Cyclicality of Add-on Pricing

Branko Boskovic*
University of Alberta

Sacha Kapoor†
Erasmus University

Agnieszka Markiewicz‡
Erasmus University

Barry Scholnick§
University of Alberta

December 31, 2017

Abstract

We use the universe of transactions data from a nationwide retailer of household durables to study the importance of add-on goods, for the firm, and for our understanding more generally of their incentives and behavior. Specifically, we examine the pricing and markups for extended warranties, a classic example of an add-on good, and compare it with the pricing and markups for the base durable. We use the markup data to first show that the firm has much more room to adjust the extended warranty price, and barely any to adjust the price of the base durable. We show next that the extended warranty price exhibits much more sensitivity to economic activity, demand shocks, and cost shocks. Finally, we show that estimates of demand elasticities and inflation exhibit severe bias if they are based on base good prices alone. The findings imply add-on price data is critical for understanding incentives and behavior in the growing number of markets where add-on goods are sold.

JEL:
Keywords: Add-on pricing, price discrimination, durable goods, extended warranties, business cycle

1 Introduction

The rise of the add-on good is one of the most salient developments in the retail sector. The development is driven by the idea that prices for base goods act simply as a tool for drawing the customer in. Low or loss-leader prices on base goods encourage customers to incur the time or

*Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta T6G 2R6, Canada (email: boskovic@ualberta.ca).

†Erasmus University Rotterdam, Erasmus School of Economics H09-22, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands (email: kapoor@ese.eur.nl).

‡Erasmus University Rotterdam, Erasmus School of Economics H09-20, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands (email: markiewicz@ese.eur.nl).

§Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta T6G 2R6, Canada (email: barry.scholnick@ualberta.ca).

[1]We are grateful to Maarten Bosker, Casper de Vries, Robert Dur, ..., and ... for helpful comments and discussions.

travel cost of visiting a business. These sunk costs (or the sunk cost fallacy) make it costly for the customer to visit competing businesses, and allow for higher markups on (add-on) goods that are only sold at the point of sale. These markups can provide businesses with the incentives to continue to produce even though they operate in markets where competition is notoriously fierce. During this development, there has been an explosion of retail price data, and of its use for learning about the incentives and behavior of the firm. This usually includes information on base prices only, or excludes information that allows one to identify add-on goods. In this paper we make the very basic observation that add-on price critical for understanding the incentives and behavior of a firm that operates in a market where add-on goods are sold.

To do this, we conduct an econometric case study of a nationwide Canadian retailer of household durable goods (*e.g.* appliances, electronics, furnishings). The retailer has provided us with data on every one of the more than 6 million purchases that took place between 1999 and 2010. The data includes detailed information on the prices of extended warranties, a classic example of an add-on good. They make the base durable better because they insure it beyond the lifetime of the manufacturer's warranty. Moreover, the extended warranty price is almost never advertised. In our data, for example, they are advertised for 0.09% of all transactions. We take information on prices of extended warranties, as well as information on prices of the durables themselves, and show that the prices for this add-on are critical for understanding the incentives and behavior of the retailer. We do this in several steps.

We take raw data on prices and costs to show first that there is more room for the retailer to adjust extended warrranty prices. Figure 1 presents scatters of prices against costs for the base durable (top figure) and extended warranty.[1] The top figure shows prices for durables are almost perfectly explained by their costs. One dollar more in costs is associated with just over one dollar more in price. The R^2 from a regression of the base durable price on cost exceeds 95 percent. Alternatively, the extended warranty cost has no explanatory power. A one dollar increase in cost increases the extended warranty price by four cents. The R^2 is 1.3 percent. Together, the figures also imply that extended warranty markups are more than three times the size of markups on the base durable.

We show extended warranty prices are more sensitive to economic activity broadly defined. Our information includes detailed information from hundreds of stores from across the country, as well as finer details on customers, including precise information about where they live. We exploit this to estimate the impact of local unemployment rates on prices and show that a one standard deviation increase in unemployment lowers the base durable price by a little more than 1 percent, and the extended warranty price by about 11 percent. The drop in the warranty price lets the

[1] Costs for the base good are generally made up of the manufacturer's price, commissions to salespersons for base good sales, and the costs of keeping inventories. For the extended warranty, costs are generally based on repair costs, as well as commissions for extended warranty sales. We discuss the construction of costs in greater detail later in the paper.

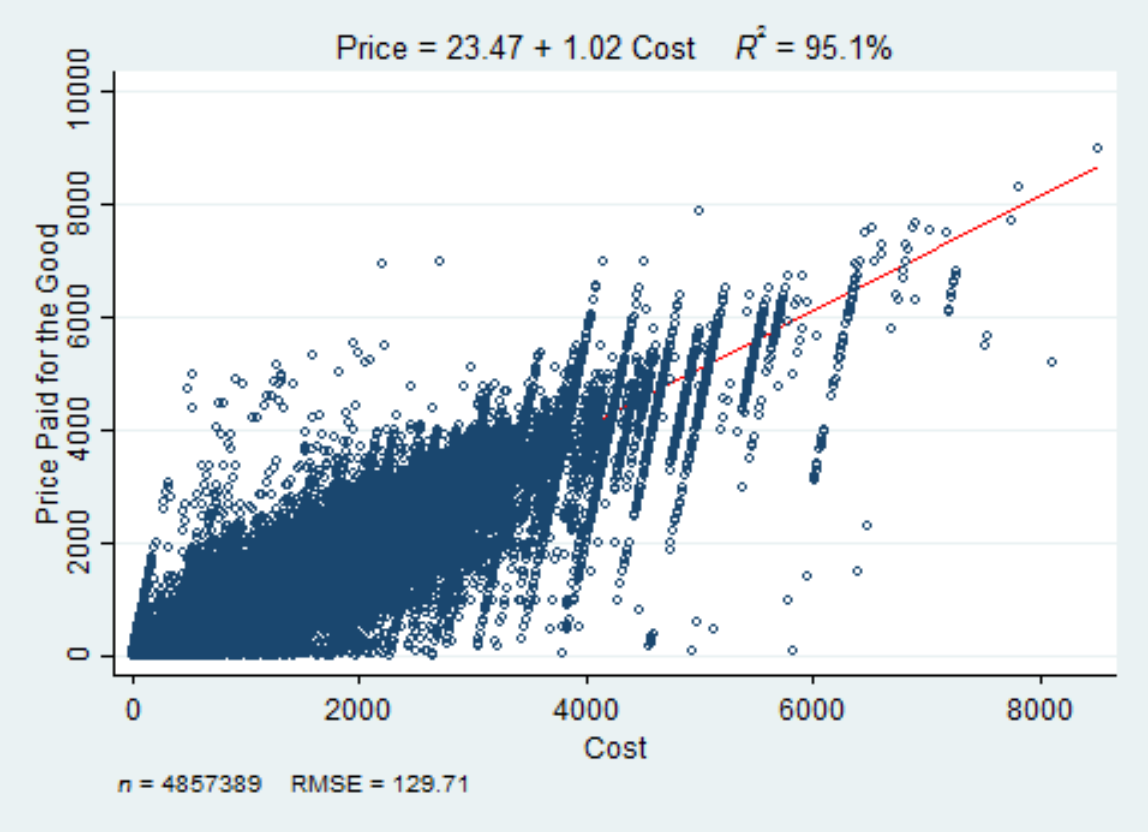

(a) Base Good

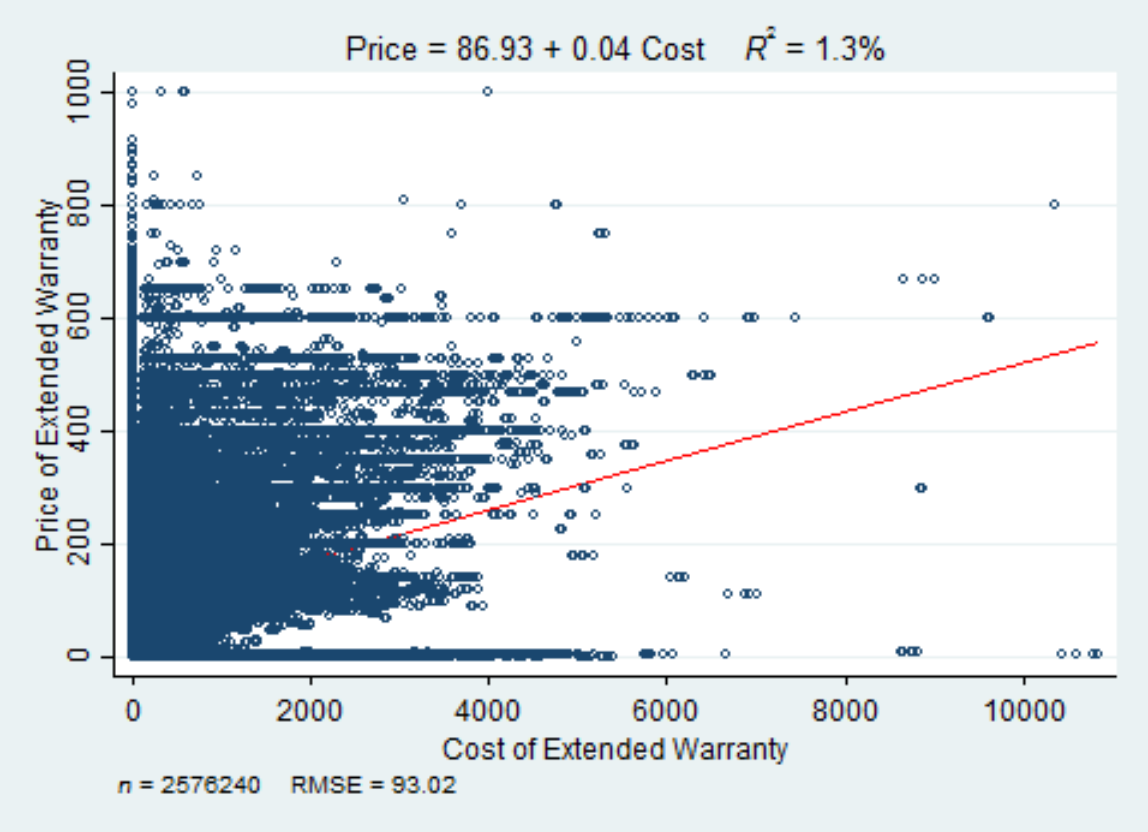

(b) Extended Warranty

Figure 1: Markups on the Base Durable and Extended Warranty.

retailer maintain, and sometimes increase, extended warranty take up. This is what one might expect if the base durable was being priced at or close to cost, while the extended warranty was being used to earn a profit.

We show extended warranty prices are also more sensitive to demand shocks. We exploit Canada's sheer size and substantial regional differences in industrial focus to explore the impact of consumer income, their risk preference (which governs the preference for extended warranties), the sunk costs that allow for markups of the extended warranty, and their propensity to make warranty claims. Consumer income has no impact on the price of the base good, but increases the price of the extended warranty to the point where there is a slight decrease in extended warranty take up. Risk preference leads to a lower price for the base good, a higher price for the extended warranty, and more take up. As is implied by the theory of add-on pricing, where sunk costs are high markups are high. Finally, it seems that the retailer is cognizant of the presence of bad types (consumers with a high claim propensity). In regions with lots of these types the retailer charges higher extended warranty prices, gets less take up, and earns the same price for the durable itself. The evidence implies the retailer's response to demand shocks mostly operates through the extended warranty price.

We also show that extended warranty prices are more sensitive to cost shocks. We first analyze the extent of pass through of cost shocks to the consumer. We show that adverse shocks are cost-absorbing for the durable, and cost-amplifying for the extended warranty. That is, prices for the durable decrease 1-for-1 with their cost. Prices for the extended warranty decrease by the more than their cost. We dig further into the data and analyze the impact of an organizational change that weakened salespersons' incentives to extend warranties. As with our other results, the organizational change had a larger relative impact on the extended warranty price. We go further and analyze the impact of costs from a couple angles. All of these analyses reinforce the conclusion that the margin of adjustment for the retailer is the extended warranty price.

We estimate own and cross-price demand elasticities for the base durable and a bundle that includes the extended warranty. We compare these with naive own-price elasticities for the base durable, elasticities that ignore the presence of the extended warranty. We show that the naive elasticity 0.16 times the elasticity which takes account of the extended warranty price.

Our data lets us demonstrate directly that the incentive to sell extended warranties is strong. Revenue and profits from the base good were respectively 4.1 billion and 156.2 million dollars. Revenues and profits from the extended warranty were respectively 230.3 and 83.6 million dollars. Thus, extended warranties generate 5 percent of total revenue, but 35 percent of total profit. This number, while certainly enough to want to sell extended warranties, is somewhat smaller than what has been suggested by several industry analysts. The Wall Street Journal, for example, reported in 2007 that at several retailers, while extended warranties generate about 4 percent of revenue, they generate somewhere between 50 and 100 percent of profit. Either way, it is worth noting

that our analysis is the first to rigorously pin down this number within the context of a systematic large-scale empirical analysis.

We use the conclusion to discuss the broader implications of our basic argument for price inflation, which is typically used to draw conclusions about firms' pricing responses to economic conditions. We show that an inflation measure which accounts for the extended warranty price is, in fact, less volatile than one that does not. We also show that the measurement error between our adjusted and naive inflation rates depends, to some extent, on the level of economic activity in the economy. This, in turn, implies biases in empirical analyses that rely on rates that ignore add-on prices. We end the paper with a discussion of the implications for electronic retail and the external validity of our study.

1.1. Related Literature. The paper builds a bridge between two strands of the economics literature, while at the same time making a substantial contribution to each. First, our paper contributes to the literature that studies add-on pricing. The literature has competing schools of thought about the importance of add-on goods. The Chicago school argues that any profit on add-on goods will translate into an equal loss on the base good. Ellison [2005] has by contrast argued that this need not be the case. That a seller of add-on goods props up the base good price, because they would rather not attract cheapskates who have a low propensity to buy the add-on. Our results favor the Ellison [2005] school of thought. Like Ellison and Ellison [2009], we show that on average the firm earns a small positive markup on the base good, and a large markup on the extended warranty. We go further and show that for gross markups, an economic downturn leads to higher markups on the extended warranty, lower markups on the base good, but a net increase in overall profit.

Second, we contribute to a literature whose interest is in the impact of demand on prices in the aggregate. Traditionally, this literature has drawn on price data from statistical agencies, or scanner price data, to speak examine whether prices adjust, by how much, and how quickly. Small or slow adjustments are taken to imply nonresponse on the part of firms. This observed nonresponse has, in turn, motivated a large number of behavioral explanations for the phenomenon. Our study emphasizes the idea that observed and actual nonresponse are not the same thing, that the response to demand is operating along a margin a statistical agency would not typically observe. We bridge these literatures because the unobserved margin in our case is the price of the add-on good.

TO BE DONE: Baxter and Landry [2012], Bils and Klenow [2004], Gagnon [2009], Esteban and Shum [2007], Carlsson and Skans [2012], Gowrisankaran and Rysman [2012], Nakamura and Steinsson [2013], Gabaix and Laibson [2006], Barsky et al. [2007] (inventory paper)

2 Context

Our study is based on the data of a nationwide Canadian retail chain that specializes in the sale of household durables (furniture, appliances, and electronics). For most goods, the chain offers

the consumer the option to extend the lifetime of the warranty beyond what the manufacturer offers. The extension covers almost 100 percent of the costs of a repair, including the costs of parts and labour, services that require a home visit by a technician, and in some cases the costs of replacement. The extension is often found at chains that specialize in the sale of durable goods. Most chains call the extension an extended warranty. We will do the same.

The context is such that the retail chain has *ex post* monopoly power over the extended warranty price. There are a couple of reasons for this. First, the price of the extended warranty is almost never posted. Our data has detailed information on the universe of transactions between January 1 1999 and December 31 2009, involving more than 6 million transactions, more than 3 million consumers, and nearly 35,000 products. During this time, the chain had 18 advertised promotions of extended warranties. The advertised promotions cover 6700 transactions, amounting to less than one tenth of one percent (0.09% to be precise) of the total number of transactions. To learn the price of the extended warranty, the consumer must speak with a salesperson at the store.[2]

Second, once the consumer has visited one of the 200 or so stores in the chain, it becomes costly for them to visit the store of competing chain. The stores are usually located in stand-alone buildings, in regions with sprawl (where the consumers usually needs a car to visit a store), and in places that are somewhat isolated from competing retailers. This, and the enormous land mass of Canada, implies the consumer would have to travel far to learn the extended warranty prices at competing retailers. These sunk travel and time costs, together with the hidden nature of the price, allow for markups over the costs of extending the additional warranty.[3]

The commission structure at retail chain reinforces the notion that the chain has market power over the extended warranty price. The chain pays salespersons commissions for the sales of base goods and extended warranties. The commission on the extended warranty is 15 percent,[4], whereas the commission on the base good is 4 percent (on average, depending on the product).

The chain gives stores and salespersons substantial discretion to set the extended warranty price. Their discretion, along with fact that the price is hidden will define the value of rich data at level of the transaction. It is unlikely that data on list prices, commonly used in the construction of consumer price indices, will include the prices of add ons like extended warranties. By definition this is information is difficult to find. It is even more unlikely that the list price will reflect the adjustments the salesperson makes in order to sell either the base good or the warranty.

[2]Customers can interact with salesperson in a couple of different ways. The salesperson can help the customer settle on a good. They can help the customer process the good after they have settled on what they want. It thus difficult to know when precisely the salesperson makes the offer of the extended warranty. (Jindal, 2015) has show that it is difficult to distinguish which is which with most data sets on extended warranties. Later we will see that our data lets us speak to the possibility that the salesperson is using the extended warranty to sell the base good.

[3]The monopoly power is ex post because the consumer paid this travel and time cost to visit the retailer in the first place. *Ex ante*, the consumer can avoid the cost by not visiting the store at all. See Ellison (2005) and Ellison and Ellison (2009) for more details.

[4]The commission was 15 percent for almost the entirety of our sample, up until May of 2009 when it was reduced to 10 percent. We will exploit the change in the empirical analysis that is to come.

The context warrants a few additional comments. First, the retail chain gives salespersons less discretion over the base good price. They have less discretion because the chain guarantees that they meet or beat the prices other retailers are asking for the base good. Presumably, the chain is compelled to do this by the highly competitive nature of retail markets in Canada.[5] Salespersons also have less discretion because the chain guarantees that their online price for the base good will exactly equal their brick-and-mortar price.[6]

Second, the same add on is applied to essentially every good. This facilitates a large-scale analysis of the cyclicality of the price of particular add-on. A large-scale analysis is more difficult in other contexts because the same good can have several different add ons, or because different goods can have different add-ons.

Third, to prevent resale and arbitrage, and thus maintains its market power over the extended warranty, the retailer maintains an extraordinary database on every one of its customers. The extraordinary nature of this database should become clear in the next subsection and the remainder of the paper.

3 Baseline Analysis

3.1. Data. The data includes the price paid for the durable, whether an extended warranty was purchased, the price paid for the warranty, and the cost of servicing claims that were made under the extended warranty. Summary statistics for these outcomes are found in Table 1. On average, the consumer pays 611 (Canadian) dollars for the durable, extends the warranty about 37 percent of the time, and pays about 89 dollars to do so. The expected cost of servicing an extended warranty claim is less than 24 dollars.

The last two columns of Table 1 present the markups over the expected service cost of the extended warranty. Column (5) presents the absolute mark up, the difference between the price and expected service cost of the extended warranty. Column (6) presents the relative markup, the absolute markup over the extended warranty price.

The last two columns imply the room for adjusting the extended warranty price is substantial. The absolute markup is just over 65 dollars ($p < 0.01$), amounting to just under 75 percent of the extended warranty price.

Figure 2 depicts the relationship between the expected markup on the extended warranty and the price of the base good. Each point corresponds to an ordered pair where the first entry is the base good price, and the second is the expected markup on the extended warranty. A more detailed breakdown, by the more than 140 product categories, is found at the end of the paper in Table 22.

[5]The four-firm concentration ratio in the industry is less than 35 percent. Economists would take this as an indication that the market is highly competitive.

[6]We will discuss the role of the online store in the conclusion. A more elaboration exploration is outside the scope of the point we want to make in this paper.

Figure 2 has several salient features. First, markups are large in general. Second, consumers never make claims on inexpensive goods, even though they might extend the warranty on these goods. This is implied by the 100 percent markup over the extended warranty price. Third, markups are lower for more expensive goods, sometimes falling below 20 percent.

In addition to prices, take up, and costs, the data includes the purchase date, the store where the purchase was made, and the home address of the consumer. This lets us link outcomes to economic activity around the purchase date, as well as the neighbourhood of the consumer.

Our baseline measure of economic activity is the unemployment rate. In Canada, the government tabulates the (seasonally-adjusted) monthly unemployment rate, for each of 60 predefined regions. These are referred to as Employment Insuranec (EI) regions. Each contains several cities, and is almost always smaller than any one province.[7] Our regression analysis will thus make use of the unemployment rate that prevails in the month when the consumer purchased the good, and in the employment insurance region of a customer's residence. Note that we have left a discussion of what the unemployment measures for Section **??**.

3.1.1. Unemployment and the Add on over Space. Figure 3a presents the spatial distribution of the unemployment rate. Figures 3b and 3c present the spatial distributions of the prices for the base good and extended warranty. Figure 3d does the same for take-up of the extended warranty.

The figures are generated as follows. We first grouped consumers by the first three characters of their postal code.[8] For each group, we averaged the unemployment rate over the 130 plus months in our sample. We averaged prices and take-up over the 4000 plus days in our sample. We weighted prices by the population share of the group.[9]

The raw data suggests prices and take up run opposite to the unemployment rate. In the northern territories, northern Manitoba, parts of Quebec and Eastern Canada, the unemployment rate hovers between 19 and 26 percent. Consumers in these areas pay (relatively) low prices for the base good and extended warranty. There is less take up of the extended warranty. In Alberta and southern Ontario (particularly around Toronto), the unemployment rate hovers around 4.5 to 7.3 percent. Consumers in these areas pay higher prices. There is more take up of the extended warranty.

3.1.2. Unemployment and the Add on over Time. Figure 4 shows how our data varies over time. Unlike Figure 3a, it presents the quantity of the base good, as well as the price (the dashed

[7]The regional unemployment is used to determine benefits an unemployed worker can receive from the employment insurance program.

[8]Postal codes in Canada have 6 characters. The first three characters refer to the forward sortation area. The first defines the province or a city in cases where the city has a large population. The second indicates whether the area is urban or rural. The third points to a specific rural region, city of medium size, or to a segment of a large metropolitan area. In all, there are about 1600 of these areas. In what follows, to keep things simple, we refer to these geographic identifiers as the postal code.

[9]See Figure 10 for the spatial distributions of the unweighted prices.

red lines in Figures 4a and 4b). Figures 4c and 4d depict the extended warranty price and the probability that one is sold (again, the dashed red line). Each graph compares the evolution of the outcome with that of the unemployment rate (the solid blue line).

Figure 4 shows that our sample includes two periods of sharp recessionary pressure. Initially, the unemployment rate exceeds 11 percent. It declines sharply and then gradually until the 2008 crisis starts. It moves from just under 7 to just over 9 percent thereafter.

Figure 4 also implies prices for the base good and extended warranty (Figures 4a and 4c) run somewhat opposite to the unemployment rate. Prices go up when the unemployment goes down. At the same time, more of the base good (Figure 4b) is sold. Surprisingly, low unemployment is associated with less take up of the extended warranty (Figure 4d).

3.2. Baseline Specification and Estimation. We investigate whether the patterns in the raw data hold up against further scrutiny. We estimate specifications which assume y_{cgsd}, for customer c, base good g, in store s on calendar date d, is generated according to

$$y_{cgsd} = \beta U_{er(c),m(d)} + postal(c) + m(g) + p(g) + s + d + \varepsilon_{cgsd}.$$

y_{cpsd} is either the price the customer pays or a binary variable that indicates whether they extend the warranty, and $U_{er(c),m(d)}$ is the unemployment rate in the employment region where the customer lives in the calendar month $m(d)$. $postal(c)$ are fixed effects for the (first three characters of the) postal code where the customer lives. Figure 3 implies these should correlate with the unemployment rate that tends to prevail locally. $m(g)$ is a fixed effect for the manufacturer of good g. $p(g)$ is a fixed effect for the chain's in-house categorization of products (see in Table 22). s is a fixed effects for the store. d is a fixed effect for the calendar date. Figure 4 implies the date-specific differences should correlate with the unemployment rate that prevails at that time. Importantly, note that as controls are added to the specification we get closer and closer to isolating the response of the firm to economic activity.

We will include $m(g)$ and $p(g)$ in some specifications because prices and take up should vary with the product and manufacturer (they will in particular pick up vertical preferences), and to show that they have little impact on the estimate of interest. We will exclude these variables from our preferred specifications. Because the product and manufacturer are choices, they can easily be interpreted as bad controls. Note that we will also experiment with richer specifications for time fixed effects, including some specifications that include the product $p(g) \times s \times d$.

Our preferred specifications will also exclude fixed effects for the store. While it makes little difference either way, store fixed effects are a deterministic function of important unobservables such as the current and future prospects of various regions. Presumably, the retail chain takes these unobservables into account when deciding where they should open a store. The deterministic nature of the relationship can bias the estimates of interest. A second advantage of excluding store

fixed effects is that some of our transactions are registered at one store when in fact they took place at another. This tends happen when the base good has sold out.[10] In this latter regard, by excluding store fixed effects, we are in effect treating the head office and stores a single decision maker.

Estimates of the baseline specification are found in Table 2. The dependent variables are (in order) the price of the base good, the price of the extended warranty, and the probability that it is sold (purchased). Moving left to right depicts how the point estimate differs depending on the controls in the regression. The first column presents the estimate from a regression without controls. The second from a regression with controls for the date and FSA of the consumer. The third controls for the product category. The fourth for the manufacturer. Note that the last column depicts our preferred specifications (for these outcomes) for the remainder of the paper.[11]

Columns (1) and (2) of the two bottom panels suggests nonneglible regional differences in unemployment and the price and take up of the extended warranty. One interpretation of the comparisons is that they reflect regional differences in the value consumers derive from a little extra income. A lot of these consumers reside in regions with high unemployment. They also have a propensity for purchasing lower quality goods. Regional differences of this sort can justify the (statistically significant) underestimates in the regression coefficients for the extended warranty price (0.25 versus 0.34) and take up probability (-0.008 versus 0.047).[12]

Columns (1) and (2) of the top panel suggests that there are also regional differences in the price of the base good. But the differences are inconsistent with the interpretation that the regional fixed effects reflect differences in the extra value derived from extra income. Unlike the extended warranty, the exclusion of the regional fixed effects leads to overestimates of the impact of the unemployment rate on the base good price. The coefficient on the unemployment decreases from -0.023 to -0.043. The difference, moreover, is negligible in the statistical sense.

Columns (2) and (3) suggests more pronounced differences emerge across time. Excluding fixed effects for the calendar date yields large overestimates of the impact on the price and take up of the extended warranty. The point estimates for take up and price are -0.19 and 0.008 when the fixed effects are included. They are 0.33 and 0.047 when they are not. One explanation for the overestimate is that it reflects the propensity of richer consumers to continuing to buy even when times are bad. These consumers derive less value from extra income, and are more likely to purchase the higher quality good. This explanation is, again, a bit inconsistent with the change in

[10] We deal with this measurement issue somewhat incidentally in Table 23.

[11] We should include time and province fixed effects, or product time fixed effects because of the introduction of new models.

[12] Our tests for statistical significance were based on seemingly unrelated regression. We drew a random sample of 500,000 observations (to reduce the computational burden). We stacked the data so that the X matrix was a block diagonal matrix, where the upper left block had only the unemployment rate, and the bottom right block had the unemployment rate as well as fixed effects for the consumer's neighbourhood. Using a simplified specification, we estimated a coefficient for the unemployment rate in the upper left block and one for the unemployment rate in the lower left block. We tested and rejected the equality of these coefficients.

the estimate for the base good.

The remaining columns consider how the estimates change when we add fixed effects for the product, manufacturer, and store. As noted earlier, we prefer not to use these specifications because one can interpret these fixed effects as bad controls. With that in mind, we feel good about the fact that they have little impact on the point estimates we care about.

Ultimately, Column (3) has the estimates from our preferred specification. The estimates are consistent with the retailer using the add on to deal with downturns in economic activity. A one standard deviation increase in the unemployment rate decreases the price of the extended warranty by 10.8 percent (middle panel), and increases take up by 1.2 percent (bottom panel). The estimates are also consistent with the base good being less useful as an instrument for dealing with downturns. The unemployment rate decreases the base good price by about 1.6 percent (top panel). In the next section we show the impact on quantities for the base good are statistically negligible.

In all the estimates, and the sign flips in particular, are consistent with what is presumed in Ellison's 2005 add-on pricing paper. His model assumes that consumers who are less sensitive to price differences across stores are also less sensitive to price differences across qualities. If this is true, then a smalll increase in the price of the add on will not deter high value consumers from buying the add-on good. Our results show that (without) controls for the firm charges a higher price for extended warranty. This is consistent with the firm being able to more easily markup add-on goods to high value consumers.

3.3. Unemployment and the Base Good Quantity. Ideally, we would know all the potential customers of the store, including the ones who purchase nothing. Knowing the potential customers would allow us to apply our baseline specification to the decision to buy at least one thing. Because such detail is impossible for us to obtain, we will use aggregate quantities to assess the impact of the unemployment rate.

Specifically, we will assume that the total base good quantity $base_{egm}$ purchased by consumers in employment region e in calendar month m ($m \in \{1999m1, 1999m2, ...\}$) is generated in accordance with

$$base_{egm} = \beta U_{em} + e + m(g) + p(g) + m + \varepsilon_{egm}$$

where is the natural logarithm of the base good quantity, e is a fixed effect for the employment region, m is a fixed effect for the calendar month, and everything else is as before. The identifying assumptions and interpretations of the fixed effects are similar to those of our specification for prices and take up. Estimates are found in Table 4.

Unsurprisingly, the chain sells less of the base good when unemployment is high. In the absence of controls, a one standard deviation increase in the unemployment leads the chain to sell 46 percent fewer units of the base good. With controls for the employment region and calendar date,

this number drops to about 7 percent. These results suggest that the unemployment rate may have an impact on the preference to visit the retail chain.

3.4. Causality. We explore the assumptions that allow for a causal interpretation. One issue relates to the influence, if any, the retail chain has over local unemployment rates. Our view is that this is implausible. While the retail chain is large nationally, it is small locally. A given store employs a small fraction of the labour force in the region. Further to this point, our unit of analysis helps us deal with any influence the retailer might have. The unit of analysis - the transaction - is highly atomistic. It is unlikely that the outcome of any one transaction will have more than an infinitesimal impact on the unemployment rate that prevails.

3.4.1. Independent Markets. Another issue relates to whether economic activity spills over from one province to another, and whether this spillover affects prices, quantities, and take up. The context we study has two major advantages in this regard. The first is that economic activity runs north to south rather than east to west. The biggest trading partner for provinces is not usually another province. It is the United States. Put another way, most of the variation in our unemployment rates is being generated by expansions and contractions in the US economy. The second advantage is that the employment insurance regions are demarcated in part based on the idea that they make up separate markets. These advantages limit the need to account for spillovers in empirical specification.

3.4.2. Loyalty of the Consumer. A third issue relates to the loyalty of the consumer to a particular manufacturer. A loyal consumer will always buy from the same manufacturer, even though they might downgrade the quality of the base good, perhaps by purchasing an older model. The presence of these consumers can bias downwards estimates of our price regressions, particularly if downgrading becomes more common in high unemployment regimes.

The most obvious way to deal with consumer loyalty is the inclusion of fixed effects for the model of the base good in our specifications. We excluded model fixed effects for a couple of reasons. First, and most importantly, there are 35,000 models in our data. Accordingly, fixed effects would weigh heavily on the time it takes to estimate our specifications. Second, like the product and manufacturer, the model is a bad control. It is a choice variable for the consumer, and likely a deterministic of economic activity.

We approached consumer loyalty from a different angle. Specifically, we examined how the estimated impact on the base good price varies with the age of the model. We estimated the base regression for goods that aged 15 days or less, 30 or less, 60, and so on. If the unemployment rate brings about switches to older models, and newer models have higher prices, then the unemployment rate should have a positive or less negative effect on the newest models. The estimates in Table 5

show that the estimate is closest to zero for the youngest models.[13]

The estimates in Table 5 also implies that switchers ultimately have no bearing on the average effect of the unemployment rate. Moving left to right we see the estimated relationship quickly approaches what we obtain when we use the sample as a whole.

3.4.3. Selection of Purchasers of Extended Warranty. A fourth issue concerns the applicability of estimates from regressions where the dependent variable is the extended warranty price. These regressions use information from consumers who purchased the extended warranty. It is unclear that the estimates would apply to consumers who bought the base good only. One way to tackle this issue is to try to adjust our estimates for the process that generates purchases of the extended warranty. An adjustment of this sort is problematic because the decision to purchase the extended warranty will depend on the extended warranty price. Thus, by making such an adjustment, we would be substituting one type of bias for another (selection for reverse causality).

Instead of an explicit adjustment for the process that generates extended warranty purchases, we cut the data in a way that lets us learn how the sample of purchasers differs from the sample of nonpurchasers. We examined the impact of the unemployment rate on the base good price for consumers who only bought the base good, and separately for consumers who bought the base good and the extended warranty. The estimates are found in Table 3. The estimates in the top panel use the sample of extended warranty purchasers. The bottom uses the sample who purchase the base good only.

The estimates in Table 3 imply that the unemployment has essentially the same impact on the base good price, regardless of whether the customer purchased an extended warranty. Once we account for the region the consumer lives in, and the date the purchase made, the point estimate for consumers who buy both us roughly the same (economically) as the point estimate for consumers who buy the extended warranty only.

3.4.4. Ex Post Price Discrimination. Table 3 also lets us study the role of forms of other forms of price discrimination (other than add-on pricing). The retailer can engage in other forms of price discrimination because they have market power at the point of sale, and can prevent the consumer from transferring extended warranties to other consumers.[14] They can be active about it, by haggling for the purposes of learning the value the consumer places on the base good and extended warranty. They can be passive about it, using observables to stereotype the consumer, and by letting the price depend on the stereotype. Alternatively, they could offer the base good for a lower price as long as the consumer buys the extended warranty as well.

[13] We tested equality of coefficients across the first two columns, and then from the second to third. The tests implied strongly that we should reject the coefficients are the same.

[14] In all but one of the ten Canadian provinces the retailer can stop customers from transferring extended warranties amongst themselves. In the province of Quebec they cannot.

Table 3 allow us to rule out the last of these mechanisms, namely whether the retailer sells the base good at a lower price when the warranty is extended. There are limits on what our data lets us say about the roles of the other forms. To say something concrete, we would need information on what the salesperson knew before speaking with the consumer, what they learned from speaking with the consumer, as well as what they spoke about. The retail chain do not track this sort of information in stores.

3.4.5. Demand and Supply. We have so far remained agnostic on whether changes in the unemployment rate measure a shock to demand or to costs. It is critical to know this if we wish to understand the behavior of the firm. It is, in particular, important for understanding whether the firm is adjusting prices in accordance with shifts or rotations in demand, or whether the unemployment rate is simply the firm more willing to accept a lower price for each extended warranty they sell (i.e. an outward shift in supply). The next sections will consider the extent to which the results reflect an indirect response to changes in demand or a direct response to the unemployment rate itself. With this in mind, and knowing that we will properly estimate the price elasticity of demand in the last sections of the paper, we use the baseline results to compute back-of-the-envelope (equilibrium) elasticities.

The estimates in Column 1 of Table 4 and the top panel of Table 2 imply that the equilibrium elasticity of the base good price to the base good quantity is 33.[15] The first column (bottom two panels) of Table 2 and 4 imply the equilibrium elasticity of take up with respect to price is $\frac{-0.009/0.395}{0.126} = -0.18$. The last column implies an elasticity of -0.23. This is in the range of what one would expect if the retail chain had power over the extended warranty price.

4 Demand

We investigate the retailer's adjustments to changes in demand. We analyze the impacts on prices of factors that might compel the consumer to extend the warranty, including their income, their preference for risk, the time and travel costs of physically visiting a store, and how well they understand the value of the extended warranty. We also investigate whether prices adjust depending on the claim propensity of local populations. We do this in order to understand whether the retailer takes account of relevant populations of good (low claim propensity) and bad types.

4.1. Income. Our context emits a natural proxy for the income of the consumer. In Canada, there are 10 provinces. For a variety reasons, including the enormity of the land mass and substantial geographic variation in natural resource endowments, the provinces differ considerably in the activ-

[15] We use Column 1 because it yield an elasticity that incorporates the impact of the unemployment rate on the sorting decisions of the consumer. The elasticity was calculated at a one standard deviation change in the unemployment rate $-0.462/-0.014$.

ities used to generate income. The province of Alberta, for example, generates income through the production and export of oil and natural gas. The province of Ontario, on the other hand, generates income through the production and export of manufactured goods. These provincial differences in industrial activity facilitate proxies for the income of the consumer.

Our specific proxy makes use of the endowment of oil in Alberta and the world spot price of oil. For Albertans, higher spot prices translate into higher gas prices and more income. For residents of other provinces, higher spot prices translate into higher gas prices only. The difference gives a proxy for the income of Albertans.

We recognize the imperfectness of our proxy. Changes in the spot price of oil could, in principle, have an direct impact on the people who work at the chain's stores. We have a couple things to say about this. First, having data at the transactions level helps with this. The oil price will tend impact behavior over time horizon that is much longer than that at which a transaction occurs. Second, in the next subsection we will show that we have a really pure measure of changes in demand.

Figure 5 depicts the evolution of the world spot price of oil over the duration of our sample (the blue line). Figures 5a and 5b compares it to the evolutions of the price and take up of the extended warranty. The dotted red lines describes the evolution in Alberta. The dashed-dotted green lines describes the evolution in the rest of Canada.

Figure 5a shows the Alberta and Rest-of-Canada prices for the extended warranty diverge right when the spot oil price started increasing steadily. The divergence starts around quarter 30. Their differential comovement thereafter supports the notion that the spot oil price yields a proxy for consumer income.

Figure 5b implies that propensity to take up the extended warranty is unrelated to where the consumer lives. Consumers from Alberta are just as likely to take up the extended warranty as consumers from the rest of Canada. This propensity, moreover, is stable over time, and seemingly unaffected by the increase in the Alberta price of the extended warranty. Put another way, consumers from Alberta take up the extended warranty at the same rate despite the differential rise in price. In our next steps we investigate how well the raw patterns fare against more rigorous empirical scrutiny.

We estimate

$$y_{cgsd} = \beta OilPrice_d \times Alberta_c + postal(c) + m(g) + p(g) + s + d + \varepsilon_{cgsd}$$

where $OilPrice_d$ is the crude oil price on date d,[16] and $Alberta_c$ is a binary variable that indicates whether consumer c lives in Alberta. The remaining variables are defined as before. Note that the $\mathbf{X}_c$ and γ_d encapsulate intercept differences generated by $OilPrice_d$ and $Alberta_c$. Both variables

[16]The raw data for our crude oil prices is in dollars per barrel and comes from the West Texas Intermediary - Cushing, Oklahoma.

are thus exogenous by definition. Estimates of the specification are found in the top panel of Table 6.

The impact of the spot oil price runs opposite to the impact of the unemployment rate. The spot price lowers the price of the extended warranty. Take up decreases by a small amount. Specifically, a one standard deviation increase in the spot oil lowers the price Albertans pay for the extended warranty by 16 percent. The same increase in the spot oil price decreases take up by 0.5 percentage points. Both estimates are nonzero from a statistical perspective.

Table 6 implies an equilibrium elasticity (of take up to price) that is bit smaller than the elasticity the unemployment rate implies. Here the implied elasticity is approximately -0.08. There it was between -0.23 and -0.18. The implied elasticity here is consistent with Figure 5, which shows that the propensity for extending the warranty is unchanged in periods when the spot price of oil is rising.

The estimate in the first column (of Table 6) implies that the spot price of oil has no impact on the base price of the durable good. The point estimate is quite a bit smaller than the impact of the unemployment rate. As with the unemployment rate, this is consistent with the retailer using the base good to draw in customers for the purposes of extended the warranty.

Table 7 investigates the impact of income on the quantities of the base good. The left panel examines the differential impact of the oil price on the number of consumers who visit the firm. The right panel examines the differential impact on the number of purchases per consumer. For both panels moving left to right shows how the coefficients change with consumers for the region and quarter-year combination. WE WILL REVISIT THIS.

Table 7 implies that consumers buy more of the base good when they have more income, but that the impact is economically small. Column 3 shows that a 47 percent increase in the oil price, which is equivalent to a one standard deviation increase, leads to a 0.0002 percent increase ($p < 0.01$) in the number of consumers who purchase at least one good. Column 3 shows that the oil price has no statistically significant impact on the number of purchases per consumer.

4.2. Risk and Uncertainty. We also use the spot price of oil in Alberta to build a proxy of a consumer preference for insurance. The logic behind our proxy is as follows. If the spot price measures current income, then futures price measures future income, and the variance of futures prices measures uncertainty about future income. Because in general risk preference will depend on income uncertainty, the variance of futures prices proxies for risk preference, and thus for a consumer preference for insurance.

Our empirical specification is basically the same as the specification as the last section. However, instead of interacting the Alberta dummy with the spot price of oil, we now interact it with the standard deviation of the future prices of oil.[17] Estimates of the specification are found in the

[17]The standard deviation is taken over the prices x months ahead, where $x \in \{1, 2, 3, 4, 5, 6, 12, 18, 24, 36, 48, 60\}$.

middle panel of Table 6. For the purposes of robustness, the bottom panel shows that neither the estimate for income or the insurance preference changes much when both interactions are included in the same specification. Accordingly, we will focus our discussion on the middle panel.

Income uncertainty leads to more demand for the extended warranty. A one standard deviation increase in the standard deviation of oil price futures increases the extended warranty price by 2.9 percent (Column 2). It increases take up by 0.004 percentage points (Column 3). These results are consistent with the assumption that the variance of oil price futures measures the risk preference of the consumer.

Income uncertainty leads to less demand for the base good. A one standard deviation increase in the standard deviation of oil price futures decreases the base price of the durable by 1.3 percent. It decreases quantities by X percent. We view this as further validation of our assumption that the variance of oil price futures measures the risk preference of the consumer. For most goods, other than insurance goods, demand should decrease when the consumer becomes more risk averse.[18]

4.3. Sunk Costs. Sunk costs have a prominent role in the Ellison model of add-on pricing. There is a cost to the consumer to visiting the store, and to browsing through items after having decided to visit. Once the consumer reaches the point of sale, they cannot recover these costs. Thus, at this point of the interaction with the retail chain, for the consumer to switch to another item or to another store, they must pay this sunk cost one more time. The switching cost, or the prospect of the switching cost, allows for markups of add on goods.[19]

We investigate the role of such costs in the pricing and markups of the extended warranty. To investigate this, we calculate the minimum distance between the home address of the customer and the closest store. We consider the impact of this minimum distance as well its interaction with gas prices (as measured by the spot price of oil). The estimates are found in Table 8. The first three columns use data from all consumers. The last three uses data from consumers outside Alberta. We did this because, as already noted, for these consumers changes in oil prices also imply changes in income.

The estimates in the top row show that distant customers pay lower prices, and take up the extended warranty less often. These baseline estimates are consistent with several ideas, including the idea that consumers in more remote areas have less money to spend.

The estimates in the bottom row are consistent with sunk costs affecting markups on add-on goods. Distant consumer who face high gas prices pay more for the extended warranty.

[18]In this paper we will not study all the primitives that could potentially drive the extended warranty decision, as a full blown analysis would a require a full paper by itself. We simply note that others have studied the issue in greater detail, usually focusing on the roles of loss aversion and overweighting of the probability of a breakdown in the durable, in addition to simple risk preference approaches like the one here. See Jindal (2015) for a more detailed discussion.

[19]The issue of sunk costs is of particular relevance in our case, because the cost learning prices elsewhere should vary with the value of time spent shopping, and because the value of shopping time can vary with the unemployment rate (cite).

4.4. Sophisticated Consumers and Bad Types. We investigate whether the benefits to extended warranties come from the purchases of sophisticated or naive consumers. We do this in several ways. We analyze the proclivitiies of repeat customers and of customers who have made warranty claims in the past. We then investigate whether the chain adjusts prices depending on the average claim propensity of the average consumer in a market. .

Table 12 shows correlations between our main outcomes and a dummy that indicates whether the current visit is a repeat visit. The top panel uses a sample that includes all customers. The bottom uses a sample that only includes customers who visit more than once. The dummy in the bottom panel therefore indicates whether the current visit is one of the customer's later visits. The table shows that regardless of the specification we use, repeat customers always pay lower prices and have less of a proclivity to take up the extended warranty. This is what one would expect for customers with some level of sophistication.

Table 9 shows correlations between our main outcomes and the number of claims made prior to this visit. The table shows that these customers buy more expensive goods are also more likely to take up the extended warranty. At the same time, these customers pay the same price for the extended warranty. This result and the results for repeat customers implies that there are two types of sophisticates. The first type avoids high prices and extended warranties. The second type buys extended warranties together with more expensive goods. The latter are types the retail chain itself wants to avoid.[20]

In particular, we investigate whether the impact of the unemployment rate differs depending on the propensity of a consumer from the FSA of the customer to make a claim on the extended warranty (Table 10). Looking at the averages is because it is, in theory, what the retail chain will know about the FSA. A comparison of these regressions with regressions based on individual types will be informative about the supply and demand sides with regards to problems of asymmetric information.

Table 10 shows that there is less of reduction in the extended warranty price in neighbourhoods where consumers have a higher propensity to make claims (Column 2). In these neighbourhoods, there is a smaller increase in the take up of the extended warranty (Column 3). This is what one would expect if the retail chain was cognizant of the type of consumer they face (and if they are price discriminating), and was setting prices so as to deter them from purchasing the extended warranty.[21]

Column 1 shows the interaction has an negligible impact on the price of the base good. This is

[20] It is worth noting that in our sample the total cost of warranty claims is generated by less than 5 percent of all customers. These costs can either reflect adverse selection, where customers who are prone to breaking things purchase the extended warranty, or moral hazard, where having the warranty induces the customer to take less care of the product. While it would be interesting to use our data to distinguish adverse selection from moral hazard, it is outside the scope of the main ideas we study.

[21] Are these guys more or less sensitive to price changes? These guys usually buy more expensive goods. They are the types who you don't want to buy the extended warranty, but who you do want to visit your store.

sensible if claim propensities are not all that relevant for the purchases of the base good.

5 Supply

We study the retailer's adjustments to cost shocks. We first examine whether the retailer passes cost shocks through to the consumer. We then investigate the impact of an organizational change that weakens salespersons' incentives to the extend the warranty. Finally, we investigate the extent that price changes reflect the responses of the head office of the retail chain.

5.1. Pass Through of Costs. To analyze pass through, we estimate the impact of economic activity on the cost of the base good. The estimates allow for conclusions about the extent to which changes in the base good price reflect changes in the base good cost. We do the same for the price and cost of the extended warranty.

We have information on the prices at which the chain transfers base goods to franchise stores. The transfer price is the sum of the price the chain pays to the manufacturer and the cost of storing the base good in one of the chain's warehouses. We use these prices to build costs for all stores. It is reasonable to do this because the retail chain guarantees franchisees (in writing) that the price they pay will be the same as any other store in the chain. There is no additional markup on the price the manufacturer charges and the cost of holding inventory.

Specifically, we have information on the transfer price of the more than 9000 models that were sold at franchise stores. After assigning these transfer prices to the same models at corporate stores, we obtained basic costs for more than 4.5 million of the 6 million goods sold sold in our sample.

We adjusted these basic costs to reflect the commissions that are paid to workers. In particular, we multiplied the base price by 4 percent and add this to the transfer price of the base good. In the case of franchise stores, costs are also adjusted to reflect the franchising fees they pay to the retail chain. These fees amount to 2.5 percent of gross revenue. Costs are also adjusted for the amount franchises are expected to spend on local advertising. The spending requirement amounts to 4 percent of gross revenue.[22]

We estimated the impact of the unemployment on the cost of the base good. We present estimates just from the sample where we did not impute costs (in the appendix we will show the estimates from the full sample). We did this because we wanted to show that it matters little if we use the full sample or the same with imputed costs.

The estimates are found in Table 14. The impact on the base good is found in the top panel. The impact on the base good price (for sample where we have costs) is found in the bottom panel. The left panel (the first three columns) depicts the impact on the natural logarithm of the dependent

[22]Most other costs are fixed. These fixed costs include the costs of advertising and marketing, the costs of renting or leasing the space where the firm sells the goods, and base pay for workers.

variable. The right panel (the last three columns) depicts the impact on the level of the dependent variable.

Our discussion will focus on the estimates in Columns 3 and 6, as these include our preferred specifications. For the estimates in the remaining columns, we will simply note that they are support the idea that the fixed effects capture customer sorting effects (based on income) that we discussed earlier.

The price and cost of the base good move in lockstep. Using Column 6 (top panel) as a baseline, the cost falls by \$11.38 when the unemployment rate increases by one standard deviation. The price falls by \$11.48. The comovement of price and cost are consistent with highly competitive market for the base good. Similar conclusions when the dependent variable is in natural logarithms (Column 3).[23]

For the purposes of comparison, Table 15 shows the impact of the unemployment rate on the costs of the warranty. The warranty cost is comprised of the average claim cost and commission paid to sales staff.[24] In the case of franchise stores, like the base good, it also includes the franchise fee and the spending requirement for local advertising.[25] The first three columns present the impact of the unemployment rate on the cost of the extended warranty. The last three present the impact on the price of the extended warranty. The table does not include a panel where the dependent variable is in logarithms because the warranty cost is often 0 (this is not true if we include the commission cost).

Unlike the base good, the extended warranty price decreases by more than the cost. Focusing on Columns 3 and 6, the warranty cost decreases by just over 5 dollars. The extended warranty price decreases by almost 9 dollars. Unlike the base good, cost shocks get ampiflied in the extended warranty price.

5.2. Commissions for Adding On. The unemployment rate can shift the effectiveness of commissions as incentive device. It can affectt the perceived job security of workers. If unemployment is high, workers might put in more effort because they worry more about their jobs. In our case, this translates into selling more extended warranties, and thus into a lower marginal cost when unemployment is high. The assumption is supported by recent evidence, which has shown that workers work harder when the unemployment rate is high (cite lazear and shaw).

Our context lets us investigate the extent to which is the case. In May of 2009, because of the residual effects of the Great Recession in the United States, the retail chain cut the commission on

[23]This not wholly consistent with a couple of recent papers that show the prices of inputs are sticky (Goldberg, and AER paper using the Swedish data).

[24]Note that this expected cost depends on the joint probability of a breakdown and a warranty claim when the good breaks down.

[25]One disadvantage to using the average cost of servicing the warranty. One is that we would have to assume that the average cost is a good approximation for the marginal cost. The extent to which the assumption fails depends on the size of the fixed costs of servicing warranties.

the extended warranty from 15 to 10 percent. The cut applied to all stores an employees of the chain. The commission for the durable good was left untouched.

The drop in the extended warranty commission rate allows for estimates of specifications of the form

$$y_{cgsd} = d + \alpha U_{er(c),m(d)} + \beta U_{er(c),m(d)} \times Drop_d + postal(c) + m(g) + p(g) + s + \varepsilon_{cgsd}.$$

where $Drop_d$ indicates whether calendar date d is in May of 2009 or later. At face value, the coefficient on the interaction tells us how the marginal impact of the unemployment rate differs depending on the commission rate for the extended warranty. More broadly, as long as the unemployment rate measures the outside opportunities of sellers, we can interpret β as measuring the interaction effect of the commission drop with those outside opportunities. Note that the base variable $Drop_d$ is excluded from the regression. Its impact is wholly captured by the fixed effects for the calendar date d.

Estimates are found in Table 13. Moving left to right shows how the estimates differ with the dependent variable (prices and take up). The first row presents the estimated coefficient on the unemployment rate. The second row presents the estimated coefficient on its interaction with the drop in the extended warranty commission rate.

The estimates line up with the basic idea behind multitask agency problems. Workers substitute away from selling the extended warranty. They sell fewer extended warranties at lower prices. They sell base goods at higher prices, presumably because they focus more on upselling the base good. Specifically, before the drop in the commission, a one standard deviation increase lowers the price of the base good by 3.4 percent, the price of the extended warranty by 17.6 percent, and increases take up by a little under 1 percentage point. After the drop, the price of the base good goes up by 4.8 percent, the extended warranty price goes down by more than 45 percent. Take up decreases by less than it otherwise would have.

5.3. Deviating from the Recommended Price. Our data includes information on the extended warranty price the retail chain recommends to its stores. We use this information to analyze deviations from the recommended, to learn about whether the actual sellers of the goods are themselves affected by the unemployment rate, and whether they ultimately have discretion over the price the extended warranty fetches. The basic idea is that if the unemployment is indeed inducing a direct response from the people who sell the goods, then it should generate deviations from the extended warranty price the retailer recommends.

Estimates of the impact on deviations from the recommended price are found in Table 16. The first three columns present estimates from the specifications where the main regressor is the unemployment rate. The last three present estimates from specifications that include the unemployment rate and its interaction with a dummy indicating the reduction in the commission rate. We will

again focus our discussion on the estimates from Columns 3 and 6.

The extended warranty price falls below the recommended price when unemployment is high. Column 3 implies a one standard deviation increase in the unemployment rate lowers the extended warranty by 6 dollars relative to the recommended warranty price. Column 6 implies a similar relative change. It also shows, however, that the reduction in the sales commissions would generate a 16 dollar relative reduction in periods of high unemployment. Thus, while the people who sell goods generally lower the price of the extended warranty when unemployment is high, they do this moreso when the incentive sell warranties is diminished.[26]

5.4. Loss Leaders. TO BE DONE.

5.5. Role of Centralized Decisions. The head office of the retail chain can exert a centralized influence on price discrimination through the recommended price for the extended warranty. We investigate the impact of our economic shocks on the price the chain of the extended warranty recommends.[27] This allows us to investigate whether local economic conditions affect local pricing, or prices that are posted before the consumer ever enters the store.

Estimates of the impact of the unemployment rate on economic conditions are found in Table 17. The first three columns present the impact of the unemployment rate alone. The second three columns present the impact of the unemployment rate and its interaction with a dummy that indicates a drop in the extended warranty commission rate. The last three shows how the recommended price differs depending on the spot price of oil in Alberta, as well as on the standard deviation of the futures price. We focus our discussion on the estimates in Columns 3, 6, and 9. The other columns show how the estimates differ if we ignore self-selection on the part of the consumer.

Most notably, Table 17 shows that the impact on the recommended price has a similar magnitude to the impact on the price of the base good. This is notable because both prices are set centrally by the head office. Specifically, Column 3 shows the head office drops the recommended price of the extended by about one and a half percent (check this) when and where the unemployment rate is high. The result shows up with the raw specification, with no interaction, and in the specification with an interaction with the drop in the commission rate. The table also shows that the interaction of the commission and the unemployment rate has no statistical impact on the pricing of the extended warranty. In all the estimates imply the retailer is cognizant of local economic conditions, but that the impact is not large.

[26] One mechanism we cannot speak to with our data relates to the hiring of workers. It could be that high unemployment reduces customer volume, and that because of this fewer workers are hired. The workers who remain focus more on the customers they serve. Because there are few customers the worker can spend a little extra time trying to get them to buy the good or warranty.

[27] We will elaborate on this more in later drafts of the paper.

6 Markups

Markups are critical to understanding the incentives of the firm. We analyze the impacts of economic activity on raw and implied markups, paying special attention to a comparison of raw and implied markups that take account of the extended warranty and raw and implied markups that ignore it.

6.1. Raw Markups. Table 18 examines the impact of the unemployment rate on markups of the base good and extended warranty. The impact on absolute markups (price minus cost) are found in the top panel. The impact on relative markups are found in the middle panel. The bottom panel shows relative markups when we exclude 0 prices. These are prices of goods and warranties the retailer throws in with other goods. We excude them because they make relative markups less meaningful as a statistic.[28] To this end, we will focus our discussion on the top and bottom panel, leaving the middle panel as a point of comparison for the reader.

Absolute markups for the base good are lower when unemployment is high. A one standard deviation increase in the unemployment rate decreases absolute markups by between 5 and 6 dollars. The estimates are all statistically significant at the one percent level. The fragility of the extended warranty estimates are consistent with our interpretation that the retailer adjusts the extended warranty enough to not take a loss on sales of the extended warranty.

Columns 1 and 5 are also of interest, particularly when compared with the estimates in Columns 2-4 and 6-8. For extended warranties, the coefficient without controls is statistically significant. For the base good, the coefficient is not. This is consistent with the idea that the firm gets better customers (with a low marginal utility of income) during downturns. On average they pay more for the extended warranty. Once we control for the type of the consumer, there is no impact on the extended warranty markup, and a substantial decline in the base good markup.

The relative markups in the bottom panel is as expected given the other results in the paper. When unemployment is high, the relative markup for the base good declines. The retailer is able to maintain relative markups on the extended warranty.

6.2. Price Elasticities. We estimated own- and cross-price elasticities for the base good alone and for a bundle that includes the extended warranty. Using the elasticities together with common formulae for monopoly-pricing models, we derive the markups implied by a model that accounts for the extended warranty and one that ignores it.

To obtain elasticities, we had to estimate two different specifications. We will discuss each in turn, explaining why we need to estimate two along the way. The first specification is given by

$$bundle_{cgsd} = \beta_1 ln(p_{gd}^{base}) + \beta_2 ln(p_{gd}^{bundle}) + X_{cd}\Gamma + f(c) + m(g) + p(g) \times s \times d + \varepsilon_{csgd}$$

[28] At these prices the relative markups tend to negative infinity.

where the dependent variable indicates whether a customer c purchased an extended warranty with base good g in store s on calendar date d.[29] $f(c), m(g), p(g)$ are fixed effects for the postal code of the consumer, the manufacturer of good g, and the product category of good g. s and d are fixed effects for the store and calendar date. X_{cd} includes the unemployment rate, the spot price of oil, and the futures price of oil. These controls help us soak up cost and demand shocks that might bias estimation of the price elasticities.

We include fixed effects for the manufacturer to account for differences in the durability of the good. Not accounting for durability can generate biases in the estimation of demand elasticities (Esteban and Shum, 2007).

The fixed effects $p(g) \times s \times d$ soak up several sources of unobserved variation that can comprise the estimation of such a demand system. One concern relates to the multiproduct nature of our firm. Similar goods are typically located in the same area of a store. It will be easy for a consumer to compare the price of the good they purchase with other prices of similar goods. They can do the same with the warranty, after engaging a member of the sales staff. In these regards, the fixed effects $p(g) \times s \times d$ soak up the average price (or any other moment of the other price distribution) of models and bundles from the same product category, at the same store, on the same calendar date.

A second concern relates to the prices rivals are charging. The price at our retailer should correlate with the prices at rivals (for example, if retailers are acting strategically). Competing prices should affect the demand for the base good, and perhaps indirectly, the demand for the bundle. The store and calendard date part of $p(g) \times s \times d$ should help with this to some extent, as it should encapsulate the presence and absence of competing firms.

As a precaution, we mostly rely on instrumental variable estimates of our specifications. We have two instruments. The first is the extended warranty price the retail chain recommends to stores. The second is the cost per unit of the base good.

We require a different specification to uncover own- and cross-price elasticities in the demand for the base good. We require this because each customer in our database has purchased at least one base good. We do not observe customers who do not or who visit but leave without purchasing anything at all. To obtain variation in the base good quanties, we built a slightly aggregated panel from the raw transactions data. Specifically, we sum quantities to the point where the unit of observation is made up of the employment insurance region (recall there are 60 of these) of the customer, the manufacturer and model of the product (more than 30,000 of these) they bought, and the calendar month (for example, 1998, month 1) they bought it in. The aggregation generates zeros for observations where no base good was purchased. In all it yields a panel with just under 9 million observations.

For the base good quantity, we estimated specifications of the sort

[29] We can interpret the dependent variable as the probability that a bundle is purchased or as the probability that an extended warranty is purchased.

$$base_{egd} = \beta_1 ln(p_{gd}^{base}) + \beta_2 ln(p_{gd}^{bundle}) + X_{cd}\Gamma + e + m(g) + p(g) \times d + \varepsilon_{cgd}$$

where the dependent variable is the quantity of the base good g in employment insurance region e on calendar month d. e are fixed effects for the employment region. $X_{cd}, m(g), p(g), d$ are as before, except that now d is the calendar month instead of the calendar date.

The assumptions required for identification of the specification for the base good quantity are stronger than the ones required for identification of the specification for the bundle. $p(g) \times d$ soaks up the average price (again, or any other moment of the other price distribution) of all the other goods the retail chain as a whole sells. This might yield imprecise measures of the average price of other goods (at the chain) that are sold in the specific employment region. In addition, the specification for the bundle does better at capturing the impacts of prices (or presence) of rival retailers.

Estimates of the specification for bundled good are found in Table 19. The right panel presents estimates from the instrumental variables specification. The left panel presents OLS estimates as a point of comparison. Moving left to right in each panel shows how the estimates change as controls are added. We concentrate our discussion on the estimates in Columns 5 and 10.

The base and the bundled goods are substitutes in consumer demand for the bundled good. The IV estimates show that there is more take up of the bundled good at higher prices for the base good. A 10 percent increase the price of the base good increases take up by 8.5 percentage points. The increase is equivalent to a 21 percent increase over the mean take up.

Note that price of the base good has a stronger impact on consumer demand for bundled good. A higher price for the base good increases take up by 8.5 percentage points. A higher price for bundled good decreases take up 7.1 percentage points. The 7.1 percentage point decrease is equivalent to an 18 percent decrease under the mean take up. That the price of the base good has a stronger impact on bundle demand is unsurprising. The price of the base good is usually 7 times the price of the extended warranty.

Estimates of the specification for the base good quantity are found in Table 20 confirm patterns found in estimates of the specification for bundled demand. The estimate in Column 10 of the table shows that 10 percent higher bundle prices leads to 0.43 fewer purchases of the base good. Here, as with demand for the bundle, the impact of the base good price dominates the impact of the bundle price.

The substitutability of the base and bundled good imply that profit margins should be higher when the goods are sold together. We use the raw data to investigate whether this is the case. Figure 6 plots total markups against the probability that a model is bundled with an extended warranty. The unit of observation in the figure is the model number. The total markup equals the markup on the base good when just the base good is sold. It equals the sum of the markups on the

base good and extended warranty when both sold together. The bundle probability is calculated over the entire sample. The figure shows that the total markup is generally larger for models that are usually together with the extended warranty.

6.3. Implied Markups. TO BE DONE.

7 Consumer Welfare

We will eventually use a simple to calculate consumer surplus. For now we will simply analyze the role of durability, as along with price has an impact on the utility the consumer derives from the purchase of the base good.

7.1. Consumer Surplus. TO BE DONE.

7.2. Durability. We analyze the impact of economic activity on the expected durability of the goods consumers purchase. We have several measures of durability. All come from the warranty of the manufacturer. The manufacturer offers a warranty on parts and labour. The coverage can vary within each product. The manufacturer will usually cover all parts and labor costs for the first couple of years. After that it selectively covers labor and parts for some parts of the durable. Eventually the manufacturers offers no coverage at all.

The manufacturer should, in principle, use their expectations over the durability of the good to set the length and coverage of their warranty. The manfacturer's warranty should thus let us identify the impact of economic activity on the expected durability of the products consumers purchase.[30][31]

Estimates of the impact of the unemployment rate on expected durability are found in Table 21. The dependent variable in all these regressions is the number of days the manufacturer covers (in natural logarithms). The upper left panel presents the estimates for the minimum coverage for labor. The upper right the estimates for the maximum labor coverage. The bottom panels do the same for coverage of parts.

The pattern of estimates is similar to the pattern for the base price of the good. In the absence of controls, durability improves when unemployment is high. In the presence of controls, durability tends to decline when unemployment is high. Like our other results, this supports the notion that the retail chain gets more visits from consumers who value an additional dollar a little bit less.

[30] We will look into the laws in Canada that govern this, in particular whether are minimum warranty requirements (we know there in Quebec).

[31] In the economics literature, the challenges that come with measuring durability are well known. An ideal measure would, at least for economics, require information on the value of the good in secondary markets. We tried to scrape information from secondary markets but were unable to construct a data set that we could comprehensively match to the models in our sample.

Overall, the estimates imply that economic activity has very little impact on the durability of purchases of the base good. In the specifications with controls the point estimates are statistically precise zeros. The point estimates are statistically significant but economically small. For example, the estimate in Column 8 implies that one standard deviation increase in the unemployment rate reduces the maximum coverage for labour by 0.25 percent.[32]

8 Conclusion

A large body of empirical research uses list or average prices to learn about the behavior of the firm and the welfare of the consumer. Our most basic observation is that there are limitations to inferences that use these data in cases where the firm pushes unadvertised goods at the point of the sale. We have shown that having no information on the prices unadvertised goods can lead to biases in estimates of the response to economic activity, the response to demand and cost shocks, and the price elasticities of demand. Ultimately, this lends itself to biased inferences about the behavior of the firm and consumer welfare.

We conclude the paper with a discussion of the broader implications of our study. We will first illustrate the impact of not accounting for the prices of unadvertised goods on measured price inflation. We next discuss how our results might apply given the dramatic advances in electronic retail since the time frame of our sample. Finally, we revisit the external validity of our study.

8.1. Measurement Error in Price Inflation. As far as we know, it is uncommon for statistical agencies to collect add-on prices. This is for good reason, as by definition these prices are costly to collect. Because they are hidden, a price collector would have to visit a store, and speak with someone who works there, in order to learn the add-on price. They would have to do this for every good in the price basket. The advent of online price collection has not necessarily made this easier. Retailers that operate online often require various pieces of personal information before the customer can learn the prices of add on goods (think about airline tickets for example). If someone were to try to scrape the add-on price, they would have invent personal information (including a credit card number in some cases) in order to book at ticket and learn the add-on price.

Perhaps the best bet for statistical agencies is the use of transactions data like the one we study.[33] In accordance with this, we use our transactions data to consider how price inflation differs depending on whether or not it accounts for the price of the extended warranty. To do this, we use a Jevons (geometric mean) price index. We use the Jevons index because it is simple to

[32]Note that the results imply not having a direct measure of durability in our other regressions implies a minimal level of bias.

[33]Statistical agencies are well aware of the advantages of scanner data. Their use, however, has not yet been adopted widely.

understand, and because other indices yield the same basic conclusions.[34] The starting point for our index is the (geometric) mean price P_d over all transactions in calendar month d

$$P_d = (\prod^{N_d} p_{bd}^I (p_{bd} + p_{wd})^{1-I})^{\frac{1}{N_d}}$$

where p_{bd} is the price of the base good, p_{wd} is the price of the extended warranty, I is an indicator for whether only the base good was purchased, N_d is the number of purchases in calendar month d, and where the counting index is omitted. Note that P_d weighs each price by the number of purchases. It can be written as a weighted average of the prices across regions, or as a weighted average of the prices across model, or both.[35] Ultimately, our adjusted price index is given by

$$\frac{P_d}{P_0}$$

where P_0 is the average price in the first month of 1999. By the same token our naive price index is the price relative for

$$P_d = (\prod^{N_d} p_{bd})^{\frac{1}{N_d}}.$$

Figure 7 compares the price indices over our sample time frame. Figure 7a compares our adjusted and naive inflation rates. Figure 7b examines the cyclicality of the difference between the rates. It tells us whether the differences depend on the level of economic activity in the economy. Figure 7c compares the rates after we adjust them for seasonality.[36] Figure 7d shows how the difference in seasonally-adjusted rates depends on the level of economic activity.

A couple of things are implied by the figures. First, Figures 7a and 7d are suggestive of some dependence between the difference in the rates and the level of economic activity. The naive rate approaches the adjusted rate when the economy is doing well.[37] Second, once the rates are adjusted for seasonality, the naive inflation rate exhibits more volality than our adjusted rate. The extent to which this implies that inflation rates are mismeasured depends on the weights placed on goods, like ours, cars, or hotels, in official measures of inflation.

We can compare inflation rates based on our transactions data with the official CPI for durables

[34] Another reason is that we observe exact prices. Many of the advances in price index measurement are a consequence of the difficulties that come with measuring exact prices.

[35] To see more clearly, one can take the term inside the outer bracket and break it up into the products of prices for each model and geographic region. Doing so shows that the exponent on the base price is the sum of the I's for the model and geographic region. The exponent on the base plus extended warranty price is the sum of the $1 - I$'s.

[36] Our adjustment for seaonality takes P_d, regresses it on month dummies, and uses the residual to calculate the inflation measure.

[37] The figures are consistent with estimates of unconditional specifications in the body of the paper. This is not inconsistent with the results in the body of the paper. The main results in the body condition on various fixed effects in order to deal with the impact of self-selection by the consumer, and to more cleanly identify the behavior of the firm. The statistics in Figure 7 are, of course, unconditional.

(Figure 8). The official CPI is most comparable with our adjusted inflation rate in Figure 7c (the one that is also adjusted for seasonality). Both are fairly flat for most our sample time frame, but decline quickly thereafter.

8.2. Cost of Retailing Online. Our study highlights an important advantage to physical (brick-and-mortar) stores. The advantage of stores is that they can use employees to help them push add ons like extended warranties. These employees can, in turn, better help the stores cope with recessionary pressure. The loss in this kind of flexibility is thus a cost to operating online.

Our data allows us to examine this issue empirically. During the period under study, the retailer was experimenting with its online arm. Our data includes all the transactions that occured during this experimentation phase.

Figure 9 lends further support to the idea that the loss in flexibility is a cost to going online. Figure 9a depicts the evolution of the extended warranty price depending on whether it was sold at a physical store or onine (green dash-dot line for brick-and mortar, red dashed line for online). Figure 9b does the same for extended warranty take up.

Figure 9a shows the in-store price of the extended warranty is usually lower than the online price. Figure 9b shows the in-store take-up is usually much higher tha online take up. This should be the case if the sales staff have a role in sales of the extended warranty. In these regards, an important question relates to whether the overhead expenses or fixed costs of operating a store outweigh the benefits from getting more take up of the extended warranty.

8.3. External Validity. One of the limits of our study is that our sample comes from a period where the retail sector as whole was in the early stages of a mass transition. Because of the rise of electronic retail, the landscape in the years to come might look very different from the landscape that existed during our sample period. Having said that, and given the recent developments in the retail sector, we would find it surprising if add ons became less important in the coming years.

Our study should have greater applicability to settings where the risk preference of the customer affects the demand for add on good, such as insurance for the home, motor vehicles, mobile phones, or a wide range of financial products. It should have less applicability to add ons such as leather interiors in cars, minibar items in a hotel room, cartridges for a computer printer, goods where the importance of risk preference is not so clear.

References

Robert B. Barsky, Christopher L. House, and Miles S. Kimball. Sticky-price models and durable goods. *American Economic Review*, 97(3):984–998, June 2007.

Marianne Baxter and Anthony Landry. Ikea: Product, pricing, and pass-through. *Federal Reserve Bank of Dallas Working Paper Series*, (132), November 2012.

Mark Bils and Peter J. Klenow. Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5):947–985, 2004.

Mikael Carlsson and Oskar Nordstrom Skans. Evaluating microfoundations for aggregate price rigidities: Evidence from matched firm-level data on product prices and unit labor cost. *American Economic Review*, 102(4):1571–1595, June 2012.

Glenn Ellison. A model of add-on pricing. *Quarterly Journal of Economics*, 120:585–637, 2005.

Glenn Ellison and Sarah Ellison. Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77:427–452, 2009.

Susanna Esteban and Matthew Shum. Durable-goods oligopoly with secondary markets: the case of automobiles. *Rand Journal of Economics*, 38(2):332–354, Summer 2007.

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121:505–540, 2006.

Etienne Gagnon. Price setting during low and high inflation: Evidence from mexico. *Quarterly Journal of Economics*, 124(3):1221–1263, August 2009.

Gautam Gowrisankaran and Marc Rysman. Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219, 2012.

Emi Nakamura and Jon Steinsson. Price rigidity: Microeconomic evidence and macroeconomic implications. *Annual Review of Economics*, 5:133–163, 2013.

Table 1: Summary Statistics for Base Good and Add on. Column 5 uses a t-statistic that allows for unequal variances to test for a difference between the mean price and expected cost of the extended warranty. The standard error of the difference is in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good	Extended Warranty				
	Price Paid	Take Up	Price Paid	Expected Cost	Price - Cost	Expected Markup
	(1)	(2)	(3)	(4)	(3) - (4)	(6)
Sample Mean and	610.90	0.37	88.65	23.49	65.16***	0.74
Standard Deviation	(1727.80)	(0.48)	(93.40)	(164.34)	(0.12)	
Observations	7,039,654	7,039,654	2,626,335	2,626,335	2,626,335	2,626,335

Figure 2: Add on Markup and Base Good Price. The figure excludes (2 out of more than 140) products with expected markups that are less than zero.

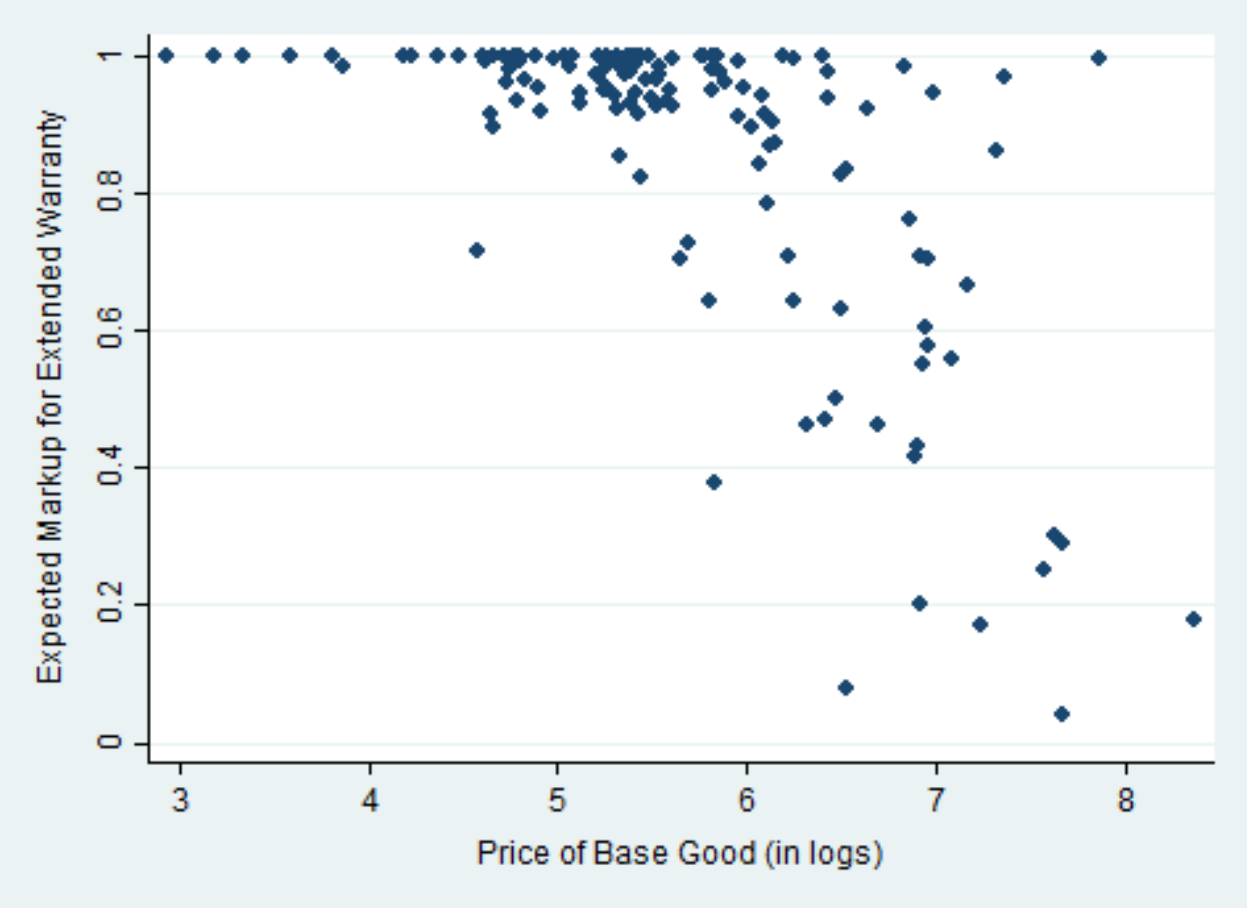

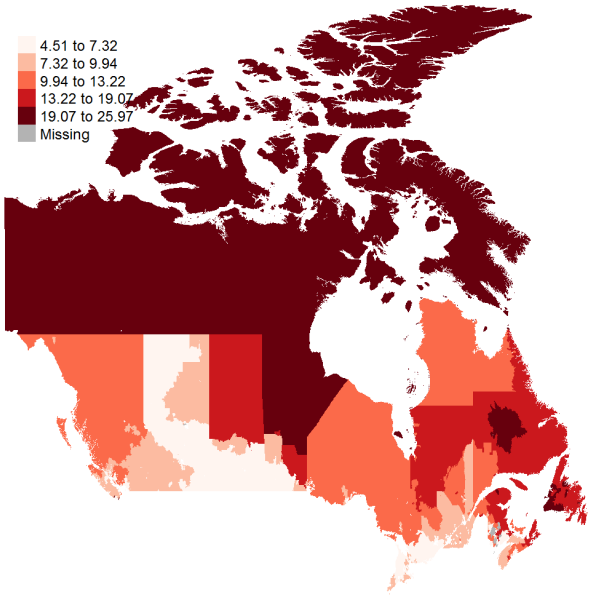

(a) Unemployment Rate

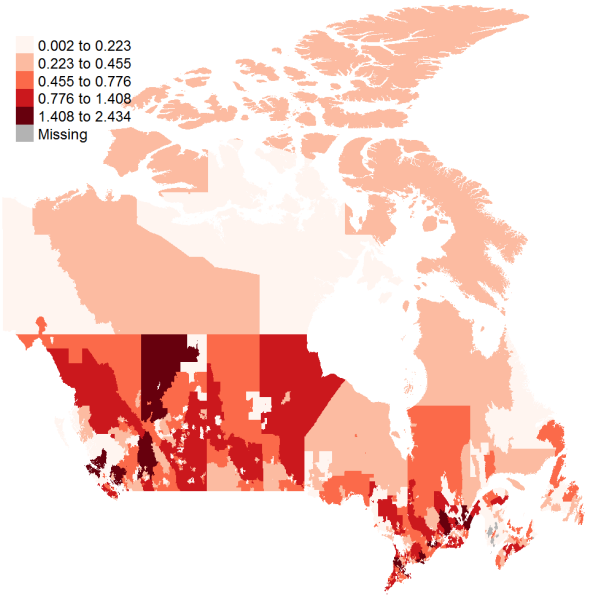

(b) Base Good Price

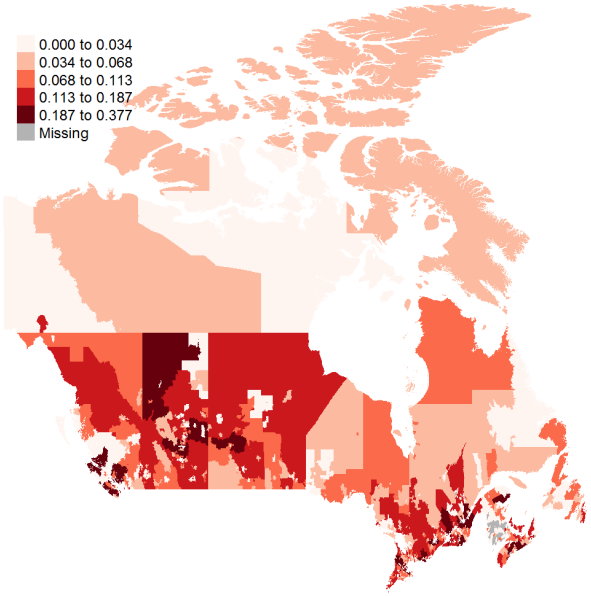

(c) Extended Warranty Price

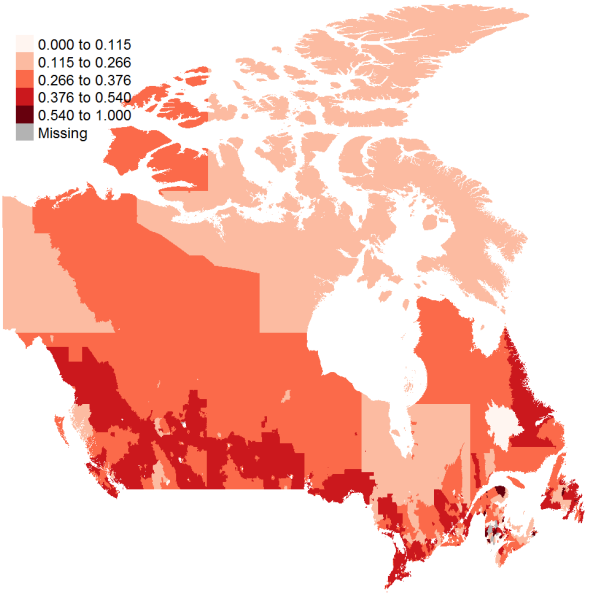

(d) Take Up of Extended Warranty

Figure 3: Adding on, over Space. Prices are weighted by the population share of the forward sortation area. All figures are based on the k-means ($k = 5$) clustering algorithm.

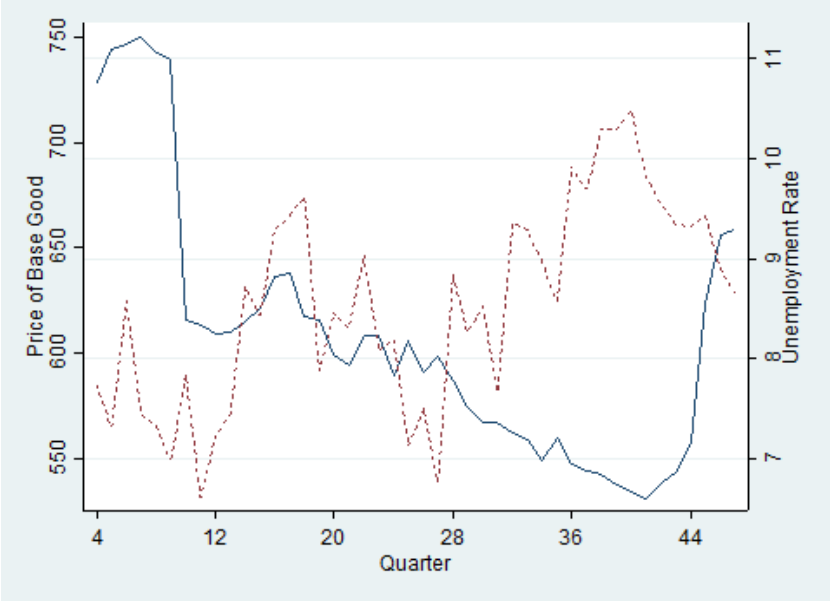

(a) Base Good Price

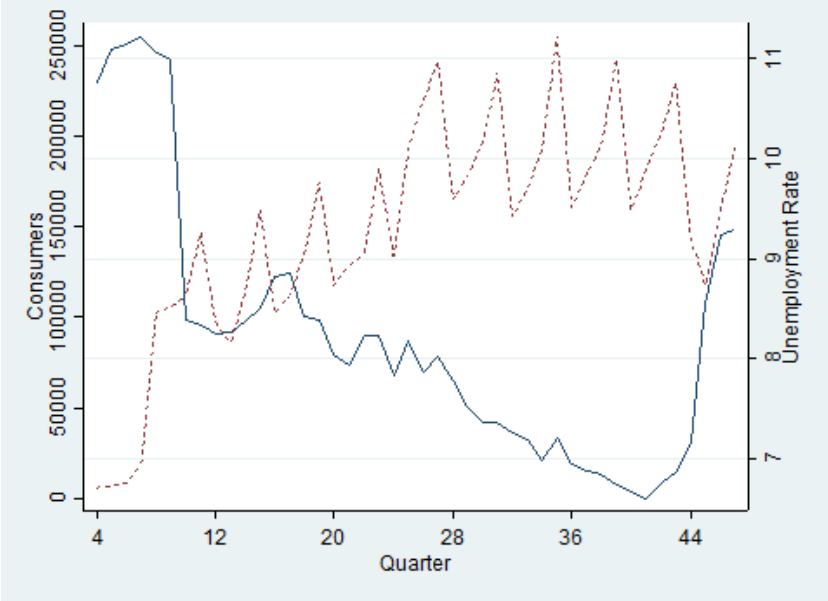

(b) Base Good Quantity

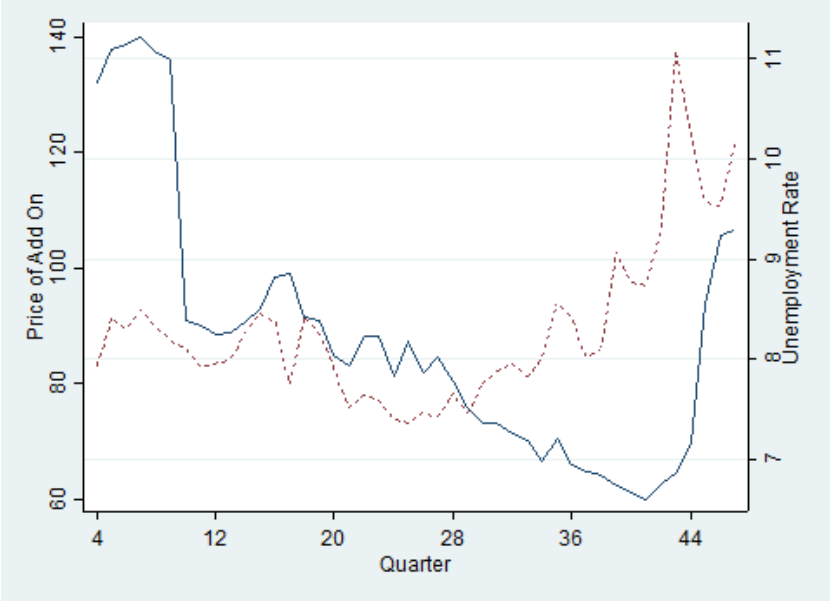

(c) Extended Warranty Price

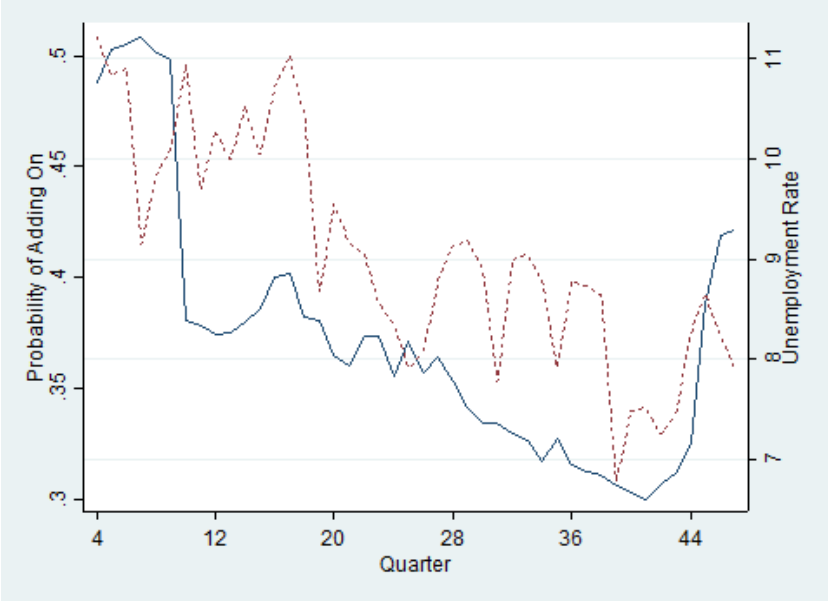

(d) Take Up of Extended Warranty

Figure 4: Adding on, over Time.

Table 2: Unemployment and the Add on. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Price (in logs)				
Unemployment Rate (logs) in	-0.028	-0.048**	-0.031	-0.028***	-0.023***
Consumer's Neighbourhood	(0.020)	(0.019)	(0.020)	(0.009)	(0.006)
Observations	6771481	6771469	6771443	6771443	6771443
R^2	0.000	0.034	0.074	0.704	0.786
	Price of the Extended Warranty(in logs)				
Unemployment Rate (logs) in	0.252***	0.339***	-0.186***	-0.181***	-0.180***
Consumer's Neighbourhood	(0.029)	(0.029)	(0.025)	(0.024)	(0.025)
Observations	2486234	2486203	2486183	2486183	2486183
R^2	0.002	0.024	0.065	0.158	0.161
	Take-up of Extended Warranty				
Unemployment Rate (logs) in	-0.009	0.047***	0.008	0.009**	0.008**
Consumer's Neighbourhood	(0.010)	(0.006)	(0.005)	(0.004)	(0.004)
Observations	6784076	6784064	6784038	6784038	6784038
R^2	0.000	0.024	0.039	0.128	0.157
Customer FSA		✓	✓	✓	✓
Date			✓	✓	✓
Product Category				✓	✓
Manufacturer					✓

Table 3: Bundling the Base Good and Add on. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

		Price (in logs) of Base Good		
		(1)	(2)	(3)
Base Good and Warranty	Unemployment Rate (logs)	-0.067***	-0.161***	-0.030***
	in Consumer's Neighbourhood	(0.013)	(0.010)	(0.006)
	Observations	2534382	2534351	2534331
	R^2	0.001	0.033	0.064
Just Base Good	Unemployment Rate (logs)	0.003	-0.026	-0.032
	in Consumer's Neighbourhood	(0.026)	(0.024)	(0.024)
	Observations	4237099	4237082	4237060
	R^2	0.000	0.042	0.093
	Consumer's Neighbourhood		✓	✓
	Calendar Date			✓

Table 4: Base Good Qty and the Unemployment Rate. Quantities are based on local customers. These are customers residing within 50 miles of the store. We get similar patterns using other definitions of a local customer. The sum of the two dependent variables gives the effect on the number of purchases overall. The unit of analysis is the store-quarter-year. Standard errors are clustered along two dimensions, the quarter-year and the store. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Quantity		
Unemployment Rate	-0.923***	-1.157***	-0.137***
(in natural logarithms)	(0.011)	(0.013)	(0.011)
Observations	147447	147424	147423
R^2	0.053	0.602	0.794

Table 5: Downgrading during Downturns. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Price (logs) for Base Goods Aged Less Than					
	15 days	30 days	60 days	90 days	180 days	360 days
	(1)	(2)	(3)	(4)	(5)	(6)
Unemployment	-0.004	-0.009	-0.023	-0.021	-0.026	-0.027
Rate (logs)	(0.025)	(0.020)	(0.020)	(0.020)	(0.019)	(0.018)
Observations	99140	196070	433887	692976	1635985	3566606
R^2	0.468	0.408	0.313	0.232	0.141	0.118

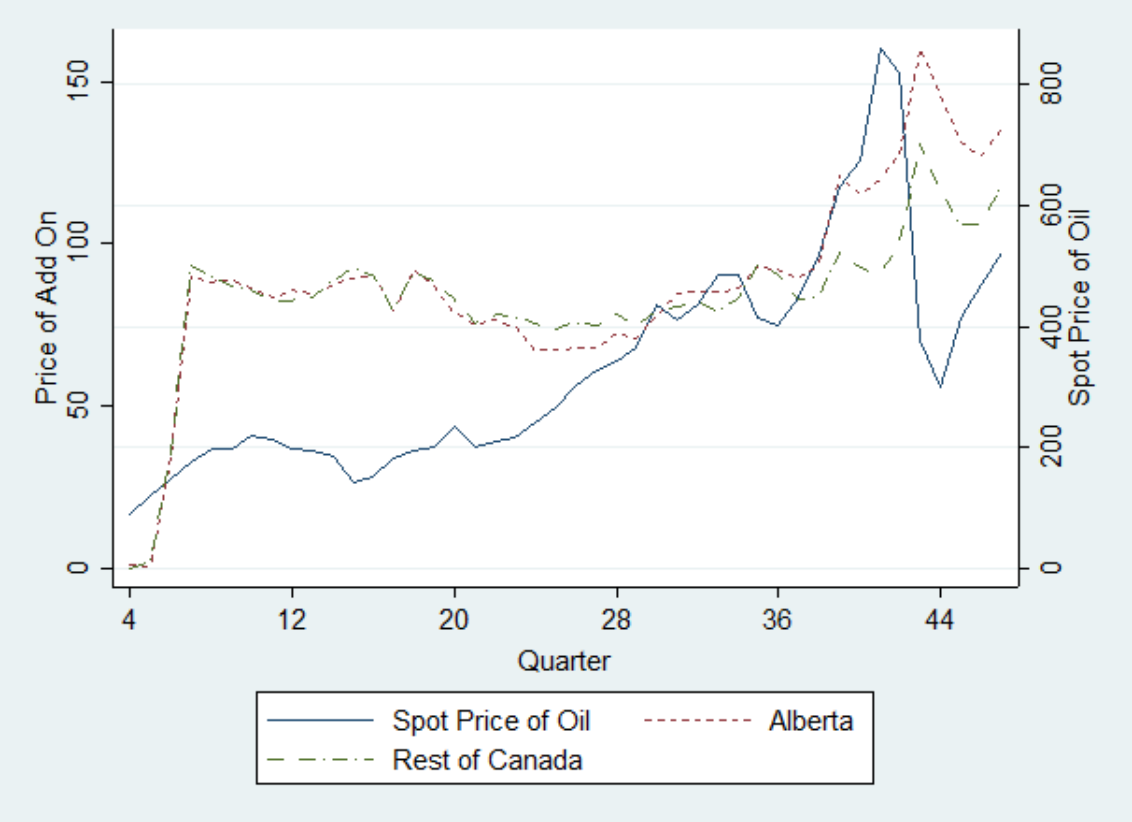

(a) Add-on Price

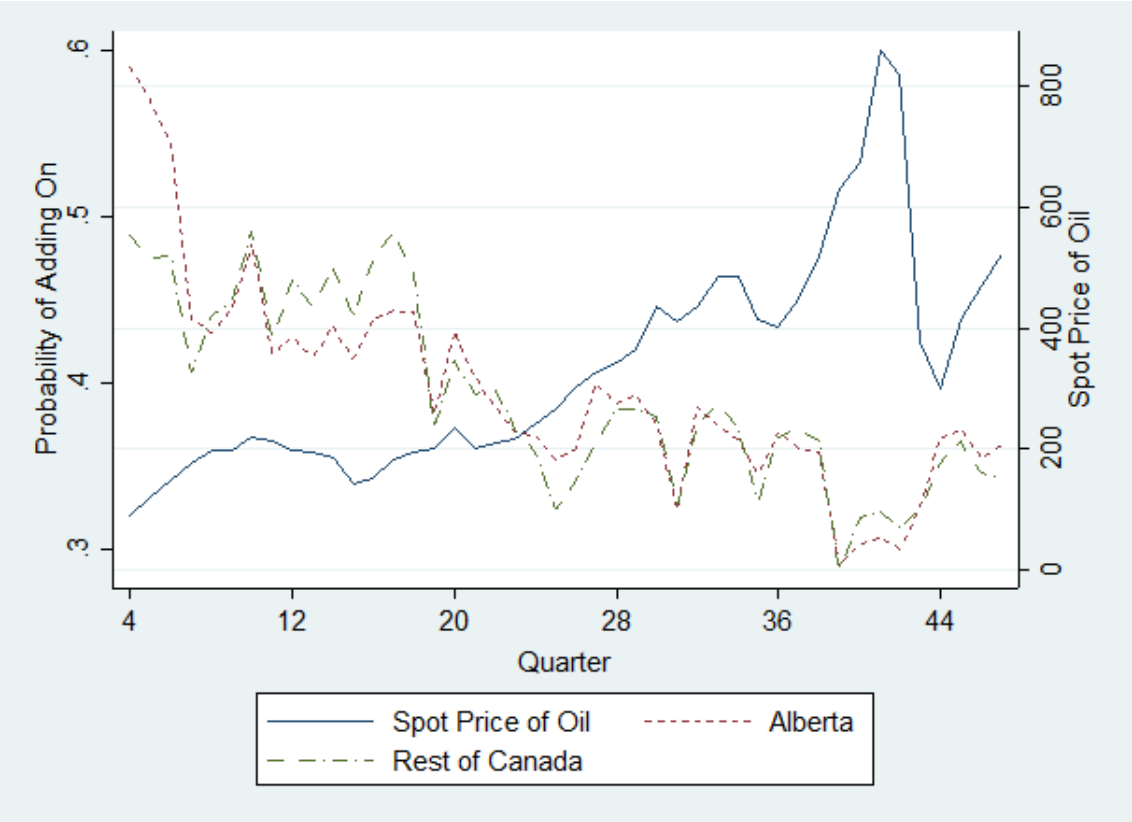

(b) Warranty Probability

Figure 5: Income Shocks and the Prices and Quantities of the Base Good and Extended Warranty.

Table 6: Income, Uncertainty, and the Price of the Add on. Estimates of specifications that include fixed effects for the neighbourhood of the consumer, as well as for the calendar date. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Prices are in logarithms.

	Dependent Variable		
Interaction of Alberta dummy with	Base Good Price	Warranty Price	Warranty Take Up
Spot Price of Oil	-0.004	0.160***	-0.005*
	(0.005)	(0.015)	(0.003)
Observations	6771510	2486245	6784105
R^2	0.074	0.066	0.039
Standard Deviation of Oil Price Futures	-0.013***	0.029***	0.004***
	(0.003)	(0.011)	(0.001)
Observations	6771510	2486245	6784105
R^2	0.074	0.065	0.039
Spot Price of Oil and with	-0.005	0.160***	-0.005*
	(0.005)	(0.015)	(0.003)
Standard Deviation of Oil Price Futures	-0.013***	0.032***	0.004**
	(0.003)	(0.012)	(0.001)
Observations	6771510	2486245	6784105
R^2	0.074	0.066	0.039

Table 7: Base Good Qty and the Mean Oil Price. Quantities are based on local customers. These are customers residing within 50 miles of the store. We get similar patterns using other definitions of a local customer. The unit of analysis is the store-quarter-year. The results imply there are more consumers (purchasing at least one good) when their income increases. The results imply income has a negigligible impact on the number of base goods the consumer purchases. Not that the oil price is the mean observed oil price in the quarter. The mean in the full sample is 58.74. The standard deviation is 27.66. Standard errors are clustered along two dimensions, the quarter-year and the store. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Consumers Purchasing One Good or More			Purchases per Consumer		
	(1)	(2)	(3)	(4)	(5)	(6)
Interaction of Alberta Dummy	0.0004**	0.0008***	0.0004***	-0.0002	0.0011**	0.0002
and the Spot Price of Oil	(0.0002)	(0.0003)	(0.0001)	(0.0003)	(0.0005)	(0.0003)
Observations	5613	5613	5613	5613	5613	5613
R^2	0.0031	0.5210	0.5905	0.0004	0.6548	0.7065
Store		✓	✓		✓	✓
Quarter-Year			✓		✓	✓

Table 8: Sunk Costs. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	All Consumers			All Consumers outside Alberta		
	Base Good Price	Warranty Price	Warranty Take Up	Base Good Price	Warranty Price	Warranty Take Up
	(1)	(2)	(3)	(4)	(5)	(6)
Distance Between Home of the Consumer and the Closest Store in the Chain	-0.057***	-0.032***	-0.011***	-0.067***	-0.036**	-0.013***
	(0.013)	(0.011)	(0.003)	(0.016)	(0.014)	(0.003)
Interaction with the Spot Price of Oil	-0.011***	0.023***	-0.005***	-0.012***	0.017**	-0.004***
	(0.003)	(0.007)	(0.001)	(0.004)	(0.008)	(0.001)
Observations	6770334	2485879	6782919	5374548	1973802	5384268
R^2	0.075	0.065	0.039	0.082	0.064	0.043

Table 9: Bad Types. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the customer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Price	Warranty Price	Warranty Take Up
	(1)	(2)	(3)
Number of claims prior to this visit	0.012***	-0.010	0.005*
	(0.004)	(0.008)	(0.003)
Observations	2420431	821964	2426048
R^2	0.115	0.070	0.061

Table 10: Price Adjustments and the Average Type. This table has fixed effects for the manufacturer. These fixed effects kill the interaction effect between claim propensity and the unemployment rate. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Price (1)	Warranty Price (2)	Warranty Take Up (3)
Unemployment Rate	-0.011***	-0.197***	0.011***
	(0.004)	(0.024)	(0.003)
Interaction with Claim Propensity in the FSA where customer lives	-0.004	0.078***	-0.015***
	(0.003)	(0.028)	(0.004)
Observations	6288040	2446965	6299735
R^2	0.420	0.086	0.114

Table 11: Price Adjustments and Individual Type. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Price (1)	Warranty Price (2)	Warranty Take Up (3)
Claim Propensity of the Customer	0.159***	-0.006	0.120***
	(0.005)	(0.010)	(0.002)
Unemployment Rate	0.005	-0.196***	0.015***
	(0.005)	(0.024)	(0.003)
Interaction	-0.000	0.020***	-0.002
	(0.003)	(0.005)	(0.001)
Observations	6288045	2446974	6299740
R^2	0.077	0.064	0.084

Table 12: Repeat Customers. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Base Good Price		Warranty Price		Warranty Take Up	
	(1)	(2)	(3)	(4)	(5)	(6)
Repeat Visit	-0.227***	-0.175***	-0.047***	-0.208***	-0.064***	-0.025***
(dummy variable)	(0.014)	(0.011)	(0.006)	(0.009)	(0.006)	(0.002)
Observations	6771510	2486245	6784105	6771472	2486195	6784067
R^2	0.010	0.002	0.002	0.082	0.065	0.039
Repeat visit of a customer	-0.282***	-0.192***	-0.091***	-0.268***	-0.042***	-0.057***
who visits more than once	(0.014)	(0.012)	(0.006)	(0.009)	(0.006)	(0.002)
Observations	3790829	1402979	3799137	3790813	1402920	3799121
R^2	0.015	0.002	0.008	0.103	0.072	0.056
Customer FSA		✓		✓		✓
Calendar Date		✓		✓		✓

Table 13: Commissions for Adding On. Estimates of specifications that include fixed effects for the neighbourhood of the consumer, as well as for the calendar date. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Prices are in logarithms.

	Base Good Price	Warranty Price	Warranty Take Up
Unemployment Rate (logs) in	-0.034*	-0.176***	0.008
Consumer's Neighbourhood	(0.020)	(0.024)	(0.005)
Interaction with Drop in Commission	0.082***	-0.279***	-0.005
on Extended Warranties (15% to 10%)	(0.020)	(0.066)	(0.009)
Observations	6771481	2486234	6784076
R^2	0.074	0.065	0.039

Table 14: Cost and Pass Through of the Base Good. The base good cost is the transfer price from the retail chain to the store, scaled up by the commission rate for sales of the base good. Estimates of specifications that include fixed effects for the neighbourhood of the consumer, as well as for the calendar date. Standard errors are clustered on a grouping of the employment region and month, the unit of observation for the unemployment rate. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Cost of the Base Good					
	Logs in First Three Columns, Dollars in Last Three					
Unemployment Rate	0.04	-0.08**	-0.06***	9.88	-99.08***	-22.76***
	(0.02)	(0.04)	(0.02)	(9.04)	(12.26)	(8.65)
Observations	409167	403645	403614	409169	403647	403616
R^2	0.00	0.04	0.14	0.00	0.03	0.07
	Price of the Base Good					
	Logs in First Three Columns, Dollars in Last Three					
Unemployment Rate	0.04	-0.11***	-0.06***	12.35	-125.77***	-22.96**
	(0.03)	(0.04)	(0.02)	(10.25)	(13.21)	(9.14)
Observations	408584	403077	403045	409169	403647	403616
R^2	0.00	0.04	0.12	0.00	0.03	0.07
Customer FSA		✓	✓		✓	✓
Calendar Date			✓			✓

Table 15: Cost and Pass Through of the Warranty. The warranty cost is comprised of the average cost of a claim and the commission for the salesperson. For franchise stores, it also includes the franchise fee and the cost of advertising locally (both of which are in proportion to gross revenues). Standard errors are clustered on the employment insurance region, the cross-sectional unit of observation for the unemployment rate. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Warranty Cost			Warranty Price		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemployment Rate	-14.55***	1.68	-5.13***	-1.28	-6.54***	-8.94***
	(1.45)	(2.33)	(1.96)	(1.10)	(1.15)	(1.18)
Observations	2583061	2497727	2497708	2576234	2491318	2491296
R^2	0.00	0.01	0.03	0.00	0.02	0.05
Customer FSA		✓	✓		✓	✓
Calendar Date			✓			✓

Table 16: Deviation from the Suggested Price and the Unemployment Rate, and the Commission. Standard errors are clustered on the employment region (employment benefits are governed by the circumstances in teh region you live in) and the month. This is the most disaggregated level for the unemployment rate. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Difference between Observed and Recommended Extended Warranty Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Unemployment Rate	13.08***	28.78***	-6.19***	16.09***	38.17***	-5.83***
	(0.98)	(2.64)	(0.94)	(1.10)	(2.70)	(0.92)
Interaction with Reduced Commissions For Sales of Extended Warranties (from 15 to 10 percent)				-13.75***	-16.00***	-10.63***
				(0.66)	(0.70)	(2.72)
Observations	2617291	2531229	2531209	2617291	2531229	2531209
R^2	0.00	0.02	0.10	0.01	0.03	0.10
Customer FSA		✓	✓		✓	✓
Calendar Date			✓			✓

Table 17: Recommended Price, Unemployment, and the Commission. Standard errors are clustered on the employment region (employment benefits are governed by the circumstances in teh region you live in) and the month. This is the most disaggregated level for the unemployment rate. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Recommended Price of the Extended Warranty (natural logarithms)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Unemployment Rate	-0.09***	-0.22***	-0.03***	-0.14***	-0.35***	-0.03***			
	(0.01)	(0.02)	(0.01)	(0.01)	(0.03)	(0.01)			
Interaction with a				0.22***	0.23***	-0.00			
Reduced Commissions				(0.00)	(0.01)	(0.01)			
For Sales of Extended									
Warranties (from 15									
to 10 percent)									
Interaction of Alberta dummy									
with Spot Price of Oil							0.21***	0.20***	0.01***
							(0.00)	(0.00)	(0.00)
with Standard Deviation							0.04***	0.04***	-0.01***
of Futures Price of Oil							(0.01)	(0.01)	(0.00)
Observations	2625276	2538733	2538714	2625276	2538733	2538714	2537859	2537829	2537810
R^2	0.00	0.03	0.13	0.02	0.05	0.13	0.01	0.03	0.13
Customer FSA		✓	✓		✓	✓		✓	✓
Calendar Date			✓			✓			✓

Table 18: Markups and the UnemploymentRate. Standard errors are clustered on the FSA (the first three characters of the postal code) of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Markup on the Base Good				Markup on the Extended Warranty			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Absolute Markups				
Unemployment Rate	-3.57	-8.22***	-5.18***	-5.67***	12.98***	-3.79	-3.71*	-3.71
	(2.43)	(1.36)	(1.17)	(1.19)	(1.89)	(2.32)	(2.23)	(2.34)
Observations	4865371	4724571	4724568	4724568	2576233	2491295	2491283	2491280
R^2	0.00	0.07	0.18	0.18	0.00	0.05	0.07	0.08
Customer FSA		✓	✓	✓		✓	✓	✓
Calendar Date		✓	✓	✓		✓	✓	✓
Manufacturer			✓	✓			✓	✓
Product Category			✓	✓			✓	✓
Store			✓	✓			✓	✓

Table 19: Demand for Bundle. OLS estimates in Columns 1 to 5. IV estimates are in Columns 6 to 10. Instrumental variables are the wholesale cost of the base good and the recommended price of the extended warranty. Standard errors are clustered on the employment insurance region. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Take Up of the Extended Warranty									
	OLS (first five columns)					Instrumental Variables (last five columns)				
Base Price	0.16***	0.15***	0.15***	0.15***	0.15***	1.15***	1.13***	1.03***	0.86***	0.85***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Bundle Price	0.00	0.01***	-0.01***	0.04***	0.04***	-1.03***	-1.00***	-0.93***	-0.71***	-0.71***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Unemployment Rate					0.01***					0.02***
					(0.00)					(0.00)
Spot Price of Oil (for Alberta)					-0.02***					-0.01***
					(0.00)					(0.00)
Futures Oil Prices (Standard Deviation, for Alberta)					0.01***					0.01***
					(0.00)					(0.00)
Employment Region		✓	✓	✓	✓		✓	✓	✓	✓
Manufacturer			✓	✓	✓			✓	✓	✓
Product Category × Store × Calendar Date				✓	✓				✓	✓
Observations	5914646	5697251	5697247	5679320	5679320	4433237	4303319	4303318	4288105	4288105
R^2	0.10	0.11	0.13	0.18	0.18	-0.20	-0.18	-0.10	0.04	0.04

Table 20: Demand for Base Good. OLS estimates in Columns 1 to 5. IV estimates are in Columns 6 to 10. Instrumental variables are the wholesale cost of the base good and the recommended price of the extended warranty. Standard errors are clustered on the employment insurance region. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Quantity of Base Good									
	OLS (first five columns)					Instrumental Variables (last five columns)				
Base Price	0.20***	0.29***	0.42***	0.30***	0.31***	-6.94***	-7.63***	-6.89***	-4.75***	-4.74***
	(0.05)	(0.08)	(0.11)	(0.09)	(0.10)	(1.25)	(1.44)	(1.27)	(0.97)	(0.96)
Bundle Price	-0.43***	-0.57***	-0.55***	-0.77***	-0.79***	7.28***	8.02***	7.30***	4.27***	4.26***
	(0.10)	(0.14)	(0.14)	(0.17)	(0.17)	(1.33)	(1.54)	(1.36)	(0.91)	(0.90)
Unemployment Rate					-0.19**					-0.25
					(0.07)					(0.17)
Spot Price of Oil					0.03					0.01
(for Alberta)					(0.05)					(0.12)
Futures Oil Prices					-0.00					0.00
(Standard Deviation,					(0.02)					(0.05)
(for Alberta)										
Employment Region		✓	✓	✓	✓		✓	✓	✓	✓
Manufacturer			✓	✓	✓			✓	✓	✓
Product Category × Store × Calendar Date				✓	✓				✓	✓
Observations	7602387	7602387	7602385	7601909	7437442	4367509	4367509	4367509	4367002	4349601
R^2	0.00	0.03	0.04	0.06	0.06	-0.18	-0.18	-0.13	-0.00	-0.00
Naive Own-Price	-0.11***	-0.14***	-0.01	-0.18***	-0.18***	-0.20***	-0.22***	-0.10***	-0.77***	-0.77***
Estimates	(0.02)	(0.03)	(0.01)	(0.02)	(0.02)	(0.04)	(0.05)	(0.02)	(0.14)	(0.14)

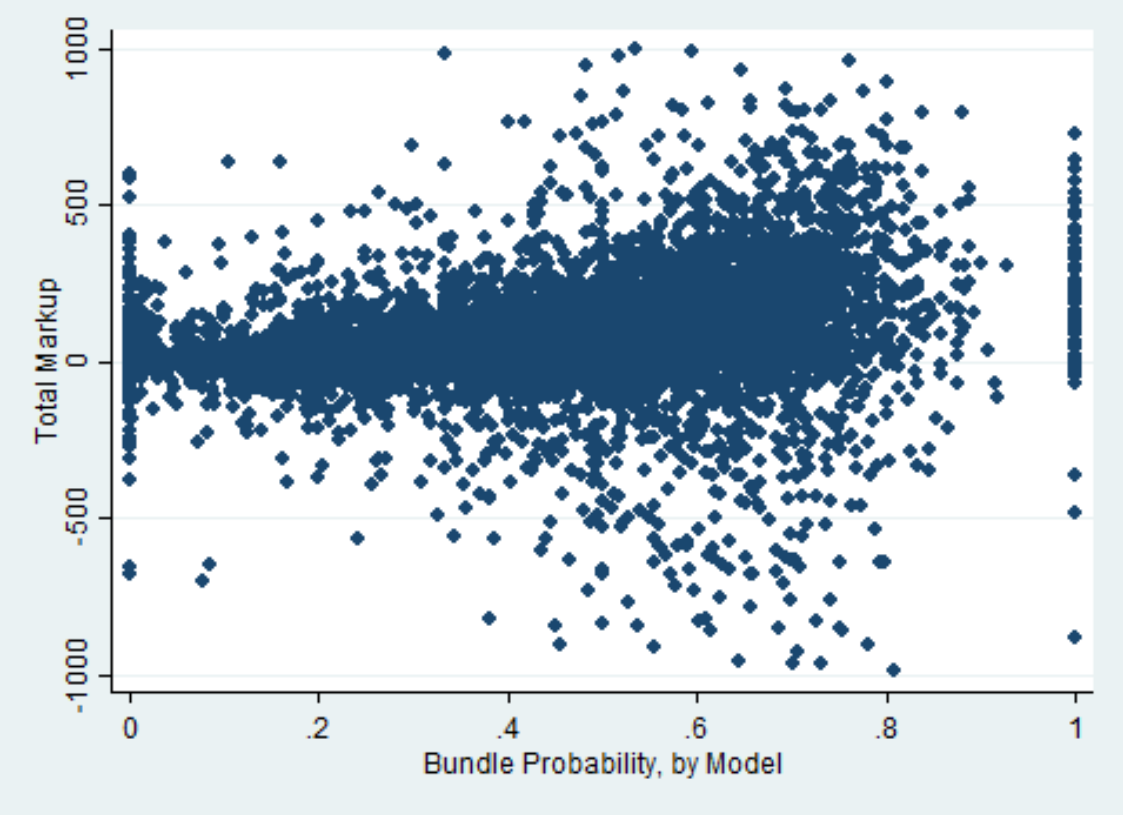

Figure 6: Total Markup and Bundle Probabilities. The unit of observation is the model of the base good. Total Markup is the markup on the base good when just the base good is sold. It is the markup up on the base good and the extended warranty when they are sold together.

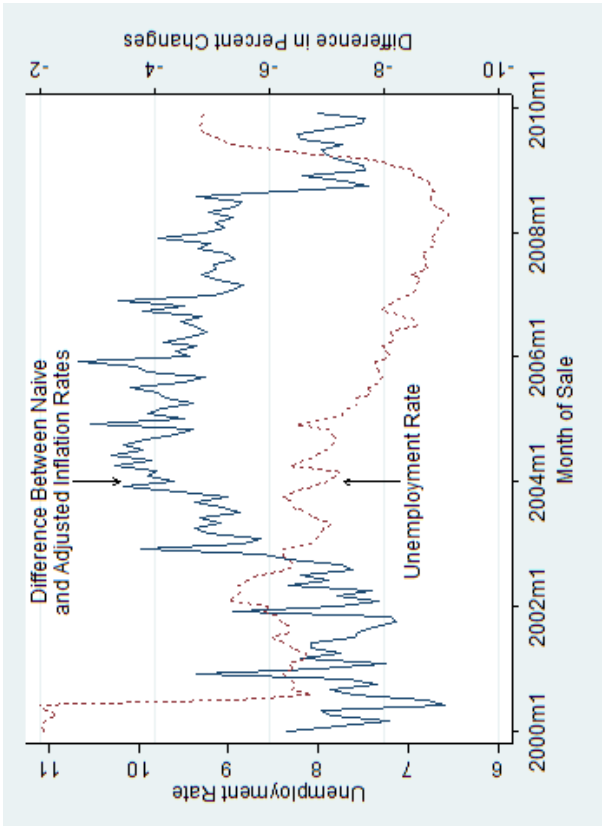

(a) Base Good Price

(b) Base Good Quantity

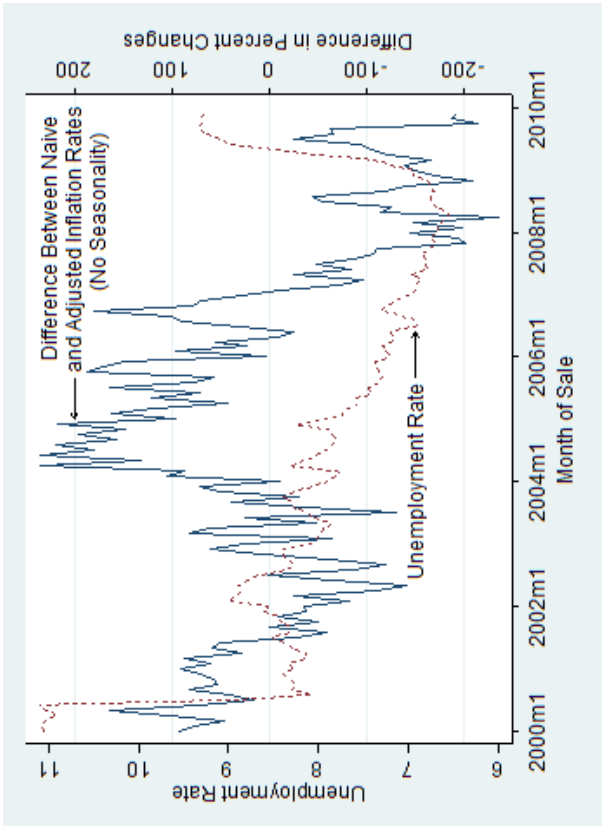

(c) Extended Warranty Price

(d) Take Up of Extended Warranty

Figure 7: Naive Inflation Rates.

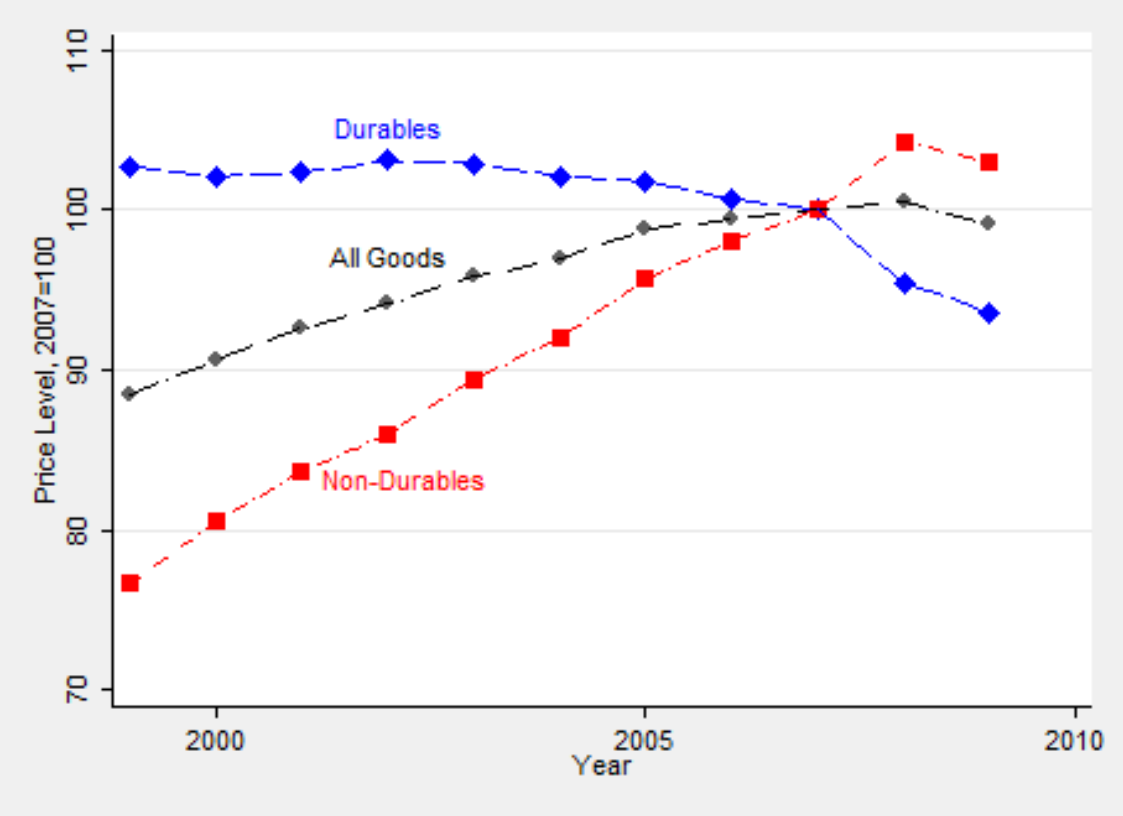

Figure 8: Price Movements and the Business Cycle

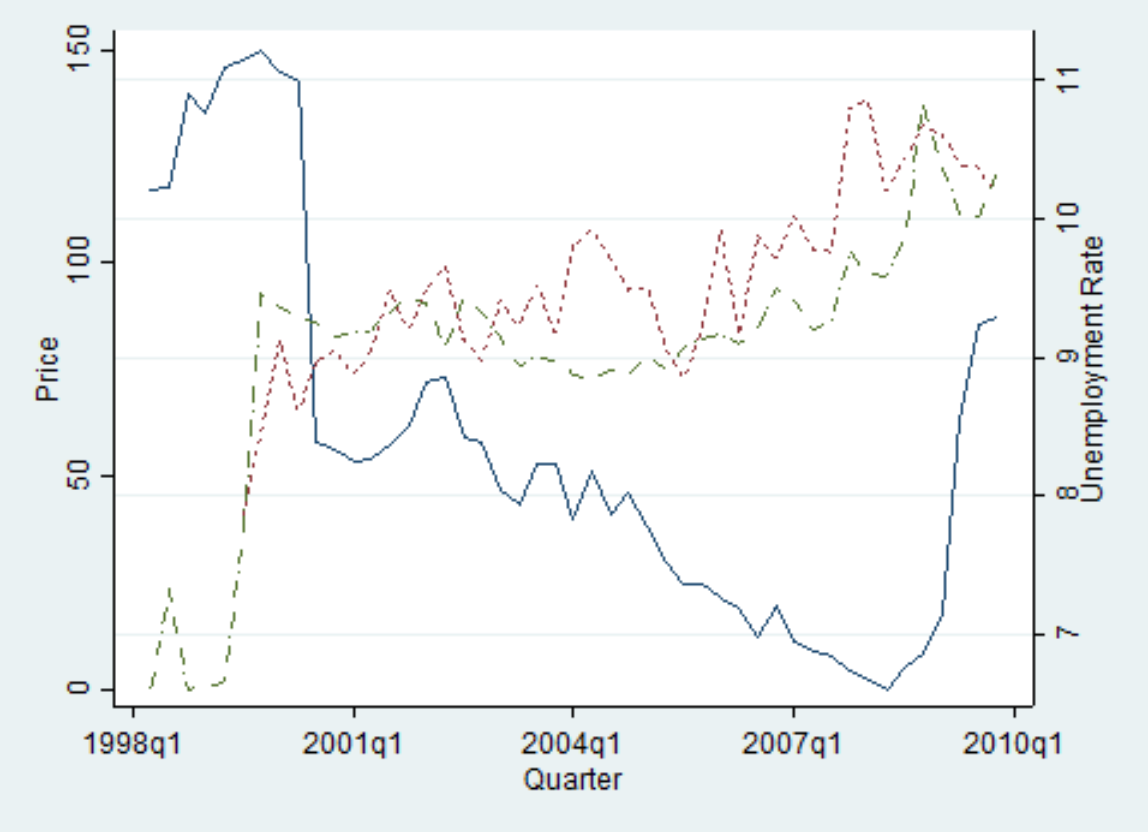

(a) Extended Warranty Price

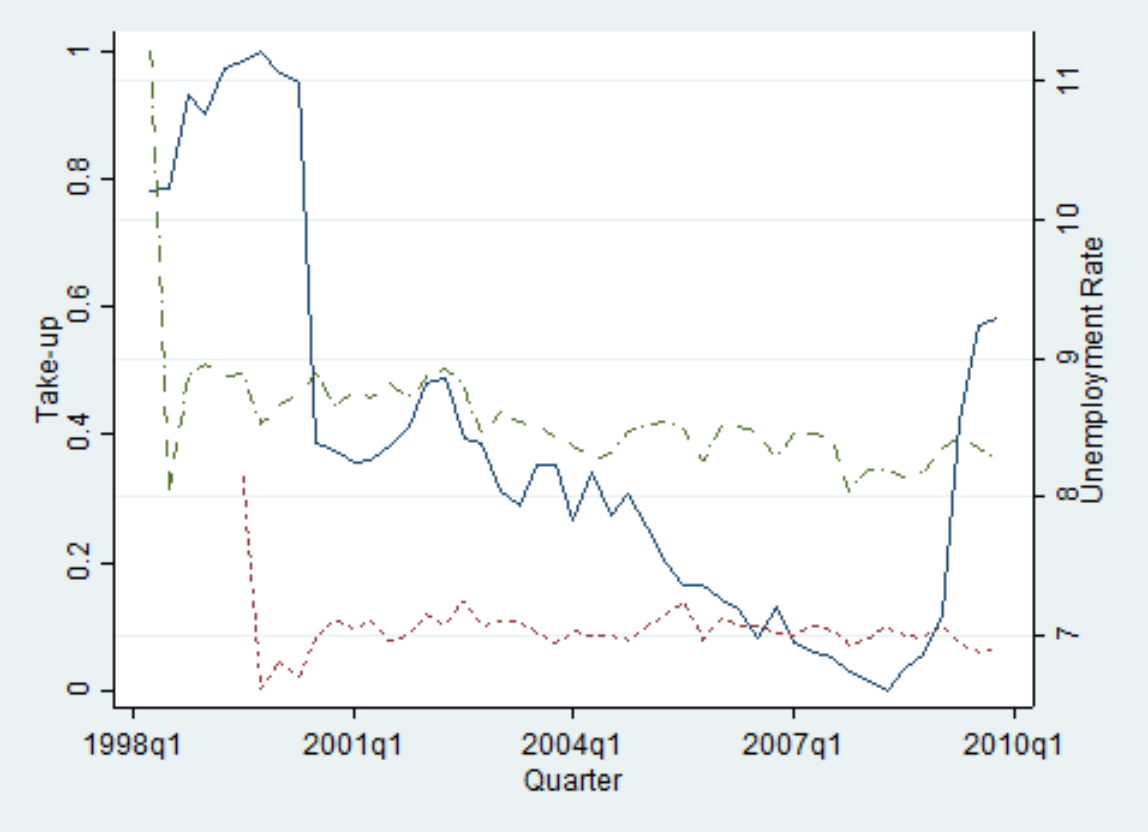

(b) Take Up of Extended Warranty

Figure 9: Online versus In-Store. Green dash-dot line is for purchases made at brick-and-mortar stores. Red dashed lines for purchases made online. Blue line for the unemployment rate.

Table 21: Expected Durability of the Base Good. The dependent variable is the number of days (in logs) covered by the manufacturer's warranty. The dependent variable in the top panel is the number of days where labor costs are covered by the manufacturer. The dependent variable in the bottom is the number of days where the manufacturer covers parts. The left panel presents estimates for the minimum days the manufacturer covers (the coverage period for some parts of the good is longer than for the others). The right panel presents estimates for the maximum days covered. Standard errors are clustered on the employment insurance region. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Warranty Coverage from Manufacturer									
	Minimum Coverage for Labor					Maximum Coverage for Labor				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Unemployment Rate (logs)	0.0127***	-0.0028**	-0.0030***	-0.0027***	-0.0002	0.0376***	-0.0003	-0.0049***	-0.0030***	-0.0001
	(0.0028)	(0.0012)	(0.0007)	(0.0007)	(0.0005)	(0.0053)	(0.0017)	(0.0012)	(0.0011)	(0.0007)
Observations	6471900	6235828	6235822	6235822	6114079	6471900	6235828	6235822	6235822	6114079
R^2	0.0004	0.0428	0.6574	0.6575	0.7848	0.0011	0.0616	0.5524	0.5526	0.7974
	Minimum Coverage for Parts					Maximum Coverage for Parts				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Unemployment Rate (logs)	0.0197***	-0.0027**	-0.0028***	-0.0026***	-0.0009*	0.0864***	0.0051	-0.0041***	-0.0029*	-0.0020*
	(0.0031)	(0.0012)	(0.0007)	(0.0007)	(0.0005)	(0.0126)	(0.0034)	(0.0015)	(0.0015)	(0.0010)
Observations	6508821	6271653	6271647	6271647	6149316	6508821	6271653	6271647	6271647	6149316
R^2	0.0008	0.0608	0.6594	0.6595	0.7826	0.0023	0.1100	0.6719	0.6721	0.8270
Consumer FSA		✓	✓	✓	✓		✓	✓	✓	✓
Calendar Date		✓	✓	✓			✓	✓	✓	
Manufacturer			✓	✓	✓			✓	✓	✓
Product Category			✓	✓				✓	✓	
Store				✓				✓	✓	
Category × Store × Month × Year					✓			✓	✓	✓

A Appendix

A.1. Proximity to the Change in Economic Activity. We examine the impact of the economic activity around the neighbourhood of each store. If economic activity has a direct impact on stores and salespersons, then it should be the case that more local shocks have a larger impact on than ones from far away. To do this, we consider restricted samples with consumers who live within x miles (x in 5,10,20,50,100) of the store where they bought the good. Estimates based on these samples are found in Table 23.

The estimates reinforce the conclusion that economic activity elicits a direct response from the supply side. They show that prices and take up are more responsive to localized measures of economic activity.[38] It helps that we can condition on fixed effects for the FSA of the customer. This accounts for the differences in consumer sensitivity to prices.

[38]Table 23 has a second advantage. Namely it excludes consumers from far away, a sale might be assigned to a store that is far away because the store consumer visited ran out of inventory.

B Appendix - Tables and Figures

Table 22: Expected Markups on the Add on, by Product. Products are defined by the retailer. The statistics are ordered by expected markups.

		Base Good		Extended Warranty			
Base Good	Count	Price Paid	Take Up	Price Paid	Expected Cost	Price-Cost	Expected Markup
External DR/Burner	77	215.61	0.08	29.17	0.00	29.17	1.00
WebTV Receiver	71	222.89	0.27	29.62	0.00	29.62	1.00
Copier	71	204.56	0.30	48.09	0.00	48.09	1.00
Cell Phone	2606	88.16	0.01	17.99	0.00	17.99	1.00
Modem	4	491.54	0.50	24.99	0.00	24.99	1.00
Black and White Mini-TV	2	28.00	1.00	9.99	0.00	9.99	1.00
Micro-Cassette	1288	99.64	0.11	30.94	0.00	30.94	1.00
Clock Radio	11149	24.16	0.01	20.75	0.00	20.75	1.00
Vacuum Electric Hose	112	131.37	0.01	39.99	0.00	39.99	1.00
Vacuum Powerhead	331	239.12	0.01	39.99	0.00	39.99	1.00
WebTV Keyboard	64	78.21	0.06	23.10	0.00	23.10	1.00
Icemaker	5401	118.30	0.00	19.99	0.00	19.99	1.00
Amplifier	29	122.12	0.17	35.96	0.00	35.96	1.00
Sewing Machine	1669	219.29	0.01	29.10	0.00	29.10	1.00
Hi-Fi VCR/Receiver	21	321.01	0.38	42.49	0.00	42.49	1.00
Humidifier	388	65.93	0.06	32.90	0.00	32.90	1.00
C.A.S. Equalizer	246	230.70	0.35	33.87	0.00	33.87	1.00
Garburator	686	192.10	0.08	31.51	0.00	31.51	1.00
Converter	138	319.14	0.27	39.88	0.00	39.88	1.00
Satellite Radio	1172	45.12	0.06	29.56	0.00	29.56	1.00
C.A.S. Recorder/Amp/Tuner	11	346.40	0.27	79.32	0.00	79.32	1.00
Small Appliance	1587	153.02	0.00	24.28	0.00	24.28	1.00
Headphone	3093	36.24	0.01	25.87	0.00	25.87	1.00
DVD Player/Receiver	154	599.24	0.20	43.62	0.00	43.62	1.00
Coffee Maker	453	111.36	0.01	11.49	0.00	11.49	1.00
Computer Speakers	983	18.81	0.01	31.22	0.00	31.22	1.00
N/A	303	337.49	0.39	63.41	0.00	63.41	1.00
Digital Recorder	98	183.92	0.32	25.51	0.00	25.51	1.00
Breadmaker	30	105.25	0.17	9.99	0.00	9.99	1.00
FRS Radio	1094	68.49	0.04	14.57	0.00	14.57	1.00
Equalizer	753	161.29	0.23	32.68	0.08	32.60	1.00
Cooktop Cartridge	1631	271.19	0.10	38.41	0.15	38.26	1.00
Gas Range	306	2614.89	0.67	144.37	0.63	143.74	1.00
MP3/IPOD Player	31100	118.58	0.07	20.83	0.10	20.74	1.00
Colour TV, 15-17 inch	1802	146.16	0.23	29.18	0.14	29.05	1.00
Fireplace	21389	521.22	0.17	48.41	0.33	48.08	0.99
Speaker	75741	388.31	0.28	41.15	0.33	40.82	0.99
Colour TV, 14 inch or less	39484	120.68	0.18	25.36	0.22	25.14	0.99
Digital Picture Frame	2952	101.82	0.04	30.48	0.28	30.20	0.99
Colour TV, 20-21 inch	72801	209.17	0.10	29.45	0.27	29.18	0.99
Digital Satellite Receiver	85122	203.27	0.12	35.78	0.43	35.35	0.99
Colour 20 inch Monitor	48265	223.74	0.11	44.78	0.56	44.21	0.99
Black and White TV/Radio/Deck	5122	47.68	0.15	12.20	0.21	11.99	0.98
Gas Wall Oven	62	933.29	0.39	71.78	1.24	70.54	0.98

Continued on next page

Table 22 – continued from previous page

		Base Good		Extended Warranty			
Base Good	Observations	Price Paid	Take Up	Price Paid	Expected Cost	Price-Cost	Expected Markup
Turntable	497	157.31	0.24	31.53	0.57	30.96	0.98
Fax Machine	3286	193.26	0.22	55.18	0.99	54.19	0.98
Colour TV, 20-21 inch	31974	255.67	0.29	46.41	0.85	45.57	0.98
Scanner	8469	114.45	0.17	34.90	0.67	34.23	0.98
Sealed Speaker	7483	333.90	0.29	33.67	0.72	32.95	0.98
Bookshelf Speaker	3622	217.11	0.23	41.33	0.95	40.38	0.98
Colour TV, 21-29 inch	83789	349.09	0.13	75.35	1.78	73.57	0.98
Speaker with Amplifier	390	619.15	0.21	66.88	1.72	65.16	0.97
Compact Refrigerator	57361	210.78	0.20	30.57	0.81	29.76	0.97
Satellite Dish Antenna	8319	181.05	0.14	37.78	1.00	36.78	0.97
N/A	438	252.76	0.20	68.87	2.05	66.82	0.97
Video Projector	64	1578.27	0.48	165.05	5.05	159.99	0.97
Colour TV, 6 inch or less	582	184.69	0.07	24.33	0.78	23.55	0.97
Printer	48259	124.62	0.19	35.83	1.28	34.55	0.96
Hoodfan/Vent	33314	250.82	0.21	18.55	0.70	17.86	0.96
GPS Unit	6488	236.03	0.15	40.45	1.53	38.92	0.96
Freezer	210451	361.63	0.36	50.78	2.03	48.75	0.96
Walkman	3103	113.84	0.23	23.66	0.97	22.69	0.96
Portable Air Conditioner	11651	395.33	0.33	57.08	2.63	54.44	0.95
C.A.S. Cassette	2141	192.08	0.53	33.54	1.57	31.97	0.95
DVD Player	627496	133.09	0.19	30.45	1.50	28.96	0.95
TV/VCR Combination	29072	333.74	0.28	66.82	3.43	63.40	0.95
Digital Camera	220612	189.11	0.19	40.95	2.12	38.83	0.95
Window Air Conditioner	50773	269.59	0.29	50.93	2.71	48.23	0.95
Car Stereo	3752	166.52	0.34	30.43	1.63	28.80	0.95
Cassette Deck	6064	223.33	0.48	35.89	1.96	33.94	0.95
LCD TV	496076	1079.36	0.38	176.37	9.96	166.41	0.94
Digital Camcorder	13025	440.22	0.38	62.08	3.58	58.49	0.94
CD Player	27007	199.41	0.42	33.44	2.03	31.41	0.94
Barbeque	54983	243.47	0.27	25.85	1.57	24.28	0.94
Coin-Op Dryer	737	620.18	0.12	55.75	3.42	52.33	0.94
Telephone	29568	119.84	0.19	26.21	1.73	24.48	0.93
SubWoofer	1936	265.62	0.29	35.14	2.33	32.81	0.93
DVD/VCR Combination	73826	166.25	0.29	29.03	2.00	27.04	0.93
Vacuum Cleaner	81216	218.63	0.19	36.58	2.54	34.04	0.93
Wine Cooler	5790	271.73	0.24	30.80	2.26	28.54	0.93
Steam Cleaner	3824	249.43	0.34	35.81	2.66	33.15	0.93
Coin-Op Washer	887	760.21	0.11	86.19	6.65	79.54	0.92
Dehumidifier	1819	203.71	0.24	31.26	2.48	28.79	0.92
VCR	133897	135.71	0.30	29.16	2.43	26.72	0.92
Printer/Fax/Copier	97	227.14	0.39	43.48	3.68	39.80	0.92
Portable CD Player	26126	103.67	0.21	23.88	2.02	21.86	0.92
Home Theater Receiver/Speaker	255000	446.46	0.31	55.79	4.73	51.05	0.92
Receiver	53860	387.55	0.48	39.06	3.52	35.54	0.91
Colour 21-29 inch Monitor	198087	463.94	0.39	80.18	7.79	72.38	0.90
Portable Stereo	31414	105.33	0.24	25.22	2.60	22.62	0.90
TV/DVD/VCR Combination	30197	412.60	0.28	72.82	7.69	65.14	0.89
Wringer Washer	100	467.11	0.36	80.26	10.12	70.14	0.87
LCD Monitor	5709	459.74	0.36	25.34	3.33	22.00	0.87
Plasma TV	128752	1505.09	0.44	206.44	28.81	177.63	0.86

Continued on next page

Table 22 – continued from previous page

		Base Good		Extended Warranty			
Base Good	Observations	Price Paid	Take Up	Price Paid	Expected Cost	Price-Cost	Expected Markup
Microwave Oven	329525	207.42	0.23	34.06	5.00	29.06	0.85
Music System	31779	432.98	0.39	53.22	8.34	44.88	0.84
Camcorder	83382	681.87	0.57	81.70	13.63	68.07	0.83
Compact Washer	524	667.64	0.52	78.43	13.54	64.89	0.83
Compact Audio System	54086	229.16	0.40	39.39	6.99	32.40	0.82
Trash Compactor	610	450.17	0.30	41.53	9.05	32.48	0.78
Front Load Washer	192379	960.10	0.63	140.75	33.80	106.95	0.76
Palm Organizer	836	297.93	0.19	44.53	12.25	32.28	0.72
Answering Machine	457	97.67	0.22	22.41	6.41	16.00	0.71
Portable Dishwasher	20590	500.76	0.48	81.74	23.85	57.89	0.71
Colour TV, 30-37 inch	143536	1014.22	0.55	120.33	35.43	84.90	0.71
Mini CD Recorder	472	284.38	0.46	52.21	15.55	36.66	0.70
Gas Cooktop	3778	1054.40	0.46	60.33	17.97	42.36	0.70
Noteboook	8147	1295.25	0.30	220.84	74.28	146.56	0.66
C.A.S. CD Player	1136	331.11	0.60	42.76	15.34	27.42	0.64
Electric Dryer	475560	516.52	0.49	37.29	13.41	23.88	0.64
Electric Cooktop	8492	664.13	0.44	54.87	20.34	34.53	0.63
Hard Drive CPU	31883	1041.07	0.58	158.70	62.64	96.05	0.61
Stack Washer/Dryer	21352	1050.77	0.41	130.81	55.09	75.73	0.58
Computer/CPU	18443	1188.18	0.60	174.86	77.32	97.54	0.56
Cooktop-Downdraft	1576	1031.78	0.51	74.24	33.29	40.96	0.55
Gas Dryer	15920	649.66	0.54	45.10	22.48	22.62	0.50
Automatic Washer	370333	611.27	0.49	107.27	57.07	50.19	0.47
Electric Range	412605	808.25	0.46	71.98	38.61	33.37	0.46
Built-in Dishwasher	331419	555.79	0.43	70.31	37.86	32.45	0.46
Refrigerator	579798	1000.48	0.49	80.69	45.86	34.83	0.43
Cooktop-Solid Element	817	989.68	0.62	61.28	35.82	25.46	0.42
Colour Monitor	40023	338.98	0.52	4.67	2.91	1.76	0.38
Range-Mod/Downdraft	2042	2067.16	0.65	88.13	61.73	26.39	0.30
Projection Colour TV	259836	2140.08	0.64	212.79	151.51	61.28	0.29
LCD Projector	477	1953.86	0.44	183.32	137.36	45.96	0.25
Digital Colour TV	274	1015.03	0.68	121.65	97.22	24.43	0.20
Colour TV, 40-42 inch	204	4336.01	0.56	353.78	290.05	63.73	0.18
Electric Wall Oven	12674	1393.49	0.58	73.62	61.15	12.47	0.17
Compact Dryer	6775	685.97	0.59	29.27	27.00	2.27	0.08
DLP Projection TV	11978	2136.66	0.68	209.49	200.87	8.62	0.04
Range-Solid Element	31955	1291.44	0.58	73.14	79.13	-5.99	-0.08
Wall Oven/Microwave	498	3541.03	0.77	90.78	270.16	-179.39	-1.98
Video Player Only	3	48.67	0.00	.	0.00	.	.
Espresso Maker	102	73.98	0.00	.	0.00	.	.
Microphone	3	154.50	0.00	.	0.00	.	.
Warming Drawer	21	795.09	0.00	.	0.00	.	.
Range-Modular	3	1650.99	0.00	.	0.00	.	.
Vacuum Accessory	27	67.08	0.00	.	0.00	.	.
Video Printer	38	192.14	0.00	.	0.00	.	.
20-29 inch TV for Hotel	16	436.60	0.00	.	0.00	.	.
Black and White TV with Radio	26	40.46	0.00	.	0.00	.	.

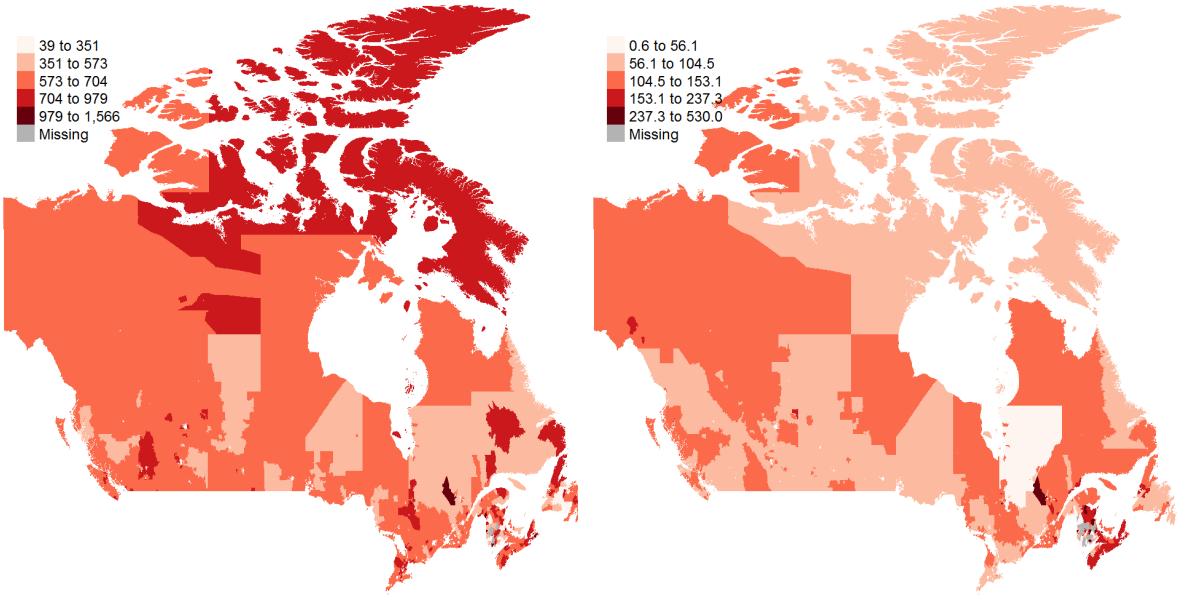

(a) Base Good Price **(b)** Extended Warranty Price

Figure 10: Prices (Unweighted) over Space.

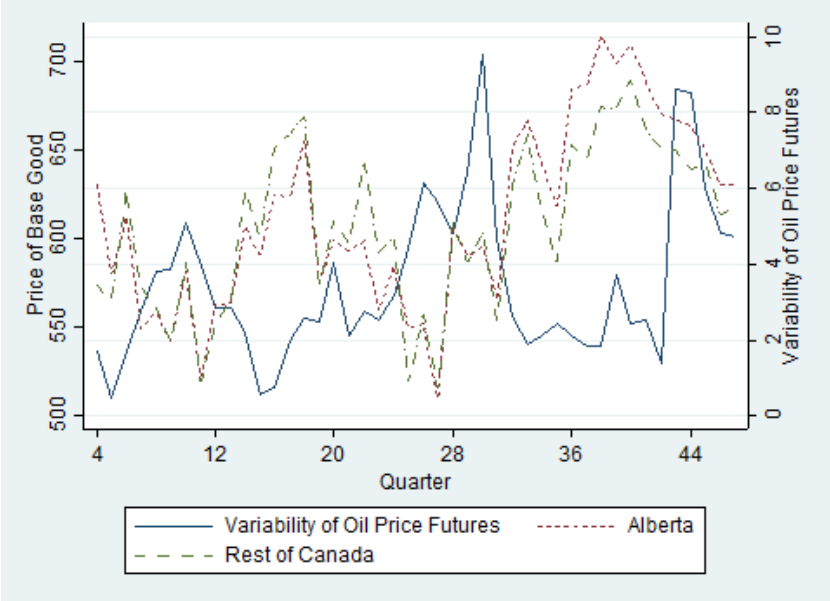

(a) Base Good Price

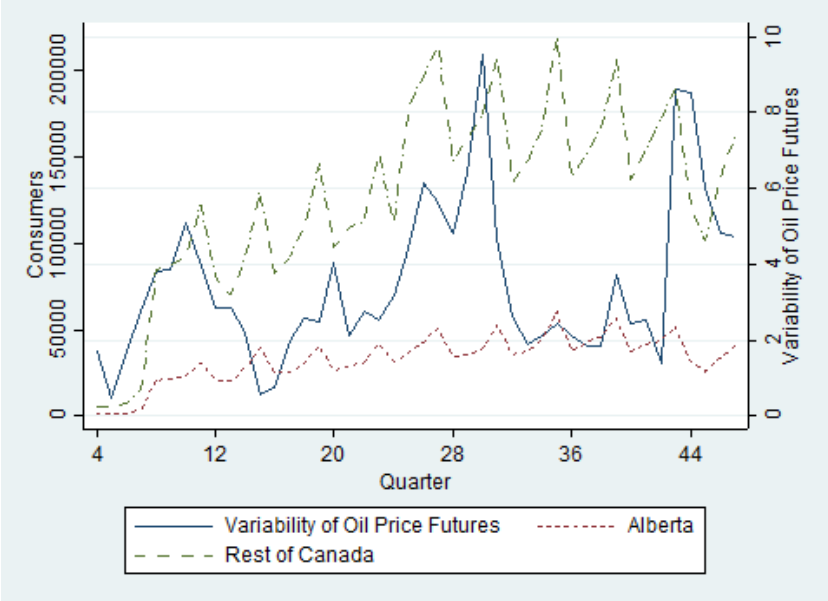

(b) Base Good Quantity

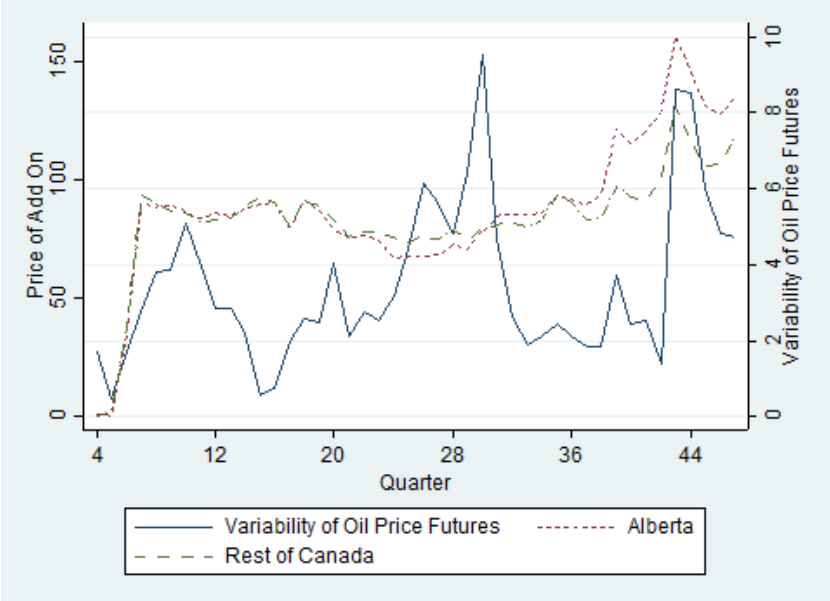

(c) Add-on Price

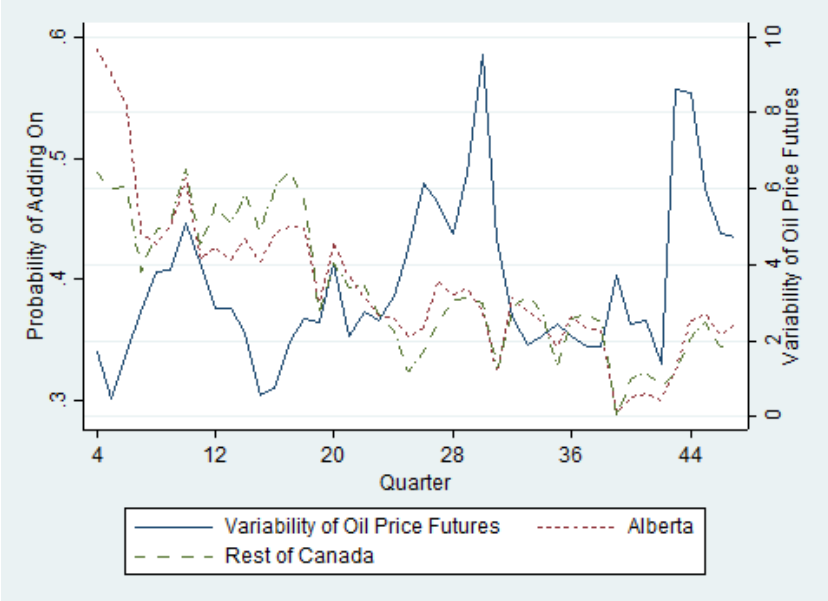

(d) Warranty Probability

Figure 11: Income Shocks and the Prices and Quantities of the Base and Add-on Good.

Table 23: Stores and Locations. The point estimates are from regressions of the outcome on the unemployment rate for consumers who live within x miles (x in 5,10,20,50,100) of a store. Prices and the unemployment rate are in logarithms. Standard errors are clustered along two dimensions, the calendar date and the neighbourhood of the consumer. They are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Radius around Store Location				
	5 miles	10 miles	20 miles	50 miles	100 miles
Impact of Base Good Price	-0.032***	-0.024***	-0.020***	-0.018***	-0.018***
	(0.006)	(0.005)	(0.004)	(0.004)	(0.004)
Observations	3321562	4694060	5493340	6042149	6250409
R^2	0.789	0.790	0.789	0.789	0.788
Impact of Extended Warranty Price	-0.240***	-0.255***	-0.232***	-0.197***	-0.190***
	(0.037)	(0.029)	(0.026)	(0.024)	(0.024)
Observations	1248781	1781337	2099189	2312096	2387977
R^2	0.166	0.160	0.159	0.159	0.160
Impact of Extended Warranty Probability	0.012**	0.015***	0.016***	0.012***	0.012***
	(0.006)	(0.005)	(0.004)	(0.004)	(0.004)
Observations	3327310	4702458	5503343	6053233	6261892
R^2	0.150	0.150	0.151	0.151	0.151