TI 2020-060/VI
Tinbergen Institute Discussion Paper

Cyclicality of Add-on Pricing: Evidence from Extended Warranties

Branko Boskovic[1]
Sacha Kapoor[2]
Agnieszka Markiewicz[2]
Barry Scholnick[1]

[1] University of Alberta
[2] Erasmus University Rotterdam

Add-on Pricing over Regional Business Cycles: Evidence from Extended Warranties

Branko Bošković Sacha Kapoor Agnieszka Markiewicz Barry Scholnick[a]

July 17, 2020

Abstract

What explains apparent price rigidity over the business cycle? This paper shows that price fluctuations may be hidden in prices of add ons. Using 10 years of extended warranty data from a nationwide Canadian retailer, we show that extended warranty prices respond strongly to changes in local economic activity whereas prices of underlying durable goods do not. The procyclicality is driven by a shift in price setting behavior, where local stores use extended warranty discounts to make base durable demand less price elastic and, as a result, increase sales of durable goods. Discounts on extended warranties were especially sharp during the Great Recession.

JEL: E30, L81

Keywords: Add-on pricing, business cycles

[a]Bošković (boskovic@ualberta.ca) and Scholnick (barry.scholnick@ualberta.ca): Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta T6G 2R6, Canada. Kapoor (kapoor@ese.eur.nl) and Markiewicz (markiewicz@ese.eur.nl): Erasmus University Rotterdam, Erasmus School of Economics H09-22, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands. We thank Mark Bils, Maarten Bosker, Robert Dur, Etienne Gagnon, Robin Lumsdaine, Andreas Pick, Lorenzo Pozzi, Giorgio Primiceri, Valerie Ramey, Vincent Rebeyrol, as well as participants at ASSA Meetings in Philadelphia, 49th Konstanz Seminar on Monetary Theory and Monetary Policy and seminar participants at various universities for comments that helped improve the paper.

1 Introduction

Add ons improve the quality of a basic good or service vertically and "their prices are not advertised and would be costly or difficult to learn before one arrives at the point of sale" [Ellison, 2005]. They are pervasive in retail. Airlines offer add-on or ancillary services such as check-in baggage, or extra-leg room above a basic seat where add-on prices are revealed only after the plane ticket was purchased. Hotels propose breakfast or internet access over a standard room and prices of these ancillary services are known only after the hotel room is occupied. Retail banks offer overdraft credit beyond basic deposit services and make the overdraft price clear only after the overdraft option is used. The hidden nature of add-on prices and the associated inconvenience cost of learning them elsewhere enable firms to earn positive profits in equilibrium.[1] This capacity to use add-on prices to extract rents on the marginal transaction provides retailers with a natural means for adjusting to business cycle fluctuations. The hidden nature of add-on prices has made it difficult however to learn whether this is the case.

We conduct the first large-scale investigation of add-on price setting and dynamics using 10 years of confidential transactions data from a nationwide Canadian retailer of household durables. The data facilitates a large-scale investigation for several reasons. First, it includes detailed price information on a textbook example of an add on, extended warranty, a service that can be added on to most of the 35000 base durable goods the retailer sells.[2] Second, the data covers every one of the more than 6 million customer purchases, by more than 3 million customers (approximately 10 percent of the population in 2001 and 26 percent of the number of households), that took place all across Canada between December of 1999

[1] The theoretical literature on add-on prices tries to explain why add-ons and their hidden prices exist. See [Ellison, 2005] and [Gabaix and Laibson, 2006] for two prominent examples.

[2] In other contexts, a large-scale analysis is more difficult because the same good can have several different add ons, or because different goods can have different add ons, making it extremely difficult to identify add-on goods and to link them with the associated base good.

and December 2009. This lets us exploit comprehensive regional and time variation in economic activity to measure durable and extended warranty price fluctuations over regional business cycles. Third, the data includes rich cost information and the extended warranty prices headquarters suggest to stores. This information facilitates an examination into the microeconomic foundations of price fluctuations.

We show that a one percentage point increase in the regional (sub-provincial) unemployment rate is followed by a $2.15 decrease in the extended warranty price in the next month. It is followed by a cumulative decrease of $6.89, 7.7 percent of the mean extended warranty price, over the following year. We find marginal impacts on the durable goods' prices over similar time horizons. Sharp discounts on extended warranty prices and durable price rigidity support the idea that extended warranty prices are an important but hidden margin of adjustment to local business cycle fluctuations. We exploit heterogeneity in economic activity across Canadian provinces and use variation in the world oil prices to show that these patterns are robust to alternative measures of economic activity, and in particular to a clean shifter of consumer demand.

The patterns of our baseline estimates suggest that extended warranty discounts may be a vehicle for stores to boost durable goods sales in response to depressed consumer demand. To check this explicitly, we estimate demand systems for durables and extended warranties that incorporate in-store discounts on extended warranty prices relative to headquarter suggested prices. We show that 100 percent discounts boost associated base good sales by 3.6 percent. Base good sales increase because the discounts make consumer demand for durables substantially less elastic to the transaction price. Specifically, a 100 percent discount on extended warranties reduces the transaction price elasticity of durables demand from 0.271 to 0.218.

We show that extreme discounts were especially common during the Great Recession, consistent with add ons being an important margin of adjustment during periods of severe

economic distress. We show further that the argument that extended warranty prices are an important but hidden margin of adjustment is consistent with trends in aggregate markups for both base durables and extended warranties. Finally, we use the data from the U.S. airlines sector to show that add ons were also critical in other industries during the Great Recession.

Similar to Chevalier, Kashyap, and Rossi [2003], Bils and Klenow [2004], Nakamura and Steinsson [2008], Hosken and Reiffen [2004], Kehoe and Midrigan [2015], Anderson et al. [2017], we study the role of temporary sales for characterizing the degree of price flexibility. Our case highlights the importance of unadvertised and transaction-to-transaction sales. To the extent that the extended warranty price variation reflects costly negotiations with consumers, this result further complements the findings of Zbaracki et al. [2004] who investigate price setting behavior of a large U.S. industrial manufacturer. They find that customer costs related to communication and negotiation exceed price adjustments or menu costs by a factor of 20.

Our paper contributes to a broader literature that examines the speed with which prices adjust to economic shocks. Similar to our paper, recent work exploits regional variation to assess the speed of price adjustment. Stroebel and Vavra [2019] use detailed micro dataset to document the causal response of local retail prices to changes in local house prices, finding elasticities of 15-20 percent across housing booms and busts. Beraja, Hurst, and Ospina [2018] combine household and retail scanner data to document a strong relationship between the unemployment rate and real wage growth.

Mounting empirical evidence indicates that households change their shopping behavior over the business cycle. Aguiar, Hurst, and Karabarbounis [2013] argue that time spent on shopping increases during recessions, and Krueger and Mueller [2010] and Nevo and Wong [2019] document a rise in several measures of shopping intensity during the Great Recession. Coibon, Gorodnichenko, and Hong [2019] show that, during recessions, consumers switch to

low-cost stores.[3] Our focus instead is on price setting behavior and on demonstrating how, during the Great Recession, the retailer was exploiting extended warranties to adjust prices. To this end, we propose a new mechanism of adjustment by the retailer and also show that consumers responded to this mechanism by increasing purchases of goods with discounted extended warranties.

The remainder of the paper is organized as follows. Section 2 describes our dataset. Section 3 studies the pricing responses to changes in local unemployment rates, our baseline measure of local economic activity. Section 4 shows that our baseline estimates are qualitatively robust to a relatively pure measure of a demand shock. In Section 5 we estimate a demand system to identify the mechanism driving the observed procyclicality of extended warranty prices. Section 6 examines the role of add ons in the airlines sector. Section 7 concludes with a summary of the findings and robustness checks found in the online appendix.

2 Data

2.1. Base Durables and Extended Warranties. Our analysis is based on the data of a nationwide Canadian retail chain which specializes in the sale of household durables, especially home appliances and consumer electronics. The retail chain is among the top 4 in terms of market share in its relative subsector [Industry Canada, 2013]. Almost all goods are offered with the option to extend the lifetime of the warranty beyond what the manufacturer offers. We follow the retailer, and other retailers, in calling this extension an extended warranty.

The data covers the universe of transactions between January 1 1999 and December 31 2009, involving more than 6 million transactions, more than 3 million consumers, nearly

[3]See Gagnon, López-Salido, and Sockin [2017] for a direct counterargument to this hypothesis.

35,000 products, and around 270 stores. The data includes transaction prices for durables and extended warranties, whether an extended warranty was purchased, the suggested warranty price, as well as the cost of servicing claims made under the extended warranty. The suggested warranty price is the benchmark price headquarters sets for stores.

The chain has corporate and franchise stores. Franchises purchase durables at cost (the price of the manufacturer) plus the cost of keeping the durable good in inventory. We observe these transfer prices and use them to construct costs for all goods sold at both corporate and franchise stores. Our dataset also includes the claim costs for all claims made within our sample. We observe claims and thus *ex post* extended warranty costs until about 2007. Later we use these claims to construct *ex ante* extended warranty costs for each manufacturer and product.

Basic summary statistics are found in Table 1. The top panel shows average price and cost for the base durable. Customers pay 628.9 dollars for a durable that costs the retailer 549.7 dollars. These costs include the manufacturer's price, salesperson commissions for base durable sales (4 percent), royalty costs to the chain if the store is a franchise (4 percent), inventory and marketing and advertisement costs (2.5 percent). Extended warranty costs include commissions on extended warranty sales (15 percent), potential royalties to the retailer (4 percent), and the costs from servicing claims on the extended warranty. Servicing covers 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. Although the prices of extended warranties are almost never advertised, the retailer calculates marketing and advertisement costs of 2.5 percent of total revenue, including that from the sales of extended warranties. We do the same.

The lower panels of Table 1 reports statistics for extended warranties. The suggested price exceeds the realized warranty price by about 50 percent. Customers extend the warranty 40 percent of the time and pay 89.4 dollars on average to do so. The costs of the claimed

warranties amount to 61.6 dollars on average implying a large average markup of 31 percent relative to the base-good markup of 13 percent. Note that this large markup is driven by the low frequency of claims (12 percent). The average cost of the realized claims is quite high, amounting to 488.6 dollars.

To learn the extended warranty price, consumers must speak directly with store representatives. Moreover, once a consumer has visited one of the stores in the chain, it becomes costly for them to visit a competing retailer. The stores are usually housed in stand-alone buildings and located in regions with sprawl so that consumers must travel by car to learn 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 warranty.

The commission structure 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, whereas the commission on the base good is 4 percent.[4] The commission structure also reinforces that notion that both extended warranty and base durable pricing are decentralized.

Note that the average realized extended warranty price is more than 45 dollars lower than the suggested price in Table 1. This price wedge highlights the importance of transactions data for our purposes. Most other data sets, like list-price data collected by representatives of a statistical agency or by web-scraping, exclude add-on prices (*c.f.* Cavallo [2018]). The inherent unobservablity of add-on prices therefore makes it generally impossible to draw reliable conclusions about their cyclical properties.[5]

The low claim rate in Table 1 raises questions about whether an extended warranty can

[4]The commission was 15 percent for almost the entirety of our sample, up until April of 2009 when it was reduced to 10 percent.

[5]This argument was supported by our conversations with Statistics Canada and U.S. Bureau of Labour Statistics, wherein both noted that their respective consumer price indices (CPIs) do not include most add ons and extended warranty prices specifically.

Table 1: Summary statistics. Unit of observation is the transaction. All prices and costs are in Canadian (CAD) dollars. If the extended warranty was purchased but no claim was made, the cost is set to zero. The average claim cost is calculated over realized claims and the frequency indicates the number of claims made relative to the total number of purchased extended warranties in our sample. Standard deviations are in parentheses.

Base durable		
	Price	Average cost
	628.91	549.71
	(613.89)	(526.98)
Observations	6538033	4865375

Extended Warranty				
	Price	Average cost	Suggested price	Take up
	89.41	61.57	133.94	0.40
	(93.62)	(261.70)	(109.01)	(0.49)
Observations	2576246	2583078	2585128	6538033

Claims on Extended Warranties		
	Average cost	Claim made
	519.46	0.12
	(583.05)	(0.32)
Observations	306177	2585128

be interpreted as an add on, as it improves base durable quality rarely. It is important to note that add ons are defined by their perceived quality at the point of sale rather than by their realized quality after the purchase is made, however, and that by this token extended warranties are add ons. We discuss this issue further in Online Appendix OA.1, where we use census data to provide direct evidence that extended warranties satisfy the assumptions Ellison [2005] uses to characterize add ons theoretically.

2.2. Raw Price Dynamics. The top left panel of Figure 1 depicts base good price dynamics for Frigidaire freezers in Kingston, Ontario. The bottom left panel shows base good price dynamics for Whirlpool automatic washers in Oshawa, Ontario. The panels on the right plot extended warranty price dynamics for these same product categories. The dotted and solid lines show price dynamics for suggested and realized prices, respectively. Because suggested prices are set by headquarters, the difference between suggested and realized describes the extent to which stores exploit discretion over prices.

The panels on the right show suggested extended warranty prices are higher than realized extended warranty prices most of the time. In fact, in our entire sample, the realized exceed the suggested prices 96.5 % of the time, in line with the mean differences for the entire sample (Table 1). There are several dips in realized extended warranty prices, most notably during the Great Recession, when realized prices were often close or equal to zero. There were no obvious and systematic dips in base-good prices during the recession, although over the entire sample, they exhibit greater month-to-month variation than realized extended warranty prices.

2.3. Local Economic Activity. Our baseline measure of local economic activity is the unemployment rate from Statistics Canada, which is seasonally-adjusted for each of 58 pre-

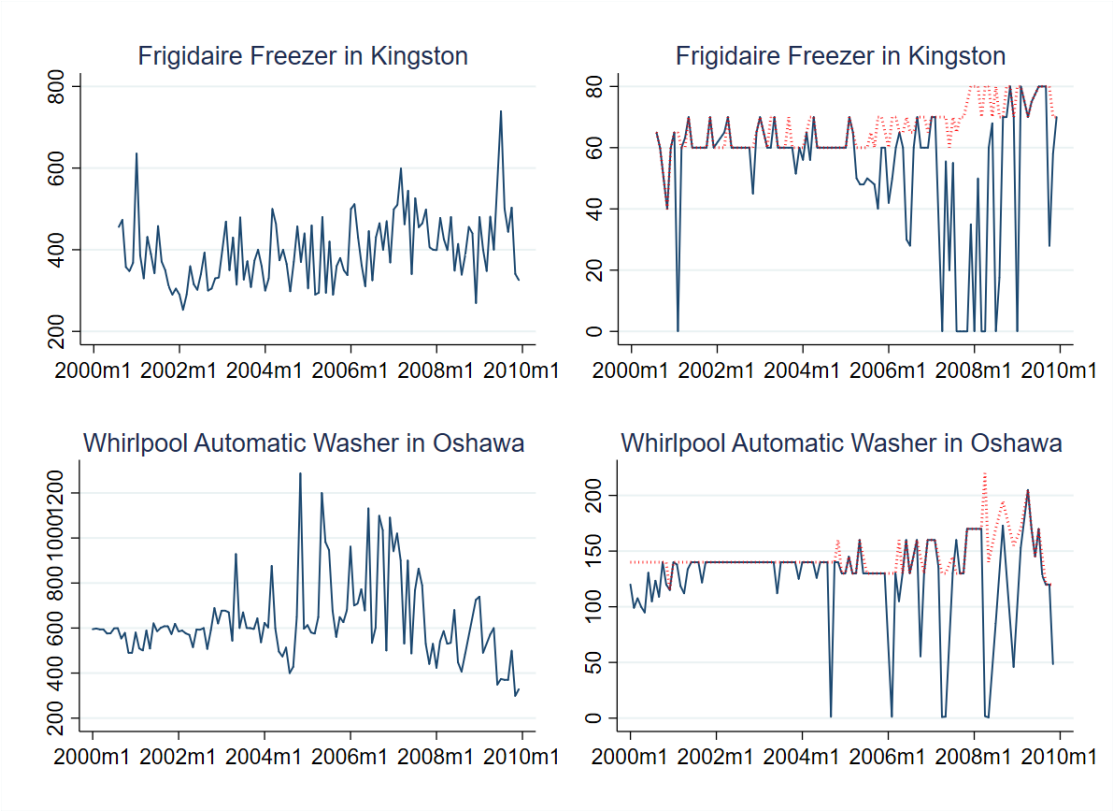

Figure 1: Realized Prices and Suggested Extended Warranty Prices. The left panels plot base durable prices. The right panels plot realized extended warranty prices. The dotted red lines in the right panels plot suggested extended warranty price.

defined administrative areas, referred to as Employment Insurance (EI) regions.[6] Each region contains several cities, and is almost always smaller than any one province. The regional unemployment rates are provided at the monthly frequency and computed as 3-month moving averages. We use the data between January of 2000 (2000M1) and December of 2009 (2009M12) for 55 regions. We exclude Yukon, Northwest Territories, and Nunavut because these territories are sparsely populated and because their unemployment is constant at a rate of 25 percent.

2.4. Unit of Analysis. We aggregate the transaction data to the month, product category and manufacturer for each store in the sample. Product categories are defined by the retailer. There are 140 categories: 15-17 inch color tv, 17-19 inch color tv, fireplace, digital picture frame, gas wall oven, among others. Aggregation produces series that have the same frequency as the unemployment rate data and lets us define lags in the analysis naturally.[7] It also lets us estimate a demand system that includes base durable quantities as well as extended warranties, which is impossible to do at the transactions level because transactions are only observed when a base durable is sold.

3 Cyclicality of Prices

We estimate

$$p_{tscm} = \beta u_{t-\ell,r(s)} + \rho p_{t-\kappa,scm} + \alpha_{r(s)} + \gamma_{tcm} + \varepsilon_{tscm} \tag{1}$$

where p_{tscm} is the median price of the base durable or extended warranty in calendar month t at store s for product category c and manufacturer m. $u_{t-\ell,r(s)}$ is the unemployment rate at time $t-\ell$ in the EI region r, to which store s belongs. We consider alternative lags ℓ of the

[6] The unemployment rate for the EI region is used to determine the EI benefits an unemployed worker is eligible for.

[7] Note that not all the products are sold every day in every store. The aggregation allows us to have an observation every month and to control for persistence in prices.

unemployment rate, at 1, 6, and 12 months, because shifts in the economic environment may affect prices with lags of several months. κ is the lag length on prices. We also compute the long-run (cumulative) impact of the local unemployment rate on prices using 12 lags for both prices and unemployment rates. $\alpha_{r(s)}$ and γ_{tcm} are fixed effects for region and the month-category-manufacturer combination. ε_{tscm} is a random variable reflecting idiosyncratic price changes.

Our interest is in β, which measures the response of the median price to the lagged unemployment rate for the region. Estimates can be interpreted causally if $E[\varepsilon_{tscm}|u_{t-\ell,r(s)}, p_{t-\kappa,scm}, r(s), tcm] = 0$. It is unlikely that lagged warranty prices or base good prices at the level of product category, store, and manufacturer influence local unemployment rates. Unobserved heterogeneity generated, for instance, by the sector of production is captured by EI region fixed effects. Month-category-manufacturer dummies help with unobserved time-varying heterogeneity in the category-manufacturer combination, such as differential propensities for obsolescence or shifting demands across products and manufacturers. The month-category-manufacturer fixed effects capture movements in the national business cycle, implying, among other things, that our cyclicality estimates reflect response to the movements in local economic conditions and partial out the effects of the aggregate business cycle.

Estimates of the baseline specification are reported in Columns (1)-(6) of Table 2. The leftmost and middle panels report estimates of β for extended warranty and base good prices, respectively. Moving left to right within each panel shows how the estimates differ depending on the lag length of the unemployment rate. The lower half of the top panel reports long-run effects, in regressions with 12 lags of the unemployment rate and 12 lags of the dependent variable.

The column (1) estimate implies that a one percentage point increase in the local unemployment rate is followed by a \$2.15 drop in the extended warranty price in the next month. The long-run effect on the extended warranty price is large: a one percentage point increase

Table 2: Pricing and Local Economic Activity. Unit of observation is defined by calendar month, store, product category, and manufacturer. Regressions include fixed effects for month-category-manufacturer combination and employment insurance region, as well as lags of dependent variable. Warranty Price Discount equals the absolute value of the difference between the realized and suggested extended warranty price. Standard errors are in parentheses and clustered on employment insurance region. $***$, $**$, and $*$ denote statistical significance at the 1, 5, and 10 percent levels.

	Warranty Price			Base Good Price			Warranty Price Discount		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Local Unemployment Rate at									
t - 1	-2.152***			-0.711			1.795**		
	(0.595)			(0.641)			(0.565)		
t - 6		-2.161***			-0.627			1.810**	
		(0.604)			(0.769)			(0.570)	
t - 12			-2.186***			-0.462			1.789**
			(0.584)			(0.812)			(0.547)
Long-Run Effect	-6.894***	-6.894***	-6.894***	-0.155	-0.155	-0.155	5.694***	5.694***	5.694***
	(1.950)	(1.950)	(1.950)	(1.380)	(1.380)	(1.380)	(1.926)	(1.926)	(1.926)
Observations	530615	511142	490949	1156770	1123359	1085902	530615	511142	490949
R^2	0.565	0.560	0.554	0.879	0.878	0.877	0.457	0.453	0.450

in the unemployment rate reduces the price of the extended warranty over the following year by $6.89, or 7.7 percent of the mean extended warranty price. The effect sizes are similar if we use 6 and 12 month lags of the unemployment rate. The estimates are all highly significant statistically. By contrast, the analog estimates in Columns (4)-(6) show that the cyclicality coefficients for base-good prices are economically small relatively to the base good price and statistically insignificant at this level of aggregation.

One concern with specifications at this level of aggregation relates to whether our estimates reflect recession-induced consumption shifts from luxury to more basic quality goods within the same manufacturer (*c.f.* Coibon, Gorodnichenko, and Hong [2019]). Such shifts would also generate a negative relationship between prices and unemployment rates. To ensure that our estimates do capture responses of prices, rather than shifting consumption patterns, we estimate the analogous specification at the transaction level while controlling for unobserved heterogeneity for each manufacturer's model in Online Appendix Table OA.4.1. The estimates therein are similar, showing extended warranty price effects that are large relative to the price mean, and base durable price effects that are small relative to the price mean.

The result that the warranty prices are more cyclical than base-good prices could simply reflect decentralized versus centralized pricing of the retailer. We explore this alternative interpretation in Online Appendix OA.2. We exploit the reduction in the commission on extended warranties from 15 to 10 percent that took place in April 2009. We show that this change resulted in a differentiated response of base good and extended warranty prices. Specifically, we find that the base good prices significantly changed after the shift in the commission structure, demonstrating that the stores are, in fact, able to adjust base good prices when needed.

4 Evidence from a "Pure" Demand Shock

There is some ambiguity as to whether the unemployment rate measures shocks to demand or supply and thus as to whether β measures demand or supply responses. While the distinction does not compromise characterizations of the cyclicality of base durable and extended warranty prices, it is informative for the microfoundations of our estimates. Accordingly, we construct a relatively clean shifter of demand to estimate β.

Canada has 10 provinces where for various reasons, including the enormity of the land mass and substantial geographic variation in natural resource endowments, the provinces differ considerably in their economic activities. The province of Alberta, in particular, is especially dependent on income from sectors relating to the production, distribution, and export of oil and natural gas, with it making up 27.8 percent of GDP on average from 1999 to 2010, compared to 9.5 percent on average for the other provinces and territories. This differential dependence on oil and gas facilitates creation of proxies of incomes shifters for consumers.

Our specific proxy makes use of the high endowment of oil in Alberta relative to the rest of Canada and hence the differential impact of the oil price on its consumers' income. To see this formally, let p^{oil} denote the world spot price and suppose that the price has recently increased to $\Delta p^{oil} > 0$, where Δ denotes a first difference. For Albertans, the increase in the oil price translates into higher consumer prices q (gas in particular) and additional income I

$$\Delta p^{oil}_{AB} \rightarrow \Delta q + \Delta I \quad (Alberta)$$

where $\Delta q > 0$ and $\Delta I > 0$. For residents of other provinces, higher oil prices translate into

an increase in consumer prices only:

$$\Delta p^{oil}_{ROC} \rightarrow \Delta q \quad (Rest\ of\ Canada)$$

The difference $\Delta p^{oil}_{AB} - \Delta p^{oil}_{ROC}$ gives ΔI, a proxy for a change in consumer income in Alberta resulting from a shift in oil prices.

To take the aforementioned logic to the data, we estimate

$$p_{tscm} = \beta AB_s \times p^{oil}_t + \phi p_{0scm} \times f(t) + \alpha_{r(s)} + \gamma_{tcm} + \varepsilon_{tscm} \tag{2}$$

where AB_s indicates whether store s is in the province of Alberta, p^{oil}_t is the median crude price of oil in calendar month t, p_{0scm} is the initial value of the dependent variable, and $f(t)$ is a second order polynomial in the calendar month. Monthly crude oil prices are obtained from West Texas Intermediary (Cushing, Oklahoma) and are measured in US dollars. We standardize p^{oil}_t by its full sample mean and standard deviation.

The parameter of interest is β. It measures the effect of the oil price on durable good and extended warranty prices in Alberta, relative to the rest of Canada. Based on our argument above, we expect $\beta > 0$ for all prices because an increase in oil prices generates an outward shift in consumer demand in Alberta. We expect a larger β for extended warranty prices if this is the margin the retailer is using to adapt to demand shocks.

One concern with identification relates to whether a positive $\beta > 0$ reflects supply as well as demand side effects. Increases in the world price of oil typically raise labor demand and wages in oil and gas industry in Alberta, which in turn improve the outside opportunities of retail workers in Alberta, disproportionately. Improved outside opportunities can shift the selling effort of employees via: (i) an employee turnover channel that trades off higher for lower quality personnel; (ii) an efficiency wage channel where employees have weaker incentives to sell durable goods and extended warranties at the margin. Either channel will

Table 3: Pricing and a "Pure" Demand Shock. Unit of observation is defined by calendar month, store, product category, and manufacturer. Warranty Price Discount equals the absolute value of the difference between the realized and suggested extended warranty price. Standard errors are in parentheses and clustered on employment insurance region. $***$ and $**$ denote statistical significance at the 1 and 5 percent levels.

	Warranty Price		Base Good Price		Warranty Price Discount	
	(1)	(2)	(3)	(4)	(5)	(6)
Alberta × World Oil Price	6.925***	7.496***	1.287	3.658	-4.758***	-5.884***
	(1.940)	(1.948)	(4.629)	(2.767)	(1.703)	(1.613)
World Oil Price (in US $)	6.416		0.504		0.601***	
	(4.767)		(0.102)		(0.053)	
Alberta (0/1)	3.187**		7.466		-4.455	
	(1.340)		(11.198)		(3.055)	
Constant	72.317***		553.599***		5.047***	
	(1.491)		(10.709)		(1.508)	
Controls						
Store-Specific Price Trends	N	Y	N	Y	N	Y
Month-Category-Manufacturer Fixed Effects	N	Y	N	Y	N	Y
Employment Insurance Region Fixed Effects	N	Y	N	Y	N	Y
Observations	771942	546069	1277824	1271011	771942	546069
R^2	0.006	0.556	0.001	0.878	0.048	0.429

generate a β that reflects a shift in consumers' income coming from the supply as well as the demand side. To address the concern that β also reflects supply side shifts, we include the interaction $p_{0scm} \times f(t)$ that captures this type of unobserved trends, which can track both retail prices and consumer incomes in Alberta. Later we provide direct evidence of oil prices generating a relative shift in consumer demand in Alberta.

Estimates are found in Table 3. The extended warranty price estimates in Column (1) are from a simple difference-in-difference specification, without interaction term. The estimates in Column (2) are from the specification described by Equation (2). The analog estimates for impact on base durable prices and extended warranty discounts are found in Columns

(3)-(4) and Columns (5)-(6).

Column (2) shows a one standard deviation increase in the median crude oil price increases the extended warranty price by $7.496 dollars in Alberta, relative to the rest of Canada. The estimate is commensurate with our baseline estimates as it amounts to a 8.4 percent increase over the mean extended warranty price. Column (4) shows a relative increase of $3.658 in the base good price, though the estimate is statistically insignificant. Column (6) shows a relative decrease of $5.884 in warranty price discounts. This estimate is smaller in absolute value than its analog in Column (2), demonstrating that the response to increased consumer demand is partially reflecting headquarters decisions. These patterns are highly consistent with the baseline estimates of price responses to the unemployment rate, reported in Table 2. Columns (1),(3), and (5) show the patterns are similar even in the case where we use the simplest difference-in-difference specification.

Our interpretation of the differential effects of world oil prices on durable and extended warranty prices is based on two primary assumptions. The first is that consumer gas prices follow similar trends across provinces, that is $\Delta q = \Delta q_{AB} = \Delta q_{ROC}$. The second is that the remainder $\Delta p^{oil}_{AB} - \Delta p^{oil}_{ROC}$ measures a shock to consumer demand via the income difference ΔI. We assess these two assumptions in turn.

Figure 2 plots time series' for consumer gas prices in Alberta and the rest of Canada. It shows consumer gas prices are always lower in Alberta, 9 cents less per litre on average relative to the rest of Canada. The figure also suggests consumer gas prices follow similar trends across provinces. Online Appendix OA.4.2 checks this further, showing substantial overlap between the two time series after they are detrended by a second order polynomial in the calendar month. Similar across-province trends support the assumption that is $\Delta q = \Delta q_{AB} = \Delta q_{ROC}$.

To assess the assumption that $\Delta p^{oil}_{AB} - \Delta p^{oil}_{ROC}$ measures a shock to consumer demand via ΔI, we compare the impact of oil prices on wages and employment in Alberta versus the

Figure 2: Provincial Trends in Consumer Gas Prices. Data comes from Statistics Canada table 18-10-0001-01, which reports the monthly average retail consumer prices for gasoline by province. The dashed blue line denotes the average monthly unleaded gas price in Alberta. The red solid line denotes the mean of the average monthly unleaded gas price across Canadian provinces. The overall mean in Alberta is $0.81 per litre with a standard deviation of 0.18. The overall mean in other provinces is $0.90 per litre with a standard deviation of 0.17.

rest of Canada and we estimate:

$$w_{tri} = \beta AB_r \times p_t^{oil} + \phi w_{0ri} \times f(t) + \gamma_t + \alpha_r + \delta_i + \varepsilon_{tri} \tag{3}$$

where w_{tri} is either the annual wage in current (as of 2020) Canadian dollars or number of workers employed in industry i of province r at year t. w_{0ri} is the value of the dependent variable in the base year of the sample and $f(t)$ includes linear and quadratic polynomials in t. The $w_{0ri} \times f(t)$ interaction lets us control for a differential trend in p_t^{oil} across provinces. γ_t, α_r, δ_i are fixed effects for the year, province, and industry, respectively and ε_{tri} is an idiosyncratic error term. Estimates based on the sample period for our main data set are found in Table 4. Note that we use robust standard errors because of the small number of provinces and industries and because finite sample adjusted clustered standard errors gave cause for concern.[8]

Column (1) shows that a one standard deviation increase in the price of oil increases the relative annual wage in Alberta by 1956 dollars and the relative number of employed workers by 6006. The estimates amount to 5.5 and 5.9 percent increases over their respective means and they support the idea that the world price of oil shocks consumers' income in Alberta and that $\Delta p_{AB}^{oil} - \Delta p_{ROC}^{oil}$ ultimately measures a shock to consumer demand via ΔI.

We end this section by revisiting the assumption that β reflects demand rather than supply side effects. In order to investigate whether the shifts in oil prices have spillover effects into other than natural recourses sectors, we estimate their impact on wages and employment in two industries: (i) oil and gas and (ii) the wholesale and retail sector, representative for workers of our retailer. We again estimate the specification in 3 and report results in in Columns (2) and (3) of Table 4. Column (2) reports estimates of the differential effects of the crude oil price on annual wages and total employment in the wholesale and retail

[8]For instance, we were unable to estimate wild two-way clustered bootstrap standard errors for wages because the variance-covariance matrix was not positive definite.

Table 4: Do oil prices generate a demand shock in Alberta? Unit of observation is defined by calendar year, province, and sector. Alberta is a binary indicator variable. ROC is an acronym for Rest of Canada. World Oil Price is the price of crude oil in US $ standardized by its mean and standard deviation for the full sample. Retail sample refers to the North American Industry Classifications (NAIC) "Wholesale and retail trade [41, 44-45]." Oil & Gas refers to "Forestry, fishing, mining, quarrying, oil and gas [21, 113-114, 1153, 2100]." The numbers are the relevant NAIC codes. The groupings of NAIC codes are determined by Statistics Canada. The annual wage data comes from Statistics Canada table 14-10-0063. The employment data comes from table 14-10-0092. Robust standard errors are in parentheses. $***$ and $**$ denote statistical significance at the 1 and 5 percent levels.

	Full Sample (1)	Wholesale and Retail Trade (2)	Natural Resources (3)
	Dependent Variable = Annual Wages (2020 CAD $)		
Alberta × World Oil Price	1956***	1925***	2744**
	(255)	(237)	(1123)
Mean Annual Wage (Alberta)	40554	29789	62794
Mean Annual Wage (ROC)	34858	24428	48292
Observations	1920	120	120
R^2	0.940	0.981	0.947
	Dependent Variable = Employment (Persons)		
Alberta × World Oil Price	6006***	15587***	20630***
	(1836)	(1992)	(5308)
Mean Employment (Alberta)	116473	286390	126620
Mean Employment (ROC)	99482	252533	20108
Observations	1590	100	100
R^2	0.973	0.999	0.993
Controls			
Province Fixed Effects	Y	Y	Y
Year Fixed Effects	Y	Y	Y
Province and Industry Specific Trend	Y	Y	Y
NAIC Fixed Effects	Y	N	N

sector and Column (3) reports estimates of the differential effects on the natural resources sector. Column (2) shows that a one standard deviation increase in oil prices increases relative annual wages in wholesale and retail by $1925 and relative annual wages in natural resources by $2744. The $819 difference further increases the annual wage gap between the sectors. Column (3) shows a one standard deviation increase in oil prices increases relative employment in wholesale and retail by 15587 persons and relative employment in natural resources by 20630. The estimates are supportive of a dominant demand side effect in Alberta.

A dominant demand side effect is consistent with several basic facts about oil and gas sector in Alberta. There are two types of jobs in oil and gas industry. The first involves physically demanding "blue collar work relating to mining and extraction. This work is typically undertaken by young unmarried men who have the freedom to live and work in isolated camps for extended periods and can weather extreme working conditions. The second involves cognitively demanding "white collar work. These job types are impractical or unattainable outside options for many retail workers. This impracticality or unattainability moderates the effects of world oil prices on supply side behavioral changes in retail. Moreover, high world oil prices generate positive externalities on incomes in other sectors in Alberta. High world oil prices comove with booms in other sectors, and specifically in Construction. Booms in other sectors increases the chances of a dominant effect on consumer demand for household durables, appliances, and extended warranties.

5 Discounts, Demand, and the Great Recession

The patterns in Figure 1 suggest local stores may be "throwing in" extended warranties to boost sales of durable goods during recessions. Online Appendix Figure OA.4.1 shows further that a disproportionate share of extended warranties are sold with a price close to

zero. These patterns arise outside of advertised promotions on extended warranties, which happen rarely in approximately 0.1 percent of all transactions.

Accordingly, we construct a measure of disount at the store:

$$d_{tscm} = |wp_{tscm} - swp_{tscm}|$$

where $|\cdot|$ denotes the absolute value, wp_{tscm} is the median of the realized extended warranty price, and swp_{tscm} is the median of the suggested extended warranty price in month t, store s, product category c, and the manufacturer m. Because the suggested price is larger than the realized price in 96.5 percent of all transactions, d_{tscm} generally measures extended warranty price discounts, the extent to which stores decrease the warranty price relative to the suggested price.

Estimates of the effect of the unemployment rate on warranty price discounts are found in the rightmost panel of Table 2. The effect is smaller than the effect on realized extended warranty prices in the leftmost panel. This suggests that part of the regional cyclicality of these discounts can be attributed to store-specific changes in headquarter suggested prices. The differences between the coefficients in the leftmost and rightmost panels are also small, however, implying that most of the price adjustment is done in store. We find similar patterns in Columns (5) and (6) of Table 3 when we use the world oil price to measure demand shocks in Alberta.

We estimate demand systems to present direct evidence of warranty price discounts helping stores boost sales of base goods. To motivate the construction of the demand system and to facilitate interpretations of elasticities, we note that the "list" price for the bundle is effectively equal to

$$listp_{tscm} = swp_{tscm} + bgp_{tscm},$$

which is equivalent to

$$listp_{tscm} = \underbrace{|swp_{tscm} - wp_{tscm}|}_{\text{discount effort } d_{tscm}} + \underbrace{(bgp_{tscm} + wp_{tscm})}_{\text{transaction price } tp_{tscm}}.$$

We estimate a demand system that allows us to disentangle the contribution of "throwing in the warranty for free" from the contribution of the transaction price itself.

In particular, we estimate:

$$Q_{tscm} = \beta_1 d_{tscm} + \beta_2 tp_{tscm} + \mathbf{X}_{tscm}\mathbf{\Gamma} + \alpha_s + \gamma_{tcm} + \varepsilon_{tscm} \tag{4}$$

where Q_{tscm} is either the number of extended warranties or base goods sold, tp_{tscm} is the sum of the median prices of extended warranties and base goods, and $\mathbf{X}_{tscm}$ is a vector of controls that includes the unemployment rate lagged one month, lagged number of extended warranties sold, and lagged number of base goods sold. To be able to interpret the estimated coefficients in terms of elasticities, we add 1 to our key explanatory variables and convert them into natural logarithms.

Conditional on the impact of the transaction price tp_{tscm}, β_1 measures the percentage change in quantities caused by a 100 percent discount. β_2 captures the percentage change in quantities generated by a percentage change in the transaction price, once the influence of the discount is accounted for. A regression weighted average of β_1 and β_2 measures the list price elasticity of demand. We refer to β_2 as the transaction price elasticity of demand and focus our elasticity discussion on this parameter. Later we will consider how β_2 coefficient varies with the discount.

Identification relies on the assumption that there is no unobservable variable that correlates with unadvertised discounts and prices and that tracks variation across months, stores, and manufacturers, and product categories. The assumption may not hold if local stores en-

courage sales staff to push a particular set of goods in a given month, for example, because of an excess of these goods in inventory or because of the activities of local rivals. To address it, we use an instrumental variables strategy where our instruments' set consists of lagged discounts and the sum of headquarter suggested extended warranty prices and the costs of base durables to stores. We do not use warranty costs as part of the instrument set because, unsurprisingly, they lack predictive power for extended warranty prices.

Basic demand system estimates are found in the top panel of Table 5. Columns (1) and (2) report IV estimates of the effects of warranty price discounts and the transaction price. Columns (3) and (4) report first stage estimates and OLS estimates are found in Online Appendix Table OA.4.2.

The first stage displays notable patterns. The third row of Column (3) shows some persistence in warranty price discounts. Positive and significant coefficient in the fourth row of Column (3) shows that the retailer has greater incentive to cut extended warranty prices for high priced durable goods. Column (4) shows that large lagged warranty price discounts correlate negatively with the transaction price. Columns (3) and (4) show that the sum of headquarter suggested prices and base durable costs correlates strongly and positively with warranty price discounts and the transaction price. The estimate in fourth row of Column (4) is close to one suggesting that the market for the durable goods is highly competitive.[9] The F-statistics at the bottom of the top panel imply the first stage is relevant.

The IV estimates show a 100 percent warranty price discount increases the number of extended warranties sold by 5.9 percent. A 13.5 percent increase in the expected transaction price, roughly equivalent to the mean extended warranty price, increases extended warranty sales quantities by $100 \times 0.135 \times 0.038 = 0.51$ percent. This estimate is statistically insignificant at the 10 percent level, however. The same 100 percent warranty price discount increases base

[9] It also provides an explanation for why the differences between the IV and OLS estimates are smallest for the base good price coefficients.

Table 5: Warranty Discounts and Consumer Demand. Unit of observation is the calendar month, store, product category, and manufacturer. Reported dependent and independent variables are in natural logarithms. Regressions include fixed effects for the month-category-manufacturer combination and employment insurance region, as well as lags of the unemployment, extended warranty quantity, and base good quantity. Standard errors are clustered at the level of the employment insurance region and are in parentheses. $***$ denotes statistical significance at the 1 percent level.

	IV Estimates		First Stage		
	Extended Warranty Quantity	Base Good Quantity	Warranty Price Discount	Total Price	Transaction Price × Discount
	(1)	(2)	(3)	(4)	(5)
Warranty Price Discount	0.059***	0.036***			
	(0.011)	(0.009)			
Transaction Price	0.038	-0.185***			
	(0.025)	(0.023)			
Warranty Price Discount $(t-1)$			0.180***	-0.005***	
			(0.013)	(0.000)	
Suggested Warranty Price + Base Good Cost			0.633***	0.822***	
			(0.045)	(0.016)	
$F(2, 57)$			576.96	1287.23	
Sanderson-Windmeijer $F(1, 57)$			420.10	1657.20	
Kleibergen-Paap Wald F			101.20		
Warranty Price Discount	-0.288***	-0.272***			
	(0.032)	(0.039)			
Transaction Price	-0.059***	-0.271***			
	(0.016)	(0.018)			
Transaction Price × Discount	0.060***	0.053***			
	(0.006)	(0.008)			
Warranty Price Discount $(t-1)$			0.710***	-0.216***	1.912***
			(0.045)	(0.011)	(0.310)
Suggested Warranty Price + Base Good Cost			0.772***	0.767***	5.216***
			(0.062)	(0.018)	(0.470)
Transaction Price IV × $t-1$ Discount			-0.080***	0.032***	-0.133**
			(0.008)	(0.002)	(0.060)
$F(3, 57)$			497.19	6201.38	658.48
Sanderson-Windmeijer $F(1, 57)$			534.61	200.63	414.89
Kleibergen-Paap Wald F				16.07	
Observations	434632	434632	434632	434632	434632

durable sales quantities by 3.6 percent. A 13.5 percent increase in the expected transaction price decreases base durable sales by 2.5 percent. We infer, from these results, that warranty discounts are used to boost durable goods demand.

Why would the transaction price increase extended warranty sales quantities? Extended warranties are a type of insurance and consumers are more likely to insure expensive durable good. We use warranty price discounts to investigate the contribution of stores to this propensity in Column (1) of the bottom panel of Table 5. Extended warranty demand is more price elastic at small warranty price discounts relative to the constant elasticity estimates in the top panel. A 13.5 percent increase in the transaction price with no warranty price discount decreases extended warranty sales quantities by 0.80 percent. Demand becomes less price elastic as warranty prices are discounted relative to suggested prices. A 100 percent discount decreases the transaction price elasticity for extended warranty demand by 0.06. The estimate shows stores have a strong influence on consumer propensities for insuring expensive durables.

In Column (2) of the bottom panel, we investigate whether the transaction price elasticity of durable demand depends on warranty price discounts. Base durable demand is also more price elastic at small discounts relative to the constant elasticity estimates in the top panel. A 13.5 percent increase in the transaction price with no warranty price discount decreases base durable quantities by 3.66 percent. Durable demand also becomes less price elastic as warranty prices are discounted. A 100 percent discount decreases the transaction price elasticity for base durable demand by 0.053. The estimate implies warranty discounts boost durable demand because they make consumer demand for durables less price elastic.

The left panel of Figure 3 plots quarterly time series for warranty price discounts as a percentage of suggested prices. The figure shows a sharp discounting starting in 2007Q3, peaking during the Great Recession at around 45 percent in 2008Q1, before returning to pre-recession levels towards Q4 of 2008. These patterns together with the demand system

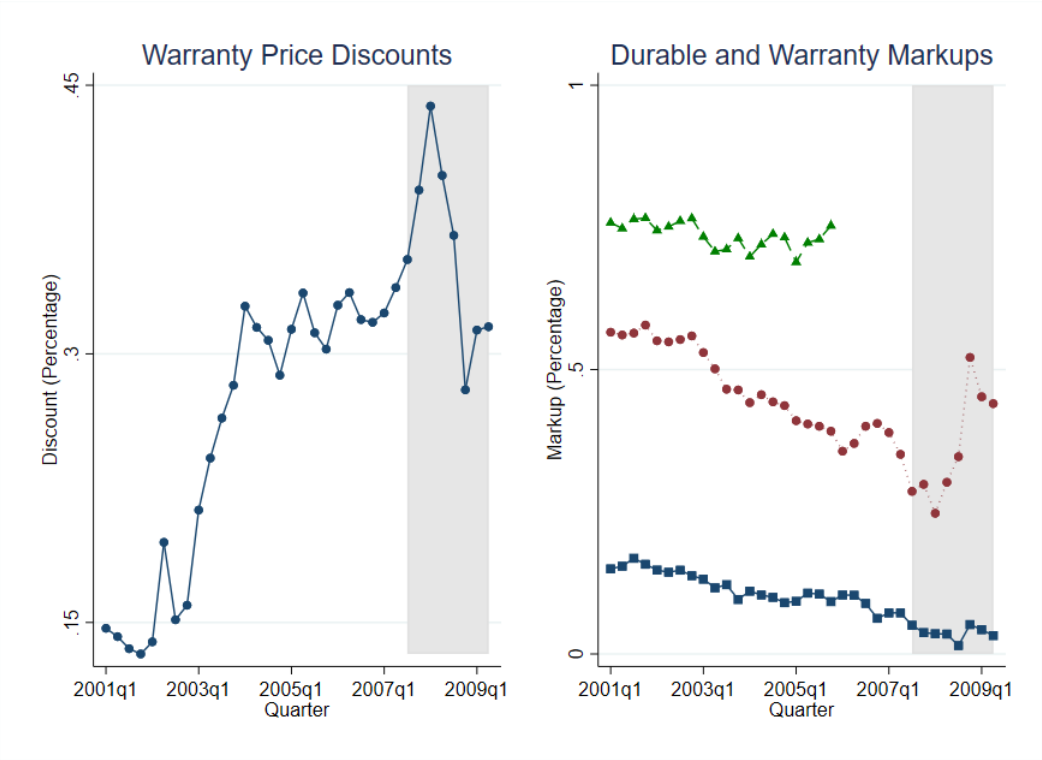

Figure 3: Warranty Discounts and Markups during the Great Recession. The left panel plots the quarterly time series for warranty price discounts as a percentage of the headquarter suggested price. The right panel plots quarterly time series for the ex post or realized markup for extended warranties (green triangles, dashed line), an ex ante or expected markup for extended warranties (red dots, dotted line), and the markup for durable goods (blue squares, solid line). The ex ante markup is based on the average cost for a particular manufacturer and product category. Ex post markups are right censored at 2007 because we do not observe extended warranty claims made after 2009. The shaded gray area denotes the Great Recession in Canada (2007Q3 to 2009Q2).

estimates suggest that stores were using warranty price discounts to boost durable goods demand during the Great Recession specifically.

The right panel of Figure 3 plots time series for ex post extended warranty markups until 2007, ex ante extended warranty price markups, and durable goods price markups. Ex post extended warranty markups are based on the realized cost of the extended warranty after a claim has or has not been made. Ex post extended warranty markups are right censored at 2007 because we do not observe extended warranty claims made after 2009. For ex ante extended warranty markups, we first computed the average extended warranty cost for each manufacturer and product category combination prior to 2007. We then assigned the average

cost to extended warranties sold from 2007 onwards. The ex ante warranty markup is just the difference between the extended warranty price and this average cost.

Ex post markups are stable between 69 to 77 percent before 2007. The high ex post markup suggests that the cost of "throwing in" the extended warranty is very low. Ex ante markups decrease sharply entering the Great Recession, reflecting the spike in warranty price discounts carried by the retailer and plotted in the left panel of Figure 3. Towards the end of the Great Recession, ex ante markups bounce back, coinciding with the reduction in discounts on extended warranties. Finally, we observe very small fluctuations in the base good markups, which continued their decline through the Great Recession. This is most likely due to the fact that markups on the base durable were already close to 0 at the onset of the Great Recession.

6 Add-on revenue dynamics in U.S. airline sector

We examine the dynamics of add on revenues using the publicly available data from U.S. airline sector. The left panel of Figure 4 plots the revenue share derived from ancillary activities such as baggage fees and cancellation fees. The add on revenue share increased gradually from 2 to 4 % prior to the Great Recession. It increased sharply from a little over 4% to about 8% during the Great Recession. Overall it increased from 1.8% in 1999 to 11% in 2018.

The right panel of Figure 4 shows that both add-on and passenger revenue changed sharply during the Great Recession. The growth rate for add-on revenue increased sharply from below 25 percent to above 50 percent. Passenger revenue contracted sharply by contrast. We cannot say anything about why add ons became more important to airlines during the Great Recession because, unlike with our extended warranty data, the airline data is aggregated and, perhaps more importantly, add on or ancillary prices are hidden.

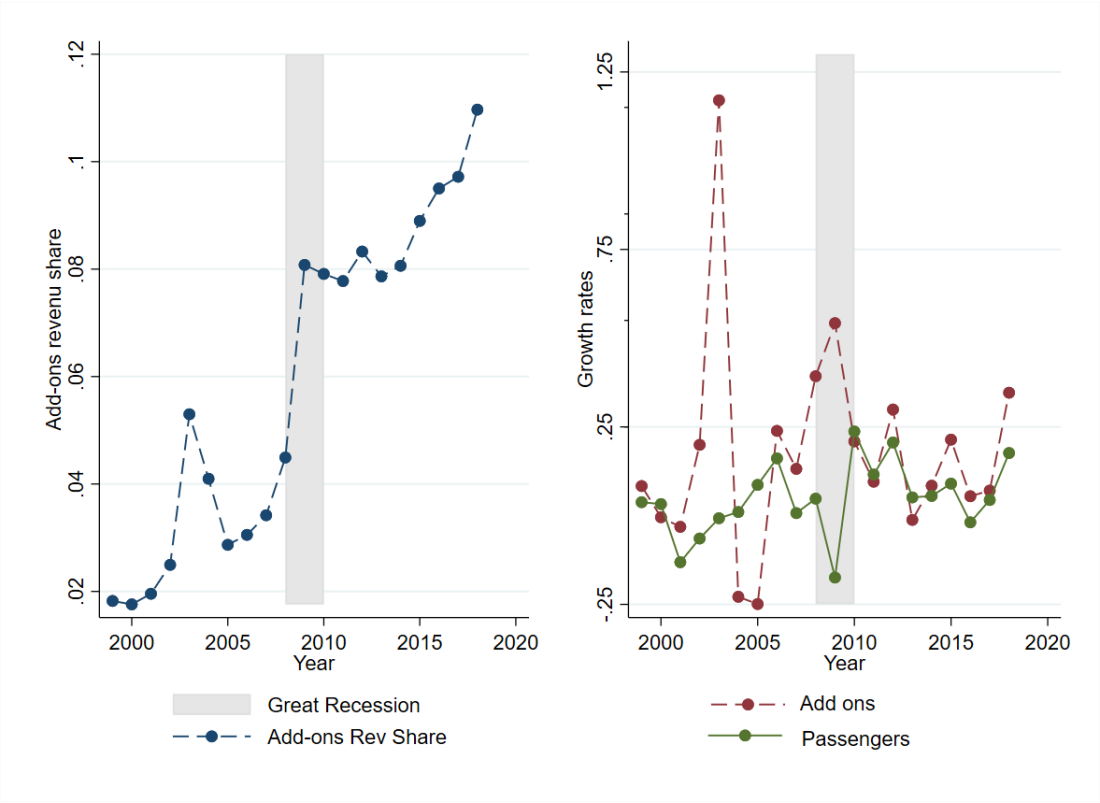

Figure 4: U.S. Airline Revenues between 1999 and 2018. Grey area covers the Great Recession. Add-on revenue share is the ratio of ancillary to the sum of ancillary and passenger revenues (blue dashed line in the left panel). Ancillary revenue is sum of revenues from baggage fees, cancellation fees, and miscellaneous operating revenues. The right panel depicts annual growth rates in passenger revenue (solid green line) and in add-on revenues (dashed red line). Data was retrieved from the Airline Data Project at MIT.

7 Conclusion

Studies of price dynamics generally show substantial price rigidity over the business cycle. These studies are however unable to account for the prices of add ons, as they are typically hidden and thus difficult to measure by definition. To help fill the gap, we draw on 10 years (1999-2009) of detailed data from a nationwide Canadian retailer of household durables to study the cyclicality of extended warranties, a customary add-on service. We show that changes in the regional unemployment rate are followed by sharp decreases in extended warranty prices. By contrast, base durable prices are rigid over regional business cycles. Our evidence suggests price fluctuations may be hidden in the prices of add ons.

We study the mechanisms underlying the relative flexibility of extended warranty prices. We estimate a demand system and show that a 100 percent warranty price discount increases base durable sales quantities by 3.6 percent. We show that warranty price discounts boost base durable demand because they make consumer durables' demand less elastic to the transaction price. A 100 percent warranty price discount decreases the transaction price elasticity of durable demand from 0.271 to 0.218. We show further that a disproportionate share of extreme warranty discounts takes place during the Great Recession. Our narrative is consistent with raw patterns in markups on base durables and extended warranties.

We conduct several robustness checks in the online appendix. We use census data to provide evidence in support of the characterization that extended warranties are add ons. We estimate transaction level regressions that control more precisely for product quality, further supporting the interpretation that our findings reflect the cyclicality of price adjustments for a given product quality by stores, rather than the cyclicality of quality adjustments by consumers. We leverage a change in the commission to salespersons for extended warranty sales to support the idea that our findings should be interpreted through the lens of add on rather than decentralized pricing.

References

Mark Aguiar, Eric Hurst, and Loukas Karabarbounis. Time use during the great recession. *American Economic Review*, 103(5):1664–96, 2013.

Eric Anderson, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jn Steinsson. Informational rigidities and the stickiness of temporary sales. *Journal of Monetary Economics*, 90:64 – 83, 2017.

Martin Beraja, Erik Hurst, and Juan Ospina. The aggregate implications of regional business cycles. *Manuscript*, 2018.

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

Alberto Cavallo. Scraped data and sticky prices. *Review of Economics and Statistics*, 100 (1):105–119, 2018.

Judith A. Chevalier, Anil K. Kashyap, and Peter E. Rossi. Why don't prices rise during periods of peak demand? evidence from scanner data. *The American Economic Review*, 93(1):15–37, 2003.

Olivier Coibon, Yuriy Gorodnichenko, and Gee Hee Hong. The cyclicality of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review*, 105(3):993–1029, 2019.

Glenn Ellison. A model of add-on pricing. *The Quarterly Journal of Economics*, 120(2): 585–637, 2005.

Xavier Gabaix and David Laibson. Shrouded Attributes, Consumer Myopia, and Information

Suppression in Competitive Markets. *The Quarterly Journal of Economics*, 121(2):505–540, 05 2006.

Etienne Gagnon, David López-Salido, and Jason Sockin. The cyclicality of sales, regular, and effective prices: Business cycle and policy implications: Comment. *American Economic Review*, 107(10):3229–42, 2017.

Daniel Hosken and David Reiffen. Patterns of retail price variation. *The RAND Journal of Economics*, 35(1):128–146, 2004.

Industry Canada. Canada's changing retail market. 2013.

Patrick Kehoe and Virgiliu Midrigan. Prices are sticky after all. *Journal of Monetary Economics*, 75:35 – 53, 2015.

Alan B. Krueger and Andreas Mueller. Job search and unemployment insurance: New evidence from time use data. *Journal of Public Economics*, 94(3):298–307, 2010.

Emi Nakamura and Jón Steinsson. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4):1415–1464, 2008.

Aviv Nevo and Arlene Wong. The elasticity of substitution between time and market goods: Evidence from the great recession. *International Economic Review*, 60(1):25–51, 2019.

Johannes Stroebel and Joseph Vavra. House prices, local demand, and retail prices. *Journal of Political Economy*, Forthcoming, 2019.

Mark J. Zbaracki, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics*, 86(2):514–533, 2004.

Add-on Pricing over Regional Business Cycles: Evidence from Extended Warranties

Online Appendix

Branko Bošković Sacha Kapoor Agnieszka Markiewicz Barry Scholnick

July 17, 2020

OA.1 Are Extended Warranties Add ons?

The low claim rate in our summary statistics (Table 1 of the main text) raises questions about whether an extended warranty can be interpreted as an add on, as it only improves base durable quality in a handful of cases. It is important to stress, as we do in the main text, that add ons are defined by their perceived quality at the point of sale rather than by their realized quality after the purchase is made. As long as the base durable is perceived to be better with the extended warranty at the point of sale, such that all consumers would choose the bundle (base durable + extended warranty) over the base durable alone if the extended warranty were thrown in for free, extended warranty can be interpreted as an add on. The low claim rate can be rationalized by biases in consumer perceptions, such as upward distortions of the failure probability of the base durable [Abito and Salant, 2018].

Two assumptions characterize add ons in the add-on pricing game developed in [Ellison, 2005]. The first is that "high type" consumers are more likely to purchase an add on. The second is that high types have less of a propensity for switching stores to take advantage of small price differences. We investigate the validity of these assumptions to extended warranties to further reinforce the notion that they are add ons. To do this, we draw on our transaction-level data and information from the Canadian Census of 2006 and compute the distance between the purchase store and home of the consumer, which gives us a rough proxy of the propensity to switch stores. The census data includes fine grained information relating to median household income in the postal code of the consumer. We use income information as a rough proxy for the consumer type, with high types having high incomes and lower marginal utility of one additional dollar of income.

Figure OA.1.1(a) summarizes the relationship between the frequency of extended warranty purchases and the median household income. The extended warranty purchase probability (y axis) is computed as the frequency of warranty sales in each EI region over the

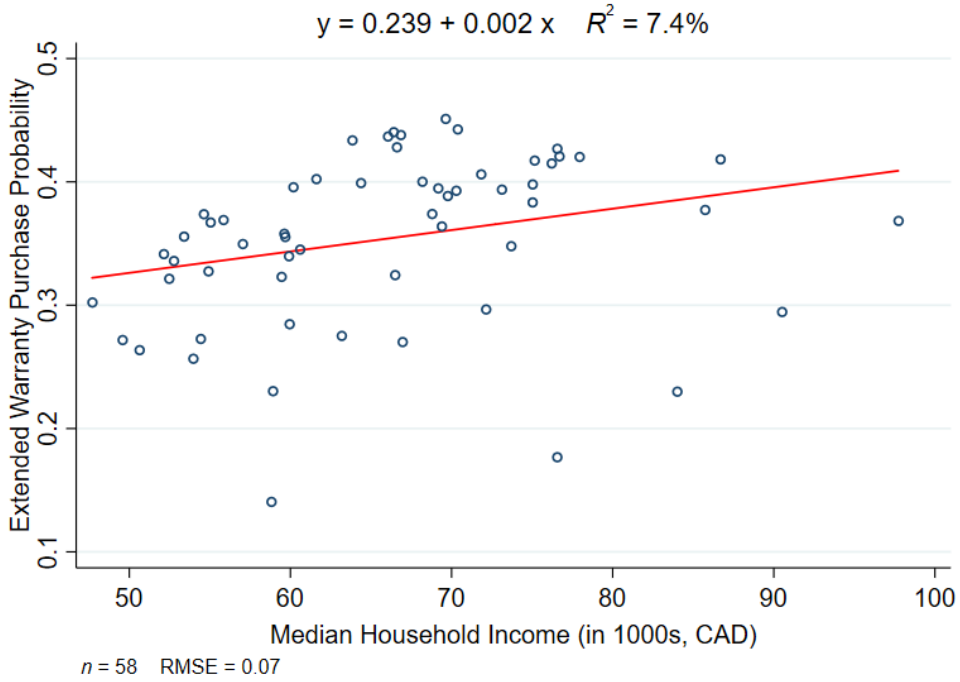

(a) Extension-Income Gradient

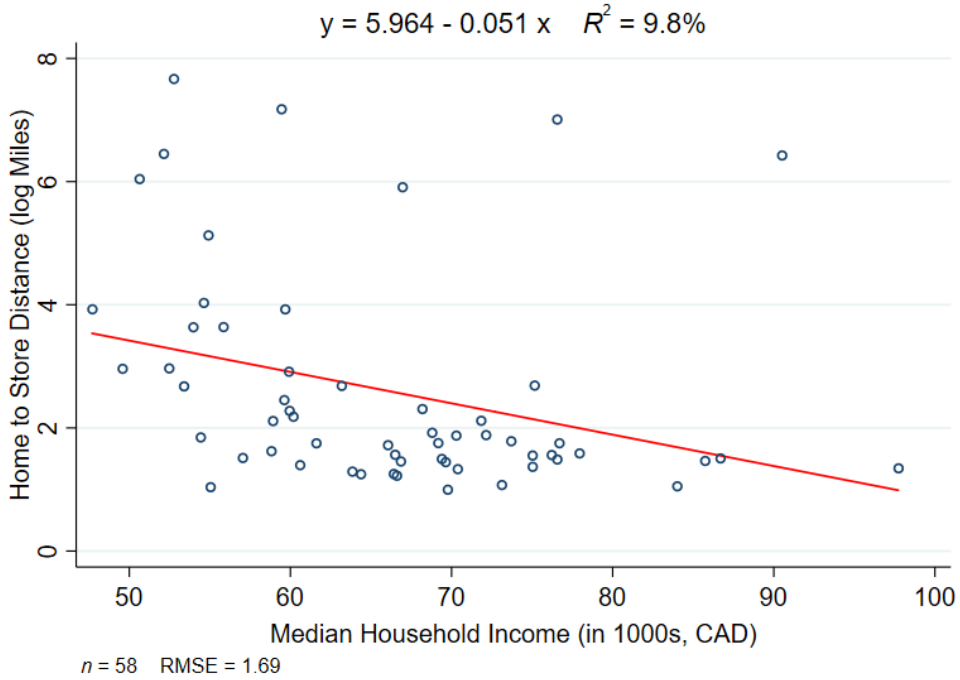

(b) Distance-Income Gradient

Figure OA.1.1: Income Gradients. Unit of analysis is EI (Employment Insurance) region. Extended warranty and distance statistics are calculated using transactions data for the entire sample period. Extended warranty purchase probability is frequency of transactions where an extended warranty was purchased. Home to Store distance is log of the average distance between customer postal code and purchase-store postal code. Income statistic is median of the median incomes across all postal codes in the EI region. Median income is constructed this way because the 2006 Canadian Census only provides summary income statistics at the level of the postal code. CAD = Canadian Dollars.

entire sample period. We constructed median income in several steps. We obtained median household income for each postal code in each EI region from the 2006 Canadian census. We calculated the median of this median for the EI region and we refer to it as median household income (x axis). Figure OA.1.1(b) plots the relationship between the distance between home and store of the purchase and median income. Home to store distance is calculated as the log of the average distance from the postal code of the consumer to the postal code of the store where the good was purchased.

Figure OA.1.1(a) shows a positive correlation between the extended warranty purchase probability and median household income. A 10000 dollar increase in median household income is associated with a 2 percentage point increase in the extended warranty purchase probability. This estimate is statistically significant at the 5 percent level with robust standard errors. Figure OA.1.1(b) depicts a negative correlation between the home-store distance and median household income. A 10000 dollar increase in median household income is associated with a 0.51 log point decrease in the extended warranty purchase probability, which amounts to a 28 percent decrease in the median log distance to the purchase store. This estimate is statistically significant at the 5.5 percent level with robust standard errors. The correlations depicted in Figures OA.1.1(a) and OA.1.1(b) are therefore consistent with the basic assumptions of the Ellison [2005] add-on pricing game because high income consumers have a greater propensity for purchasing extended warranties and a lower propensity to travel far to make a purchase. These findings also lend further support that extended warranty is an add on.

OA.2 Add On versus Base Good Pricing or Decentralized versus Centralized Pricing?

We investigate whether our findings should be viewed through the lens of centralized versus decentralized pricing rather than the lens of base good versus add on pricing, as both views would predict greater regional cyclicality for extended warranty prices. We leverage a change in the commission structure for sales employees to show that our findings should be viewed through the lens of base good versus add on pricing. In April of 2009, towards the end of the Great Recession, the retail chain cut the extended warranty commission 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. We use the reduced pay for performance to show that, like extended warranty prices, base durable prices are decentralized.

To show this, we compare price differences from 3 months before and after the commission reduction with the same price differences in 2008. Specifically, we estimate

$$p_{tscm} = \beta_0 + \beta_1 2009_t + \beta_2 AMJ_t + \beta_3 2009_t \times AMJ_t + \alpha_{r(s)} + \gamma_{cm} + \varepsilon_{tscm} \tag{5}$$

where 2009_t indicates whether the year for t is 2009 and AMJ_t indicates whether the month for t is April, May, or June. Our interest is in β_3, which measures yearly differences in price differences across the months of the commission reduction. We use this coefficient to draw inferences about whether base durable prices are in fact decentralized.

Estimates are found in Table OA.2.1. Note that we consider specifications that exclude category-manufacturer fixed effects in Columns (1) and (3). We do this because the new commission structure can affect prices via two channels. First, it reduces the marginal incentive to put sales effort into obtaining high prices on extended warranties. Second, it increases the relative incentive to push high priced durables. In specifications without category-manufacturer

fixed effects, our estimates of β_3 reflect the effect of the new commission structure on both marginal and relative incentives. In specifications with category-manufacturer fixed effects, our estimates of β_3 reflect only the effect on the marginal incentive to obtain high extended warranty prices. This latter effect on base durable prices is informative about whether base durable prices are decentralized.

Column (1) shows a decrease of \$3.56 in the relative price difference for extended warranties across years. Column (3) shows an increase of \$27.89 in the relative price difference for extended warranties across years. The extended warranty price decrease and base durable price increase are consistent with how the commission structured changed in 2009.

Column (2) shows a more pronounced extended warranty price effect when category-manufacturer fixed effects are included, suggesting that the effects on marginal incentives are more pronounced once the specification accounts for the increased relative incentive to push high priced base durables. Column (4) shows a less pronounced based durable price effect when category-manufacturer fixed effects are included. This shows that the relative incentive to push high priced durables is quite important, and that the marginal incentive effect survives category-manufacturer fixed effects. This further supports the idea that base durable prices are decentralized along the intensive margin of store effort.

OA.3 References

Jose Miguel Abito and Yuval Salant. The Effect of Product Misperception on Economic Outcomes: Evidence from the Extended Warranty Market. *The Review of Economic Studies*, 86(6): 2285-2318, 2018.

Glenn Ellison. A model of add-on pricing. *The Quarterly Journal of Economics*, 120(2): 585-637, 2005.

Table OA.2.1: Do Stores Control Base Durable Prices? Difference-in-differences estimates of the effect of a reduced commission on extended warranty sales on extended warranty and base durable prices. The estimates uses 6 months around the commission reduction month (January to June) from 2009 and the same 6 months from 2008. The retail chain reduced the extended warranty commission from 15 to 10 percent in April 2009 for all employees. The commission for the durable good was unchanged. Standard errors are clustered on the employment insurance region and are in parentheses, with *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Extended Warranty Price		Base Good Price	
	(1)	(2)	(3)	(4)
(Month≥April)×(Year = 2009)	-3.557**	-5.908***	27.888***	7.178**
	(1.391)	(1.199)	(5.559)	(3.014)
(Year = 2009)	12.851***	14.847***	-37.383***	-38.617***
	(1.551)	(1.520)	(5.366)	(3.169)
(Month≥April)	1.082	3.851***	-11.283***	-8.894***
	(1.094)	(0.888)	(3.718)	(1.420)
Constant	89.028***	87.218***	614.495***	618.748***
	(0.821)	(0.741)	(3.178)	(1.349)
Fixed Effects				
Category-Manufacturer	N	Y	N	Y
Employment Insurance Region	Y	Y	Y	Y
Observations	81297	81250	155553	155503

OA.4 Additional Figures and Tables

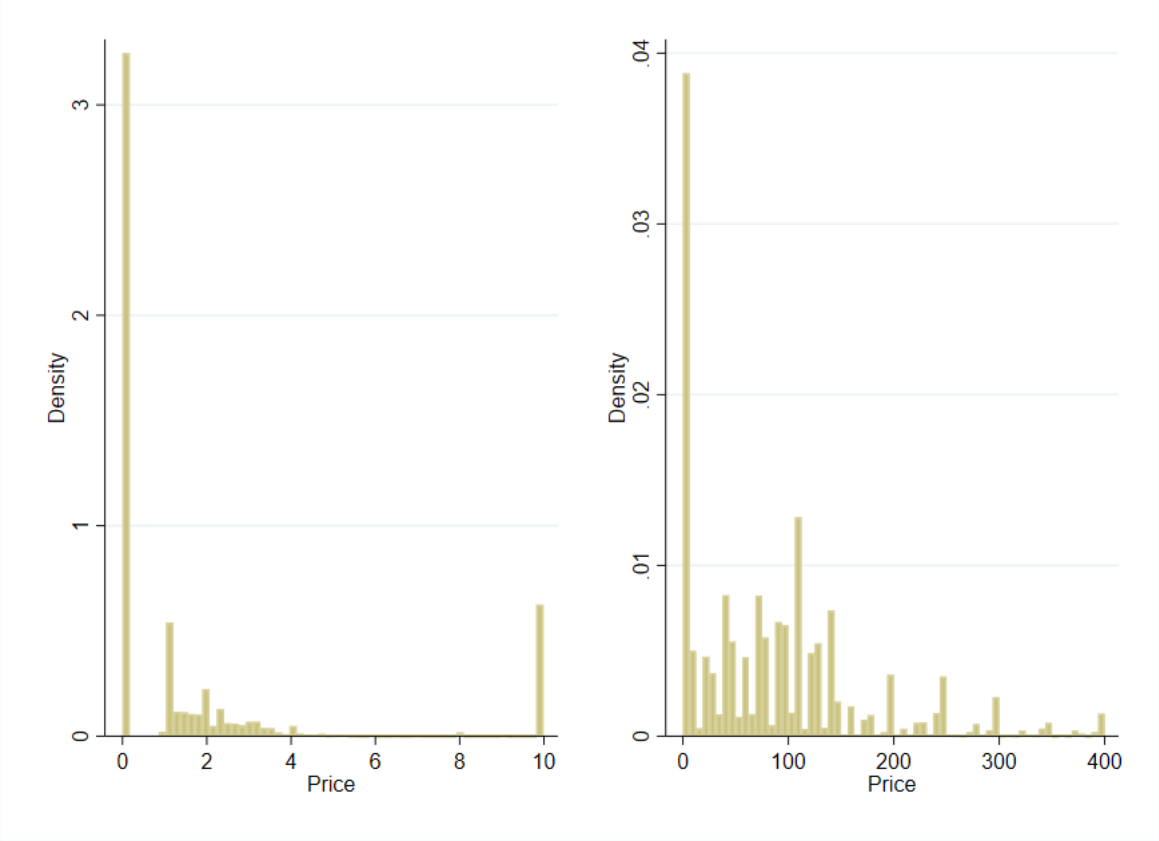

Figure OA.4.1: Distribution of Extended Warranty Prices. The histograms show that a disproportionate share of extended warranties are sold with a price which is close to zero. The left panel plots the histogram for prices between 0 and 10 dollars. The right panel plots the histogram for prices between 0 and 400 dollars, which is essentially the full support, less outliers.

Table OA.4.1: Price adjustments for a fixed product quality by stores or quality adjustments by consumers? The unit of observation is defined by transaction. Regressions include fixed effects for the month-category-manufacturer-model and employment insurance region, as well as lags of the unemployment rate. Long run effects are cumulative effects in regressions with 12 lags of the unemployment rate and 12 lags of the dependent variable. Standard errors for long run effects are computed via the Delta Method. Standard errors are clustered at the level of the employment insurance region and are in parentheses. $***$, $**$, $*$ denote statistical significance at the 1, 5 and 10 percent levels.

	Warranty Price			Base Good Price			Warranty Price Discount		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Local Unemployment Rate at									
t - 1	-2.536***			-0.708*			2.028***		
	(0.810)			(0.369)			(0.730)		
t - 6		-2.490***			-0.657*			1.948**	
		(0.819)			(0.362)			(0.745)	
t - 12			-2.579***			-0.662**			1.993***
			(0.745)			(0.318)			(0.682)
Implied effect of a 1σ	-6.364***	-6.200***	-6.406***	-1.777*	-1.635*	-1.644**	5.088***	4.851**	4.951***
increase in Lagged UR Rate	(2.033)	(2.040)	(1.851)	(0.926)	(0.903)	(0.790)	(1.831)	(1.855)	(1.694)
Long-Run Effect		-3.669***			-0.854**			2.914**	
		(1.126)			(0.434)			(1.036)	
Mean of Dependent Variable		79.869			582.176			36.456	
Observations	2,281,497	2,190,674	2,115,712	5,900,342	5,699,837	5,535,294	2,281,497	2,190,674	2,115,712
R^2	0.497	0.486	0.482	0.982	0.982	0.981	0.453	0.449	0.447

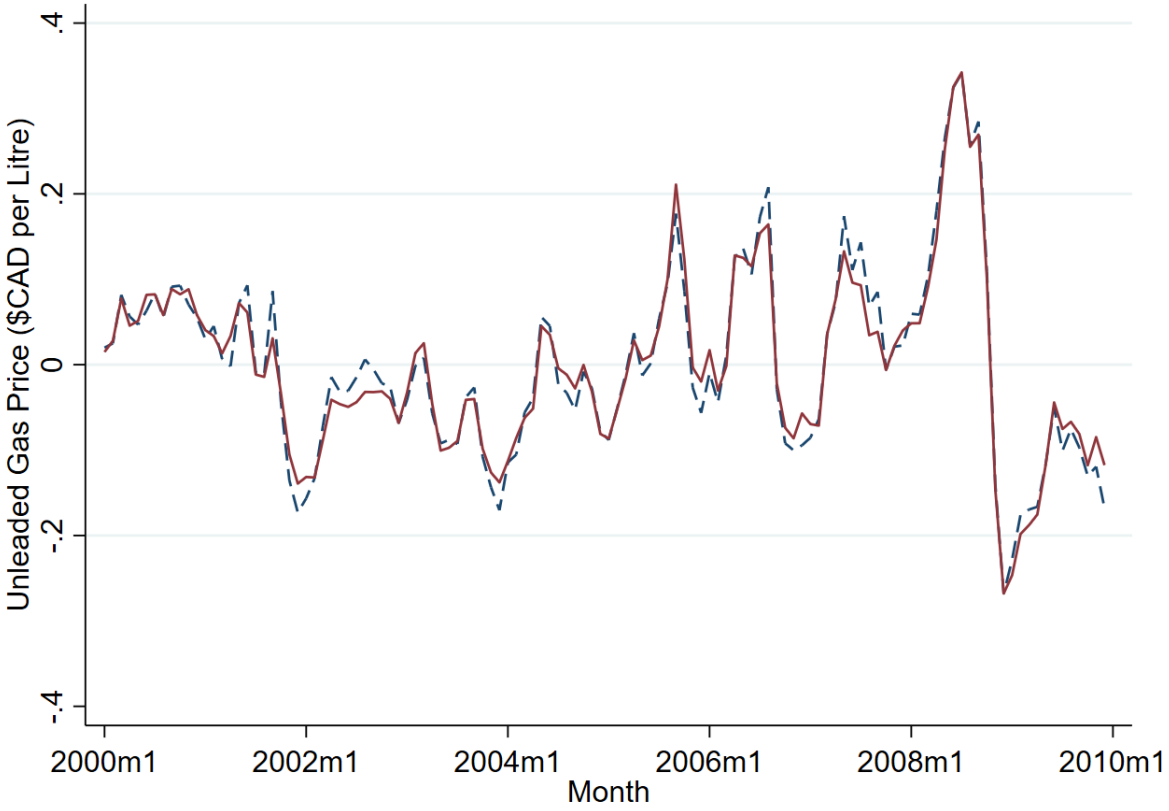

Figure OA.4.2: Detrended Time Series for Consumer Gas Prices. Data comes from Statistics Canada table 18-10-0001-01, which reports the monthly average retail consumer prices for gasoline by province. The dashed blue line denotes the average monthly unleaded gas price in Alberta. The red solid line denotes the mean of the average monthly unleaded gas price across Canadian provinces. Both series are detrended by a quadratic polynomial trend.

Table OA.4.2: Warranty Discounts and Consumer Demand (OLS Estimates). Unit of observation is the calendar month, store, product category, and manufacturer. Reported dependent and independent variables. Regressions include fixed effects for the month-category-manufacturer combination and employment insurance region, as well as lags of the unemployment, extended warranty quantity, and base good quantity. Standard errors are clustered at the level of the employment insurance region and are in parentheses. $***$ denotes statistical significance at the 1 percent level.

	OLS Estimates	
	Extended Warranty Quantity	Base Good Quantity
	(1)	(2)
Warranty Price Discount	0.026***	0.014***
	(0.002)	(0.002)
Total Price	0.144***	-0.077***
	(0.018)	(0.016)
Warranty Price Discount	-0.150***	-0.107***
	(0.005)	(0.004)
Total Price	0.056***	-0.139***
	(0.015)	(0.014)
Total Price × Discount	0.035***	0.024***
	(0.000)	(0.001)
Observations	434632	434632