of consumer protection. *The Review of Economic Studies*, **80** (3), 1002–1026.

Jindal, P. (2014). Risk preferences and demand drivers of extended warranties. *Marketing Science, forthcoming*.

Jolson, M., Wiener, J. and Rosecky, R. (1987). Correlates of rebate proneness. *Journal of Advertising Research*, **27** (1), 33–43.

Ko, K. J. and Williams, J. (2013). The effects of regulating hidden add-on costs. *mimeo*.

Kőszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, **97** (4), 1047–1073.

— and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly Journal of Economics*, **128** (1), 53–104.

Loewenstein, G., O'Donoghue, T. and Rabin, M. (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics*, **118** (4), 1209–1248.

Lutz, N. and Padmanabhan, V. (1998). Warranties, extended warranties, and product quality. *International Journal of Industrial Organization*, **16** (4), 463–493.

Mann, D. and Wissink, J. (1990). Money-back warranties vs. replacement warranties: A simple comparison. *American Economic Review*, **80** (2), 432–436.

Miao, C.-H. (2010). Consumer myopia, standardization and aftermarket monopolization. *European Economic Review*, **54** (7), 931–946.

Michel, C. (2016). Contractual structures and consumer misperceptions. *mimeo*.

OFT (2002). Extended warranties on domestic electrical goods. *Office of Fair Trading market study report OFT387*.

— (2012). Extended warranties on domestic electrical goods. *Office of Fair Trading market study report OFT1403*.

Petrikaite, V. (2017). Reversed search for the best product: a model of costly product returns. *mimeo*.

Rogers, D. (2003a). NOP Report of Quantitative Research into Consumer Experiences of Purchasing Extended Warranties: Annex 1. *Research Report prepared for the Competition Commission*.

— (2003b). NOP Report of Quantitative Research into Consumer Experiences of Purchasing Extended Warranties: Annex 2. *Research Report prepared for the Competition Commission*.

Sandroni, A. and Squintani, F. (2007). Overconfidence, insurance, and paternalism. *American Economic Review*, **97** (5), 1994–2004.

Sydnor, J. (2010). (over) insuring modest risks. *American Economic Journal: Applied Economics*, **2** (4), 177–199.

Warranty Week (2005). Extended warranty administrators: While auto and pc manufacturers have the top spots, insurance companies and third party administrators grab the bulk of the pie. *http://www.warrantyweek.com, January 19*.

Westra, B. (2002). Report on the consumer survey about extended warranties. *Report prepared for the Office of Fair Trading by FDS International Ltd*.

A Proofs

Proof of Proposition 1

The proof goes in several steps. 1. As shown in the main text, from $q \geq \frac{1}{2}$ it follows that no firm has an incentive to set extended warranty (henceforth EW) prices such that consumers buy multiple EW contracts even if they buy multiple base contracts. We thus restrict to the single EW purchase case. 1a. Assume first that $\theta = 0$. For a given warranty payment $w > 0$, sophisticated consumers are willing to pay a maximum warranty price $x = (1-q)F(w)E[w-r|w>r] \leq 0$ such that selling an EW to sophisticated consumers never leads to positive firm profits. 1b. Assume next that $\theta = 1$. We show that no firm can do better than to offer a positive EW contract that maximizes the difference between a naive consumer's predicted rent from the EW contract and the expected claim costs next to a 0 coverage default contract $\gamma_0 = (x, w) = (0, 0)$. Because naive consumers do not pay attention to the EW contracts before making a base product choice, EW contracts will not affect their base product choices. Because each firm has a quasi-monopoly at the point of sale, it can charge a naive consumer's full predicted willingness to pay for such a warranty, which leads to the EW contract in equation (3). Given a 0 coverage default contract, any warranty level for a positive EW contract that does not maximize this difference leads to a lower profit for any firm. 1c. We show that the menu of EW contracts $\Gamma = \{(0,0), (\tilde{\pi}(q,\tilde{w}), \tilde{w})\}$ is also optimal whenever both consumer types are present, i.e. when $\theta \in (0,1)$. For the same reasons as when $\theta = 1$, a firm's expected warranty profits are weakly decreasing in the default EW coverage. For the same reasons as when $\theta = 1$, each firm selling a base product to both consumer types can do no better than offering a contract $(\tilde{x}, \tilde{w})$ that is purchased by naive consumers. It thus follows that there is no menu of EW contracts leading to higher EW profits than $\Gamma = \{(0,0), (\tilde{\pi}(q,\tilde{w}), \tilde{w})\}$.

2. If a firm has positive demand from consumers, it has again a monopoly for selling EW contracts. Therefore, as shown in 1., no firm can do better than offering a menu of EW contracts $\Gamma = \{(0,0), (\tilde{\pi}(q,\tilde{w}), \tilde{w})\}$. It thus also follows that this is the optimal menu of EW contracts when setting a base product price such that consumers buy multiple quantities of the base product. 3. As shown in the main text, if the price is at least $\underline{b} = q(1-q)I$, each consumer cannot do better than buying at most one base product. 4. We next prove the first part of the proposition. Each consumer chooses one of the firms with the lowest base product price equally likely, given that this price is equal or below the maximum willingness to pay Iq. Therefore, each firm with a positive market share faces both consumer types with the same proportion as in the whole population, which yields an expected profit $\theta\tilde{\pi}(q,\tilde{w})$ from selling

a positive EW contract per consumer, as shown in 1. If $c - \theta\tilde{\pi}(q, \tilde{w}) > \underline{b}$, in a symmetric equilibrium firms set a base product price $p_i = c - \theta\tilde{\pi}(q, \tilde{w})$. Suppose not and that all firms have a price above this level. Then at least one firm can deviate by slightly undercutting the lowest price of its competitors to increase its profits. Next, suppose not and a firm sets a price below this level. This leads to negative profits for this firm which cannot be optimal. Therefore, the only equilibrium candidate that survives has all firms with positive market shares setting a price $p_i = c - \theta\tilde{\pi}(q, \tilde{w})$. Suppose next that $c - \theta\tilde{\pi}(q, \tilde{w}) < \underline{b}$, such that $\underline{b}$ binds as base product price. Any firm deviating to a higher base product price would again lead to zero demand. Because of point 2. and because of the base product being sold as a loss leader, the optimal deviation when lowering the price would be to slightly undercut the base product price. Any lower price would lead to full demand for all consumers, but also to consumers buying multiple base products. In this case, because the base product being a loss leader and from 2. it follows again that the optimal deviation would be to undercut the base product price to only slightly undercut the base product bound $\underline{b}$ while maintaining the same menu of EW contracts. This deviation results in all consumers buying exactly two base products from this firm. This is less profitable than keeping the base product price at $\underline{b}$ if $\frac{1}{M}(\underline{b} - c + \theta\tilde{\pi}(q, \tilde{w})) > 2(\underline{b} - c) + \theta\tilde{\pi}(q, \tilde{w})$. But this can exactly be rewritten as the condition in the proposition, such that no profitable deviation is possible when the condition holds. 5. We next prove the second part of the proposition, i.e. uniqueness of the equilibrium if $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$. With the same reasoning as in step 4, there cannot be an equilibrium in which firms set prices higher than $\underline{b}$, because at least one firm would always have an incentive to undercut the other firms. Any firm setting a base product price below $\underline{b}$ and having positive demand always experiences losses: the firm sells $k \geq 2$ base products to consumers that are sold as a loss-leader. From 2. and the product being sold as a loss leader it follows that a firm cannot achieve a higher profit than $2\frac{\underline{b}-c}{k} + \theta\tilde{\pi}(q, \tilde{w}) < 0$ because of $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$. This completes the proof.

Proof of Proposition 2

We first show that at least two independent EW providers offer an EW contract $((1 - q)F(\tilde{w})\tilde{w}, \tilde{w})$. This follows from a simple Bertrand pricing argument. Each EW provider only offers contracts that lead at least to zero expected profits. Suppose now that all EW providers set contracts such that at least one firm makes positive profits while setting an EW contract (x, w), $w \notin \arg\max_{w' \in [0,\overline{w}]} \tilde{\pi}(q, w')$. Then at least one firm can always profitably deviate by setting an EW contract $(x - \epsilon, \tilde{w})$, for any arbitrarily small and positive ϵ. But any profitable EW price with warranty level $\tilde{w}$ will be undercut by a different firm. Therefore,

if at least two firms set $((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$ and no other firm sets a contract that generates expected losses if it were the only firm in the market, no firm has an incentive to deviate. This leads to the contract $((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$. This implies zero expected profits from the EW contracts for each retail firm. Therefore, price competition for the base product leads to base product prices equal to marginal costs, i.e. $p_i = c$ for each retail firm i in the industry. This completes the proof.

A.1 Proof Proposition 3

1. Given the minimum default EW level w^{min}, the new default contract turns to $(x, w) = (0, w^{min})$ under this regulation. Compared to the case without minimum default warranty level the firm now has to give a rent $(1-q)\tilde{F}(w^{min})\tilde{E}[w^{min} - r | w^{min} > r]$ to naive consumers for any positive EW contract. The profits for a firm selling a positive contract with coverage $w > w^{min}$ to naive consumers, given that it faces both consumer types are thus $\theta\tilde{\pi}(w, q) - \theta(1-q)\tilde{F}(\tilde{E}[w^{min} - r | w^{min} > r] - (1-\theta)(1-q)F(w^{min})w^{min}$. But this results in the problem $\max_{w \in [w^{min}, \overline{w}]} \theta[\tilde{\pi}(q, w) - \tilde{\pi}(q, w^{min})] + (1-q)F(w^{min})w^{min}$, which then results in the same optimal warranty level $\tilde{w}$ as in Proposition 1. The same optimal warranty level $\tilde{w}$ arises when no sophisticated consumer or only a part of the sophisticated consumers demands the base product. Similar to the proof of Proposition 1, a firm's profit is weakly decreasing in the default minimum warranty level, which sophisticated consumers anticipate. Therefore, each firm in equilibrium offers a menu of EW contracts $\Gamma^{min} = \{(0, w^{min}), (\tilde{\pi}(q, \tilde{w}) - (1-q)\tilde{F}(\tilde{w})\tilde{E}[w^{min} - r | w^{min} > r])\}$. 2. Suppose both consumer types still demand the product after the policy. An equilibrium in which both types only purchase a single quantity always remains feasible, because firms have weakly lower EW profits due to the policy. With an identical argument as in the proof of Proposition 1, in this case the price turns to $p^{min} = \max[0, c - \theta[\tilde{\pi}(q, \tilde{w}) - \tilde{\pi}(q, w^{min})] + (1-q)F(w^{min})w^{min}]$. It follows that as long as $F(w^{min}) > 0$, naive consumers are strictly better off because of this policy. Sophisticated consumers' change in utility compared to no intervention is $-\theta\tilde{\pi}(q, w^{min}) + \pi(q, w^{min}) < 0$. Thus, if in absence of any intervention the equilibrium price is not binding at zero, it is sufficient for sophisticated consumers to receive a lower utility due to the minimum warranty level. The change in overall consumer surplus in this case is $-(1-\theta)(1-q)F(w)(w^{min} - E[w^{min} - r | w^{min} > r]) \leq 0$. If the price without intervention is binding at zero and firms make positive profits, there are cases in which an intervention can decrease the expected warranty profits while still not be able to increase product prices. This does not affect sophisticated consumers' surplus while it increases naive consumers' surplus. However, if the base product price increases due to the intervention, the change in consumer surplus is negative for sophisticated consumers. This

shows that a minimum warranty level weakly decreases sophisticated consumers' surplus in this case. When the price is binding at 0 without regulation, a minimum expected warranty level increases consumer surplus whenever the decrease in firms' profit is bigger in absolute magnitude than the increase in the expected return costs that consumers face. This already shows that the effects on overall consumer surplus are ambiguous. The change in overall welfare due to the intervention is $-(1-\theta)(1-q)F(w)(w^{min} - E[w^{min} - r|w^{min} > r]) \leq 0$, which implies that the policy weakly decreases overall welfare. 3. Suppose that both consumer types demand multiple products after the introduction of the policy. This implies that the total price per product a naive consumer pays weakly decreases, which implies that naive consumers benefit from the policy. We next show that in this case the policy has ambiguous effects on sophisticated consumers. As an example, suppose there are 2 firms, and $\underline{b} = I[q(1-q)] = c - \frac{\theta\tilde{(q,\tilde{w})}}{3}$, while $I[1-(1-q)^3] < c - \frac{\theta\tilde{(q,\tilde{w})}}{2}$. In absence of the policy, there exists an equilibrium in which each consumer buys a single base product at price $p = c - \frac{\theta\tilde{(q,\tilde{w})}}{3}$, such that each sophisticated consumer has an expected surplus $Iq - c + \frac{\theta\tilde{(q,\tilde{w})}}{3}$. Suppose further that with positive probability consumers incur 0 return costs. Then, for any positive minimum default warranty level $w^{min} = \epsilon$, the above equilibrium is not feasible anymore. For ϵ close to zero, the new equilibrium price can thus be approximated as $p = c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$, such that the expected surplus of a sophisticated consumer turns to $I[q(1-q)] - 2c + \theta\tilde{\pi}(q,\tilde{w})$. This is higher than in the former case if $c < I[q(1-q)] + \frac{2}{3}\tilde{\pi}(q,\tilde{w})$, which holds because $c = I[q(1-q)] + \frac{\theta\tilde{(q,\tilde{w})}}{3}$. Because $c > I[q(1-q)]$, selling multiple products to consumers further reduces overall welfare. 4. Suppose that only a fraction α, $0 \leq \alpha < 1$ of sophisticated consumers buys a positive quantity of the base product. Then it follows that sophisticated consumers derive at most a zero consumption utility from the base product including minimum default warranty after the policy, such that they are worse off than without the policy, which also implies an overall welfare loss. The fraction of (naive) consumers that buys an EW contract compared to those that only buy a product in the population is thus $\frac{\theta}{\alpha(1-\theta)+\theta} > \theta$. Therefore, with a similar argument as in Proposition 1, the base product price weakly decreases in this case, while the EW prices also decrease. This is sufficient for naive consumers to be at least weakly better off. This completes the proof.

Proof of Proposition 4

When $\tilde{w} = 0$ this holds trivially. We next look at the case when $\tilde{w} > 0$. Because $\underline{b} < c - \frac{\theta\tilde{\pi}(q,\tilde{w})}{2}$, from Proposition 1 it follows that the single base product quantity equilibrium is unique. 1. We first focus on the case in which $\beta < 1$. We first show that given that there is at least one retail firm that only offers positive EW contracts $\gamma = (x, w)$ that are not equal

$\gamma^{ind} \equiv ((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$ to consumers, and $\tilde{\pi}(w,q) \geq x$, for all of these contracts - i.e. they make no losses with these contracts - all independent EW providers that sell a positive market share will always offer an EW contract $\gamma = ((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$. This follows from the fact that a fraction $\beta > 0$ of naive consumers become aware of the different independent EW contracts joint with a Bertrand pricing argument which is identical to that in the proof of Proposition 2. Therefore, given a firm sets an EW contract $(x,w) \neq \gamma^{ind}$ with $\tilde{\pi}(q,w) \geq x$ that gets initially chosen by a naive consumer with positive probability, a fraction of β of these consumers cancel their contract because of becoming aware of the independent warranty providers' contracts and switching to these. Offering a contract (x,w) with $\tilde{\pi}(q,w) < x$ that gets chosen by naive consumers with positive probability is strictly dominated by offering only a default contract $(0,0)$. Offering an EW contract γ^{ind} next to a $(0,0)$ default contract will yield to zero EW profits. However, a fraction of $\theta(1-\beta)$ consumers is not aware of the option to switch EW contracts. But then it follows that firms maximize their EW profits by offering the menu of EW contracts $\Gamma = \{(0,0),(\tilde{x},\tilde{w})\}$ as in Proposition 1, which sophisticated consumers correctly anticipate. With a similar argument as in that Proposition it follows that all retail firms who have a positive market shares set a base product price $p = \max[0, c - \beta\theta\tilde{\pi}(q,\tilde{w})]$. 2. Suppose $\beta = 1$. All naive consumers always buy a contract $\gamma = ((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$. Given a firm sets a positive EW contract $(x,w) \neq \gamma^{ind}$, it will make a non-positive profit from it. Therefore, from a Bertrand pricing argument, base product competition will lead to $p = c$ for each product. From above it still holds that at least two independent EW providers always offer an EW contract γ^{ind} if at least one retail firm does not offer this contract and makes non-negative EW profits. Therefore all naive consumers end up buying an EW contract γ^{ind} in equilibrium whenever $\beta = 1$. This completes the proof.

Proof of Corollary 3 We first prove point 1 of the corollary. If a firm unshrouds the menus of EW contracts Γ given that all of its rivals do not unshroud the warranty contracts, consumers become aware of a competitive EW market. From Proposition 3 it follows that in this market independent insurance firms offer an EW contract $\gamma_c = ((1-q)F(\tilde{w})\tilde{w}, \tilde{w})$, such that all firms make zero profits, such that unshrouding the EW contracts is not profitable for any retail firm. We next prove point 2 of the corollary. If naive consumers become sophisticated via unshrouding of the menus of EW contracts, their willingness to pay for an EW goes down to the sophisticated consumers' willingness to pay. From step 1 in Proposition 1 it follows that naive consumers who become sophisticated do not generate positive profits anymore, which leads to at most zero profits for a firm that unshrouds the EW contracts. Therefore, given that none of its rival firms unshrouds the warranty contracts, a firm can never increase its profit by unshrouding the costs itself. We next prove point 3 of the corollary.

If consumers do not anticipate buying an EW contract before going to a store even if EW contracts are disclosed, they always buy a warranty from a retailer that sets the lowest base product price available, irrespective of the EW contracts offered. Suppose firm i unshrouds all extended warranty contracts in the industry. Because naive consumers' demand and valuation of i's base product are irrespective of the EW levels offered in i's menu of the EW contract, Γ_i, it follows that no menu of EW contracts can lead to higher EW profits for firm i than the menu $\{(0,0),(\tilde{x},\tilde{w})\}$, where $\tilde{w}$ and $\tilde{x}$ are defined by equation (3) by the definition of this menu to be profit maximizing. But this implies that given at least one other firm j sets a base product price $p_j = \max[\underline{b}, c-\theta\tilde{\pi}(q,\tilde{w})]$, firm i cannot make a higher profit from undercutting firm j's price compared to setting base product price $p_i = \max[\underline{b}, c-\theta\tilde{\pi}(q,\tilde{w})]$. If $p_j = \underline{b}$, undercutting it would lead to lower profits as shown in Proposition 1. If $p_j = c-\theta\tilde{\pi}(q,\tilde{w}) > \underline{b}$, for any lower price, a firm makes losses. This completes the proof.

B Further extensions

B.1 Characterization of equilibria without parametric restrictions

In this section we remove several parametric restrictions in the main text and characterize how to find all equilibria of the game in this case. In particular, we now consider $q > 0$ and $c > 0$. In case of multiplicity of optimal warranty levels, there exist equilibria that are asymmetric in the extended warranty level offered but symmetric with respect to the base product price. A consumer rather purchases $k-1$ instead of $k \geq 2$ base products if the base product price is sufficiently high. This occurs if $[1-(1-q)^{k-1}]I-(k-1)p > [1-(1-q)^k]I-kp$, which can be rewritten as $p > [(1-q)^{k-1}q]I \equiv \underline{b}(k)$. Therefore, given p is the lowest base product price offered by any firm in the industry, each consumer purchases $k-1$ products if $\underline{b}(k) \leq p < \underline{b}(k-1)$. In any equilibrium in which consumers buy $k-1$ base products, competition drives prices again either down to the price floor $\underline{b}(k)$, or in case this base floor does not bind, down to 0 firm profits. With a similar reasoning as for establishing the lower price bound when buying one base product in the proof of Proposition 1, it follows that firms set the base product price $p^*(k-1)$ which can be written as

$$p^*(k-1) = \max\left[\underline{b}(k), c - \frac{\tilde{a}(k-1)\theta\tilde{\pi}(q,\tilde{w})}{k-1}\right], \tag{4}$$

where $\tilde{a}(k)$ is a scale factor that accounts for a potential increase in the maximum extended warranty profits when selling multiple extended warranty contracts instead. Note that for any integer $k \geq 2$, it holds that $k > \tilde{a}(k) \geq \tilde{a}(k-1) \geq 1$. The condition $\tilde{a}(k) \geq \tilde{a}(k-1)$ follows from the extended warranty profits being weakly increasing in the number of base

products sold. Firms still have the possibility to offer the same extended warranty contracts when offering the base product at a higher price, which would lead to the same extended warranty profits per consumer. When increasing the number of extended warranty contracts sold per consumer, this increases the extended warranty profits per consumer always less than linearly. When selling $k-1$ extended warranty contracts to consumers, the profit becomes $(k-1)(1-q)^{k-1}\tilde{F}(w)\tilde{E}[w-r|w>r]-[\sum_{i=1}^{k-1}(1-q)^i]F(w)w$. For any integer $k\geq 2$, an increase in k, i.e. for any purchase of multiple extended warranty contracts, there is never a proportional increase in the per-product extended warranty profits because of the exponential term $(1-q)^{k-1}$ in the above expression, from which it follows that $k>\tilde{a}(k)$.

For a price $p^*(k-1)$, a firm has no incentive to deviate to a lower base product price if there does not exist any positive integer $l\geq k$ such that $l(p^*(l-1)-c)+\tilde{a}(l)\theta\tilde{\pi}(q,\tilde{w})\geq \frac{1}{M}[(k-1)(p^*(k-1)-c)+\tilde{a}(k-1)\theta\tilde{\pi}(q,\tilde{w})]$. This leads to the condition

$$c\geq \frac{Mlp^*(l-1)-(k-1)p^*(k-1)+[M\tilde{a}(l)-\tilde{a}(k-1)]\theta\tilde{\pi}(q,\tilde{w})}{Ml-k+1}, \quad (*)$$

both for $l=k$ and for all integers $l>k$ for which $[\tilde{a}(l)-\tilde{a}(l-1)]\theta\tilde{\pi}(q,\tilde{w})>\underline{b}(l)-\underline{b}(l+1)$. Denote the highest of these values for l by $\bar{l}_{k-1}$. If $\bar{l}_{k-1}=k$, this turns to

$$c\geq p^*(k-1)+\frac{[M\tilde{a}(k)-\tilde{a}(k-1)]\theta\tilde{\pi}(q,\tilde{w})}{(M-1)k+1}.$$

To determine all equilibrium base product prices of the game one can use the following steps iteratively.

1. Start with $k=2$. This is because $Iq>c$, which ensures that all consumers buy at least one base product in equilibrium.

2. Compute $\bar{l}_{k-1}$, $\underline{b}(k)$, $\tilde{a}(k-1)$, and for each integer $l=k,..,\bar{l}_{k-1}$, compute $\underline{b}(l+1)$, $\tilde{a}(l)$.

3. For each integer $l=k,..,\bar{l}_{k-1}$, compute $p^*(l-1)$ in equation (4) and then check whether condition $(*)$ is satisfied.

4. If condition $(*)$ is satisfied, $p^*(k-1)$ is an equilibrium base product price, and $\tilde{a}(k-1)\theta\tilde{\pi}(q,\tilde{w})$ is the equilibrium industry extended warranty profit.

   (a) If for no integer $l\in\{k,..,\bar{l}_{k-1}\}$ it holds that $\underline{b}(l)\geq c-\frac{\tilde{a}(l)\theta\tilde{\pi}(q,\tilde{w})}{l}>\underline{b}(l+1)$, the algorithm stops. This is because no lower base product price can be achieved as to generate a higher demand for base products per consumer and at the same time leading to non-negative firm profits.

   (b) If there is an integer $l\in\{k,..,\bar{l}_{k-1}\}$ such that $\underline{b}(l)\geq c-\frac{\tilde{a}(l)\theta\tilde{\pi}(q,\tilde{w})}{l}>\underline{b}(l+1)$ holds, start from point 2. with $k=k+1$.

5. If condition $(*)$ is not satisfied, start from point 2. with $k = k + 1$.

When $c < q(1-q)I$, it follows immediately that in each equilibrium consumers always purchase multiple base product quantities. Qualitatively, such equilibria do not differ from other multiproduct equilibria for which this assumption does not hold.

B.2 Alternative forms of consumer behavior and naivete

Hyperbolic discounting and projection bias Consumers having time-inconsistent preferences is one possible explanation for underestimating return costs in our model. If consumers have (β, δ) preferences under hyperbolic discounting, as in for example DellaVigna and Malmendier (2004), an overestimation of their short-term patience β yields an underestimation of the opportunity costs of returning a product. A second possibility is projection bias as in Loewenstein *et al.* (2003). If consumers' return costs are low in a motivated state and high in an unmotivated state, under projection bias they underestimate the return costs they face in an unmotivated state while making a decision in a motivated state.

Risk aversion and loss aversion So far we have abstracted from risk aversion in our analysis. If consumers are risk averse, this increases their willingness to pay for an extended warranty. However, for positive firm profits to persist under this assumption in equilibrium, myopia with respect to not anticipating an extended warranty purchase before buying a base product is still required. Furthermore, several empirical studies show that pure risk aversion cannot account for the high markups of extended warranties in several consumer electronics industries. Jindal (2014) sets up a survey experiment to distinguish between risk-aversion and loss aversion of consumers with respect to extended warranty sales for washing machines. His calibrated model indicates that loss aversion is significantly more important in predicting consumer behavior than risk aversion. Related to our model, however, his experimental setup does not allow to test for return costs or any other return frictions. Huysentruyt and Read (2010) present experimental evidence that the magnitude of the feeling of loss in case of product breakdown is a more important purchase predictor for extended warranties than the attributed probability of product breakdown.

We next formalize a small model variant that incorporates reference-dependent loss aversion.[21] Assume that naive consumers are initially myopic with respect to the possibility of buying an extended warranty. For simplicity, we abstract from all return costs and assume that sophisticated consumers are all risk neutral with respect to product breakdown. At the point of sale, naive consumers form an exogenous reference point after making a purchasing

[21] See Kőszegi and Rabin (2007) for a model of loss aversion and endogenous reference points.

decision with respect to the base product. This reference point is the average risk-neutral consumption utility given a base product price p and zero warranty, $Iq-p$. Define the reference dependent utility as $U^{RD}((x,w)|qI-p) = qU(I-x-p|qI-p)+(1-q)U(w-x-p|qI-p)$. This setting is consistent with the interpretation of empirical results in Chen *et al.* (2009).[22] Due to loss aversion, consumers disproportionately suffer more from losses, i.e. product breakdown, than they benefit from gains. This increases their willingness to pay for an extended warranty contract that reduces the suffering in the breakdown period. A firm thus offers an extended warranty contract (x,w) that maximizes the difference between the reference dependent utility and the expected claim costs:

$$(x,w) \in \arg\max_{x,w} \tilde{\pi}^{RD}(q,w,x) = U^{RD}((x,w|qI-p) - (1-q)w - U^{RD}((0,0)|qI-p).$$

Whenever this profit is positive, this variant leads to a similar equilibrium as in our basic framework with observable product qualities, except for no cost frictions being present. Competition at the point of sale leads to the reference-dependent utility maximizing allocation subject to the constraint that firms make zero average profits. In terms of potential policy interventions, a minimum default extended warranty level weakly decreases the willingness to pay for a positive extended warranty contract, as it already increases the utility in the loss state. In case there are no cost frictions, the effects on overall welfare are different: any minimum warranty level is weakly beneficial for consumer surplus.[23] In case of return frictions, the usual trade-off between lowering the prices of extended warranty contracts via a minimum warranty level and increasing social costs of return applies.

Exogenous quality misperceptions Another alternative assumption to the underestimation of return costs is that consumers exogenously overpredict the failure rate of a base product, and for simplicity do not incur any cost of product return. Huysentruyt and Read (2010) interpret survey evidence as consumers to initially overpredict product failure rates. In case sellers' sales techniques lead to overpredicting product failure at the point of sale, this increases the willingness to pay for extended warranties. If naive consumers do not anticipate buying an extended warranty contract before making a base product purchase, under retail competition the main results qualitatively hold compared to our baseline case.

[22] Also see Sydnor (2010) for a discussion of the choice explanations for insurances against risks.

[23] Consumer surplus frictions could easily be incorporated using a known return cost distribution. An important normative question in an extended warranty context with loss averse consumers would be whether loss-aversion permanently affects consumer utility, or whether it can be temporarily induced by salesmen at the point of sale.