Title
Essays in Natural Resource Economics

Author
Ray, Kevin Daniel

Publication Date
2020

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays in Natural Resource Economics

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Kevin Daniel Ray

Committee in charge:

Professor Richard Carson, Co-Chair
Professor Joel Watson, Co-Chair
Professor Wendy Liu
Professor Craig McIntosh
Professor Dale Squires

2020

The dissertation of Kevin Daniel Ray is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Co-Chair

University of California San Diego

2020

DEDICATION

To my wife Gretchen, for your support and encouragement as I wrote this dissertation as the second biggest step in our life together, behind only our wedding.

EPIGRAPH

The conservation of natural resources is the fundamental problem.

Unless we solve that problem it will avail us little to solve all others.

— Theodore Roosevelt

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Richard Carson for his support as the co-chair of my committee. His vast knowledge of the relevant applied and methodological literature was invaluable in honing my research questions and using the appropriate methods to answer them. My thanks as well to Professor Joel Watson who continued to attend my talks and help me as my work moved from blending theory and application to being almost strictly applied.

I would also like to acknowledge Dr. Dale Squires and Dr. James Hilger who have provided me with the data and funding for much of the work contained in this dissertation. Their wisdom and experience in the messy political arena of fisheries management has helped keep my research grounded in what is practical and useful.

Last of all, I would like to thank my friends and family for their love and encouragement, and all the joyful experiences we have shared over my time as a doctoral student.

Chapter 1, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Arizona American Water for their generous provision of the data for academic research. I am grateful for the help of Professor Gary Thompson in getting this project off the ground through the early data collection stages, and for his seed of an idea to employ quantile regressions to analyze the heterogeneity of price elasticities. Additional thanks to Professor Richard Carson and Professor Kaspar Wuthrich for help in selecting the appropriate quantile regression methods, and numerous seminar participants for their comments which helped shape this paper.

Chapter 2, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Bumble Bee and Nielsen via the Kilts Center for Marketing at Chicago Booth for their generous provision of data. I am grateful for Dr. Dale Squires who believed in my potential to address these fundamental questions about demand for environmental attributes enough to hand this project to me. Additional thanks to Professor Richard Carson for his guidance in selecting the proper scope for this research and the inspiration to use a survival model.

Chapter 3, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Dr. James Hilger who selected me to work on the Coastal SEES project of which this research is an offshoot, and whose curiosity regarding the application of big data techniques to fisheries was the spark for this paper. Thanks also to Professor Art Miller for his feedback on the relevant climate and oceanographic variables to include, Dr. Kristin Roll for comments and feedback, and to Dr. Emmanis Dorval for useful conversations regarding market squid biology and fishing.

VITA

2007	B. S. in Economics and German Studies *Magna Cum Laude*, University of Arizona
2011-2012	Research Assistant, University of Arizona
2012	M. S in Agricultural and Resource Economics, University of Arizona
2012-2019	Graduate Teaching Assistant, University of California San Diego
2020	Ph. D. in Economics, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays in Natural Resource Economics

by

Kevin Daniel Ray

Doctor of Philosophy in Economics

University of California San Diego, 2020

Professor Richard Carson, Co-Chair
Professor Joel Watson, Co-Chair

This dissertation consists of three papers on the economics of natural resources. Two of these examine the impact of price on consumption, and the third illustrates the benefits to applying big data methods to the estimation of fishery production functions.

Chapter 1 estimates the price elasticity of demand for residential water across the distribution of users and over the course of time after the change. Using monthly billing data provided by a private utility in Phoenix, Arizona, the distribution of price responses is estimated via quantile regression. The results indicate that most households reduced their usage about 200 gallons per month, while the top 20% reduced usage by 500 to 800 gallons per month. Converting this into

a percentage change to compute the elasticity inverts the results such that the lowest 20% of households have the largest percentage reduction in use. The reductions in usage take four months to emerge, after which it is apparent that summer reductions are greater than winter reductions. Chapter 2 explores the strategic behaviors exhibited by households in response to high-low pricing patterns for purchases of canned tuna, and how these behaviors differ for purchases of eco-labeled brands. Promotional price elasticities are estimated using a demand system for six retailers over five years. The promotional price elasticity for the eco-labeled brand (-2.8) is similar to Bumble Bee (-2.2) and StarKist (-2.9), implying that purchases of eco-labeled canned tuna respond similarly to sales prices. A survival model is estimated using household scanner panel data to analyze the inter-purchase timing decision, with the results showing that stockpiling behaviors are similar between eco-labeled and conventional tuna but price is a more significant factor in accelerating purchases for the eco-labeled brand. Chapter 3 illustrates the efficacy of LASSO variable and instrument selection methods for estimating production in fisheries. Production functions are estimated using traditional model selection and LASSO model selection for both ordinary least squares and instrumental variables methods. The results show that LASSO performs marginally better at ordinary least squares and significantly better for instrumental variables, which tests for endogeneity indicate should be the preferred method.

Chapter 1

Get Off My Lawn! How and when water usage changes in Response to a price increase in Phoenix Retirement Communities

1.1 Introduction

As urban areas expand world-wide, policy makers struggle to simultaneously balance three major goals of water supply: affordability for low-income households, growing scarcity of clean water, and adequate return-on-investment for water providers. Costs for water providers are continuing to grow with scarcity, increasing water pollution, and an aging infrastructure; thus higher prices are needed to keep the water providers fiscally solvent. However, rising prices are bad for consumers, particularly those with low income. Many water providers have turned to block-rate pricing schedules to help address access for low-income households while providing incentives for high-usage households to conserve. However, with complicated price schedules it becomes challenging to forecast the impact of rate increases on revenue. Correctly forecasting revenue and welfare implications requires not just understanding *how much*, but *when* the changes occur (for instance the short-run versus long-run, and seasonality) and *who* reduces their use. Properly accounting for the heterogeneity and timing of response are important factors in achieving a balance of the three goals of water policy.

The paucity of information regarding which households reduce water usage and when is primarily due to inherent limitations in available data. Early studies of residential water demand relied on cross-sectional variation or longitudinal variation, each subject to a number of confounding factors. Panel data provides some solutions to the confounds, but still requires assumptions about which price consumers observe and requires an instrumental variable for price which is "inexcludably" linked to quantity. Quasi-experimental methods have begun to address these shortcomings by exploiting features of particular rate changes to estimate "treatment effects" of price changes (Nataraj and Hanemann, 2011; Baerenklau et al., 2014; Ito, 2014; Yoo et al., 2014; Wichman, 2014). However, the identification strategies in these papers are unable to estimate BOTH who and when because of the particular construction of each "control group". This paper is able to address both who and when because of the availability of a control group of similar group of households who did not experience a price increase. Using a quantile difference-in-differences framework I estimate the heterogeneity of response as well as how the response develops over time.

Plausible arguments can be developed to suggest that households which use the most water are the least or the most responsive to price increases. It is ultimately a question that must be evaluated empirically, and the existing literature has found these high usage households have less elastic demand. This study also finds that households using the most water have the least elastic demand, but is suggestive that the elasticity does not tell the whole story. The estimated elasticity is relatively constant around -0.10 for the top 80% of households by usage, with higher elasticities around -0.25 for the bottom 20%. The results are in agreement with Baerenklau et al.

(2014) who find that the reduction in pure levels is increasing with household usage, but it is the conversion to percent changes which produces the inverse relationship with elasticities.

A number of papers about residential water have used Phoenix data, but these retirement communities differ in several key respects. The dominant landscaping is gravel with succulents and cacti, with grass being an uncommon choice. Higher water use plants such as citrus trees and rose bushes do appear with regularity though. Another cause for lower water usage is the relative scarcity of backyard pools, as the community pools are an effective substitute. As a result of these factors, the 95th percentile household in these communities uses about as much water as the 75th percentile Phoenix home[1] If the prior results that low-usage households are more price elastic that should be amplified in this sample.

The seasonality of the price response can be informative to both separating indoor reductions from outdoor reductions and to estimating how "sticky" the reductions may be. The results show the decrease in water usage to be greater in the summer, with summer elasticities greater in magnitude than winter elasticities. These results suggest a possible flaw with the assumptions of many previous water demand papers that estimate seasonal elasticities. In estimating the price elasticities from differing summer and winter prices as they have, they assume that water usage is only a function of current price without regard to the history of prices. If usage is a function of the history of prices, a phenomenon called hysteresis, then some of the reduction in winter usage should be attributed to the increased summer rates and not to the change in winter rates. Recent work in the electricity literature finds evidence for hysteresis (asymmetric price elasticities), and the fact that water and electricity are derived demands implies that purchases of more efficient durable goods and habit formation theoretically justify the hysteresis.

The study households take at least four months to show any statistically significant decrease in water usage. This could possibly be attributed to the price change taking place in December and being observed during winter when demand is lower. However, comparing the

[1] Klaiber et al. (2014) reports a median of approximately 23 ccf = 17 kgal in 2003

water usage in the first and second year after the price increase shows that even in the winter months, the price increase does lead to a decrease in usage in the second year. This delayed response suggests that other quasi-experimental estimates of the price elasticity of water demand may be underestimating even the short-run effect by looking at only the first month after the price increase.

The remainder of the paper is laid out as follows: Relevant literature is discussed in Section 1.2. The underlying theoretical model is examined in Section 1.3. The overall study context and the data set used are described in Section 1.4. Econometric estimation methods are outlined in Section 1.5, and results are analyzed in Section 1.6. Section 1.7 provides some concluding observations.

1.2 Related Literature

This section looks at the econometric techniques used in the literature to estimate the relationship between price and water usage. This is followed by a discussion of the existing results with regards to heterogeneous reactions to price.

1.2.1 Identification of Price Elasticity with Block Rate Pricing

The complexity created by increasing block rates and the fundamental endogeneity of price have made econometric identification of the price elasticity rather tenuous. The earliest studies tended to ignore the endogeneity and simply estimated the reduced form OLS demand function. Recognizing the endogeneity, instrumental variables became a staple in the residential water demand literature, but have fallen into disuse of late in part because the exclusion restriction is impossible to satisfy since price is *inseparably* a function of usage. Around the early 2000s, a few papers applied the Discrete/Continuous Choice (DCC) model which accounts for how the rate structure creates kinks in the budget curve. While representing a conceptual improvement in estimation technique in theory, it assumes perfect knowledge of price and structural knowledge of the household utility function up to the set of estimated parameters.

But as behavioral economics gained in prominence, the evidence kept mounting that consumers could not be counted upon to behave as economic theory would predict. Borenstein (2009) argues against the traditional "perfectly optimizing consumer" paradigm and even goes so far as to state:

> It seems safe to say that not only do most consumers not know how much power or water they have used since their current billing period began, most consumers don't know when their current billing period began. (p. 3)

If household behavior is not optimized in roughly the same manner as theory suggests, then the DCC model is misspecified. A test of this model's specification comes from the theoretical prediction of "bunching" in which households with demand curves that intersect vertical portions of the supply curve all demand the exact quantity associated with the discrete change in price (Castro-Rodríguez et al., 2002). Borenstein finds no such observable bunching, nor does Ito (2010, 2014) with one data set of Southern California electricity usage and one of Southern California water usage. This behavior is also absent from the data used in this study.

With the growing emphasis on identification, quasi-experimental approaches have begun to find applications to water demand within the past decade. Defining households pre-price change as the control, Klaiber et al. (2014) and Yoo et al. (2014) use a single difference approach to identify change in water usage as a function of changes in weather and price. Because the elasticity is a function of the estimated regression *constant*, the identification of the elasticity in Klaiber et al. (2014) is threatened by potential omitted variable bias, particularly if things other than price are changing over the same time period. By estimating the elasticity as a function of the change in prices, Yoo et al. (2014) lessen the dangers of omitted variable bias, but their results are still susceptible to it, since a single-difference identification strategy requires some rather strong assumptions about other variables changing at the same time in order to interpret it as identifying a true "treatment effect".

One of the best identified estimates of water elasticity in the literature comes from Nataraj and Hanemann (2011). The authors exploit the creation of a new block in the rate structure, using households just below the cutoff of 40 ccf in summer water usage as a control group for households just above 40 ccf. This identification strategy limits their analysis to one point in the usage distribution. Employing a difference in differences strategy in a regression discontinuity framework they find these high-usage households to have a short-run marginal price elasticity of -0.12. In another quasi-experimental study, Wichman (2014) defines "treatment" as gaining information about the new price structure contained in the bill, and uses the schedule of meter readings to define households that receive this information later in the month as a control for those that received it early. This limits him to a single month of analysis where he uses a diff-in-diff approach with decile dummies to consider heterogeneous treatment effects across the usage distribution. In related literature about electricity demand, Ito (2010) uses an ideal "natural experiment" data set of households within a mile of the border separating two electrical utilities with different prices to perform difference-in-difference estimation. His identification comes from comparing usage across the utility boundaries and allows considerably more freedom to explore heterogeneity across space and time. Another quasi-experimental strategy employed in Baerenklau et al. (2014) uses pre-price change usage histories to generate a predicted usage pattern, and then to use these predictions as the control group.

1.2.2 Price Elasticity Heterogeneity

The past decade has also seen significantly more interest in the heterogeneity of price elasticity across user groups. Heterogeneity is typically measured across one of two dimensions: income and usage. Empirical results tend to find that the price elasticity of water demand is decreasing with income (Agthe and Billings, 1987; Renwick and Archibald, 1998). Mansur and Olmstead (2012) analyzed detailed household data and found that the poor have the lowest indoor usage elasticity and the highest outdoor usage elasticity. These results are interpreted by Ferraro and Price (2013) as suggesting that households using the most water have the lowest price sensitivity.

Almost all water papers reporting on heterogeneity in usage have found a monotonic relationship with the lowest users being the most price sensitive. Baerenklau et al. (2014) divided the sample into low, medium, and high usage households and estimate responses for each, finding that the more a household used the more they would reduce in levels. However, when converting from levels to an elasticity, the relationship flips such that the high usage households had the least elastic demand. Analyzing aggregate water usage data from Phoenix, Arizona Klaiber et al. (2014) and Yoo et al. (2014) report that the largest users are the least price sensitive. Klaiber et al. (2014) define the dependent variable as change in the τth percentile within a Census block group (roughly 600 homes).

Yoo et al. (2014) uses quantile regression to analyze the heterogeneity in price elasticities, reporting that the price sensitivity is decreasing with higher quantiles. This is, however, a result of their specific modeling approach and is not consistent with the interpretation that larger water users are less sensitive. Quantile regression necessarily looks at the quantiles of the dependent variable, which in their case is *differenced* water usage. Thus the "10th quantile" corresponds not with bottom 10% of water users, but with the 10% of users who had the largest magnitude decrease in water usage. Because the dependent variable upon which the quantiles are based is the *change* in water usage, it follows that the results must indicate that higher quantiles have smaller price elasticities.

In the presence of Phoenix's seasonally varying water prices, Klaiber et al. (2014) and Yoo et al. (2014) both require the assumption that water usage is a function *only* of today's price. Water usage, however, is dependent on habit formation, landscaping choices, and appliance efficiency so the history of prices may play a significant role. Investment in drip irrigation, xeriscaping, or high efficiency washing machines is likely to happen when water prices are high. If prices subsequently fall, it is only rational for households to bear the fixed costs of converting back to water-intensive capital if it provides more utility *and* prices are expected to remain low. This does not seem to be a reasonable assumption since most water-saving technologies provide similar levels of functionality and water prices are generally expected to continue increasing. The finding of higher winter elasticities in Klaiber et al. (2014) may be a result of attributing the usage decreases that were spurred by summer prices to the smaller winter price differences. Such an explanation is supported by Baerenklau et al. (2014) where a price *decrease* does not lead to any increase in water usage.

The exception to the rule that the lowest usage households were found to be the most price sensitive is Wichman (2014), which finds a non-monotonic relationship. Using a variant of a triple difference strategy to identify the treatment effect across deciles, he finds that the lowest and highest deciles of households do not respond to the price increase. The price elasticity is found to be generally increasing from the 1st decile through the 8th decile, and decreasing for the 9th and 10th deciles. One concern with estimates based on binned water usage is that mean-reversion by households in the upper and lower bins may result in bias by attributing this mean-reversion to the price change. Ultimately, the heterogeneity of price elasticity as a function of usage remains an open question in need of better estimation.

1.3 Theoretical Underpinnings

Understanding consumption of water and electricity is a complex issue with many deviations from standard consumer demand models. New models were created to allow for the increasing block rate structure in Taylor (1975) and Nordin (1976) who find that an optimizing consumer will set $MB(Q) = MC(Q)$ and should ignore inframarginal prices, treating the lower prices on earlier units of consumption as an income subsidy. This elegant theoretical solution faces a number of problems in practice. The assumption that households know the utility rate structure has been found to be very unlikely, with both surveys and empirical studies finding that consumers are unaware of or not responsive to the marginal cost. A more subtle issue is that electricity and water use incorporates elements of derived demand and can be susceptible to violations of the non-satiation assumption.

1.3.1 Knowledge of Price and Quantity

Standard demand theory assumes that consumers are conscious of the price paid and the quantity consumed. But with household utilities there are a number of reasons to question this assumption. The price structure is complicated and the units of measurement are abstract, making it difficult for a household to compare marginal benefits and marginal costs.

The marginal rate is hard for many consumers to identify on a utility bill, and it is rare that utilities include on the bill/insert information about rates in tiers above the household's current highest tier. Motivated consumers can usually find this information somewhere on the utility's website. The burgeoning popularity of automatic bill payment make it all the more likely that households are not actively looking at their utility bills. Furthermore, for most households their information about usage and prices comes from the most recent bill, which could be a few weeks to months out of date depending on the billing lag and frequency. As a result, households may have accurate information about the past rate but not necessarily the current one. A major critique Borenstein (2009) levels against the use of *ex post* price as the determinant of usage is the difficulty of forecasting usage for the entire billing period. Because utility usage is often driven by weather shocks, an accurate prediction of usage (which determines price) would require at least 30 days of accurate forecasts for temperature and rainfall. Most weather sources will only report 10 to 15 day ahead forecasts, and those come with a relatively high degree of uncertainty, especially when extreme weather events are possible.

There are also likely to be significant issues regarding household understanding of the quantity consumed. The standard measure of household electric usage is the kilowatt-hour: the power consumption of one thousand watts for one hour. This is a relatively small amount of electricity, but not easily observed or understood. To assume that consumers are perfectly informed optimizers is to assume that they know how many watts all of their electrical devices use, how many hours they have been on during the billing period, and when their billing period begins and ends. Borenstein (2009) suggests that most households do not even know when their latest billing cycle began, which is the most likely to be known out of this list. Conversely with water, the units of kgals and ccf (1000 gallons and 776 gallons respectively) represent such large quantities of water that households are likely unable to identify "marginal" changes. Households are unlikely to know how many loads of laundry or minutes of watering the lawn represent a marginal increase in their bill or a jump point in a tiered rate structure. Thus even with knowledge of the marginal cost, it is tenuous to assume they are basing their consumption on the marginal benefit.

An alternative model developed in Borenstein (2009) attempts to address these uncertainties with a more realistic approach towards household behavior. He proposes that consumers have “behavioral rules” such as “set the thermostat to X during the day and Y at night” and “run the dishwasher every night”. These rules are set at the start of the billing period, ostensibly informed by prior bills and a forward looking estimate at price. Within this basic framework, I further propose that households will adjust their behavioral rules after a price increase, and this adjustment may take multiple periods as they are using a trial-and-error approach to getting the utility bill to the level they desire. This insight will be important for interpreting the results.

1.3.2 Derived Demand and Capital Investments

Another important theoretical feature of utility demand is that many of its uses are not direct but rather derived from services the utility can perform. Because households use electricity for “refrigeration services” and “laundry services” the demand for utilities is the result of the quantity of utility services demanded *and* the efficiency of appliances which provide these services. Therefore the observed price elasticity is the aggregate effect of price changes on the quantity of services and changes in efficiency. Disaggregating this effect will require detailed data that may hopefully be available for future research, but even in this setting there is still benefit to considering what impact this may have on household behavior.

Changes in the services demanded represent the standard movement along the demand curve, as households take shorter showers or adjust their thermostat to save money. These changes are available to the household immediately and at no financial cost. If this represented the entirety of changes in utility usage after a price change, then there would be no dependence on the history of prices and one would expect to see demand return to normal after temporary price changes.

On the other hand, if households invest in more efficient appliances to provide utility services, then the entire demand curve for utilities will shift. Assuming that utility services have constant or diminishing marginal benefits, improving the efficiency will result in the demand for utilities shifting inwards as the same services can be provided with less engergy or water. Because the price increase plays a role in these investment decisions, reductions in usage caused by them should be considered as part of the price elasticity. However, the price elasticity measured in these cases is decoupled from the slope of the demand curve. If the price were to return to the original price, usage would remain lower than it had been before the price increased thanks to the efficiency improvements. Two challenges induced by this feature are that the price elasticity becomes dependent on the history of prices and on the pace of technological development.

The first challenge, referred to as hysteresis or asymmetric price elasticities, has been observed in the energy demand literature (Agnolucci et al., 2004; Grubb, 2005). where temporary increases in price result in permanent changes in demand. The second challenge implies that the effects of every price change are a function of the capital stock in place at the time of the price change and the efficiency of the available replacements. This implies that empirical measures of price elasticity are dependent on both place *and* time, severely limiting the external validity of estimates.

The decision to invest in capital upgrades is based on the net present value of the upfront costs and delayed benefits. For owner-occupied residences, the decision is subject to household discounting rates. The problem is even more challenging for renter-occupied residences, since often the capital stock is provided by the owner who does not bear the utility costs and it may not be possible to fully price the benefits of efficiency into the rent. In any case, higher prices should induce increased capital investment, but because the stream of benefits is small and long-lived this decision can be highly dependent on appliance prices which follow a predictable pattern (e.g. deep discounts on Black Friday). With the price increase observed here of approximately \$0.45 per 1000 gallons, a delay to purchase expensive capital upgrades in order to take advantage of anticipated sales would be fully rational.

1.4 Data

This study considers two communities known as Sun City and Sun City West that lie in northwestern Phoenix, Arizona. Sun City is an unincorporated town, founded in the 1960s as a Del Webb retirement community, and nearby Sun City West was created about 15 years later to accommodate demand as Sun City filled up. A satellite view of the two communities is presented in Figure 1.1. That they are both retirement communities is potentially important because households living on a fixed income and choosing to retire in one of these may be more conscious of their expenditures on water. On the other hand, retired individuals are likely to spend more time in their home, and are more likely to be active gardeners. However, an additional unique factor in these communities are the "snow birds" who spend their summers in cooler climates and only reside in the Phoenix homes during the winter. To accommodate this, only households which received at least 23 bills in the 24 relevant months are considered. As long as the two communities have reasonably similar numbers of "snow birds", this should not pose any additional problems.

With data spanning the period of 2005-2010, the study focuses on the price change in Sun City West that happened on December 1, 2006, while the price in Sun City remained unchanged. This rate change was an increase of \$0.45/kgal across the entire three-tier schedule, along with a \$3.70 increase in the service charge. Figure 1.2 displays the price schedules for the two communities.

Water consumption and prices were provided by EPCOR Water (previously Arizona American Water) for the water districts of Sun City and Sun City West, covering the billing periods from January 2005 through December 2010. Because of the usual delay between meter readings and billing, this represents the water usage period between late November 2004 and early December 2010. Meters at individual street addresses are read roughly every thirty days but bill start dates vary across addresses because meter readers must physically collect usage at each

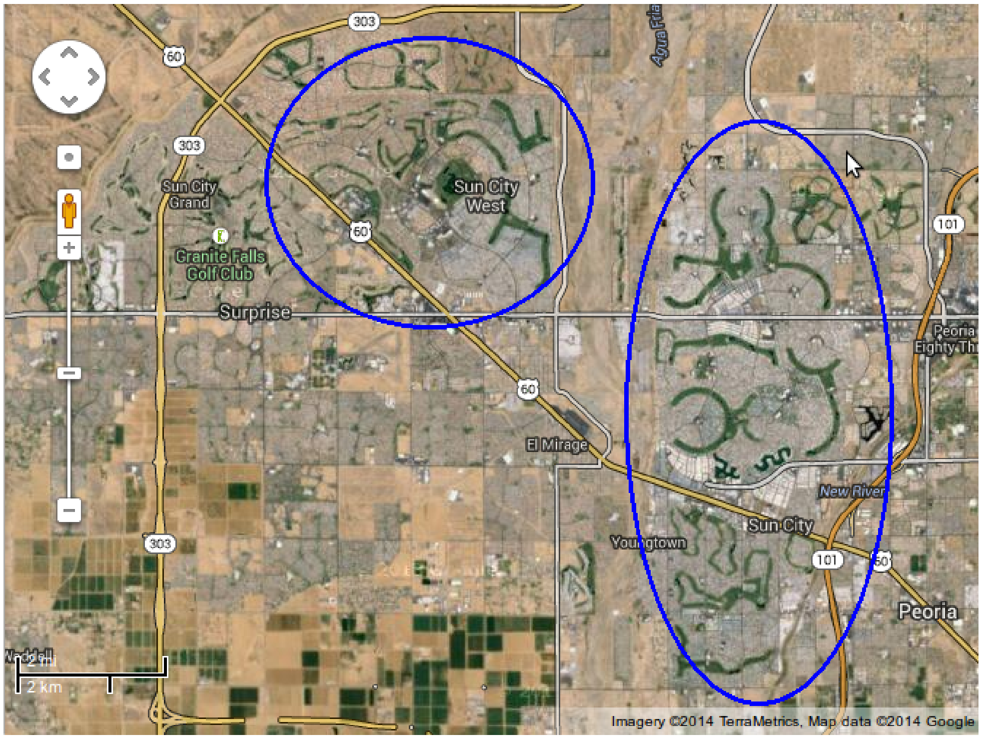

Figure 1.1: Map of Northwest Phoenix: Sun Cities Highlighted

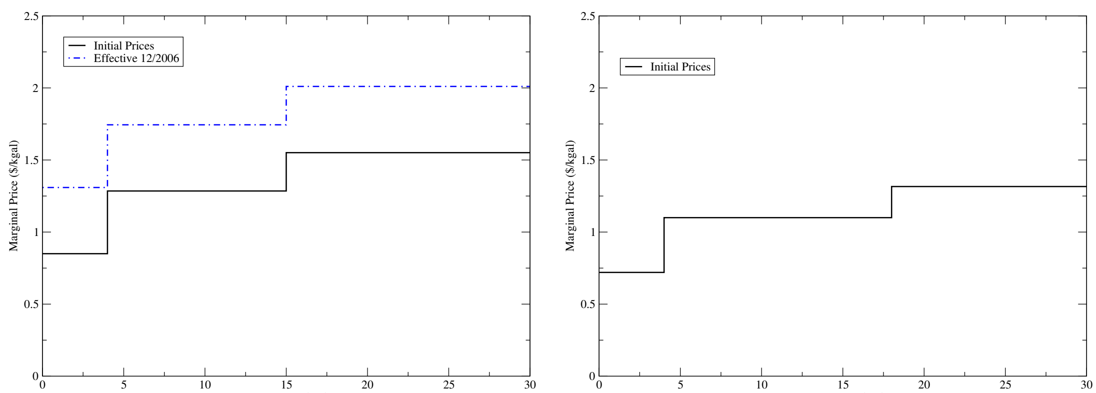

(a) Sun City West

(b) Sun City

Figure 1.2: Block Rate Price Schedule

residence. As a result, the data are asynchronous in that not all customers' bills begin on the same date each month. EPCOR Water supplied all water bill details, from which the total bill and usage amounts are calculated. The process of matching the addresses to the county assessor's database to get property characteristics effectively removed multi-family housing from the sample, since the water bill and property tax are usually not linked to the same physical address (e.g. the water bill is attached to unit 226, while the property tax assessment has no unit number). The two communities are primarily single-family housing, so relatively few households were lost.

Household characteristics data may be useful for estimating water demand at the household level, so data for residential and vacant properties (a residence recorded as vacant in 2011 may have been in use during the sample period) were collected from the Maricopa County Assessor's Office for the 2011 tax year. In estimating how much water a household may use, characteristics such as the size of the lot, and the size of a pool, if present, are naturally candidates for predictive variables.

As a measure of the irrigable area, yard size is calculated as $YardSize = LotSize - PoolSize$. The square footage of living area, was not subtracted from lot size because multistory homes resulted in negative computed yard sizes. Due to different landscaping choices, irrigable area can only proxy for irrigation demands. Landscaping in the Sun Cities ranges all the way from zero-water usage options like strictly gravel and artificial turf, up to gardens, citrus trees, and small grass lawns maintained both summer and winter (a separate winter grass seed can be spread and watered heavily in the fall). The dominant choice in both communities is low-maintenance landscaping, with gravel, and sparse vegetation which is frequently native drought-resistant plants such as cacti. A sample can be seen in Figure 1.4. Parks and small areas of grass are maintained as public property, with those water bills paid by HOAs. The large green areas seen in the satellite image of Figure 1.1 are most likely such public spaces and golf courses. Furthermore, the HOAs often pay for community pools, which is a major factor in the lower than average pool ownership (5-6% compared to 21.5% for Phoenix). These public green spaces and pools are not included in

Table 1.1: Sun City Household Statistics (N = 16,931)

	Mean	Min	Q1	Median	Q3	Max
Assessed Value ($)	104,887	25,047	85,500	102,100	120,800	602,700
Living Area (sq. ft)	1,544	388	1,236	1,536	1,827	4,382
Yard Size (sq. ft)	9,003	268	7,500	8,720	9,773	102,446
Pool Size (sq. ft)	424	111	364	418	450	960
Construction Year	1974	1954	1968	1972	1978	2009

Pool Size statistics exclude homes without pools.

Table 1.2: Sun City West Household Statistics (N = 13,282)

	Mean	Min	Q1	Median	Q3	Max
Assessed Value ($)	154,616	74,200	123,500	146,500	174,700	393,100
Living Area (sq. ft)	1,787	918	1,497	1,741	2,020	4,301
Yard Size (sq. ft)	9,615	234	8,700	9,300	10,400	28,500
Pool Size (sq. ft)	440	150	390	426	480	800
Construction Year	1988	1911	1983	1989	1994	2003

Pool Size statistics exclude homes without pools.

the sample of residential households.

Tables 1.1 and 1.2 summarize the household characteristics of Sun City and Sun City West. There are some differences due to Sun City West being newer, but what is necessary for estimation in this context is that the communities have overlap in their water usage rather than in their characteristics. There is significant overlap in the amount of water used as shown in Figure 1.3.

Weather data are collected from the Arizona Meteorological Network (AZMET) and Maricopa County Flood Control District (FCD). Potential evapotranspiration (ETo), which represents the amount of water lost to the atmosphere through soil evaporation and plant transpiration for a cool season grass 3-6” in height (Brown and Kopec, 2000), is included in place of temperature and humidity because the effect of temperature and humidity on water consumption comes from evaporation of standing water and need for increased watering as plants transpire more water.

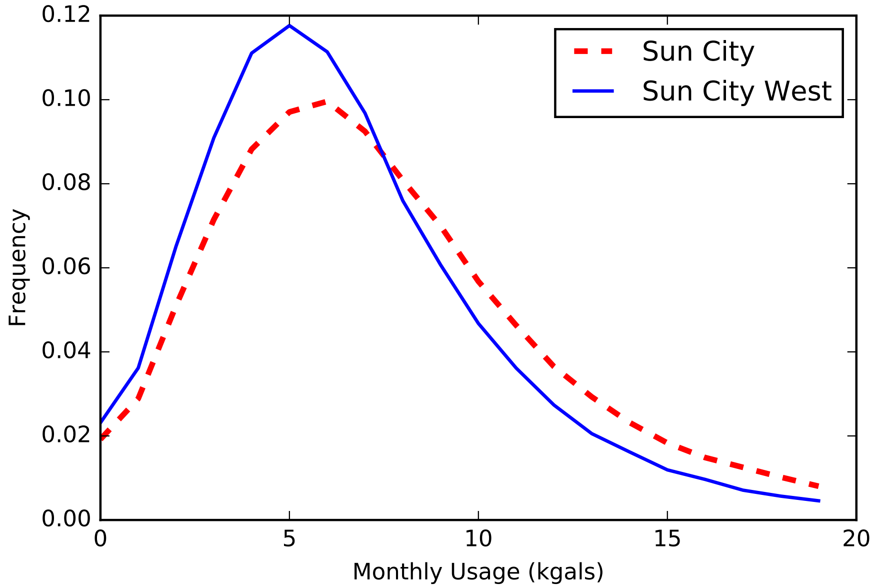

Figure 1.3: Distribution of Water Usage Prior to Price Change

Figure 1.4: Sample Sun City West Neighborhood

Both of these effects are more accurately measured through ETo. The "Original AZMET ETo" is calculated using a modification of the Penman equation proposed by (Brown, 2005). The stations report hourly ETo, as well as daily totals. Comparing Figure 1.6 with Figures 1.5 shows that seasonal movements of ETo and water usage are strongly correlated. This is not surprising since landscaping needs more water in hotter months. Households' adjustment to watering may not be perfect, but residents who desires a green lawn will increase watering after observing wilting or browning. Conversely, in the winter months they will decrease watering for the intuitive reason, or because they observe that their lawn is muddy long after watering.

AZMET also gathers daily rainfall totals; however, AZMET stations are on average less reliable for rainfall totals and more distant from the study residences than the FCD stations. Therefore the daily rainfall totals were collected from FCD monitoring stations centrally located in the study areas.

All weather data are collected on a daily basis, and in order to account for the varying starting dates and lengths of billing cycles, the total ETo, rainfall, and rainy days are calculated to coincide with the exact billing period for each customer's monthly bill.

1.4.1 Description of Data Generating Process

Underlying the water usage data is a set of irregularities imposed by the work-week schedule and billing units. Because of weekends and holidays, meter-readers will not get the reading on the same calendar day of each month, leading to irregular lengths of billing cycles. This can be easily addressed by including the length of billing cycle as an explanatory variable. More challenging is the particular measurement error structure imposed by the discrete billing units. The meter-reader reports that the meter is at 5 kgals whether the meter reads 5,000 gallons or 5,999 gallons, and then the next month reports the meter is at 10 kgals for readings of 10,000 gallons to 10,999 gallons. This means that the household could be billed for 5 kgals of usage if they had used 4,001 gallons up to 5,999 depending on the start and end points. Thus the data

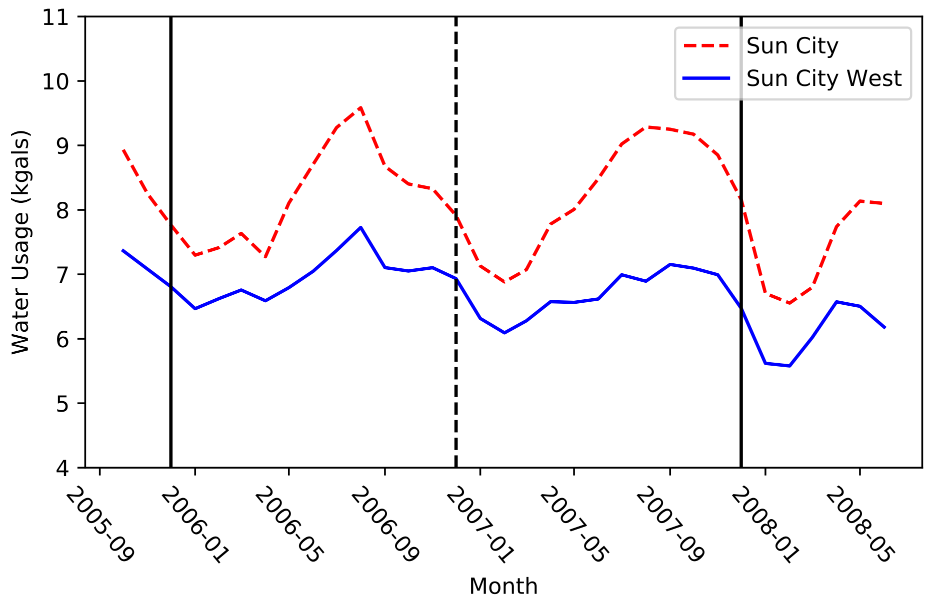

Vertical lines indicate months when the prices change (Solid = Sun City West, Dashed = Sun City)

Figure 1.5: Average Monthly Water Usage

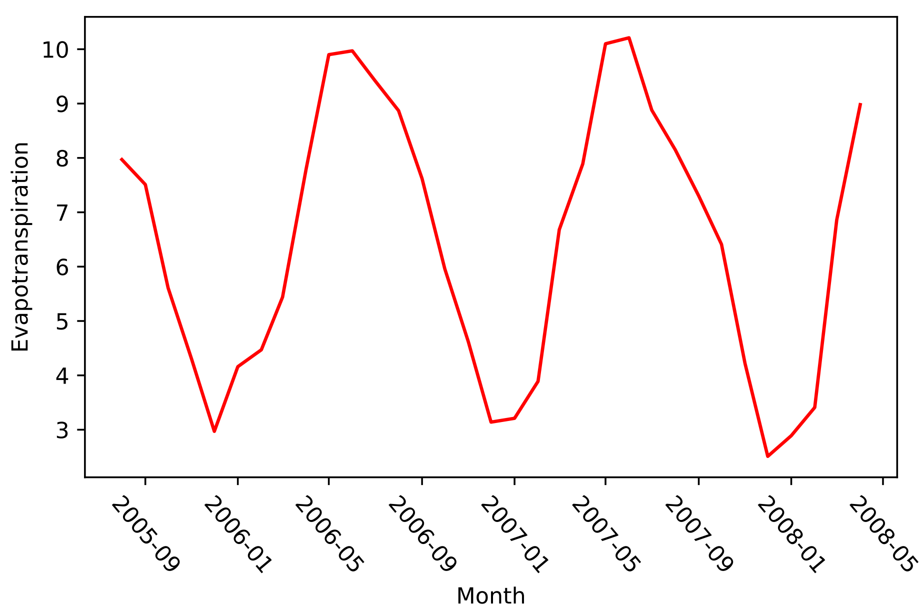

Figure 1.6: Comparison of ETo Calculation Methodologies (Shared Sun City/Sun City West Station)

could be described as having a lower bound of 0 with interval censoring[2].

However, simple interval censoring fails to consider the distinct impact this has on the error structure of the data. The mean error is still zero, since eventually households get billed for their usage. What this misses is the distinct temporal effect implied by the data. Because of the tipping point for a discrete measure of kgals, households will tend to get strings of errors that are of one sign, followed by a large error of inverse sign. For instance, a household using 1,900 gallons per month would pay for 1 kgal in one month followed by 2 kgal for nine months. Similarly, a household using 2,100 gallons per month would pay for 2 kgal for nine months followed by 3 kgal for one month. Because a single large error is likely to be followed by a small error of opposite sign this implies $E(\varepsilon_t|\varepsilon_{t-1}) \neq 0$.

The chosen estimation methods do not explicitly account for these error structures, but do ameliorate their impact on the inference. Taking the annual average of water usage means that the interval of the error is shrunk from $\pm$ 999 gallons to $\pm$ 999/12 = 83.25 gallons. The annual average also solves the issue with serial correlation as the errors between years will not be correlated in the same fashion. As for estimating monthly treatment effects, the variance of the jittering estimator increases when the conditional quantile rankings of the original usage are far from the conditional quantile rankings of the jittered usage, as would be the case when these measurement errors are large. And by taking observations that are not sequential (12 or 24 months apart) there is no longer any serial correlation in the subset of data being used for analysis.

1.5 Methods

A difference-in-differences framework is used as the general identification strategy for estimating quantile treatment effects. This paper's two goals are estimating the heterogenous treatment effects across households and estimating when and how households respond to the

[2]Models accounting for interval censoring produced similar results to standard methods, but quantile methods have limited support for interval censored data.

price increase. Each goal will require different estimating methods due to the discreteness of the dependent variable. The annual treatment effect is estimated across the distribution of households using the change-in-changes model with covariates (Athey and Imbens, 2006; Melly and Santangelo, 2015). In estimating the monthly quantile treatment effects discreteness becomes a larger issue, therefore I employ the method proposed in Machado and Silva (2005) to estimate quantile treatment effects for count data.

1.5.1 Identification Concerns for Difference-in-Differences

The typical use of difference-in-differences models for the average treatment effect requires parallel trends. Extending the identification technique to quantile treatment effects requires a different assumption: time invariance of the unobservables within groups, sometimes called rank similarity (Athey and Imbens, 2006; Melly and Santangelo, 2015). Time invariance requires that unobservable characteristics that determine an individual's rank have the same distribution over time. This is a weaker assumption than rank invariance which requires the the individuals rank in the distribution not change over time. While parallel trends are rather easy to test statistically and to demonstrate graphically, determining whether time invariance holds is more challenging. Several tests have been proposed, and the methodology proposed by Melly and Santangelo (2015) is used here. These authors suggest that with 2 pre-treatment periods, a placebo test is sufficient to validate the identifying assumption. Indeed the estimated QTE in the placebo model, shown in Figure 1.7 are almost everywhere not statistically different from zero. The identifying assumptions are valid if the placebo effect is 0 for all quantiles which cannot be rejected, thus validating the the use of a difference-in-differences methodology to estimate the quantile treatment effects.

Having established the time invariance, I need to show that the rate change was exogenous and that the communities faced similar pricing structures prior to the rate change. Both Sun City and Sun City West are served by the same water provider, with the Arizona Corporation

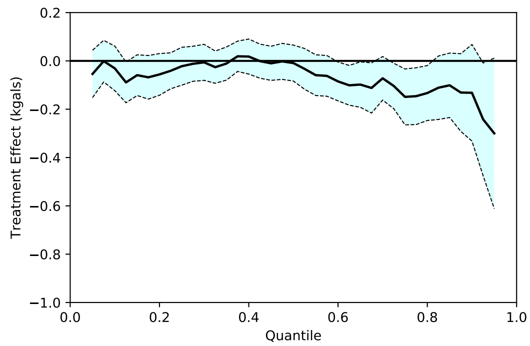

Bootstrapped 95% Confidence interval shown

Figure 1.7: Test for Time Invariance of the Unobservables (Placebo Test)

Commission governing the allowable price schedules. Rate hearings are every three years, and there is a one-and-a-half year gap between when the private water provider is allowed to lobby for rate changes in these two communities. Thus the *timing of the rate change* is purely a result of bureaucratic rules, and is suitably exogenous to both communities. The *magnitude of the rate change* is primarily a function of inflation and changes in costs of delivery. These factors are beyond the control of the two communities, so the magnitude is also plausibly exogenous.

Since prices across utilities and districts are rarely the same, one must define what "similar prices" means and to demonstrate that these requirements are met. The key consideration is that households should not have been induced to make significant technological or behavioral changes as a result in price differences. To justify a difference-in-differences methodology for electricity usage Ito (2014) claims that "Until the summer of 2000, SCE and SDG&E had nearly the same two-tier nonlinear price schedules". A closer look at this study shows that the marginal prices at Tier 1 are 22% higher for SCE, and at Tier 2 are 8% higher. Furthermore, using 21 kWh/day as the reference point with 10 kWh as the baseline, and ignoring fixed charges (none are mentioned in the article), the final bill for a 30 day month would be \$75.90 for SCE, and \$66.60 for SDG&E.

The marginal price charged in Sun City West is always higher than the price charged in Sun City as shown in Figure 1.8, but aside from the range of the 15th to 18th kgal, they are never more than 30 cents different. Since consumption for 93% of household-months were 15 kgal or less, the price difference was no more than 20 cents per kgal for the vast majority of households. At its greatest the price difference is \$0.45 per 1000 gallons.

However, Sun City faced a slightly higher service charge of \$7.41 compared to \$6.74 for Sun City West, so the total bill for Sun City is higher until 5 kgal as shown in Figure 1.9. Approximately 93% of all household-months in the two years of data prior to the first price change are between 0 and 16 kgal, a range over which the difference in the bills are negligible.

The difference in prices would be more of an issue if there was a higher level of usage in the two communities, but over 95% of households in both communities are using less than 20 kgals during the sample period. As such, households would not be expected to have adopted significant technological or behavioral changes due to existing price differences.

1.5.2 Econometric Methods

Addressing how price sensitive water users are over the range of *usage* is an ideal situation for quantile regression. The Quantile Treatment Effect is similar to a LATE, and in its most basic form is calculated as the τ-quantile of the treated distribution minus the τ-quantile of the untreated distribution. Thus, in the context of a difference-in-differences framework with a price change, it can identify the price sensitivity across the distribution of households on the basis of water usage rather than other factors such as income or yard size.

Change-in-changes with Covariates

The methodology of change-in-changes (CiC) was put forward by Athey and Imbens (2006) and improved by Melly and Santangelo (2015). It is used here to estimate the response to the price change in levels and converted into elasticities. A simple way to conceptualize the change-in-changes model is to imagine replacing the expected values in the traditional difference-in-differences model with the full inverse CDF of values. The method works to recover the whole distribution of the counterfactual outcome with a semi-parametric estimator, allowing for the computation of unconditional quantile treatment effects. The primary benefits of the method are the flexibility of the estimator and the validity of the empirical bootstraps for the estimated treatment effects. Another advantage is that time and groups are treated asymmetrically, whereas traditional difference-in-difference regressions give them symmetric weights.

Without covariates, the method works by taking an outcome y and associated quantile q from the pre-treatment period in the treatment group, then finding the quantile q' where y would fit in the pre-treatment control group distribution, finally computing the distance to q' in the counterfactual where the untreated group receive treatment. This is labeled δCIC in Figure 1.10. The model with covariates proposed by Melly and Santangelo (2015) works by performing this type of transformation for all $x \in \mathbb{X}$ and then integrating over $\mathbb{X}$ to obtain the unconditional quantile outcome.

A considerable challenge to the identifying assumptions of quantile regression in this context is the violation of the strict monotonicity requirement resulting from discrete outcome variables. Fortunately the method is flexible enough to allow for discrete data with some additional considerations, and the method performed well with simulated data[3] even with the discrete outcome variable. Some of the discreteness is alleviated by computing the annual household water usage, but this ultimately only multiplies by 12 the number of possible outcome variables.

A further reason for computing the annual average household water usage is that using monthly data would not lead to appropriate interpretation in a quantile model. Because the dependent variable would be “household-months” then the quantile treatment effect would be ranked over these household-months rather than over households alone. Using the average instead of the sum simply allows the analysis to be in terms of changes per monthly bill to aid in comparability with the month-by-month analysis that follows.

A description of the individual covariates will be useful. The first covariate is the average number of days in the billing cycles in each of the two years, because the discrete measurement

[3] Information about the simulation available by request.

of water usage can be impacted by the length of the billing period. It would also be useful to account for weather. However the two communities share a single monitoring station for evapotranspiration and a single monitoring station for rainfall, so the household average annual weather impacts are practically identical, and thus washed out by the difference-in-differences estimator. While annual weather is not of use here for estimating the treatment effect, it turns out the estimated effects of monthly weather on monthly household water consumption, a proxy for landscaping choices, is a useful covariateThis result was determined by the simulations.. The sensitivities are estimated as the coefficients on ET and rainfall from a household-level regression of monthly usage on evapotranspiration, measured rainfall, days in the billing period, and the time dummy.

As in the empirical chapter of Melly and Santangelo (2015), no estimating equation is presented because the estimating process contains multiple steps which are not made more enlightening by listing covariates and details specific to this application. The quantile treatment effects are estimated from the 5th percentile to the 95th percentile with a step of size of 2.5, and the empirical bootstrap was iterated 1000 times. To account for the panel nature of the data, the bootstrap selects both time periods for a household.

Month by Month Quantile Treatment Effects

In addition to the question of how the average household usage changes in response to the price increase, I investigate when the household usage changes. To perform this analysis the data is restricted to bills which are 12 (or 24) bills apart, with one before the price change and one after the price change. By working specifically with the bill number, I avoid problems where multiple bills occur in the same month due to short billing cycles (i.e. one bill on the 1st, one bill on the 31st). This also treats each bill after the price change as a chance for households to update their knowledge of the price and their usage, and adjust their behavioral rules accordingly.

Narrowing in to look at the usage on a month by month basis, the CiC method ceases to perform well in simulations as discreteness is a greater problem. The preferred model for the month by month analysis is quantile regression of count data as proposed in Machado and Silva (2005) and implemented in STATA with Miranda (2006). This quantile model is estimated using the exponential function much like in the traditional Poisson count model. In addition to using the exponential function, this procedure generates a set of random draws from a uniform distribution which they call "jitters," which are then added to the discrete variable before running quantile regressions. This process introduces noise for the technical purpose of artificially satisfying strict monotonicity, which could potentially alter the outcomes. For this reason, they suggest repeating this process m times and averaging across these outcomes which generates a covariance matrix

$$V^A = \frac{1}{m}\underbrace{D^{-1}AD^{-1}}_{Artificial\ Noise} + (1-\frac{1}{m})D^{-1}BD^{-1}$$

with $\frac{1}{m}$ weight on the nuissance factor, and $(1-\frac{1}{m})$ weight on the covariance matrix, suggesting that a number of repetitions will lead to reliable estimates of both coefficients and standard errors.

The principal empirical assumptions necessary for identification in their model are:

(A1) $Y \in \mathbb{N}_0$, and the conditional probability function of Y given $\mathbf{X}$ is uniformly bounded away from 0 for almost every realization of $\mathbf{X}$.

(A2) At least one of the covariates is continuously distributed

With (A1) and (A2) satisfied, one can generate a Uniform [0,1) "noise" term U, and define $Z = Y + U$ to artificially create a continuous dependent variable.

With respect to these data, the Y of water usage is from the set of non-negative integers, and for almost every realization of $\mathbf{X}$ (billing period length, household characteristics, evapotranspiration, rainfall, and time/treatment dummies) one would not predict zero water usage. The weather variables and household characteristics are continuously distributed, satisfying the second assumption.

While it would be ideal to incorporate the information inherent in panel data in estimating the quantile treatment effects, fixed effects estimators for quantile regressions can alter the interpretation of the results. This is most clearly seen by de-meaning the data, in which case the distribution from which quantiles are computed is no longer $F(Y_i)$, but $F(Y_i - \bar{Y})$[4].

To inform the discussion of the quantile treatment effects, a standard fixed effects model is estimated according to

$$w_{it} = \alpha_i + \alpha_{(t \geq 12)} + \gamma(\mathbb{I}_{i \in SCW})(\mathbb{I}_{t \geq 12}) + \beta \mathbf{X}_{it} + \varepsilon_{it}$$

where w_{it} is the number of kgals of water used by household i in month t, with γ representing the average treatment effect. For all panel models, the vector of covariates $\mathbf{X}_{it}$ includes the number of billing days for household i in bill month t is included, as either a single variable or a set of dummies. The second model includes weather variables (evapotranspiration and rainfall during household i's billing period t), and third model adds the weather variables interacted with property characteristics $\{yardsize, poolsize\} \times \{rainfall, evapotranspiration\}$. To address how the response to the price change evolves over time the same model is estimated on the subset of the data with one bill taken after the price change, and a comparison drawn by taking the twelfth (or twenty-fourth) bill prior.

Moving to the estimation of the quantile treatment effect, the model is generally specified as

$$\begin{aligned} Q_{Z_{it}}(\tau|\mathbf{x}) &= \tau + \exp(\mathbf{x}'\gamma(\tau)) \\ \mathbf{x}'\gamma(\tau) &= \alpha_{(t \geq 12)} + \alpha_{i \in SCW} + \gamma(\mathbb{I}_{i \in SCW})(\mathbb{I}_{t \geq 12}) + \beta \mathbf{X}_{it} \end{aligned}$$

where $Q_{Z_{it}}(\tau|\mathbf{x})$ is the τth quantile of the conditional distribution of $Z_{it} = w_{it} + U_{it}$ with U_{it} being

[4]A method for addressing this concern is proposed in Powell (2016) however the simulations indicate that the count data poses a serious problem to identification using this method.

generated from the uniform distribution. $\mathbf{X}_{it}$ includes the length of the billing cycle, weather variables and interactions, as well as time invariant household characteristics living area, yard area, pool size, assessed value, age, and age squared. Note that billing days are no longer included as a set of dummies, since these results are highly similar[5], but wreak havoc with standard errors at low quantiles. Furthermore due to findings in the fixed effects model, all quantile models include the weather interaction terms.

[5] Available from author upon request

1.6 Results

1.6.1 Average Treatment Effects

I start with a simple model estimating the average treatment effect as the baseline for later analyses of heterogeneity and the process of model selection. The number of days in the billing period is an important variable because meter readers come out to the household at irregular intervals and the amount of water used per bill will mechanically increase as the number of days during the billing period increases. The sample has been restricted to bills between 25 and 35 days, since bills outside that interval were rare and may be indicative of a non-billing related issue. The most flexible solution to this problem is including a set of dummy variables for each billing period length in order to shift the intercept accordingly. Analyzing the estimated coefficients on these dummy variables reveals a relationship that is highly linear, and so the plausible parsimonious alternative specification is to include the number of days in the billing period as a numeric variable.

Because differences in weather between the two sample years may drive differences in usage, the estimates should be improved by inclusion of weather covariates as controls. The weather variables used are inches of rainfall and "Potential Evapotranspiration", a combination of temperature, humidity, intensity of sunlight, and wind which measures the weather's impact on plants. Most prior studies have assumed an effect of rainfall and evapotranspiration on usage that is homogeneous across the sample. However, one may see a heterogeneous response based on the yard size as well as pool size (the data includes pool surface area in square feet). Therefore, interaction terms are added with $ET * YardSize$, $ET * PoolSize$, $Rain * YardSize$, and $Rain * PoolSize$, which for notational simplicity I will often refer to as $Weather \times PropertySize$.

So the full estimating equation in levels is given by

$$
\begin{aligned}
W_{it} =& \alpha_i + \beta_1 I_{t>12} + \beta_2 \left(I_{i \in SCW}\right)\left(I_{t>12}\right) + \beta_3 DaysInBill_{it} \\
& + \beta_4 Weather_{it} + \beta_5 \left(Weather_{it} \times PropertySize_i\right) + \varepsilon_{it}
\end{aligned}
$$

In light of the count nature of the data, specifically with a number of zeros in both dependent and independent variables, to estimate a model in percentage terms, a fixed effects Poisson model is also presented.

The results of the fixed effect model are presented in Table 1.3. They show that on average households reduced their usage about 300 gallons, or 4% in the Poisson model, per month in the year following the price change. The chosen method of handling number of days in the billing cycle makes little difference in the estimated effect, and both methods produce nearly identical R-squared values. The weather variables improve the R-squared and alter the estimated treatment effect by a small amount. However adding the weather/property interactions leads to only small changes in model fit and estimated treatment effect. Due to a strong theoretical justification and the empirical evidence, (3) will be the preferred model, including the weather interaction terms and a numerical variable for days in the billing period.

The average usage in Sun City West before the price change was 7.15 kgals, therefore the treatment effect translates to a reduction in usage of approximately 4%, matching the results in the Poisson model. Using these results it is possible to compute price elasticities for marginal price, average total price, and average variable price. For the marginal price the percent change in price is 54.1% at the first tier, 35.8% at the second tier, and 29.7% at the third tier Therefore, assuming homogenous elasticities, it comes to -0.07 at the first tier, -0.11 at the second tier and -0.14 at the third tier. For the average total price (including fixed service charge), the average usage of 7.15 corresponds with an average price before the change of \$2.01 and \$3.00 after (since there was a non-trivial increase to the fixed charge). This represents a price increase of 50%,

Table 1.3: Fixed Effects Estimates of Treatment Effect

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment Effect	-0.325***	-0.308***	-0.304***	-0.321***	-0.292***	-0.289***
(In Levels)	(0.0363)	(0.0362)	(0.0361)	(0.0363)	(0.0362)	(0.0362)
Observations	713,976	713,976	713,976	713,976	713,976	713,976
Days in Billing Period	Linear	Linear	Linear	Dummies	Dummies	Dummies
Weather Vars		YES	YES		YES	YES
Weather × Property Size			YES			YES
R-squared	0.014	0.023	0.025	0.014	0.023	0.025

Robust standard errors in parentheses, N = 29,537

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

Table 1.4: Poisson Fixed Effects Estimates of Treatment Effect

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment Effect	-0.0457***	-0.0428***	-0.0427***	-0.0451***	-0.0400***	-0.0398***
(In Percent)	(0.00480)	(0.00478)	(0.00478)	(0.00479)	(0.00478)	(0.00478)
Observations	601,115	601,115	601,115	601,115	601,115	601,115
Days in Billing Period	Linear	Linear	Linear	Dummies	Dummies	Dummies
Weather Vars		YES	YES		YES	YES
Weather × Property Size			YES			YES

Robust standard errors in parentheses, N = 24,898

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

producing an elasticity estimate of -0.07.

Estimating a Poisson model to properly handle the count data, Table 1.4 shows that treating the water usage as continuous did not bias the results for the average treatment effect. Interpreting the coefficients as the percentage change, they correspond almost exactly with the level effects divided by the average usage of 7.15 as outlined in the preceding paragraph, estimating a 4.0-4.5% decrease in usage. Because of this similarity, the elasticity calculations are not repeated here.

The response path is also of interest. To address the adjustment period question, the data has been extended to 17 months after the price change. In the 18th month after the price change in Sun City West, Sun City experienced a price change of its own, at which time the difference-in-differences analysis is no longer valid[6]. To estimate the time path, the model is estimated for individual bills (months) after the price change. Bills after the price change are compared to the 12th bill prior for the first year, and the 24th bill prior for the second year. Since "bill 0" corresponds to December, and "bill 1" corresponds to January, for purposes of seasonality it is convenient to think of the bills as months.

The preferred baseline model is estimated in levels. The results are shown in Figure 1.11. There is a reaction period of 3 bills for which the response is statistically indistinguishable from zero, after which there appears to be a drop in usage. The negative and significant treatment effects estimated for the 12th-16th bill after the price change indicate that this is not just an effect of winter usage being lower. The spike in the 8th bill is unexpected, since it corresponds to August, one of the highest usage months. Excluding this 8th bill spike, the treatment effect has curvature corresponding with the seasons, implying reductions in outdoor watering.

[6]One could argue that the role of treatment and control groups could be switched and the analysis repeated. However it is not clear that the identifying assumptions are still satisfied, especially if the "treatment" of Sun City West had not stabilized and they continue to make new adjustments during this time period.

1.6.2 Quantile Models

Annual Effect

Quantile regression has the advantage of estimating the model coefficients for each level of the dependent variable (water usage). Results in this section employ the change-in-changes model with covariates for billing days, evapotranspiration sensitivity, and rainfall sensitivity.

Figure 1.12 shows the CiC estimate of the QTE expressed as the average monthly decrease in usage, with quantiles estimated from 0.05 to 0.95 at 0.025 intervals. Even at the lowest estimated quantile ($\tau = 0.05$) the QTE is significantly different from 0, with the reductions slightly increasing in magnitude at higher quantiles until increasing sharply after the 80th quantile. This stands in contrast with the claim of Ferraro and Price (2013) from their interpretation of Mansur and Olmstead (2012) that the highest users are the least price-sensitive.

Even at the 95th percentile the estimated reduction in usage is less than 1 kgal per month. Since each bill is measured in kgals this means that, after reducing their usage, the households would not observe their savings on a monthly basis but with even lower frequency. Almost 80% of households do not even observe the savings on a bimonthly basis, since that is the point at which the treatment effect exceeds 500 gallons.

To convert the level effects to elasticities is a simple computational exercise. The elasticity is computed as $\frac{\Delta Q(\tau)}{Q^{pre}(\tau)} * \frac{p^{pre}(\tau)}{\Delta p(\tau)}$, where $\Delta Q(\tau)$ is the quantile treatment effect for the τ quantile, $Q^{pre}(\tau)$ is the amount of water used at that quantile in the pre-treatment period, $p^{pre}(\tau)$ is the price paid for that quantile of usage in the pre-treatment period, and $\Delta p(\tau)$ is the change in price for that quantile of usage. The confidence interval is constructed by taking the lower and upper bounds of the treatment effect and running that through the formula. This structure allows for the computation of an elasticity on the basis of both average price and marginal price. Looking at the elasticity estimates in Figures 1.13a and 1.13b. As was the case in Baerenklau et al. (2014), the relationship inverts after converting the QTE to an elasticity measure as shown in Figure 1.13a.

The 5th quantile has an elasticity estimate of approximately -0.3, decreasing in magnitude up to the 20th quantile where it stays relatively flat in the range of -0.15 to -0.10. The estimated elasticity does not appear to depend on which variant of price is used. This is not surprising because the average variable price and marginal price both increased by the same amount. This is a feature of this particular case and not generalizable, so no claim about which measure of price is being perceived by households can be made. If anything it is reassuring that the experimental method reveals the same elasticities regardless of which of the candidates for perceived price is used.

These elasticity results indicate that consumers would be more accurately modeled with Cobb-Douglas utility than a linear or Stone-Geary utility function. The elasticity is essentially constant for much of the range, and in particular does not show the the asymptotic behavior at low consumption that would be implied by utility functions where a $1/Q$ term remains[7]. While the Stone-Geary assumption that consumers have some minimum water demand, which is necessary and hence not price sensitive, is almost certainly true of residential water, evidence here suggests that consumers are not up against this constraint. The low income, lack of significant irrigation, and small number of residents make this an ideal data set to observe such a "subsistence" level of consumption, making its absence all the more notable.

1.6.3 Monthly Evolution

While the preceding analysis looks at the full year after the price change, it is important to understand the process of household response to the change. To ensure the sub-sample used in the model most closely matches the desired effect, bills were numbered relative to the time of the price change (0 represents the bill with a pro-rated portion billed at the old rate and a pro-rated

[7]The Stone-Geary utility function has price elasticity of demand

$$\varepsilon_i = -\frac{1}{Q_i}\left[\beta_i\gamma_i - \frac{\beta_i}{P_i}(y - \sum_{j=1}^{n}\gamma_j p_j)\right]$$

portion at the new rate), with 1 representing the first bill entirely under the new price structure. The model also selects the bill 12 prior to the bill after the price change. Because bills represent inconsistent intervals of time (25-35 days) this is not identical to saying “the bill 12 months before”, but it is relatively close and weather covariates help to adjust for slight differences. Recall that the price change went into effect December 1st, so the x-axis approximates calendar months of the same number (1 is January).

Results for the 10th, 50th, and 90th quantiles are shown in Figure 1.14. It takes 3-4 bills after the price change for households to begin to reduce usage, with the largest reduction occurring 8 bills after the price change (roughly July-August). This stands in stark contrast with the fixed effects models which estimated a positive treatment effect for the 8th billing period due to a few outliers. The treatment effect continues to be negative and significant into the following winter, indicating that the lack of change early is more likely to be due to ignorance of the price change or delay in water-saving capital purchases, rather than a seasonally driven outcome.

For the 10th quantile, the change in usage for each month is a relatively tight zero in most time periods, with an unusually strong negative effect of -500 gallons for the December following the price change. However the effect returns to zero for the following spring months. Both the 50th and 90th quantile show statistically significant treatment effects in every billing period after the fourth month. Seasonality is rather subdued for the 50th quantile, although the 90th quantile shows a more evident summer/winter difference in treatment effect. This trend provides evidence that outdoor usage reductions are a large portion of the households’ response among the high usage households. Median households are likely making minor adjustments both indoors and outdoors to achieve the roughly 500 gallon a month reductions in the Summer.

These treatment effects can be used to calculate a price elasticity estimate for each billing month after the price change using a similar formula as above. Accounting for variance in usage and thus price across months, the formula becomes $\frac{\Delta Q_m(\tau)}{Q_m^{pre}(\tau)} * \frac{p_m^{pre}(\tau)}{\Delta p_m(\tau)}$. Elasticities were calculated for both average and marginal price, but due to substantial similarities (MP has slightly larger magnitudes but wider confidence intervals) only the average price elasticity estimates are presented in 1.15. The results indicate that in contrast to Klaiber et al. (2014), demand is slightly more elastic in the summer (bills 6-9) for the 50th and 90th quantiles. The results for the 10th quantile are rather erratic with wide confidence bands, but the lack of elasticity in summer supports the notion that outdoor water usage is not a meaningful component for such households.

To analyze the hypothesis that households are slow in responding to the price change, the full distributions of quantile treatment effects are presented in Figure 1.16. The figures show the quantile treatment effects for the first year with 95% confidence interval in blue, and for the second year with the 95% confidence interval in yellow. For the first two full bills after the price change, the estimated treatment effect in the second year is significantly lower than the first year for all but the lowest quantiles, with estimates of zero and even slightly positive effects in the first year. The third monthly bill after the price change shown in part (c) is interesting because the estimate is not statistically different from zero for the first year, while the second year estimates correspond very closely with the lower bound of the first year estimates. This appears to be the beginning of a convergence which continues in the fourth and fifth monthly bills, both of which show highly similar QTE curves for the two years with overlapping confidence intervals. This gives strong evidence that households are willing and able to reduce usage in the three calendar months following the price change, but did not do so in the first year. This delayed response makes the significant treatment effects found in Nataraj and Hanemann (2011) and Wichman (2014) all the more surprising, since they only estimate responses in the first bill following the price change. These studies are perhaps much helped by the bill occurring in summer, thus raising salience of the notifications. In any case, it is likely that their estimated price elasticities are an underestimate of the short-run elasticity.

1.7 Conclusion

By exploiting a price change in only one of two highly similar neighboring retirement communities in the suburbs of Phoenix, my analysis reveals that significant heterogeneity in the response to residential water rate increases, and interesting patterns in the evolution of these responses. The results add to the mounting evidence that low-usage households have more elastic demand, but the reduction in levels highlights that this is an artifact of the "percent change" element of elasticities. Although these analyses cannot lend any evidence regarding which price consumers respond to, it highlights the advantages of quasi-experimental methods which allow for computation of price elasticities for any perceived price. As a result of the form of rate increase (across the board $0.46/kgal), the choice of price variable does not lead to large differences in the estimated elasticity, which ranges from -0.1 to -0.3 across the distribution in this sample.

The time path shows a couple of interesting features as regards time to respond and seasonality. Most intriguing is the finding that households are slow to respond to a price change, with three to four months passing before statistically significant reductions occur. This slowness could be a result of households "learning" about the price change and its effect on their bill only through increased bills over a series of months, since a single month's bill increase could be attributed to randomness. It could also be a result of households being slow to adapt technological improvements, if for instance they need to save up for a new sprinkler system, or there is a delay in availability of supplies or labor for plumbing projects. It may even be a rational decision to wait for holiday sale prices on major appliances. Lastly it could result from households being slow to change their water-use habits, and any combination of the above.

The delayed response is an important finding in bridging the gap between economists who believe there is a price elasticity and regulators who believe it to be zero. Many utility commissions[8] still operate under the assumption that utility usage has zero price elasticity when considering rate cases. It is feasible to believe that they might look at data in the months

[8]Including the Arizona Corporation Commission which regulates the water provider examined in this study.

immediately following a rate increase and estimate that the elasticity is zero. This has significant ramifications, since the rate is set to give the fair rate of return under the assumption that usage will not change. When the usage decreases after the price change, the utility company fails to meet their fair rate of return and is forced to ask for more price increases in the next rate cases.

Looking beyond this adjustment period a clear trend of seasonality emerges. Reductions are at their greatest in August, with elasticity estimates of -0.2 and -0.4 for the 50th quantile and 90th quantile respectively, while November elasticity estimates are smaller than -0.2.

This study also suggests a number of considerations which are relevant to future utility research. Quantile regression is beginning to find more usage in applied economics, but are susceptible to misuse. First differencing and de-meaning of data, common means of leveraging the panel nature of the data, result in perverse effects when applied to quantile regressions. Furthermore, this study serves as a reminder that residential water demand is a count variable, a fact which should be incorporated in estimation techniques. The same holds for electricity, but with a more granular unit of measure for consumption (kWh) the impcats are likely to be smaller. The censoring which results in water-demand being integer valued may also lend itself to interval censored regression techniques.

From a policy perspective, I show that price increases can be effective at leading to broad-based usage reductions. Furthermore, the results suggest that "water hogs" are more responsive to price changes than average households in raw terms, even if it is a smaller percentage of usage.

1.8 Acknowledgments

Chapter 1, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Arizona American Water for their generous provision of the data for academic research. I am grateful for the help of Professor Gary Thompson in getting this project off the ground through the early data collection stages, and for his seed of an idea to employ quantile regressions to analyze the heterogeneity of price elasticities. Additional thanks to Professor Richard Carson and Professor Kaspar Wuthrich for help in selecting the appropriate quantile regression methods, and numerous seminar participants for their comments which helped shape this paper.

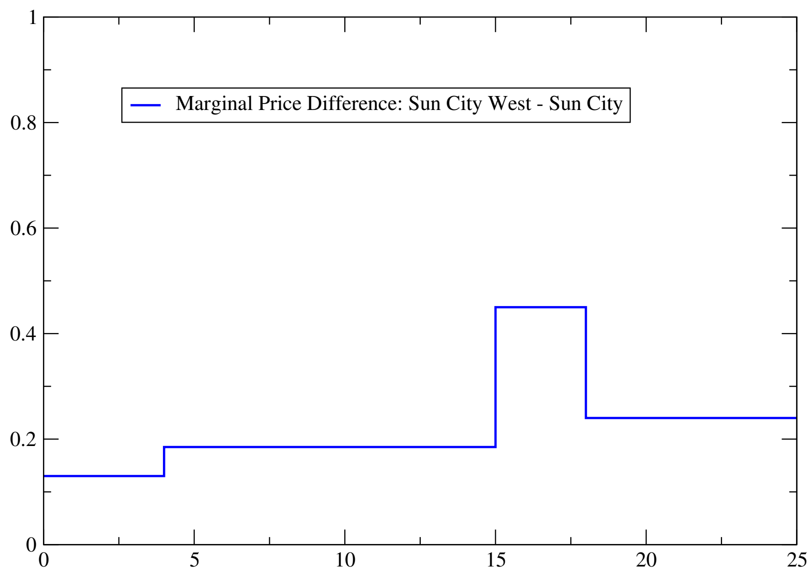

Figure 1.8: Marginal Price Difference

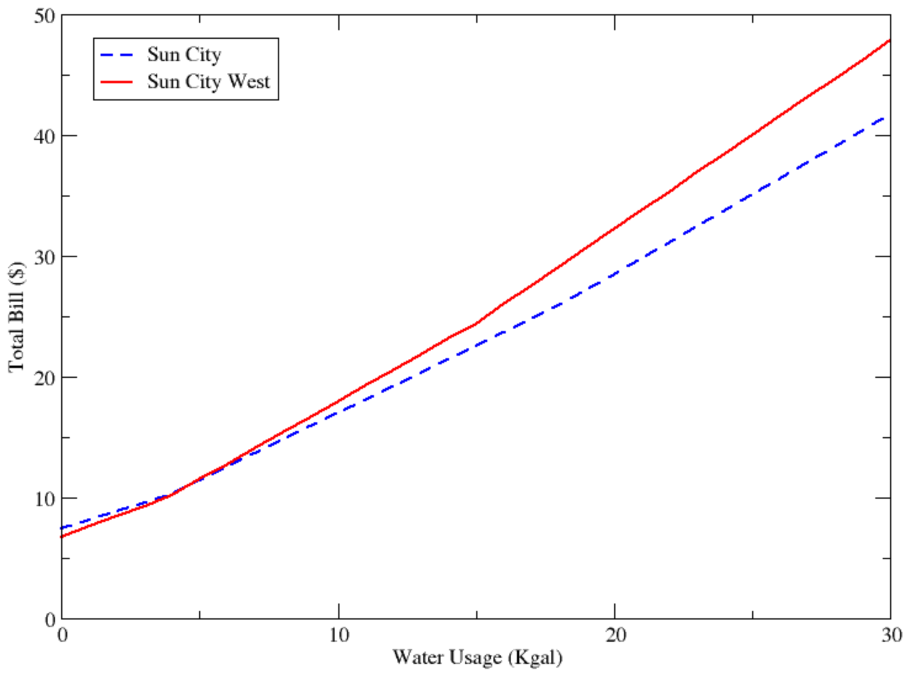

Figure 1.9: Total Bill Comparison

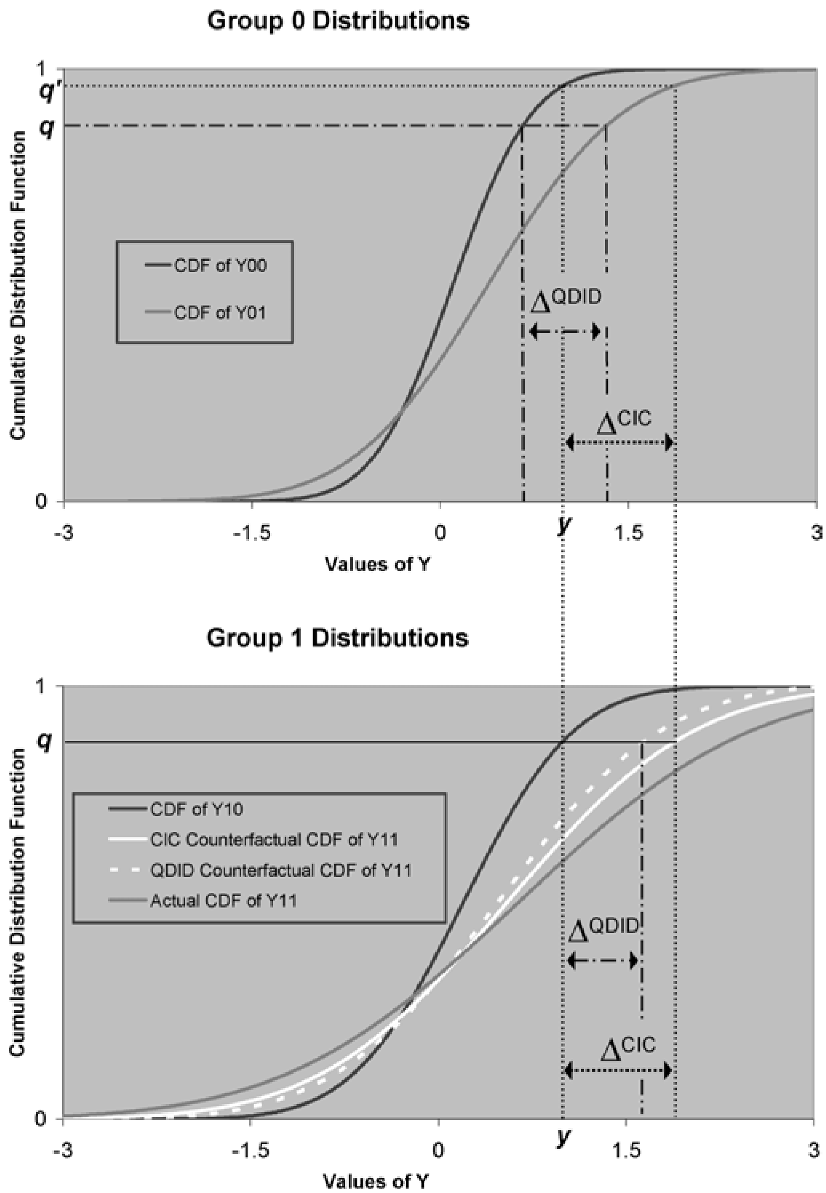

Figure 1.10: Illustration of CIC Transformation from Athey and Imbens 2006

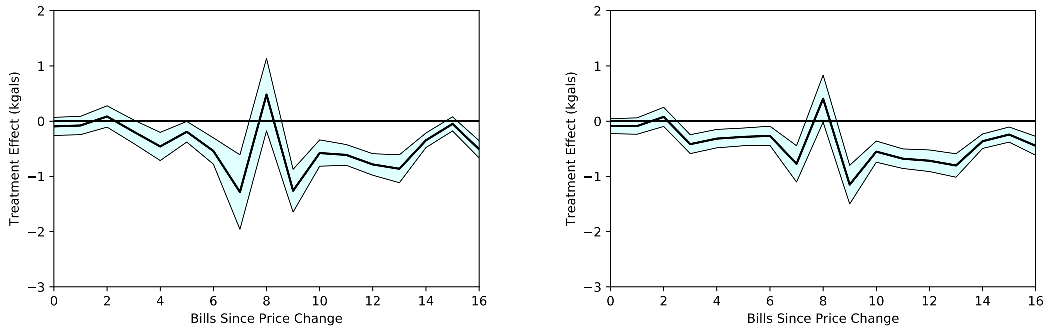

(a) Dummies for Bill Length Effect (b) Linear Billing Length Effect

Figure 1.11: Fixed Effects ATE Estimates by Bills After Price Change

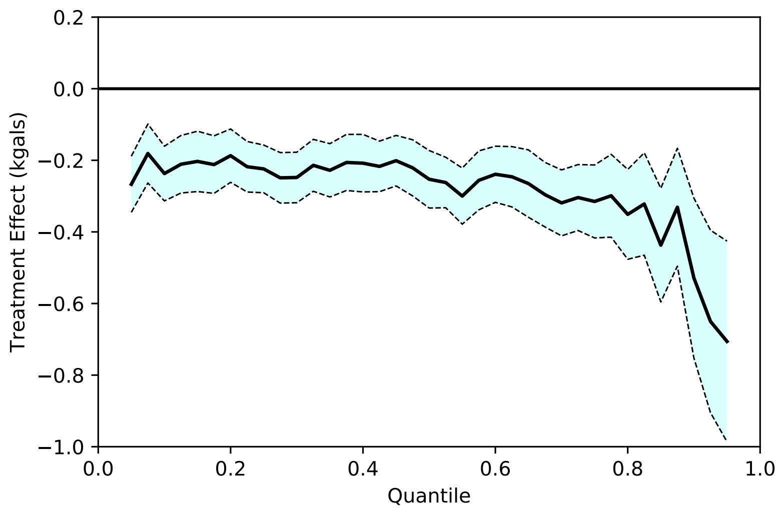

Figure 1.12: Average Annual Quantile Treatment Effects

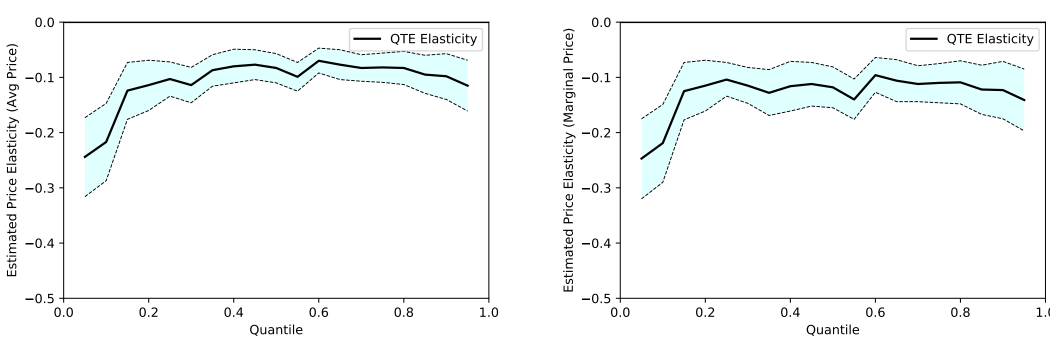

(a) Average Price (b) Marginal Price

Figure 1.13: Price Elasticities Computed from QTE

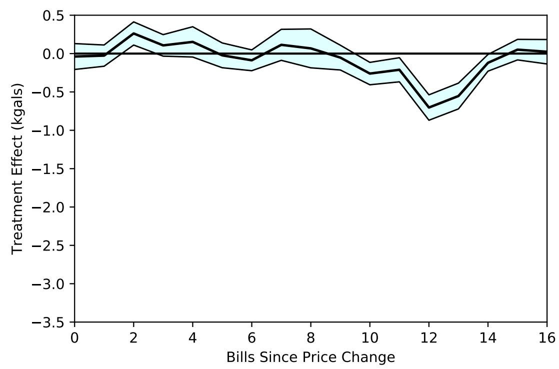

(a) 10th Quantile

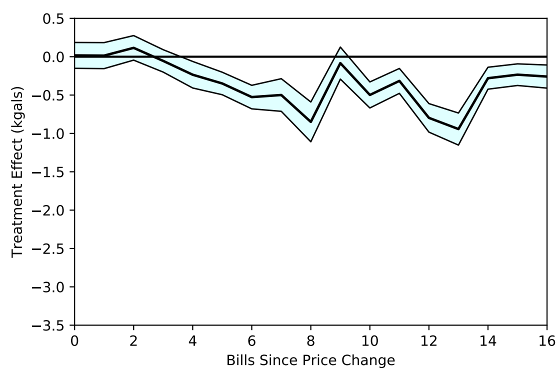

(b) 50th Quantile

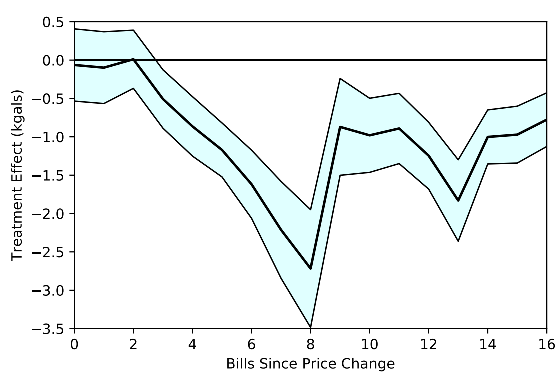

(c) 90th Quantile

Figure 1.14: Quantile Treatment Effects by Bills After Price Change

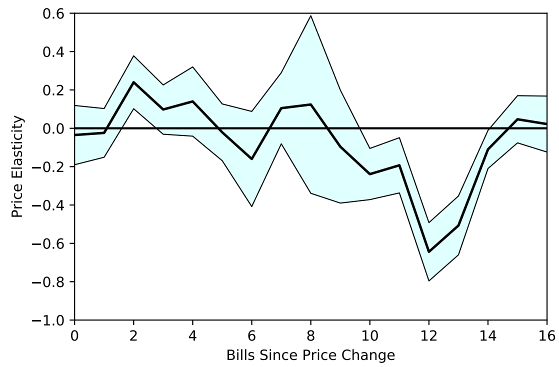

(a) 10th Quantile

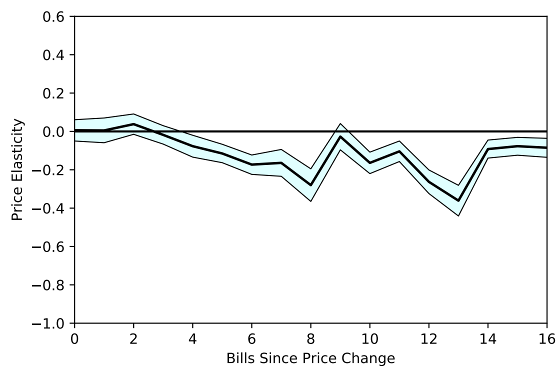

(b) 50th Quantile

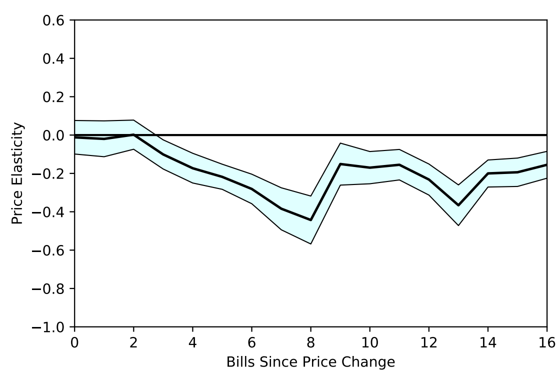

(c) 90th Quantile

Figure 1.15: Quantile Price (AP) Elasticities by Bills After Price Change

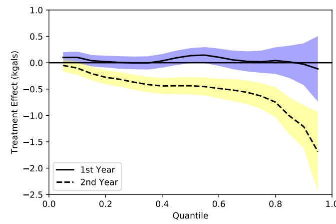

(a) 1 Bill After Price Change

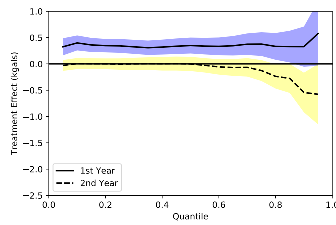

(b) 2 Bills After Price Change

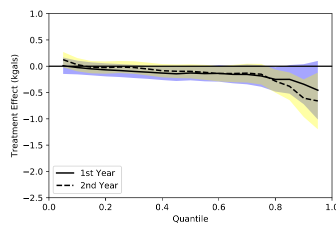

(c) 3 Bills After Price Change

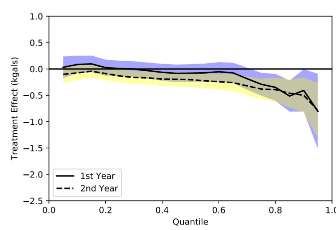

(d) 4 Bills After Price Change

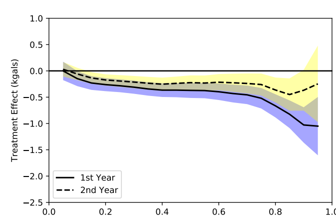

(e) 5 Bills After Price Change

Figure 1.16: Quantile Treatment Effects by Bills Since Price Change

Chapter 2

Stockpiling and Bargain-Hunting in the Demand for Eco-Labeled Canned Tuna

2.1 Introduction

The entire supply chain of eco-labeled seafood is built around the price premium that consumers pay in order to purchase seafood caught in a manner that is more sustainable for individual species and the broader ecosystem. While numerous papers have attempted to estimate this price premium, via both revealed and observed preferences, they ignore the critical fact that there is no “one true premium” in the market. Prices for food vary significantly over both time and space, with the same grocery store charging different prices week by week, and different grocery stores charging different prices in the same week. In addition to price discrimination by retailers, brand differentiation leads to a yet greater diversity of prices. Therefore this paper will seek to address how consumers respond to these ever-changing prices rather than studying the prices themselves, as the amount of revenue generated for transmission through the supply chain depends on the entire distribution of price premiums paid by consumers.

In the presence of heterogeneous willingness to pay for sustainability, this price discrimination allows the retailers to capture a greater share of the surplus, hopefully transmitting that benefit through the supply chain and leading to a better environmental outcome. In spite of this interaction between price discrimination and heterogeneous willingness to pay (WTP), existing research has largely focused on estimating the average price premium respondents claim to be willing to pay for eco-labeled and organic foods. Along with the willingness to pay in these stated preference surveys, researchers often evaluate what characteristics are predictive of that WTP. Relatively little is known about how consumers actually behave in the marketplace in light of the marketing mix employed for both conventional and eco-labeled products. Marian et al. (2014) examines the repeat purchase behavior of consumers of organic products vis-a-vis conventional products, finding some marked differences between the two. Specifically, products in the high-price tier show the greatest repeat purchase whereas organic products in the high-price tier show the least repeat purchase. Several studies have compared estimates of the price elasticity for conventional goods and organic/eco-labeled goods, such as Thompson and Glaser (2001), Glaser and Thompson (2000), Glaser and Thompson (1998), and Sun et al. (2017), finding that the demand is more price elastic in the organic/eco-labeled category. It is important, however, to note that the motivations underlying organic food purchases tend to be more health and wellness oriented, while environmental conscientiousness is more important for eco-labeled products. This paper seeks to build upon the literature regarding demand behaviors for "green" goods and gain a better appreciation for the consumers' strategic shopping behaviors and how this might influence the realized profits from greener production.

In order to estimate the promotional price elasticity, defined as the extent to which demand increases due to temporary price reductions, aggregate retailer data is used. The results show that the sustainably pole-and-line caught brand Wild Planet has a promotional price elasticity that is statistically indistinguishable from two of the three major conventional brands of solid white tuna, indicating that consumers of the sustainably caught brand are similarly strategic in buying more when prices are temporarily reduced. Cross-price elasticities reveal that the price of the sustainably caught brand is positively correlated with purchases of two of the three conventional brands, which means that either households substitute conventional brands for the sustainably caught brand when its price is high, or that they switch into the sustainably caught brand when it is on sale. However, all the prices of conventional brands are negatively correlated with the quantity sold of the sustainably caught brand. This is surprising, but it may be explained by sale markers for the conventional brands drawing consumers' attention to the tuna shelves, ultimately increasing sales for the sustainably caught brand.

In addition to the retailer data, a survival time analysis approach is applied to a panel of consumer scanner data to develop an understanding of what drives the interpurchase interval for sustainable and conventional brands of shelf-stable tuna. This analysis adds to the evidence that stockpiling is an important feature of the demand for shelf-stable tuna, with the sustainably caught brands showing an even greater delay in purchase for households which purchase larger volumes. In contrast with the promotional price elasticity results, in the case of the interpurchase timing, demand for the sustainably caught brands is significantly more price sensitive than for the conventional brands. This may be due to the difference in the dependent variables (quantity versus interpurchase interval) or due to the inclusion of more brands and quality levels of shelf-stable tuna in the consumer panel. The particular price decomposition used here provides evidence that, for the long panel of data, the consumer's reference price appears to be their previous purchase price rather than their sample average purchase price.

These results have strong implications for the use of eco-labels as a market-based solution to environmental degradation. The promotional price elasticity results suggest that retailers may be able to take advantage of heterogeneous willingness to pay for sustainably caught seafood, which could result in increased surplus being passed through the supply chain and providing stronger incentives for fishing vessels to choose sustainable methods. However, the storability of shelf-stable tuna and the stockpiling behavior allow households with higher willingness to pay to counteract the price discrimination by purchasing at a low price to consume much later. The analysis herein is hopefully a launching point for a better appreciation of how the complications of price discrimination and marketing strategies may interact with eco-labels as a market-based solution to environmental degradation.

The paper is organized as follows; Section 2.2 reviews three related branches of literature: demand for eco-labeled and organic goods, demand for storable goods, and interpurchase timing. Section 2.3 describes the data used for this paper. Section 2.4 reviews the methods and results of the empirical analyses used to estimate how promotional price sensitivity and stockpiling impact demand for eco-labeled shelf-stable tuna. Section 2.5 provides concluding remarks including areas for future research.

2.2 Literature Review

This paper contributes to the confluence of three disparate branches of literature on consumer demand. The first branch looks at the demand for eco-labeled and organic goods, and has largely focused on measuring the additional willingness to pay (WTP) and the characteristics of consumers for whom this extra WTP is positive. The second branch examines the demand for storable goods with emphasis on a household's ability to stockpile the good when prices are low. The final branch of the literature looks at the interpurchase timing for grocery products. This literature attempts to understand the regular and irregular patterns in consumer purchases in order to more accurately predict future purchase behavior in order to inform marketing decisions.

2.2.1 Demand for Eco-labeled and Organic Foods

Organic food and eco-labels are relatively new to the market. The research in this arena has focused on identifying how much more consumers are willing to pay for the goods and what demographics and attitudes are correlated with this WTP.

Eco-labels and organic labels address fundamentally different issues, and consumers' motivations for purchasing may be quite different. Eco-labels are designed to inform consumers that the production process is in some way better for the environment; therefore the consumer purchases it with an altruistic goal of preserving the natural world. On the other hand, organic foods are marketed as having less chemicals and more nutrition; thus a self-interested health goal may play a larger role, although concern for the environment is often believed to be an element in organic purchase decisions. As Yiridoe et al. (2005, p. 198) summarize,

> Some studies reported health and food safety as the number one quality attribute considered by organic produce buyers, followed by concern for the environment, suggesting that such consumers might rank private or personal benefits higher than the social benefits of organic agriculture.

The literature on organic products is large, so the review here concentrates on goods with some degree of storability, as these bear more relevance to the issues being considered. Thompson and Glaser (2001) and Glaser and Thompson (1998) examine baby food and frozen vegetables, and in both cases find that the demand for organic products is more elastic than conventional products. Furthermore, the cross-price elasticities imply that consumers are likely to switch from organic to conventional when the price gap expands but are unlikely to switch from conventional to organic when the gap decreases. Marian et al. (2014) finds opposing patterns in the effect of price on repeat purchase of conventional and organic products. This work finds conventional products in the high price tier show the most repeat purchases, while organic products in the high price tier show the least repeat purchases. However, they also find that consumers who purchase organic products in the high price tier are more likely to purchase other high-price organic products.

Relatively few papers use revealed preference methods to evaluate demand for sustainable or environmentally friendly seafood. One of the earliest of these papers and directly relevant to the good of interest here is Teisl et al. (2002). They examine the impact of the dolphin-tuna controversy and subsequent ubiquitous "dolphin-safe" labels on consumption patterns in the U.S. and find that the labels led to a statistically significant increase in purchases of canned tuna. The effect was not immediate though, suggesting that it took time for consumers to learn about the label and for them to trust the message. Roheim et al. (2011) analyze supermarket scanner data and find that the price premium on Marine Stewardship Council certified pollock products in the London market was around 14.2%. Looking at the U.S. market for canned tuna, Hallstein and Villas-Boas (2013) and Hilger et al. (2019) exploiting the same quasi-experimental data on the response to a street-light style coding of environmental quality with green representing the most environmentally friendly. They find that demand decreases only for the "yellow" coded products. The Hilger et al. paper estimates product-specific effects and finds that there is significant heterogeneity in the impact of the label across different products. Sun et al. (2017) use scanner data from natural foods grocers to analyze the price elasticity of sustainably pole-and-line caught tuna. While the "sustainably pole-and-line caught" labeling is not a traditional eco-label (there is no certifying body and no recognizable icon), it functions in much the same way as an eco-label. They find that the sustainably caught brand of albacore is more price-elastic at -3.1 than conventional brands at -0.7. Furthermore, the cross-price elasticity suggests that consumers are likely to switch from the sustainably caught brands to conventional tuna brands when the price premium is high.

As a result of the difficulty and expense of acquiring supermarket data, there are many more papers which use stated preference methods to estimate the willingness to pay for sustainably caught seafood and identify what consumer characteristics correlate with that willingness to pay. These papers include Wessells et al. (1999), Johnston et al. (2001), Jaffry et al. (2004), Johnston and Roheim (2006), Brécard et al. (2009), and Salladarré et al. (2010). In general, these papers find that willingness to pay is a function of environmental attitudes, distance from the ocean, age, education, and nationality. They also find that the WTP even varies by species being purchased. Intriguingly, Brécard et al. (2009) reports that consumers who favor an eco-labeling policy report paying more attention to price. Zander and Feucht (2018) use contingent valuation methods to ask Europeans about their WTP for sustainable seafood, estimating the premium at between 7 and 20%. From a marketing standpoint, the most important result though is the identification of three distinct market segments: 47% of respondents had no WTP for sustainability, 44% had a moderate WTP of around 20%, and 9% had a WTP of over 40%.

2.2.2 Demand for Storable Goods

The capability for households to store goods throws a wrench into traditional demand analysis. The relative storability of a consumable good is a function of its longevity and size since each household has different storage constraints. Because some households have more storage space than others, retailers can price discriminate on the basis of product size and frequency of promotions. This storability is relevant to both sides of the market, with retailers considering the impact of storability on their optimal pricing strategy (Pesendorfer, 2002; Hendel et al., 2014), and consumers accelerating their purchase and stockpiling when prices on their desired products are lower (Neslin et al., 1985). Considering the possibility for stockpiling to reduce producer surplus, Ailawadi et al. (2007) find that benefits in the form of increased consumption, brand switching, and repeat purchases more than make up for the profits lost when households that would pay the higher price purchase during a promotion. Hendel and Nevo (2006) shows that failing to account for these dynamic factors in a logit-style demand model leads to over-estimates of the own-price elasticity[1] around 30%, underestimates of the cross-price elasticities by up to a factor of 5, and overestimates the outside-option/no purchase decision by over 200 percent.

These intricacies in the supply and demand of storable goods have led to innovation in econometric modeling to account for both the inventory effects and brand preferences, such as Wang (2015). The subject of the storable goods literature has centered on four products: ketchup, yogurt, laundry detergent, and soda. Across all these products, authors find that larger packaging (e.g. 128 oz vs 64 oz or 24 pack vs 6 pack) is used to price discriminate across households with

[1]This problem is avoided here by using other models; a demand system for estimating the promotional price elasticities in which the goal is to estimate the short-run dynamics, and a proportional hazard model for estimating the impact of covariates on the interpurchase interval.

heterogeneous consumption rates and storage ability, and that stockpiling is both present in the data and important for estimating demand. The literature tends to take consumer heterogeneity as a nuisance to be accounted for in estimation, rather than a distribution to be paramaterized in terms of observables, so little is known about what variables can predict stockpiling behavior. This literature has not considered any organic or eco-labeled products.

2.2.3 Interpurchase Timing

The final strand of the demand literature reviewed here explores what factors influence the interval between purchases of a product. These irregular purchase intervals are of particular importance for retailers which are engaging in intertemporal price discrimination via the high-low pricing strategy, which leads to several questions. These questions include:

- do promotions drive consumers to purchase the product sooner than expected?
- do promotions lead consumers to switch brands?
- do promotions induce stockpiling?
- how do other marketing tools impact these purchase intervals?

Since these questions intrinsically involve a time dimension, it is not surprising that approaches based on survival analysis from the statistics literature have been adopted. A seminal paper in the marketing literature is Gupta (1991), which uses a proportional hazards model to allow for time-varying covariates to show that the inclusion of these variables leads to a significant improvement in fit over the stochastic models which had previously been the standard.

Helsen and Schmittlein (1993) show that traditional methods such as regression with inter-purchase time as the dependent variable and logistic models with a 1 for purchase during an arbitrary time window have significant issues with bias and inconsistency. Using a proportional hazard model with only time to first purchase, they find that the method improved the stability of estimates, gave more plausible estimates, and performed better at out-of-sample prediction.

Bijwaard et al. (2006) implement advances from the burgeoning recurrent event literature to look at marketing data in order to allow for the use of all purchase data, rather than the first purchase spell alone. This paper will draw heavily on the methods in this paper.

2.3 Data

2.3.1 Retailer Level Data

This aggregate data set consists of the weekly prices and sales volume of 5 ounce cans of solid white (Albacore) tuna from the three major U.S. brands Bumble Bee, Chicken of the Sea, StarKist, as well as sustainably-caught Wild Planet from 5/20/2012 through 8/13/2017 for seven retailers summed across all of their stores. Due to privacy constraints, the retailers remain anonymous, but there is one national supermarket chain (with stores in almost 20 states), a national discount retailer (a “big box store” that sells groceries in addition to clothing, sporting goods, electronics, etc.), and five regional supermarkets These regional supermarkets include two pairs of retailers for which each pair is owned by a single parent company but operate under different names with different prices. This data set does not allow for the identification of any effects of competitors’ pricing on sales because some of these retailers have a presence in many states while some operate in only one state. Further, some relevant competitive retailers are not included in the data set. While the issues preclude the identification of competition’s role in the market, it does provide the opportunity to test the co-movement of prices nationwide to determine if the price changes are strategic as opposed to changes in supply or national demand shocks. These tests indicate that the price shifts are not driven by national shocks, as can be seen from the variation across brands and retailers in Figure 2.1.

Solid white (Albacore) tuna is higher quality and typically more expensive than the other varieties of tuna which are typically sold in shelf-stable packaging, since it is a different species than "light tuna" and is a solid piece of meat rather than the many smaller pieces in "chunk" tuna. These data do not include prices or sales volume of these lower priced substitutes. However, because the included sustainably caught tuna is also a solid white albacore omitting cheaper varieties of tuna is likely to aid the comparison rather than hinder it.

The promotional pricing schedule for grocers is almost always a weekly discount. Much less common are promotional prices for shorter durations. Coupons mimic a longer duration promotion as well as providing further price discrimination. Week-long promotions generally transition from one set of sales to another on Tuesday or Wednesday. The data are collected from Friday to Friday, and the price variable is calculated as Total Revenue divided by Total Units Sold across the retailer. Therefore if prices vary across stores within the retailer, the observed price will be weighted according to the volume sold. This allows for recovery of price responsiveness, as the price variable will reflect the weighting of purchases made at the sale price versus the regular price during the same sample week. However this does mean that price is measured with error, and if the error is classical (normal and i.i.d), this implies that the coefficients estimated here may be attenuated towards zero, implying that the estimated elasticities may be conservative. Another weakness of the weekly retailer data is that it is impossible to recover the actual price paid, and therefore, unlike previous studies which show very distinct patterns of a modal regular price and periodic price reductions (e.g., Figure 1 Pesendorfer, 2002; Hosken and Reiffen, 2001), these data have an almost continuous set of prices.

All prices are in nominal terms. Over the five years of the sample, the BLS reports a 5.2% increase in CPI (132.2 to 139.1), yet the nominal price in the sample shows a downward trend. For many of the store-brand pairs the price decreases by \$0.10 - 0.20 over the sample period, which can be observed in Figure 2.1, with "regular" and/or "promotional" prices decreasing to generate this result[2]. In spite of a decrease in average nominal price, there is no significant

[2]It should also be noted that the major three brands were convicted of price fixing during the period from 2011 to 2013, which may also explain the downward trend in prices from 2014 to 2018.

increase in quantity sold. It appears that given the short window of the data, inflation was either not a salient factor in demand or almost exactly offset by income effects, so doing the analysis with real price would only inflate the rate of price decrease. Other studies in this literature have also used nominal prices.

The demand for solid white tuna shows some interesting annual trends. Figure 2.2 shows the sales volume of Bumble Bee solid white tuna from regional supermarket C, which is unique in having a nearly constant price ($1.60-$1.77) with no discernible pattern over the three years of the data, and in having no in-store competition between brands in the "regular solid white tuna" space because it only carries Bumble Bee. The price history is included in the Appendix in Figure 2.8. There are distinct negative shocks in demand that occur each year in the weeks following Thanksgiving and Christmas, with a lesser shock occurring after Easter as illustrated in Figure 2.2. The negative shocks from the holidays likely result from the large family dinners of turkey or ham and the leftovers they produce, reducing the demand for canned tuna. The time off of work and school may also be a factor in decreasing demand, as bag lunches are replaced by meals at home. Some stores and some brands do show evidence of a statistically significant increase in sales associated with Lent, but the effect is not strong in this segment of the tuna market. Nevo and Hatzitaskos (2005) reports that "the quantity sold of light tuna increases during Lent, while there is no increase, even a slight decrease, in the quantity sold of the higher quality white tuna."

Because the retailer in Figure 2.2 is a small regional grocer in New England, serious weather events are responsible for the largest spikes in demand as households stocked up in advance. The large spike in late 2012 corresponds to Hurricane Sandy, and the smaller spike in January 2016 corresponds to a Catagory 5 blizzard.

Table 2.1 and Figure 2.1 illustrate that consumers at different retailers face different prices and pricing patterns. Looking at the sales at the National Supermarket in Figure 2.3, it strongly suggests that stockpiling plays a large role in the demand since it is unlikely that households

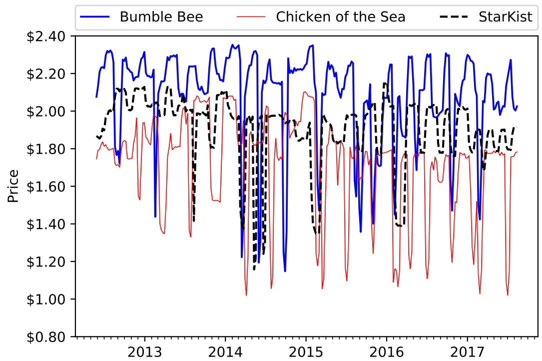

(a) Regional Supermarket A-1 Prices

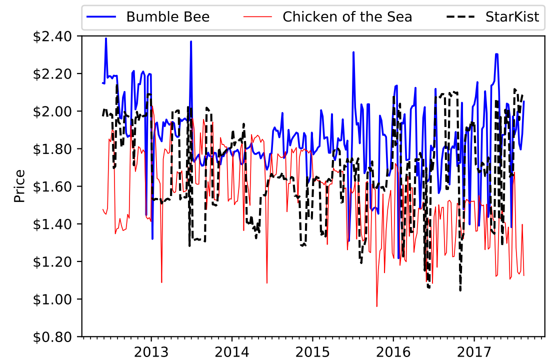

(b) Regional Supermarket A-2 Prices

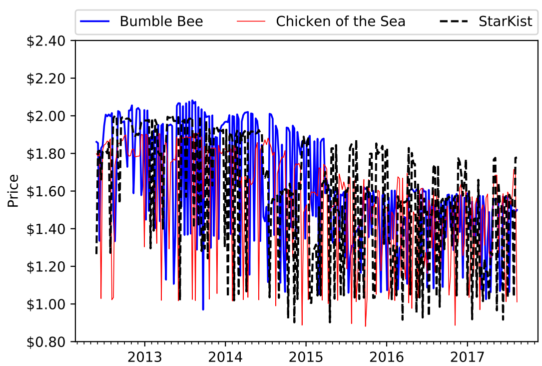

(c) Regional Supermarket B-1 Prices

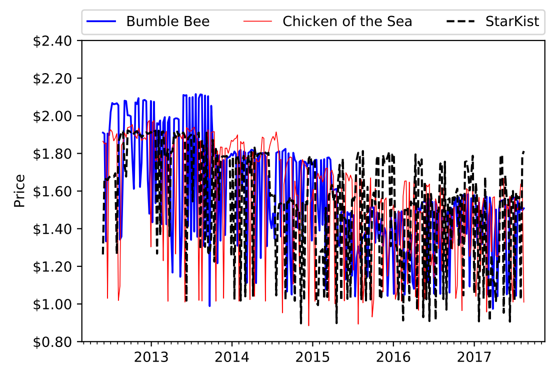

(d) Regional Supermarket B-2 Prices

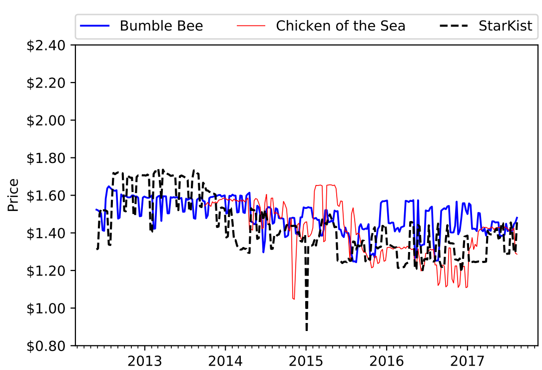

(e) National Discount Retailer Prices

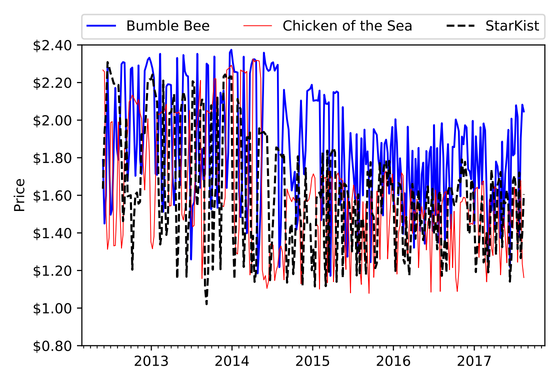

(f) National Supermarket Prices

Figure 2.1: Major Brand Prices for 5 oz. Cans of Solid White Tuna by Store

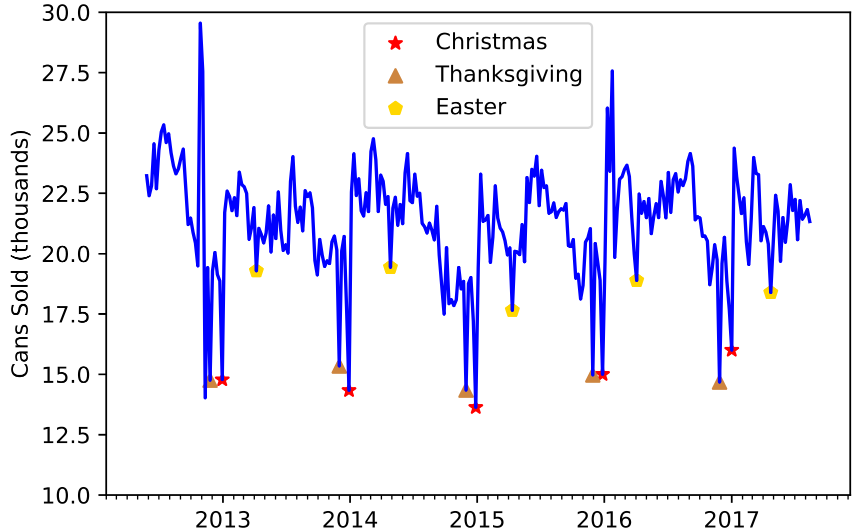

Figure 2.2: Demand for 5 oz. Cans of Bumble Bee Solid White Tuna Under Nearly Constant Price

are tripling their tuna consumption every three to five weeks given what is known about variety seeking behavior and diminishing marginal utility. Normal weekly sales are around 50,000 cans for each major brand, but during periods of promotional pricing these quantities jump to 150,000 or more.

Although it is an imperfect measure, the standard deviation of price in Table 2.1 is an informative first approximation of the strategy. It is striking that there is a high degree of consistency in the standard deviations across the conventional brands within each retailer. This suggests that whatever the pricing pattern is, it is likely being set by the retailer rather than suppliers. Since the price of Wild Planet is roughly 2-3 times that of the regular brands, a similarly scaled distribution of prices would have a 2-3 times higher standard deviation. Some retailers do have this level of difference (RSM-A-2 and NDR), while NSM actually has a lower standard deviation for Wild Planet, indicating fewer or shallower sales.

Generating a cumulative distribution of relative prices, defined as $R_{it} = p_{it}/\max_t(p_{it})$, where i is the brand and t is the week, it is possible to compare the pricing strategies used

Table 2.1: Average Price of 5 oz. Cans of Solid White Tuna by Store and Brand

		Bumble Bee	Chicken of the Sea	StarKist	Wild Planet
Regional Supermarket (RSM-A-1)	Mean	\$1.76	\$1.66	\$1.71	\$3.99
	Std Dev	0.30	0.27	0.31	0.44
Regional Supermarket (RSM-A-2)	Mean	\$1.71	\$1.70	\$1.64	\$3.55
	Std Dev	0.29	0.29	0.28	0.83
Regional Supermarket (RSM-B-1)	Mean	\$1.88	\$1.67	\$1.68	\$4.12
	Std Dev	0.16	0.17	0.23	0.30
Regional Supermarket (RSM-B-2)	Mean	\$2.14	\$1.80	\$1.92	\$4.15
	Std Dev	0.24	0.25	0.20	0.26
Regional Supermarket (RSM-C)	Mean	\$1.66	–.–	–.–	\$3.34
	Std Dev	0.03			0.28
National Supermarket (NSM)	Mean	\$1.95	\$1.72	\$1.71	\$4.23
	Std Dev	0.32	0.36	0.36	0.32
National Discount Retailer (NDR)	Mean	\$1.53	\$1.52	\$1.52	\$3.75
	Std Dev	0.07	0.11	0.15	0.39

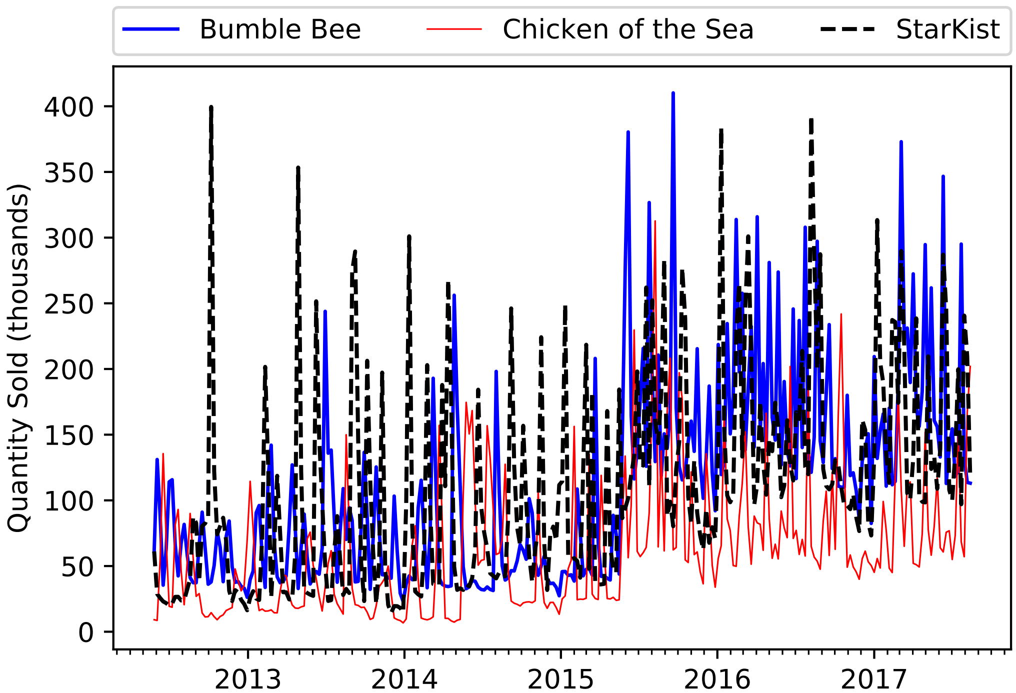

Figure 2.3: National Supermarket Sales Volume for 5 oz. Cans of Solid White Tuna

as marketing tools (as suggested in Empen and Weiss, 2011). These distributions, shown in Figure 2.4, allow for a comparison of the depth and breadth of promotional prices across the brands. Quantile regressions were used to test for significance in the point-wise differences of these distributions and are included in the Appendix Figure 2.9. In the cases where Wild Planet does not cross the lines for the conventional brand (Subfigures b, e, and f), the figure indicates that the retailer employs promotions less frequently and the sales prices are less discounted. However, even in these retailers a clear curvature exists for the sustainably caught brand, demonstrating that the retailers are using promotional prices as a frequent part of the marketing mix but not to the same extent as the conventional brands. Where the Wild Planet curve crosses the conventional brands as in Subfigure a and d, this means that the retailers offer larger percentage discounts for the sustainably caught brand than conventional brands, with approximately 30% of the observations being more discounted than one brand in panel a and all brands in panel d. Finally in Subfigure c, the promotional strategy for Wild Planet is the same as Bumble Bee and StarKist.

2.3.2 Consumer Panel Data

In addition to the retailer data, the Nielsen consumer panel data from 2004 to 2016 are used to examine the behavior of individual households. The data are trimmed to include only trips in which a product (UPC) belonging to the product module of "Shelf-stable tuna" was purchased. This category includes a variety of package sizes for both cans and pouches.

These data are a rolling panel, with households entering and exiting throughout that time. While a household is in the panel, they use a small UPC scanner at home to report what they purchased and from which store, which is then transmitted to Nielsen. The households also report their demographic information including income category, race, and ethnicity. The panel includes 40,000 households from 2004-2006 and 60,000 households from 2007 onwards. Consumers use a scanner in their home to record their purchases by UPC, and for stores which do not report prices

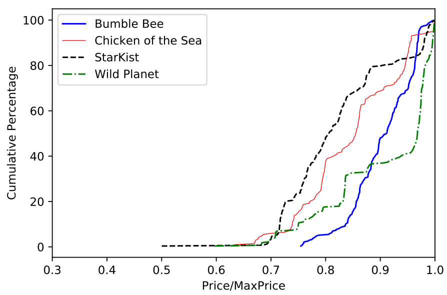

(a) NDR

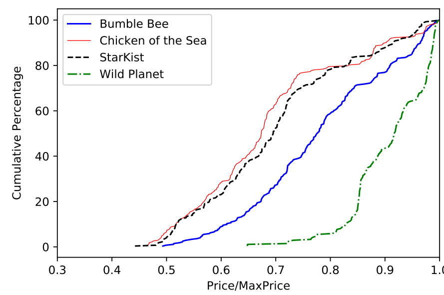

(b) NSM

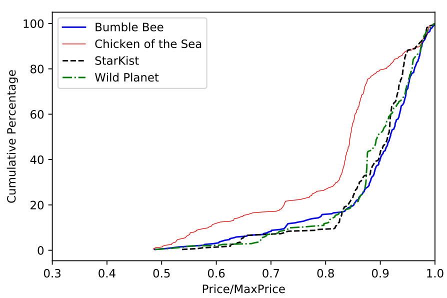

(c) RSM-A-1

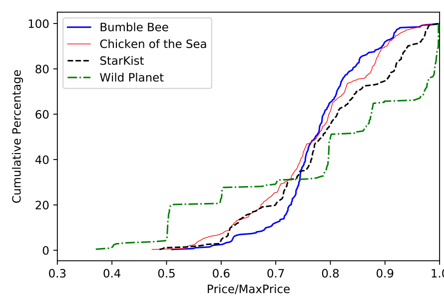

(d) RSM-A-2

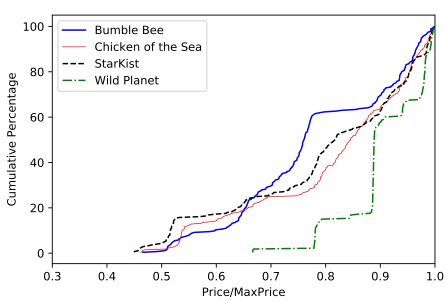

(e) RSM-B-1

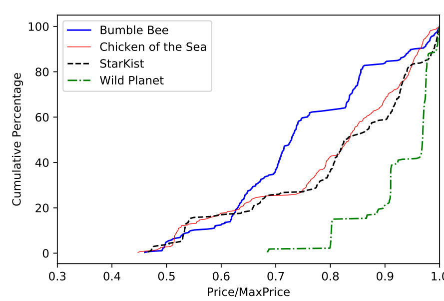

(f) RSM-B-2

Figure 2.4: Cumulative Distribution of Relative Price for 5 oz. Canned Solid White Tuna

to Nielsen, consumers also record the total price paid net of coupons[3]. With coupons, values of $0.00 for price are present and accurate.

In order to classify the brands as "sustainable" or "premium," I researched each brand's labeling and, where available, website. A brand was classified as sustainable if it contained terminology related to sustainability in the brand (e.g. Sustainable Seas) or the label (e.g. sustainably pole-and-line caught). Likewise a brand was classified as premium on the basis of name (e.g. Bumble Bee Prime Fillet), the label (e.g. fancy or premium), and an average price generally at or above $2/can. I have classified 47 of the 199 brands, a group which comprises 99.83% of purchase trips and 99.77% of revenue.

Because the panel covers 12 years of purchases, all dollar values are converted to real using the monthly Bureau of Labor and Statistics chained consumer price index data with December 1999 as the baseline. In light of the variety of sizes and packaging available, all quantities and prices are standardized to the most common sales unit, the 5 ounce can.

A subset of the data is generated with only those households with at least one purchase of "premium" tuna because households which strictly purchase "chunk light" tuna at low prices are not likely to be comparable to the target market for higher priced sustainably caught tuna. Even for those households which bought premium or sustainable tuna, these brands average around 11% of their total purchases as shown in Table 2.2. In keeping with the literature, a minimum number of purchases is required to enter into the analysis sample. The limiting factor here is the computing time to complete the shared frailty models, so the minimum of 11 purchases was chosen, as that produced a feasible computation time.

The original consumer panel data are transformed into count process input style for use in survival time modeling. This was accomplished by defining the start date as December 25th of the year prior to a household's entry into the panel, since the panel actually adds new households

[3]This leads to occasionally mis-entered data, usually via a double-press of the number button such as 1.29 as 12.99 or 11.29. I have manually cleaned the data whenever price exceeded $8 per 5 ounces in order to minimize this effect.

roughly one week prior to the new year, and their exit date as January 1 after their last year with an eligible purchase. Purchase information was broken down into household averages and deviations from this average as in Bijwaard et al. (2006), with a summary presented in Table 2.3.

This decomposition of the data allows for the testing of hypotheses regarding stockpiling and bargain seeking. The average volume measures how much tuna a household tends to purchase at once. The volume deviation measures how the amount of tuna purchased during the previous trip compares to their average purchase quantity. When included in a proportional hazard model, the coefficient on this variable will be a measure of the delay until the next purchase induced by an additional unit stockpiled. Household purchase price average serves to differentiate households by both product type and shopping strategy, with lower average prices indicating a preference for less fancy tuna, strategic purchases during sale periods, or a combination of the two. Two variables attempt to measure the impact of promotions by comparing the price of the current purchase to a reference price, with the deviation measuring the difference from the household average and the difference measuring the change from the previous purchase occasion. Including both variables allows for a comparison of reference points, which is important because households may have evolving preferences over the length of the sample. The expenditure average simply measures how much households spend on average during trips in which they purchase tuna, with the deviation identifying trips which are smaller (presumably irregular) trips and larger trips. Note that household income is a categorical variable ranging from 3 to 30, and household size is top coded at 9.

The principle effect of including averages and deviations is to change what is "held constant" in the estimation. For instance, increasing the average but holding the deviations constant looks like shifting the entire distribution of prices paid, while increasing the deviation while holding the average fixed nets out the impact that purchase has on the overall average.

Table 2.2: Percentage of All Shelf-Stable Tuna purchases Recorded as Premium or Sustainable

	Households Which Bought Sustainable	Households Which Bought Premium	All All Households
Premium	10.6%	11.4%	0.6%
Sustainable	10.7%	6.3%	0.4%

2.4 Results

2.4.1 Promotional Elasticity and the Almost-Ideal Demand System

A promotional price elasticity is the extent to which promotional price cuts increase demand for a product, and is something completely separate from the typical price elasticity which assumes the price change is more permanent. The difference arises not from the fact that the cut is temporary, but from the consumer's expectation that the price cut is only temporary but likely to be repeated. Narasimhan et al. (1996) explore the relationship between the promotional elasticity and a number of product/market characteristics. They report that the promotional elasticity is highest for those categories with a relatively small number of brands, higher category penetration, shorter interpurchase times, and higher consumer propensity to stockpile. Shelf-stable tuna would rank highly in most of these categories (on average it may have a higher interpurchase time and lower penetration than other commonly studied products), and is therefore expected to have a large promotional price elasticity.

The Almost-Ideal Demand System (AIDS) has been a staple in demand analysis since it was introduced by Deaton and Muellbauer (1980). Using expenditure shares in a system of equations, AIDS can be used to estimate a system of demand which satisfies the ideal properties of Slutsky symmetry, adding-up, and homogeneity of degree 0 in prices and expenditure. It was originally designed for use with expenditure data from an entire economy, with each equation representing one sector of the economy. However, it has also frequently been used for the estimation of smaller demand systems with just a few brands or products under the tenous

Table 2.3: Summary Statistics for Shelf-Stable Tuna Purchases from the Consumer Panel

(a) Any Households with More Than 10 Purchases (N = 1,475,657 Purchases)

	Mean	Std. Dev.	Min	Max
Volume Average	3.56	2.37	0.51	75.20
Volume Deviation	0.08	2.62	−74.00	162.13
Price Average	0.88	0.38	0.23	10.93
Price Deviation	−0.01	0.45	−9.96	82.59
Price Difference	0.00	0.57	−88.86	82.59
Expenditure Average	71.06	47.17	0.89	658.51
Expenditure Deviation	−0.86	47.42	−504.72	839.21
Household Income	19.76	5.90	3	30
Household Size	2.46	1.29	1	9

(b) Bought Premium and More Than 10 Purchases (N = 70,764 Purchases)

	Mean	Std. Dev.	Min	Max
Volume Average	3.42	2.26	0.55	28.80
Volume Deviation	0.12	2.87	−26.29	125.16
Price Average	1.25	0.70	0.39	10.93
Price Deviation	0.00	1.07	−9.96	12.42
Price Difference	0.01	1.20	−13.52	14.04
Expenditure Average	68.20	40.87	3.89	390.00
Expenditure Deviation	−0.85	47.51	−248.99	733.24
Household Income	20.59	6.20	3	30
Household Size	2.35	1.26	1	9

(c) Bought Sustainable and More Than 10 Purchases (N = 44,754 Purchases)

	Mean	Std. Dev.	Min	Max
Volume Average	3.44	2.35	0.68	28.80
Volume Deviation	0.12	2.94	−26.30	125.16
Price Average	1.34	0.82	0.45	10.93
Price Deviation	0.02	1.29	−9.96	12.42
Price Difference	0.03	1.44	−13.52	14.04
Expenditure Average	71.54	41.68	4.60	327.99
Expenditure Deviation	−1.06	48.98	−242.19	712.75
Household Income	20.88	6.14	3	30
Household Size	2.20	1.12	1	9

assumption that this small subset of the economy is separable. Of particular relevance to this paper, Thompson and Glaser (2001) and Glaser and Thompson (1998) use an AIDS model to estimate price elasticities of organic products which they compare to their conventional counterparts. They report rejecting the restrictions implied by symmetry and homogeneity. Since the demand system is not all-inclusive, an increase of all observed prices and incomes could lead to more or less consumption within the products specified in the model, thus breaking the homogeneity assumption because the sub-market is not separable. Likewise, the consumers may be choosing to substitute products outside the model for products inside the model, thus invalidating the separability assumption which leads to the symmetry requirement within the subset. The Likelihood-Ratio test statistic for Homogeneity of Degree 0 is 436.1 and for symmetry it is 475.7, with p-values less than 0.0001. In light of the evidence that the restrictions are not valid in this setting, a variant of the AIDS model is presented here without imposing homogeneity of degree 0 or symmetry. Models estimated with the restrictions do not produce substantially different results, and are included in the appendix Table 2.7.

The non-linear AIDS demand system was modeled as

$$w_{it} = \alpha_i + \sum_j \gamma_{ij} \ln p_{jt} + \beta_i \ln(x_t/P_t),$$

where w_{it} denotes the share of the i^{th} canned tuna brand in the t^{th} week, x_t represents total expenditure on canned solid white tuna ($x_t \equiv \sum_j p_{jt} q_{jt}$) with p_{jt} and q_{jt} being the price and quantity of the j^{th} canned solid white tuna brand, and P_t is defined by $\ln P_t = \sum_k \gamma_{kt} \ln p_{kt} + \sum_k \sum_\lambda \gamma_{k\lambda} \ln p_{kt} \ln p_{\lambda t}$[4]

The model was estimated for each retailer separately (included in the appendix Tables 2.9 and 2.9) and for all retailers aggregated together.

While in most cases the estimation of a demand curve would require the inclusion of an

[4]Imposed that the constant term $\alpha_0 = 0$.

instrumental variable to account for price being co-determined by demand and supply, in the case of a frequently purchased good with promotional price cuts, the observed price is plausibly exogenous. The different brands do not show significant correlation in their weekly price changes, and the magnitude of the promotional cuts dwarfs the impact of other price-shifters over the sample period. Evidence for this can be seen in scatter plots of log-price and log-quantity shown in Figure 2.5.

Using the well-established non-linear AIDS model, and estimating it via iterated Seemingly Unrelated Regression, a set of Marshallian elasticities is calculated for the whole data set in Table 2.4, with standard errors calculated using a non-parametric bootstrap with 1000 samples rather than the Delta method. In case there is coordination of weekly promotional strategies across brands influencing the results, the data are aggregated to the monthly level and the same analyses are performed as a robustness check. Results from the individual retailer level regressions are available in appendix Tables 2.8.

The results are similar in a relative sense to the doctoral thesis of Daloonpate (2002), albeit greater in magnitude which used an AIDS model to estimate the price elasticity for the three major brands of canned tuna, and reports an own-price elasticity for Starkist at -1.67, for Chicken of the Sea at -2.80, and Bumble Bee at -1.71.

The promotional elasticities estimated here, ranging across the four brands from -2.21 to -4.11. are also similar to the promotional elasticity for canned tuna calculated in Liu and Balachander (2014) at -4.14. The promotional elasticity for the sustainably caught Wild Planet is not statistically different from the elasticities for Bumble Bee and Starkist, but those three brands are significantly less elastic than Chicken of the Sea. This also holds for the monthly aggregated data, implying that whatever differences exist between the consumer-base for sustainably caught and conventional tuna, their responsiveness to promotional pricing is similar. However, such aggregate data cannot answer to what extent the response to the promotion of the sustainably caught brand is purchase acceleration or stockpiling.

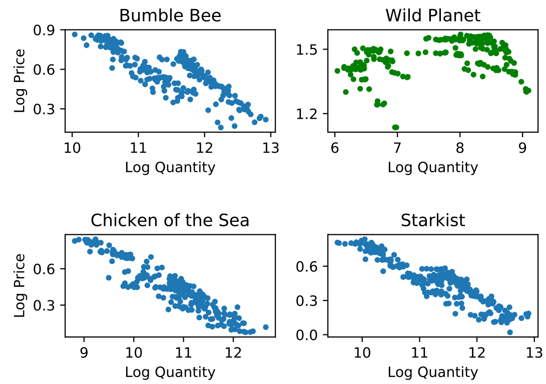

(a) National Supermarket

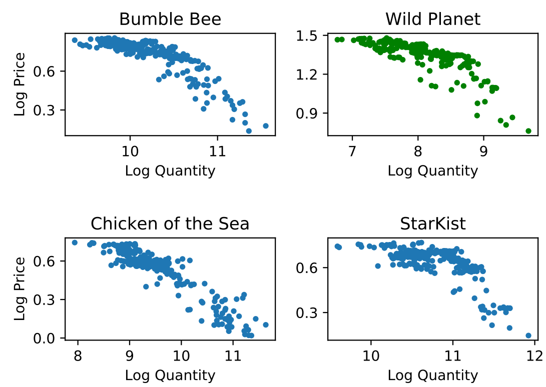

(b) Regional Supermarket A-1

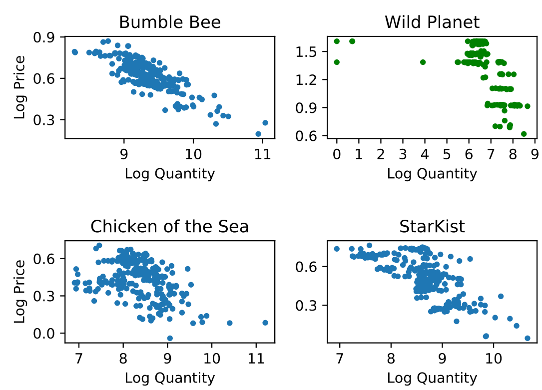

(c) Regional Supermarket A-2

Figure 2.5: Sample Scatter Plots of Log Price - Log Quantity Pairs

The data can, however, address the question of brand switching. The results show that as the price of the sustainably caught brand increases, the demand for Bumble Bee and Starkist increases slightly, indicative of switching between these brands and Wild Planet. But somewhat surprisingly, when the prices of conventional brands go up, the demand for Wild Planet decreases as well.[5] One would think that the decrease in the relative price premium would increase sales of the sustainably caught brand, but this theory seems to be counter-indicated by the results here. It may be the case that the sale price tags are responsible for drawing consumers' attention to the tuna shelf where they ultimately purchase the sustainably caught brand instead of the conventional brand which was on sale at the time. In any case, this finding casts into doubt the idea that consumer behavior may be sufficiently explained by the "price premium" as calculated by taking the difference between the brands, since something more appears to be at play in their brand choice.

The demand system results are both good and bad news for the effectiveness of eco-labels at improving environmental outcomes. The good news is that the large promotional price elasticity suggests that retailers are able to effectively target households with lower and higher willingness to pay, which should increase producer surplus. The bad news is that households' ability to stockpile is allowing some households with higher WTP to avoid the high prices and reclaim some of the surplus that they would have lost under price discrimination. The high-low pricing strategy is almost certainly a net win for the retailer (and for the entire supply chain) because the strategy would not be ubiquitous if it produced lower profits in the end. This suggests that promotional pricing of eco-labeled goods will be at least as effective at environmental protection than a single steady price, depending on how that revenue is transmitted through the supply chain..

[5]This relationship still holds in the monthly data, although many of the cross-price elasticities are no longer statistically different from zero.

Table 2.4: Marshallian Elasticities

(a) All Retailers Combined - Weekly Data

	P_{BB}	P_{CS}	P_{SK}	P_{WP}
$Q_{Bumble\ Bee}$	−2.21	0.55	0.89	0.22
	(0.06)	(0.04)	(0.03)	(0.04)
$Q_{Chicken\ of\ the\ Sea}$	1.94	−4.11	1.42	−0.57
	(0.10)	(0.11)	(0.08)	(0.11)
$Q_{StarKist}$	1.09	0.58	−2.94	0.18
	(0.07)	(0.05)	(0.04)	(0.04)
$Q_{Wild\ Planet}$	−0.34	−0.62	−0.41	−2.79
	(0.10)	(0.09)	(0.09)	(0.15)

(b) All Retailers Combined - Monthly Aggregated Data

	P_{BB}	P_{CS}	P_{SK}	P_{WP}
$Q_{Bumble\ Bee}$	−1.93	0.31	0.59	0.21
	(0.17)	(0.10)	(0.10)	(0.09)
$Q_{Chicken\ of\ the\ Sea}$	1.30	−2.82	0.97	−0.46
	(0.22)	(0.13)	(0.15)	(0.27)
$Q_{StarKist}$	0.95	0.40	−2.29	0.10
	(0.19)	(0.12)	(0.11)	(0.13)
$Q_{Wild\ Planet}$	−0.52	−0.63	−0.22	−2.53
	(0.31)	(0.20)	(0.20)	(0.33)

Standard errors in parentheses are computed via non-parametric bootstrap with 1000 replications.

2.4.2 Proportional Hazard Model for Recurrent Events

To analyze the consumer panel data, a proportional hazard model for counting processes is used. This model, originally proposed in Andersen and Gill (1982), is a variant on the semi-parametric Cox Proportional Hazard model which allows for multiple failures. The advantages of this method over using a standard proportional hazard model are many: one does not have to discard any available data, information about previous purchase decisions can be included in estimation, and allowing for calendar time effects such as seasonality. Although there are several other techniques for modeling recurrent events, the Andersen-Gill model has the useful property in the case of modeling consumer purchases of ignoring the order of purchase events (Bijwaard et al., 2006).

The model estimates the intensity process for the ith household as

$$\Pr[dN_i(t) = 1|\mathcal{H}_{it}] = Y_i(t)\lambda_0(t)exp(\beta' X_i(t))dt,$$

where $dN_i(t) = 1$ means a purchase is made by household i in the period t, $\mathcal{H}_{it}$ is the history of purchases by household i up to, but not including, t, $Y_i(t)$ is an indicator function which is 1 if that household is at risk (liable to make a purchase and under observation) in period t, $\lambda_0(t)$ is the baseline intensity process (the baseline probability across households of observing a purchase at time t), and β represents the vector of coefficients to be estimated on the vector of variables $X_i(t)$.

Following the framework of Bijwaard et al. (2006), the covariates include household average price, deviation of price from the household average, difference of price from previous purchase occasion, household average volume, deviation of previous purchase volume from household average, household income, household size, plus day-of-week and calendar month dummy variables. In addition to these variables, included here are the average spent in shopping trips in which tuna was purchased, and the deviation of the current trip from this average, as well as a set of dummy variables for race, Hispanic or non-Hispanic. A dummy variable for purchases of a sustainably caught brand is interacted with the price and volume variables to identify any differential impacts on demand for sustainably caught tuna. Ties in failure time (purchases by different households on the same day) are handled via Breslow's method, which essentially ignores the ordering during the day in which households made their purchases, for computational simplicity and because the hour in which households purchased is not relevant to these research questions.

Because households differ in their consumption and purchase rates for products, it is helpful to account for unobserved heterogeneity. This is accomplished via what is referred to as a "shared frailty" which works to allow each household to have its own baseline intensity by introducing a household-specific ν_i which is drawn from a gamma distribution and enters the intensity process multiplicatively. Accounting for this unobserved heterogeneity has been found to be critically important in estimating the coefficients on the covariates in marketing applications (Jain and Vilcassim, 1991; Bijwaard et al., 2006). In order to reduce the number of households to a computationally manageable level, the data were trimmed to include only households that had at least eleven purchases of shelf-stable tuna and at least one purchase of sustainable or premium tuna. Since the interaction terms measure how purchase timing differs within a household when purchasing a sustainably caught brand it is re-assuring that only 10-11% of purchases in this restricted sample are of sustainable brands.

The results are presented in Table 2.5. Columns 1 and 2 are the results from the subset containing households which ever purchased premium shelf-stable tuna, which likely better represents the market for the higher priced sustainably caught brands than including households which only purchase chunk light tuna. The preferred model including the shared frailty to account for unobserved heterogeneity is in column 1, with column 2 containing the results without the frailty. The statistical significance of the Gamma distribution variance is evidence that the frailty is important for the estimation. Column 3 includes the results including all households with at least 11 purchases without including shared frailty as this sample size exceeds the maximum supported size for the frailty model in STATA.

The negative and significant coefficient on average volume implies that the more tuna a household purchases at a time, the less frequently they purchase tuna, consistent with a stockpiling narrative. This is different from the result Bijwaard et al. (2006) obtain for yogurt, and may be attributed to the much shorter expiration period for yogurt which makes it harder to stockpile, thus producing a positive link between habitual consumption and quantity purchased. Reinforcing the theory that stockpiling is driving this estimate is the negative and significant coefficient on the volume deviation. It implies that the next purchase is delayed when a household purchased a greater volume on the previous trip. Turning to the interaction terms to measure the differences in demand for the sustainably caught brands, the coefficient on average volume for purchases of sustainable brands is negative and significant in the preferred model, suggesting that households either consume less of the product or endure stock-outs for longer. The coefficient for the interaction with the volume deviation is not statistically different from zero, meaning that a purchase of more or less than the average results in the same effect on the timing of the next purchase for sustainably caught and conventional tuna alike. This leads to the conclusion that for shelf-stable tuna, the more product a household purchases in an average trip the less frequently they purchase, which is indicative of the importance of stockpiling in the household's decision. This effect is stronger for purchases of sustainably caught tuna, emphasizing that stockpiling plays an even larger role in the market for the higher priced sustainably caught brands.

The coefficient on average price is positive and significant in the preferred model, suggesting that the households which purchase higher-priced varieties purchase more frequently. This would represent a significant deviation from the traditional law of demand. One possible explanation for this is the success of the high-low pricing strategy in achieving price discrimination. However, if this were the explanation one would expect the coefficient to be larger and more significant without the shared frailty, which is not the case. Because this pattern only emerges after accounting for households' individual frequency of purchase, it may be indicative that those

Table 2.5: Parameter Estimates for Proportional Hazards Models

	Bought Premium (1)	Bought Premium (2)	Full Sample (3)
Volume Deviation	−0.0328	−0.0312	−0.0383
	(0.0009)	(0.0022)	(0.0007)
Sustainable x Volume Deviation	0.0079	−0.0002	−0.0012
	(0.0050)	(0.0097)	(0.0102)
Volume Average	−0.0293	−0.0610	−0.0434
	(0.0035)	(0.0073)	(0.0022)
Sustainable x Volume Average	−0.0189	−0.0088	−0.0117
	(0.0054)	(0.0144)	(0.0108)
Price Deviation	0.0525	0.0964	0.0070
	(0.0065)	(0.0149)	(0.0052)
Sustainable x Price Deviation	−0.0638	−0.1104	−0.0511
	(0.0254)	(0.0521)	(0.0139)
Price Average	0.0994	0.0676	0.0005
	(0.0267)	(0.0487)	(0.0115)
Sustainable x Price Average	0.0796	−0.0917	0.0331
	(0.0171)	(0.0561)	(0.0224)
Price Difference	−0.0647	−0.0829	−0.0172
	(0.0050)	(0.0112)	(0.0031)
Sustainable x Price Difference	−0.1078	−0.1186	−0.0276
	(0.0192)	(0.0403)	(0.0118)
Trip Expenditure Deviation	−0.0007	−0.0005	−0.0004
	(0.0000)	(0.0001)	(0.0000)
Trip Expenditure Average	−0.0001	−0.0001	0.0000
	(0.0002)	(0.0003)	(0.0001)
Shared Frailty	Y	N	N
Est. Variance of Gamma Dist.	0.3892		
	(0.0082)		
Number of Households	4,670	4,670	61,523
Number of Purchases	216,456	216,456	1,414,305

Standard errors in parentheses are clustered at the household level.

households which purchase tuna more frequently favor a higher quality of tuna.

The negative and significant coefficient for the price difference aligns with the expectation that prices lower than the reference point should accelerate purchase, contrasting with the positive and significant sign on the deviation which implies that households are more likely to purchase when prices are high relative to the reference point. The best explanation for this puzzling result is that the reference price is more accurately described by the previous purchase price, while the positive effect of the deviation from the average may be attributed to the average price containing prices for purchases which households have not yet made and prices that may be too old to actually inform households' reference price. A moving average price may be an even better description of households' reference price, but that analysis is left for future work.

Considering the interaction terms, the coefficient on average price for sustainably caught brands is positive and significant in the preferred model, indicating that the proposed explanation of higher quality leading to more frequent purchases may be even stronger in the sustainably caught segment of the market. This may be related to green preferences, with households feeling better about purchasing shelf-stable tuna frequently when they believe it is sustainably caught. Another explanation is that some sustainably caught brands also promise a lower mercury content which could lead to them choosing to consume more.

The interaction terms with the deviation and the difference are both negative and significant, which is interesting because this means the effect is amplified for the difference from the previous purchase price but the effect of the deviation from the average price is wiped out. This zero effect for the deviation from average price is consistent with the theory that the positive coefficient without interaction is a result of switching to a higher quality, as the sustainably caught brand is one such higher quality alternative. For purchases of the sustainable brand the price effect appears to be entirely driven by the difference from the previous purchase price. Furthermore, the fact that the coefficient on the interacted price difference implies that purchase timing for sustainably caught brands is significantly more price sensitive than for conventional brands.

The impacts of trip expenditure are relatively small, with the average amount spent having no impact on purchase frequency at all. The deviation coefficient is negative and significant, meaning that trips which are smaller than average are marginally more likely to include a purchase of shelf-stable tuna.

Because household size and income are recorded as categorical variables, they are included in the model as a series of dummies rather than a continuous variable. The lowest income and household size are excluded and thus form the reference point. The estimated regression results are presented in Figures 2.6 and 2.7. Much as a fixed effect wipes out the impact of the time-invariant variables, the shared frailty model dampens the impact of household size and income. Household size has a generally increasing effect on the frequency of purchase, which is sensible given that the amount of tuna consumed in a meal would scale with household size, leading to more rapid depletion of the stockpile. From Model 2 (excluding frailty), it appears that the frequency of purchasing tuna is highest in low income households, but this impact can no longer be seen when the household-level differences are accounted for via the shared frailty and is only statistically significant in one income bracket. The results for Model 1 in fact suggest that all income brackets are about 10% more likely to purchase tuna than the excluded low-income group.

The coefficients for the demographic variables in the same models are presented in Table 2.6. The categorical variables for race and ethnicity are not significant in the models including only households which purchased premium tuna, with or without frailty, but are significant in the model with all households. While the significance is not strong, the evidence is generally suggestive that White/Caucasian and Hispanic households purchase the most frequently. Regionally, New England and Middle Atlantic households appear to purchase most frequently, with South Atlantic, and East South Central being the next most frequent purchasers.

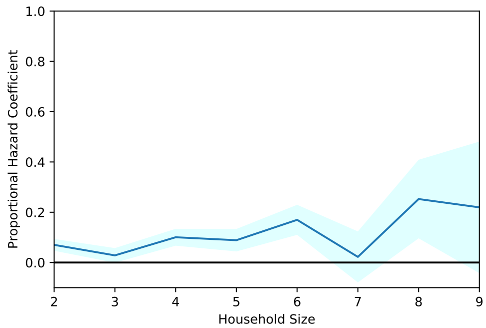

(a) Model 1

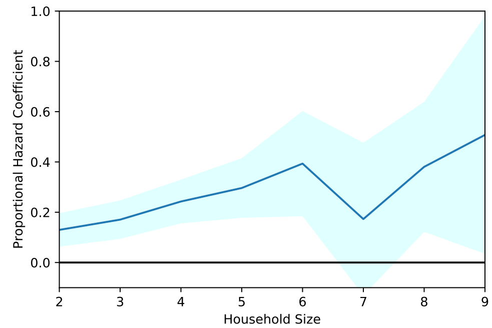

(b) Model 2

Figure 2.6: Estimated Impact of Household Size on Purchase Frequency

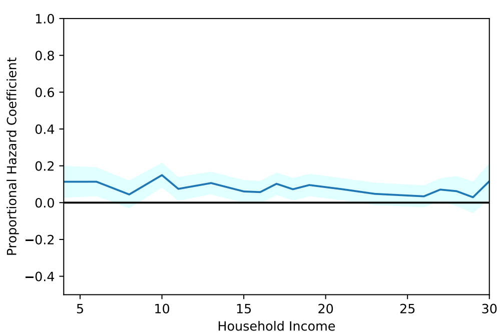

(a) Model 1

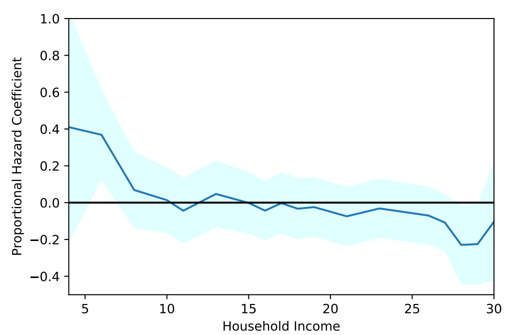

(b) Model 2

Figure 2.7: Estimated Impact of Household Income on Purchase Frequency

Table 2.6: Demographic Characteristics Results

	Bought Premium (1)	(2)	Full Sample (3)
White/Caucasian	–.–	–.–	–.–
Black/African American	−0.027	−0.106	−0.179
	(0.030)	(0.061)	(0.013)
Asian	0.049	−0.169	−0.225
	(0.047)	(0.086)	(0.025)
Other	0.011	0.028	−0.022
	(0.022)	(0.059)	(0.017)
Not Hispanic or Latino	−0.037	0.066	−0.041
	(0.027)	(0.050)	(0.016)
New England	–.–	–.–	–.–
Middle Atlantic	−0.016	−0.036	−0.038
	(0.032)	(0.049)	(0.019)
East North Central	−0.221	−0.233	−0.238
	(0.036)	(0.051)	(0.018)
West North Central	−0.260	−0.239	−0.247
	(0.048)	(0.075)	(0.020)
South Atlantic	−0.130	−0.139	−0.156
	(0.030)	(0.050)	(0.019)
East South Central	−0.121	−0.091	−0.191
	(0.047)	(0.093)	(0.023)
West South Central	−0.077	−0.128	−0.195
	(0.040)	(0.063)	(0.020)
Mountain	−0.159	−0.193	−0.220
	(0.041)	(0.061)	(0.022)
Pacific	−0.310	−0.292	−0.206
	(0.037)	(0.056)	(0.020)
Number of Households	4,670	4,670	61,523
Number of Purchases	216,456	216,456	1,414,305

Standard errors in parentheses are clustered at the household level.

2.5 Concluding Remarks

In this paper, it is established that promotional pricing is an important component of retailers' strategy in selling sustainably caught shelf-stable tuna. Both demand system and survival model analyses are estimated to better understand how consumers have responded to this strategy and how these responses compare to the patterns of demand for conventional tuna. With respect to the promotional pricing schedules retailers use to price discriminate, the aggregate demand system estimates lead to the conclusion that, as measured by the promotional price elasticity, the responsiveness to this strategy is similar in the demand for conventional brands and the sustainably caught brand. With a long supply chain, stretching from vessels to processors to distributors to retailers, the sustainably caught products have higher costs at many stages and a price premium is necessary to justify the supply. While this paper is just a beginning, more work needs to be done to understand the trade-offs involved in promotions between the revenue lost due to strategic consumers stockpiling at a discount and the gains from consumers switching from conventional brands, and how these behaviors influence the transmission of the price premium through the supply chain.

Cross-price elasticities suggest that consumers of the sustainably caught brand may switch in and out of Bumble Bee or StarKist based on the price of Wild Planet, but the demand system summarizes the data such that it cannot differentiate whether consumers switch to Wild Planet when prices are low, switch out when prices are high, or both. An investigation into the possibility of asymmetric price elasticities would be of great value for future work. Interestingly, the cross-price elasticities show that, across the board, the prices of conventional brands have a negative relationship with the demand for Wild Planet, although it is not statistically significant when averaged over monthly purchases. A possible story consistent with this is that sales on the conventional brands are catching consumers' attention enough to look at that section of the shelves, ultimately leading them to purchase the sustainably caught brand.

While the demand system leads to the conclusion that the quantity demanded of sustainably caught tuna shares a similar relationship to promotions as conventional tuna, the proportional hazard model implies that the interpurchase timing of the sustainably caught brands is more sensitive to price. Because the delay between purchases is longer on average and the quantity purchased smaller, this suggests that promotional prices for the sustainably caught brand have a greater impact on the decision to purchase at least one unit.

The evidence suggests that consumers' reference price is more accurately described by the price paid at their last purchase than their average price over the full sample. With a long panel, this could be attributed to changes in preferences for brands or quality levels causing the sample average price to be a poor representation of the reference price. It may be that a moving average price is an even better representation of the true reference price, but that investigation will be left to future work.

The survival model results indicate that stockpiling is an important feature in demand for both conventional and sustainably caught shelf-stable tuna. Households which purchase sustainably caught shelf-stable tuna appear to eat through their stockpile more slowly, since their hazard of re-purchase is significantly lower for the same volume purchased.

The interaction between "green products" and various components of the traditional marketing mix including price promotions, featured displays, and advertisement is a fertile ground for future research. In order to improve the realized environmental outcomes, it will be important for retailers and brands to know how best to persuade consumers to purchase sustainable products over the conventional alternatives. Promotional prices appear to be weakly effective at increasing consumption of sustainable brands over conventional brands, but given households which purchased at least one unit of sustainable tuna only bought the sustainable brands in 10% of their trips, more research needs to be done to understand the permanence of the shift in brands.

It is also worth noting that the "sustainably caught" labeling studied here is only a statement about the fishing gear and method, without a regulatory body or a universal symbol such as the Marine Stewardship Council or the Dolphin Safe label. While having the phrase "sustainably caught" on the package, or the word "sustainable" in the brand name communicates a clear message, consumers may not trust these labels as much as more regulated ones, and the image of charismatic mega-fauna like in the dolphin-safe label may induce a behavioral response above and beyond the mere words.

With the uncertain future resulting from the anthropogenic changes in ocean ecosystems, there is much debate about if eco-labels can sufficiently leverage market forces to protect the environment, or if greater regulation will be required. This paper cannot answer that question, but it does show that people are willing to pay a premium for sustainably caught tuna and the use of promotional pricing allows retailers to take advantage of heterogeneity in this willingness to pay for sustainability. While this price discrimination is bad for consumer surplus, the increased producer surplus should trickle through the supply chain and ultimately benefit the environment.

2.6 Acknowledgments

Chapter 2, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Bumble Bee and Nielsen via the Kilts Center for Marketing at Chicago Booth for their generous provision of data. I am grateful for Dr. Dale Squires who believed in my potential to address these fundamental questions about demand for environmental attributes enough to hand this project to me. Additional thanks to Professor Richard Carson for his guidance in selecting the proper scope for this research and the inspiration to use a survival model.

Researcher's own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Appendix

Table 2.7: Restricted AIDS Aggregate Marshallian Elasticities

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.31	0.52	0.80	0.06
	(0.06)	(0.02)	(0.03)	(0.01)
$Q_{Chicken\ of\ the\ Sea}$	1.77	−4.15	1.28	−0.03
	(0.07)	(0.10)	(0.06)	(0.03)
$Q_{StarKist}$	1.15	0.56	−2.87	0.05
	(0.04)	(0.03)	(0.03)	(0.01)
$Q_{Wild\ Planet}$	0.76	−0.01	0.48	−1.84
	(0.10)	(0.07)	(0.07)	(0.11)

Table 2.8: Unrestricted AIDS Retailer Marshallian Elasticities

(a) NDR

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−1.09	0.13	1.35	0.20
$Q_{Chicken\ of\ the\ Sea}$	−0.85	−3.21	−0.41	0.29
$Q_{StarKist}$	0.35	0.35	−2.20	0.22
$Q_{Wild\ Planet}$	−0.36	−0.21	−0.49	−2.88

(b) NSM

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.14	0.53	0.86	0.38
$Q_{Chicken\ of\ the\ Sea}$	1.16	−3.06	1.30	−0.71
$Q_{StarKist}$	0.77	0.61	−2.70	−0.16
$Q_{Wild\ Planet}$	−0.70	−1.02	−1.01	0.47

(c) RSM-A-1

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.14	0.51	0.46	0.02
$Q_{Chicken\ of\ the\ Sea}$	0.43	−3.29	0.66	−0.04
$Q_{StarKist}$	0.53	0.33	−1.54	0.22
$Q_{Wild\ Planet}$	0.44	0.17	0.34	−2.52

(d) RSM-A-2

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.04	0.31	0.32	0.25
$Q_{Chicken\ of\ the\ Sea}$	0.31	−1.33	−0.26	−0.49
$Q_{StarKist}$	1.70	−0.05	−1.57	0.23
$Q_{Wild\ Planet}$	0.62	−1.13	0.28	−2.33

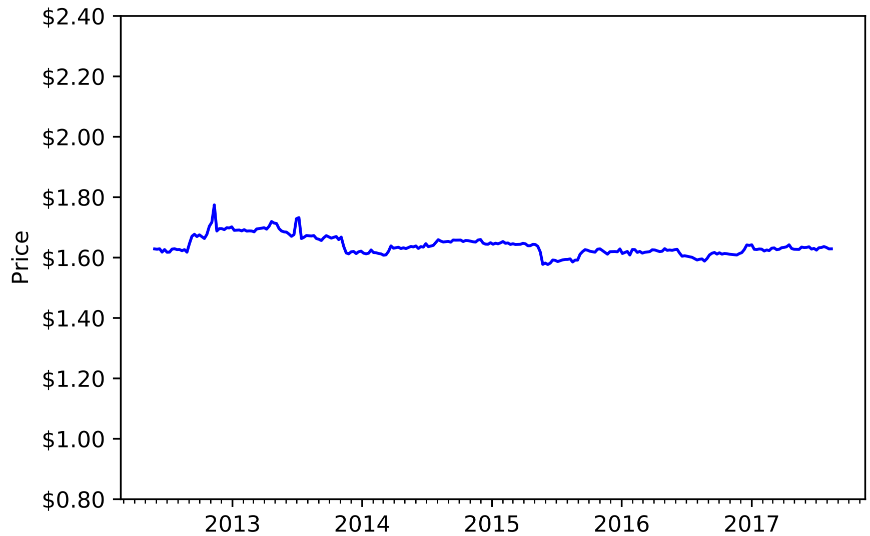

Figure 2.8: Nominal Price for Bumble Bee Solid White Tuna at Regional Supermarket C

Table 2.8: Unrestricted AIDS Retailer Marshallian Elasticities, Continued

(e) RSM-B-1

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.28	0.38	1.08	0.00
$Q_{Chicken\ of\ the\ Sea}$	3.71	−4.24	0.81	0.06
$Q_{StarKist}$	1.56	0.81	−4.48	0.10
$Q_{Wild\ Planet}$	−0.40	−1.35	0.99	−2.99

(f) RSM-B-2

	P_{BB}	P_{CotS}	P_{Star}	P_{WP}
$Q_{Bumble\ Bee}$	−2.30	0.38	1.03	−0.14
$Q_{Chicken\ of\ the\ Sea}$	4.01	−4.72	0.88	0.96
$Q_{StarKist}$	1.92	0.96	−4.92	0.08
$Q_{Wild\ Planet}$	−0.40	−1.19	0.91	−3.12

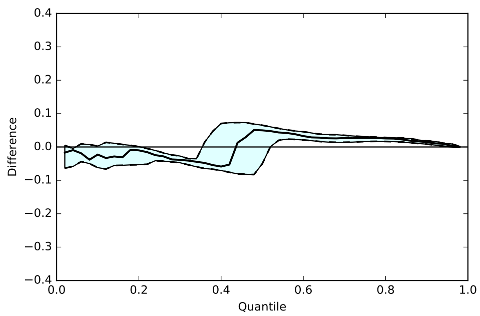

(a) NDR

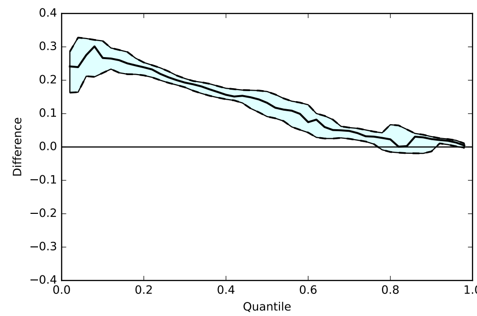

(b) NSM

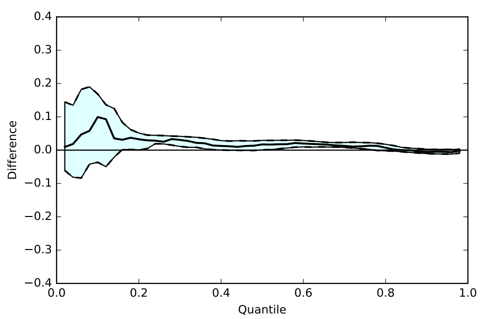

(c) RSM-A-1

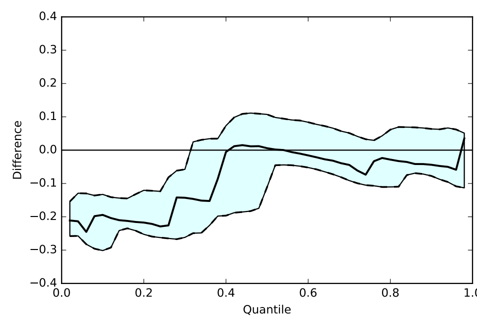

(d) RSM-A-2

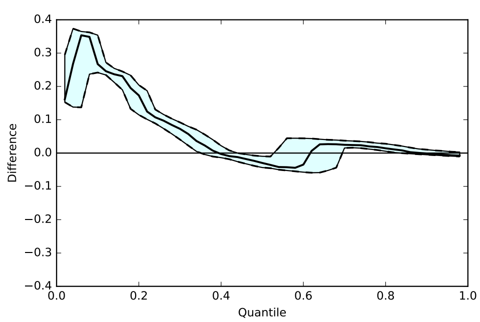

(e) RSM-B-1

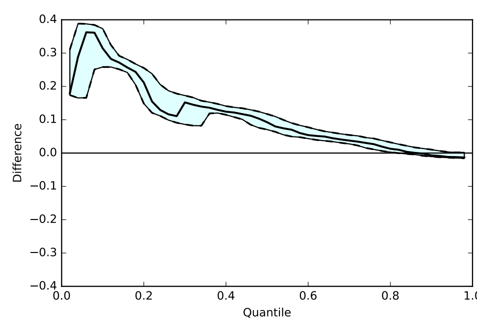

(f) RSM-B-2

Standard errors computed via non-parametric bootstrap with 1000 replications.

Figure 2.9: Difference in Distributions (Conventional minus Wild Planet) for Relative Prices of 5 oz. Canned Solid White Tuna

Table 2.9: Unrestricted AIDS Retailer Marshallian Elasticity 95% Confidence Intervals

(a) NDR

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−2.18	−1.18	−0.19	0.28	0.29	0.80	−0.09	0.27
$Q_{Chicken\ of\ the\ Sea}$	−1.63	2.42	−4.96	−0.61	0.02	2.27	−0.11	1.24
$Q_{StarKist}$	0.28	1.00	−0.02	0.33	−2.25	−1.69	0.00	0.27
$Q_{Wild\ Planet}$	−0.52	1.37	−0.11	0.75	0.00	1.13	−3.16	−1.63

(b) NSM

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−2.32	−2.16	0.51	0.61	0.77	0.89	−0.10	−0.03
$Q_{Chicken\ of\ the\ Sea}$	0.89	1.13	−3.32	−3.03	1.11	1.33	−0.23	−0.10
$Q_{StarKist}$	0.77	0.91	0.61	0.77	−2.65	−2.46	−0.13	−0.07
$Q_{Wild\ Planet}$	−1.29	−0.15	−1.97	−0.74	−1.77	−0.67	2.73	4.67

(c) RSM-A-1

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−2.39	−2.12	0.26	0.43	0.75	0.95	0.14	0.25
$Q_{Chicken\ of\ the\ Sea}$	0.34	0.73	−3.40	−2.86	0.67	1.20	0.07	0.32
$Q_{StarKist}$	0.48	0.60	0.31	0.44	−1.99	−1.79	−0.01	0.07
$Q_{Wild\ Planet}$	0.60	1.15	0.24	0.80	−0.15	0.44	−3.04	−2.20

(d) RSM-A-2

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−1.86	−1.43	0.36	0.69	−0.02	0.40	0.08	0.19
$Q_{Chicken\ of\ the\ Sea}$	0.77	1.94	−3.60	−2.11	−0.13	1.14	0.01	0.33
$Q_{StarKist}$	−0.51	0.53	−0.32	0.81	−2.59	−1.03	−0.10	0.14
$Q_{Wild\ Planet}$	0.39	1.36	−0.09	0.94	−0.17	0.56	−2.70	−2.22

Table 2.9: Unrestricted AIDS Retailer Marshallian Elasticity 95% Confidence Intervals, Continued

(e) RSM-B-1

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−2.30	−2.19	0.57	0.66	0.76	0.89	−0.00	0.01
$Q_{Chicken\ of\ the\ Sea}$	2.26	2.73	−6.01	−5.38	1.52	1.94	−0.04	0.01
$Q_{StarKist}$	2.08	2.40	1.02	1.30	−5.01	−4.57	0.02	0.04
$Q_{Wild\ Planet}$	0.13	0.60	−0.19	0.26	0.49	0.83	−1.94	−0.95

(f) RSM-B-2

	P_{BB}		P_{CS}		P_{SK}		P_{WP}	
$Q_{Bumble\ Bee}$	−2.34	−2.21	0.57	0.65	0.80	0.96	−0.00	0.01
$Q_{Chicken\ of\ the\ Sea}$	2.52	2.94	−6.26	−5.64	1.43	1.95	−0.05	0.01
$Q_{StarKist}$	2.39	2.85	0.95	1.28	−5.47	−4.87	0.02	0.05
$Q_{Wild\ Planet}$	0.10	0.62	−0.26	0.21	0.48	0.84	−1.99	−0.78

Chapter 3

Little Fish, Big Data: Demonstrating the Applicability of Big Data Techniques to Estimating Production Functions for Coastal Pelagic Fisheries

3.1 Introduction

Proper management of a fishery involves accounting for the relevant biology and oceanography in addition to economics. While much research goes into the life history of economically significant species and how oceanographic conditions affect them, there are still numerous unknowns that pose challenges to understanding a fishery in order to effectively manage it. In light of the tremendous amount of data available to researchers in the modern era, big data techniques may be brought to bear to address these challenges in understanding and managing fisheries. The LASSO (Least Absolute Shrinkage and Selection Operator) method holds particular promise for improving models of fishery production, and may provide beneficial insights for biologists as well. The method selects a parsimonious subset of predictive variables from a large pool. In addition to improving the out-of-sample predictive capacity of the model, it also provides a more transparent methodology for model selection by taking much of the model selection out of the researchers' purview.

The production function for a fishing vessel considers how catch is impacted by the primary inputs of a fishing vessel, namely capital (vessel size and gear) and "effort" which is similar to the labor measure in standard production functions. Furthermore, based on the workhorse bioeconomic model, it is assumed that the abundance of the species is also an input to production. Each of these factors of production are important for regulators to understand when making regulatory decisions, as regulatory tools such as gear restrictions and quotas will impact the effort and capital decisions of vessel owners while tools such as time-area closures and quotas will also affect the abundance of a species.

An additional complexity to be considered is whether an Instrumental Variables regression technique could improve the results. The effort level may be endogenously determined, since vessels may choose to fish more or less based on how much they are catching with each trip. Results are mixed, with Paul et al. (2009) failing to reject the null hypothesis that effort is exogenous, and Felthoven et al. (2009) rejecting the null and therefore employing instrumental variables. The null hypothesis of exogeneity of effort is also rejected in Wolff et al. (2013) and in Squires and Vestergaard (2018). The decision of whether or not to take the vessel out bears some similarities with the ubiquitous taxi driver research that shows additional layers of complexity exist in the labor supply decision. Gordon (2015) also theorized that the measurement error that likely accompanies stock assessments may violate the exogeneity assumption such that instrumental variables are needed. In these circumstances, the proper selection of instruments is another challenging task for a researcher.

The fisheries considered in this paper are Pacific Coastal Pelagic Species. These and similar species are of great importance for marine ecosystems and for human food supplies. The various species of small schooling fish and squid that comprise the Coastal Pelagic species are important food for larger fish species as well as marine mammals and seabirds. Some of the larger fish species that prey upon the coastal pelagic species are the basis of major commercial and recreational fisheries including tuna and swordfish. As such, the second most common usage for CPS landings is bait. Although the United States does not consume many of these small species, the demand for calamari makes market squid the exception. Sardine, anchovies, and mackerel are less commonly consumed in the United States and tend to be exported (Herrick, 2000). However, there is a growing push to "eat down the food chain" for the sake of sustainability (Gordinier, 2015), and there is increasing appreciation of the importance of coastal pelagic species for food security worldwide. In addition to the pressures of predators and humans, environmental pressures will play a growing role in the status of coastal pelagic species worldwide. They are known to be sensitive to water temperature, with El Niño events causing dramatic shifts in productivity and habitat preferences along the Pacific coast. The mounting impacts of climate change make it all the more important to understand the fishery.

This paper presents a proof-of-concept application of the data-driven LASSO method for estimating the production function for five fisheries. Using this method for model selection, numerous variables and transformations of these variables are considered (e.g. log, squared, cubed) in order to select relevant control variables, and to select instruments in the case of an endogenous regressor. The results indicate that the LASSO is particularly well suited to the selection of instruments, but it provides marginal improvements to the precision of coefficient estimates in the case of exogenous effort. The improvement in first stage F-statistics is dramatic, going from weak identification with the baseline instruments to strong identification with the LASSO selected instruments, along with tests for endogeneity of the effort variable providing a justification for the use of instrumental variables.

The remainder of the paper is laid out as follows. Section 3.2 provides an overview of the Pacific Coastal Pelagic Fishery including the biology and regulatory history, then proceeds into a discussion of the estimation of production functions in fishery economics. Section 3.3 details the data sources and how they are relevant to the problem at hand. Section 3.4 provides a review of the literature on big data techniques for model selection. The particulars of the application of these methods to these data and the results are included in Section 3.5. Concluding remarks are presented in Section 3.6.

3.2 Fisheries

3.2.1 Overview of Coastal Pelagic Species and the Fishery

Coastal Pelagic Species (CPS) are species that live in the water column, as opposed to near the sea floor, and tend to be found near-shore. The CPS fishery off the west coast of the United States is comprised of five species: Market Squid, Pacific Sardine, Northern Anchovy, Pacific (Chub) Mackerel, and Jack Mackerel. NOAA Fisheries and the Pacific Fisheries Management Council jointly oversee the CPS fishery off the West Coast, while Market Squid is also co-managed with the California Department of Fish and Wildlife.

Initially, only the Northern Anchovy fishery was managed, but Amendment 8 to the Northern Anchovy Fishery Management Plan (passed in December 1998) officially added sardine, squid, and mackerel to the management plan and changed the name to the Coastal Pelagic Species Fishery Management Plan. Along with this new management plan, the formerly open-access fishery was designated a limited-entry fishery south of the 39°N and permits were given to a limited number of vessels. Only Pacific Sardine and Pacific (Chub) Mackerel were added as "actively managed" species, meaning that stock assessments would be regularly performed to adjust the harvest guideline. Northern Anchovy, Market Squid, and Jack Mackerel would be "monitored," meaning that landings are tracked against annual catch limits and estimated abundance levels but without stock assessments or adjustments to the harvest guidelines.

The Pacific sardine management has undergone a number of changes during its time as

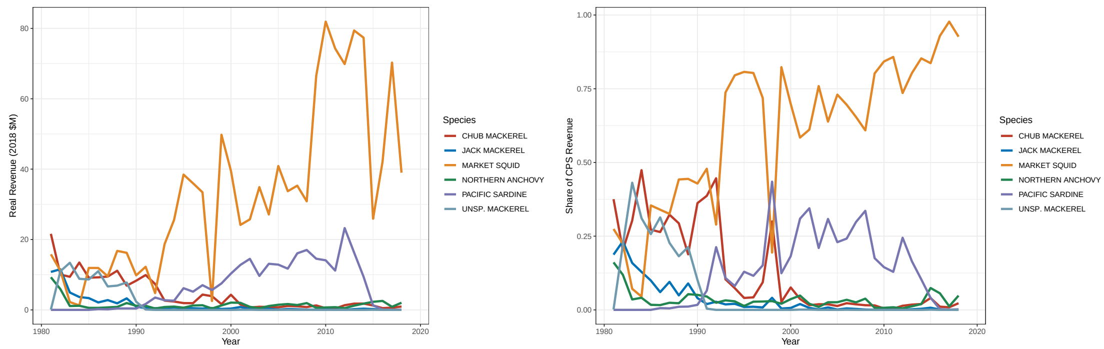

(a) Total Revenue (b) Revenue Shares

Figure 3.1: History of Revenues in the Coastal Pelagic Species Fishery

a managed species. Initially the season ran from January to December, with portions of the allowable catch being metered out on January 1, July 1, and September 15. However, the stock assessment cycle was based on a July through June calendar so the fishing season was moved to correspond with this in 2014.

The fishery operates year-round, with different species being more prevalent during certain seasons. For this reason, the limited entry program is not specific to a species, but allows vessels to target the entire set of species. The primary gear type is encircling nets, a gear type which, in conjunction with the schooling tendencies of the coastal pelagic species, allows for strong targeting with low by-catch rates and few mixed catch sets.

The fishery has significantly grown in value, largely due to increases in demand for market squid and Pacific sardine as shown in Figure 3.1a. Pacific (Chub) Mackerel was a significant piece of the fishery until 2000, after which it has never comprised more than 4% of revenue.

In spite of not being an actively managed species, the primary revenue generator among CPS is market squid (Doryteuthis opalescens). This is in large part because the demand for calamari means the exvessel price is higher at $0.27-0.50 per pound, with the other species selling for approximately $0.10 per pound (Pacific Fishery Management Council, 2019b, Table 6.3).

Squid are primarily targeted from October to February when they congregate to spawn in Southern California, while there is also a significant fishery in Northern California[1] during spring and summer. During the day, vessels hunt for squid on their own, but at night, a “light boat” attracts squid to the surface by shining bright lights onto the water’s surface while fishing vessels capture the squid and take them to port, paying the light boat a share of the take. Frequently the vessels turn back around and make another trip as soon as they are unloaded, stopping only for the weekend closure described below.

Because they are short-lived species that die shortly after spawning the entire population is replaced each year, meaning they can withstand heavy fishing pressures. Based on work by Maxwell et al. (2005) and Dorval et al. (2013) the Council adopted a definition of overfishing on the basis of egg escapement. The management goal is to ensure 30% of eggs are successfully laid before females are harvested (Pacific Fishery Management Council, 2011). This is readily provided by a total closure of the fishery from noon on Friday to noon on Sunday, as the two day closure means that there is no fishing for 28.6% of the week[2]. The species abundance is sensitive to El Niño conditions, with the warmer waters associated with El Niño leading to decreased abundance (Perretti and Sedarat, 2016).

Since 2001 Market squid has had a seasonal catch limit that does not vary from year to year, although the limit was lowered in 2005. Landings have typically fallen short of the limit, as shown in Figure 3.2a.

The abundance of Pacific sardine (Sardinops sagax) is highly variable, with regular boom and bust cycles. They are relatively short-lived, averaging about five years, but quite fertile as they reach sexual maturity at one or two years and can spawn multiple times per season (NOAA Fisheries, 2019c). There was a recent boom which peaked in 2007, but the stock declined rapidly thereafter. The directed fishery was closed down in 2015 as the population biomass fell below the

[1]A good summary of squid fishing can be found at https://www.montereyherald.com/2018/04/25/monterey-bay-fishermen-working-round-the-clock-to-pull-in-plentiful-catch/.

[2]This may only bind for some vessels, as vessel maintenance and crew personal lives may also lead to vessels not going fishing. This regulation then would address a coordination problem, allowing for undisturbed spawning for two days.

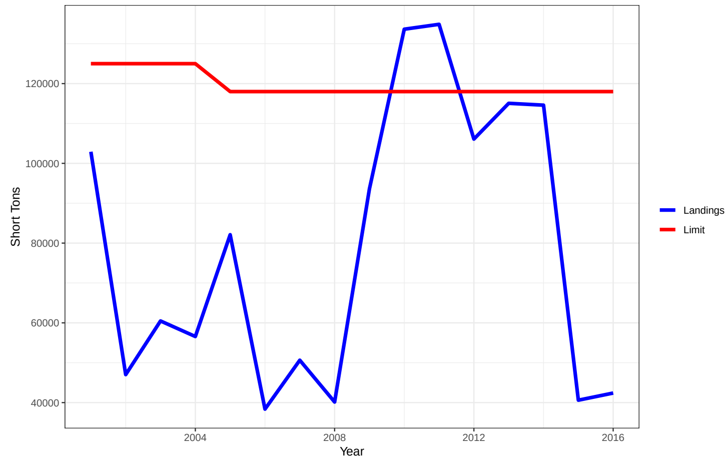

(a) Market Squid

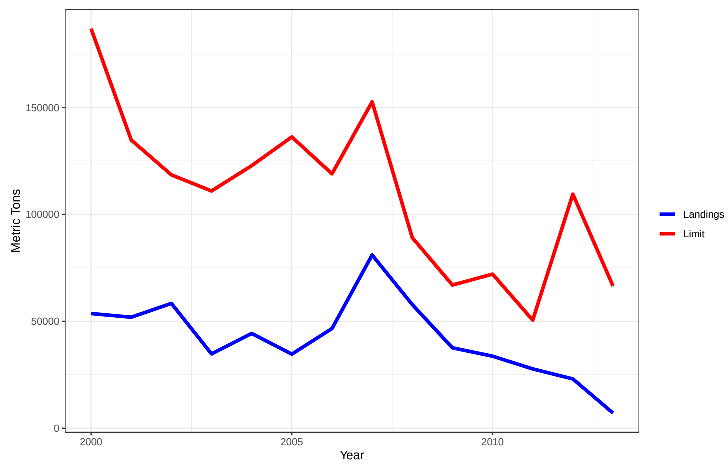

(b) Pacific Sardine

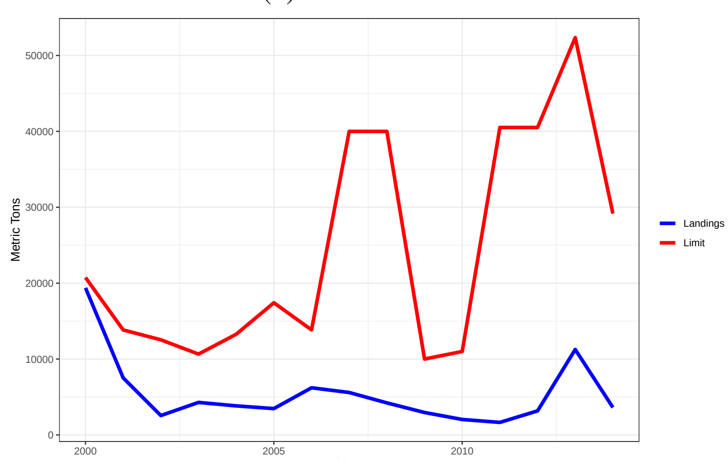

(c) Pacific Mackerel

Figure 3.2: Comparison of Landings with Catch Limits

minimum threshold of 150,000 metric tons, and has yet to re-open. In addition to facing pressures from humans, they are prey for other fish, marine mammals, and sea birds. Consequently, fishing pressure is not just for human consumption but also for bait. They are sensitive to sea surface temperature and chlorophyll-a concentrations, with their movements being predictable from a model based on these inputs (Zwolinski et al., 2011).

Pacific mackerel (Scomber japonicus) also experience boom and bust cycles. Small and fast growing, they are also a common prey for other species. They are migratory, traveling north to Washington in the Summer and south to Baja California in the winter (NOAA Fisheries, 2019b).

Jack mackerel (Trachurus symmetricus) are larger than the other species, growing up to 2 feet in length, and living as long as 35 years (Pacific Fishery Management Council, 2019a). However, they are not a significant species in the fishery (around 2% of revenue), in part because much of their range is outside of the U.S. Exclusive Economic Zone.

Northern anchovy (Engraulis mordax) are another small and fast growing species. They are used for both human consumption and bait. As with sardine and mackerel, it is a frequent prey for other species. In fact, each year 45 to 55 percent of the total stock would die of natural causes if no fishing occurred (NOAA Fisheries, 2019a).

3.2.2 Fishery Production Functions

The initial theoretical developments in the economics of fisheries, most notably Gordon (1954) and Schaefer (1957), viewed the process of harvesting fish as a function of species abundance and an input which they called "effort". Effort is an aggregate input including all of the fixed and variable inputs like labor, capital, fuel, bait, knowledge, etc. This was sufficient to derive some key theoretical insights into the economics of fisheries, but created some challenges for empirical fishery research. In order for the aggregation to form a consistent index of effort the production must have homothetic separability of inputs (Squires, 1987). In that paper Squires shows that the requirement is not satisfied for the New England otter trawl fleet, and develops a method of calculating a consistent aggregate index of effort via cost-share functions. However, for the most part subsequent papers have typically used some measure of time spent fishing as the measure of effort while including other factors of production to avoid omitted variable bias. A fair approach, considering that many of the listed factors of production such as the vessel, gear, and skipper/crew skill are more representative of fixed costs than the variable input costs like labor and fuel which effort is meant to embody.

Speaking to the challenges of empirical fisheries analysis, Gordon (2015) notes that omitted variables such as skipper skill could bias the estimates. In modern econometrics, this is less likely to be a problem as the availability of panel data methods allow for skipper skill to be account for. However other missing variables could include weather, which could impact both fish abundance and effort, economic conditions, regulatory changes, and other factors which may impact catchability. While Gordon considers the use of instrumental variables for measurement error in stock estimates, he does not consider the endogeneity of effort in the production function which may or may not be present depending on the fishery. However, endogeneity of effort is considered in Felthoven et al. (2009) and Paul et al. (2009), with the latter finding that endogeneity was a valid concern in the Alaskan Pollock fishery. In estimating a Cobb-Douglas production function for the Indian tuna fishery, Wolff et al. (2013) performed a Wald test for endogeneity in a Probit model and rejected the null hypothesis that effort was endogenous in the production function. Similarly, Squires and Vestergaard (2018) used a Durbin-Wu-Hausman test and reject the null hypothesis that days fished is exogenous in estimating production. Ignoring the endogeneity of effort would result in biased estimates of the elasticity of effort in production. Because this elasticity is used as an input to management strategy evaluations, this may lead to the implementation of regulations that may be either too onerous or not effective.

Growing out of the theory of the bioeconomic equilibrium, the production function is often measured as a Cobb-Douglas function including effort and the stock abundance. Other research including production functions in fisheries have explored issues such as technical change in Squires and Vestergaard (2013, 2018), the impact of skipper skill in Kirkley et al. (1998) and Wolff et al. (2013), and attempts to measure technical inefficiency in Hannesson (1983), Pascoe and Coglan (2000), Pascoe and Coglan (2002), and Guttormsen and Roll (2011).

3.3 Data

3.3.1 Fish Ticket Data

The data on the fishery is from the Pacific Fisheries Information Network (PacFIN) which provides information on vessels and fish tickets. A fish ticket is generated whenever a vessel brings fish to shore for sale or personal consumption, and records all the necessary regulatory details. This includes what species were landed, how many pounds were landed, the price per pound, where it was caught, and by whom it was purchased.

The data was subset to include all fish tickets from vessels that had ever participated in the CPS fishery, defined as having caught one of the CPS species using a dip net, seine, or other net gear excluding tuna seine and drift gill net.

A fish ticket is a record of exchange between a fishing vessel and a processor. Each record in the data is a single line-item from a fish ticket, where one fish ticket may have many line items. This is because a single line contains one species and one sale price, so if a vessel lands multiple species or a single species with multiple quality grades that result in differing prices.

Since the CPS fishery takes place close to shore, and most of the species do not hold up well, trips of more than one day are exceedingly rare. However, the variable in the database for number of days fished is missing for about 96% of the observations. In light of what is known about the fishery, these missing values were assumed to be single day trips.

Because the fishing seasons for which the rules and guidelines apply are different across species, a “landing season” variable is created which corresponds to the harvest guidelines and other regulatory features of that specific species and season.

The features of the CPS fishery allow vessels to take multiple trips in a day, so to provide a more accurate measure of effort a new variable was created to correspond with individual trips. Trips cannot be perfectly identified from the data because a single trip can generate multiple fish tickets depending on where and to whom the landings are sold. Therefore I developed a heuristic approach to identifying trips. Because the data are quite complicated, a conservative approach was taken, as it is more likely that boats go out only once per day, and would only take multiple trips when they are able to quickly fill the hold. In light of this, I assume that each fish ticket represents a trip when a vessel's daily landings totaled more than 105% of their observed maximum single-ticket landings. The remaining observations are labeled as one trip per day, resulting in a similar measure of effort to the summation of vessel fishing days with a correlation coefficient of .992.

For use in calculating the production function, the fish ticket data are aggregated up separately for each species. Because the species school separately incidental catch is not a major factor and the production process can be modeled as non-joint in inputs, as a trip to catch one species will not produce a meaningful catch of other species. To more accurately assess targeted fishing behavior, the target species of each fishing trip was identified as the species which generated the majority of the revenue. Only landings corresponding with the target species were maintained in the analysis data set as incidental catch would bias the estimates of the effort elasticity. The "home region" is assumed to be the region in which a vessel had the most fish tickets for that season, where California is broken into three regions and Oregon and Washington are broken into two regions. The key numeric variables are summed up: landings weight and number of trips

To get a better sense for the overall activity of vessels in the CPS, the prices and number of vessels targeting each species are computed for each season. The average fish price is computed as total revenue divided by total landings within the season for the five CPS species and the top ten revenue-generating species that are also targeted by CPS vessels: Dungeness Crab, Albacore

Table 3.1: Annual Fishery Averages by Species

Species	Number of Vessels	Exvessel Price	Number of Trips	Catch (Mtons)	Vessel Size (Gross Tons)
Market Squid	88.9	0.21	30.4	698.0	75.5
Pacific Sardine	45.9	0.06	24.9	895.2	87.2
Northern Anchovy	28.9	0.05	15.2	274.1	74.0
Pacific Mackerel	41.6	0.08	12.8	306.3	86.9
Jack Mackerel	29.0	0.08	7.3	120.3	90.3

Tuna, Yellowfin Tuna, Bluefin Tuna, Skipjack Tuna, Swordfish, Chinook Salmon, Coho Salmon, Red Sea Urchin, and Lobster. The price is converted to a real value using Bureau of Labor Statistics Consumer Price Index data with Aug. 1983 as the base. The count of vessels with a landing of the major species was also computed for each season, as the number of vessels is an indicator of abundance, economic value, and may be a measure of competition. A summary of these statistics averaged at the annual level is presented in Table 3.1.

3.3.2 Weather and Ocean Condition Data

The coastal pelagic species are rather sensitive to ocean conditions, showing significant sensitivities to temperature and chlorophyll, an indicator of phytoplankton which is a key species in the base of the food web. Because this paper focuses on aggregated measures of production, it will not employ spatially varying measures of these conditions. Instead, time-varying series of observations corresponding to a particular location are used. The monthly observations are aggregated up to the annual level, as well as the "December-January-February" level which has been shown to be the most relevant for SST along the West Coast of the United States (Alexander et al., 2002). Furthermore, three years of lagged values are collected for each of the variables.

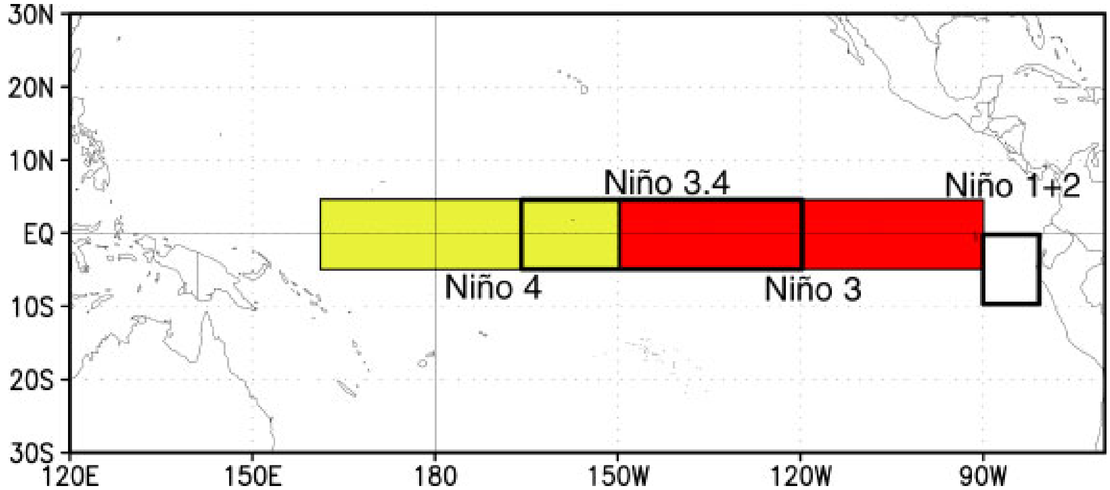

Figure 3.3: El Niño Regional Boundaries

Source: Climate Prediction Center (2019)

Periodic Oceanic Variations

The El Niño Southern Oscillation is an irregular high-frequency variation in Pacific Ocean temperatures in which periods of warmer sea temperatures are referred to as El Niño and periods of cooler sea temperatures are referred to as La Niña. These variations are measured and reported by scientists in a number of different data series, with some reporting only the sea surface temperature from certain regions and others creating indices from two or more variables to provide a more accurate picture. The data includes four regional measures of sea surface temperature as indicated in Figure 3.3. It is not obvious which region will be most relevant for the abundance of the coastal pelagic species, but the Niño 3.4 region is the basis for one of the primary indices used to monitor ENSO conditions: the Oceanic Niño Index. The Multi-variate ENSO Index version 2 (MEIv.2) uses both sea surface temperature and sea level air pressure to provide a richer view of the atmospheric and oceanic conditions of interest.

The Pacific Decadal Oscilation (PDO) is similar to ENSO in terms of its impacts, but occurs the PDO can be in a phase for 20-30 years while ENSO typically lasts 6 to 18 months. Since the two both govern warm/cold water, they can amplify or nullify one another. The

Pacific/North America Teleconnection (PNA) is a pattern of geopotential heights that appears to be correlated with ENSO and largely influences the weather of the United States. The North Pacific Gyre Oscillation (NPGO) is a periodic cycle in sea surface height which shows significant correlations with salinity, nutrients, and chlorophyll-a in California and Alaska (Di Lorenzo et al., 2008).

In addition to the tropical sea surface temperatures used to measure, daily sea surface temperature measures from Scripps Pier in La Jolla, CA were collected and aggregated to their monthly means. This measurement is of particular importance because the harvest guidelines for Pacific Sardine were explicitly formulated on the basis of the three-year average sea surface temperature from Scripps Pier. The Scripps Pier SST and sardine stock relationship was re-evaluated in McClatchie et al. (2010) which recommended removing the SST from the management guidelines. Deyle et al. (2013) takes another approach to the data, and finds that the sea surface temperature from the pier at Scripps Institution of Oceanography is a significant predictor of Pacific sardine abundance ($p = 0.04$), and also reports that measures of Pacific Decadal Oscillation ($p = 0.04$), and Northern Pacific Gyre Oscillation ($p = 0.19$) are able to improve forecasting. Other variables including measures of the El Niño index were not useful for forecasting abundance.

3.3.3 Economic Data

In addition to the prices and revenues from the PacFIN data, additional data for inflation and input prices were collected. All dollar values are converted from nominal to real using the BLS chained CPI data for all consumers. Marine diesel price data availability was limited in its time span, so national highway diesel fuel prices[3] were obtained from the U.S. Energy Information Administration to be used as a proxy. As a proxy for the cost of labor, national average hourly earnings of construction workers were obtained from Federal Reserve Economic

[3] Geographically specific fuel data is not available for years prior to 1994.

Data (FRED). Although workers on a fishing boat are usually paid a percentage of the revenue rather than an hourly rate, the construction wages correspond to an opportunity cost of their labor in a similarly physical job and should therefore be correlated with the firm's cost of labor. Lastly, as a measure of interest rates for borrowing and capital improvement, the corporate BAA bond yields were collected from FRED.

3.4 Methodology

> We hope readers will see that data mining done correctly is the opposite of "bad practice": it is an extremely useful tool that opens many doors in the analysis of interesting economic data. These tools allow researchers to add rigor and robustness to the "art" of variable or model selection in data analyses where the aim is to draw inferences about economically meaningful parameters. Belloni et al. (2014a, p. 48)

Big data techniques like the LASSO are designed to provide good predictions from data. In economics, we are often more interested in inference relative to the coefficients, but in the case of Instrumental Variables these two goals align. In the case of two-stage least squares, the first stage is essentially a prediction problem; the researcher is trying to get the best prediction of the endogenous variable subject to the exclusion restriction. The second stage allows for inference on the coefficient under certain conditions and with the proper methods.

The principal tool to be used in this paper is the Least Absolute Shrinkage and Selection Operator (LASSO), which is designed to simultaneously improve out-of-sample prediction and produce a parsimonious model to ease interpretabiltiy (Tibshirani, 1996). The LASSO regression is appropriate for situations where the function being estimated is "approximately sparse," meaning that the dependent variable can be well approximated using only a small subset of variables (Belloni et al., 2012). Zhao and Yu (2006) describes another condition for LASSO to perform well which they call "Irrepresentable Condition." This condition states that "Lasso selects the true model consistently if and (almost) only if the predictors that are not in the true model are 'irrepresentable' by predictors that are in the true model." This boils down to a question of multi-collinearity; if a variable that does not belong in the true model is a linear combination of variables in the true model, then LASSO will not consistently select the true model.

All variables should be centered and standardized prior to estimation, since what might be considered a "small" coefficient is also determined by the variance of the variables. However, as an alternative the penalty for each variable can be given a "penalty loading" which serves the same function as standardizing but requires less computation. In addition to the computational benefits, penalty loadings can also accommodate heteroskedastic and clustered errors (Belloni et al., 2012; Belloni and Chernozhukov, 2013; Belloni et al., 2014b, 2016)

The model is written as

$$\min_{\beta} \sum_{i=1}^{n} \left(y_i - \sum_{j=1}^{p} x_{ij}\beta_j \right)^2 + \lambda \sum_{j=1}^{P} |\beta_j|\gamma_j \tag{3.1}$$

where the first term is the familiar "sum of squared errors" and the second term is the "penalty level" λ times the sum of coefficients absolute values multiplied by the penalty loadings γ_j. The penalty weight λ is chosen to produce the best predictive accuracy, and can be identified based on cross-validation, information criteria, or as proposed in Belloni et al. (2012) it can be computed analytically:

$$\lambda = 2.2\sqrt{n}\Phi^{-1}\left(1 - \frac{\gamma}{2p}\right)$$

where p is the number of penalized regressors, Φ^{-1} is the inverse normal distribution, and γ is $0.1/\log(n)$ for the standard case and $0.1/\log(n_{clust})$ for the clustered case. The solution to this minimization problem will have $\beta_j = 0$ for variables which have a relatively small impact on the SSE, where the definition of "relatively small" depends on the penalty level λ chosen as described previously. A variable whose estimated coefficient is non-zero is said to have been "selected."

However, this LASSO construction relies on the assumption that σ, the standard deviation of the noise, is known or can be easily estimated. Particularly in the case with more variables in the choice set than there are observations, this estimation is non-trivial. Furthermore, the construction assumes that the noise is normality distributed. A variant on the LASSO which does not rely on the foreknowledge of σ or estimating it is the sqaure-root lasso proposed in Belloni et al. (2011), in which the SSE is replaced by the square root of the sum of squared errors. In addition to not being dependent on knowing or estimating σ, it does not rely on normality or sub-Gaussianity of the error.

The papers Belloni et al. (2012); Belloni and Chernozhukov (2013) point out a problem central to the issue of inference on coefficients with LASSO methods. While the "selection" element of LASSO works quite well, the "shrinkage" element causes the coefficients to be biased towards zero. However, performing OLS with the variables selected by LASSO recovers the unbiased coefficients. This procedure is referred to as "Post-LASSO OLS."

The LASSO is well suited to selecting instruments, particularly with data for which the number of variables is large relative to the number of observations, as demonstrated in Belloni et al. (2012). They demonstrate that the IV LASSO technique provides root-n consistency, asymptotic normality, and is robust to imperfect model selection. While that paper allows for many instruments, it assumes that any control variables to be included in the second stage are known in advance. In Belloni et al. (2014b) they develop a methodology for the case of an "exogenous treatment conditional on observables" setup called "post double selection" which is so named because there are two stages of variable selection. In the first stage, the LASSO is applied to selecting variables which are useful for predicting the dependent variable. In the second stage, the LASSO is applied to selecting variables which can predict the treatment variable. Lastly, OLS is run using the union of all the variables selected in the two stages. In Belloni et al. (2017) they extend the Post-Double Selection idea to allow for an endogenous treatment variable using orthogonal moment conditions to estimate. These methods are estimated via the STATA implementations of Ahrens et al. (2018a) and Ahrens et al. (2018b).

Throughout the series of theoretical advances in the use of LASSO for measuring treatment effects, the authors have also included applications of their method to re-evaluating well known papers and compare the results. In Belloni et al. (2012) they consider the impact of takings law on economic outcomes as in Chen and Yeh (2010). They find that the IV LASSO produces better first stage Wald statistics, implying that the LASSO is producing better first stage predictors, which reduces the standard error on the second stage estimates. In spite of this improvement in efficiency, the point-estimates are quite similar, ultimately producing the same economic interpretation. In Belloni et al. (2014b) they replicate the work of Acemoglu et al. (2001) to estimate the effect of institutions on economic performance. As in the previous exercise, they find a smaller but qualitatively similar point estimate in both the first and second stage, with a standard error on the second stage smaller than the Acemoglu et al. results. They suggest that these high-dimensional techniques may be a useful complement to existing strategies for robustness checks and sensitivity analyses.

Their choices for higher-dimensional variables depended on the problem at hand, using various mathematical transformations, splines, and interaction effects. It remains a fact that the choice of high dimensional controls and instruments is at the researcher's discretion, and the data-driven technique cannot answer whether the exclusion restriction is met. Furthermore, with the "irrepresentable condition" underlying the validity of the results, the inclusion of too many interaction terms or transformations can result in the LASSO producing spurious results. The researcher is not taken out of the process entirely, but rather than being a gatekeeper to determining which variables are included in the model, the researcher determines which variables are eligible to be included in the model.

3.5 Modeling and Results

3.5.1 Model Selection

It is necessary to select a functional form and a set of variables to be used as a baseline in order to compare the performance of the LASSO selected variables. The functional forms commonly used in the literature are the Cobb-Douglas and the translog model. The workhorse theoretical model of Schaefer (1957) assumes aggregate production is Cobb-Douglas with exponents of one on the inputs of effort and abundance. A Cobb-Douglas production model based on the model in Hannesson (1983) is chosen over the translog model for ease of interpretation and to keep the focus on establishing the viability of LASSO for fisheries. The translog functional form requires the inclusion of sets of covariates, squares, and interaction terms (e.g. K, L, K^2, L^2, & KL). It is likely that LASSO will not select every term in these sets, in which case the researcher must decide: keep only the selected variables, in which case the estimated production function would no longer be truly translog, or force the inclusion of full sets where partial sets are selected via LASSO, in which case the benefits of LASSO are not fully realized. It is therefore expedient to evaluate the benefits of LASSO for modeling fisheries with the Cobb-Douglas functional form, leaving applications to the translog for future work after having established its usefulness. Towards this end, the "overlapping group LASSO" appears the most promising for applications to translog. (Jacob et al., 2009; Yuan et al., 2011)

The model largely follows Hannesson (1983), with the estimating equation

$$\ln Y_{it} = \alpha + \beta_E \ln E_{it} + \beta_A \ln A_t + \beta_K \ln K_i + \beta_T T_t + \beta_{\mathbf{X}} \mathbf{X}_{it} + \varepsilon_{it}. \tag{3.2}$$

E_{it} is a vessel's annual number of trips as a measure of effort, A_t is the species abundance as estimated by stock assessment or catch-per-unit-effort[4] depending on available data, K_i represents capital stock as measured by the gross tonnage (size) of the vessel, and T_t represents the time trend by counting years since the start of the sample. The time trend is intended to account for Hicks neutral technological change at an increasing rate, but may also pick up other factors which have trended up over time.

For the traditional variable selection models, additional variables are included in $\mathbf{X}_{it}$ based

[4]This is possibly endogenous, but the hypothesis test fails to reject exogeneity.

on the regulatory framework of the coastal pelagic species. One set of these variables is the Scripps Pier sea surface temperature measurement for time t through $t-3$ since the three-year average temperature was included in the formula for establishing the quota for sardine, and is generally accepted as a useful measure for productivity for all of California's CPS fisheries. Furthermore, a dummy variable is included which takes the value of one after the fishery transitioned from open access to limited entry. Finally, $\mathbf{X}_{it}$ accounts for regional productivity differences via fixed effects for the region (e.g. Northern California and Central California) a vessel landed in most frequently during the season.

Using vessel fixed effects eliminates the time-invariant gross tonnage measure of capital while more effectively capturing the differences between vessels and skippers, leading to the following estimating equation

$$\ln Y_{it} = \alpha_i + \beta_E \ln E_{it} + \beta_A \ln A_t + \beta_T T_t + \beta_{\mathbf{X}} \mathbf{X}_t + \varepsilon_{it} \tag{3.3}$$

The proceeding models assume that effort is exogenously determined, but it may be that vessels are choosing to change their number of trips based on their landings of that species as in Felthoven et al. (2009), Wolff et al. (2013), and Squires and Vestergaard (2018). To account for this plausible endogeneity a Two Stage Least Squares (2SLS) model is also estimated, with the first stage regression estimated using a vector of excluded instruments $\mathbf{Z}_t$ which includes average annual ex-vessel prices of CPS and other species frequently harvested by the vessels as a measure of opportunity costs, as well as fuel prices and wage rates which determine variable costs. Felthoven et al. (2009) also used own price and substitute fish price as instruments, as well as gross tonnage and an ice-cover measure. The choice of instruments is specific to each fishery, and therefore was based on conversations with experts in the fishery.

Regarding the exclusion restriction, as long as vessels are price-takers in a global market the price of the species of interest is not determined by any individual vessel effort level, and should not be related to their catch-per-unit-effort, but would inform their decision to target that species. The price of other species is entirely separated from the production function of the species of interest, but because vessels can substitute across species these prices may influence how much effort goes towards the species of interest. Diesel fuel prices are determined by a global market, but since they represent a significant portion of operating expenses are likely to influence vessel decisions, but have no relationship to landings. The same logic applies to the wage rates.

Thus the instrumenting equation without vessel fixed effects is estimated as:

$$\ln E_{it} = \alpha + \gamma_A \ln A_t + \gamma_{GT} GT_i + \gamma_{\mathbf{X}} \mathbf{X}_{it} + \gamma_{\mathbf{Z}} \mathbf{Z}_t + \varepsilon_{it} \tag{3.4}$$

Models were estimated which also allowed for measurement error in the abundance measures, but the tests of endogeneity failed to reject the exogeneity of the variable. Because CPS are schooling species, abundance impacts the spatial extent whereas the bioeconomic model assumes fish are uniformly distributed throughout the space. In that case, the catch will be only loosely related to the abundance, and instruments will have little to no impact on the results. Since the test statistic for endogeneity is increasing in the distance between the coefficient estimated as endogenous and exogenous, when the point estimates are close to one another it is likely to fail to reject exogeneity.

For each species, the OLS and 2SLS models are estimated with and without vessel fixed effects, and using traditional variables as well as LASSO selected variables. In order to ensure the model would have the necessary variables to be a production function (and to improve comparability across models), the log abundance, log effort, and log gross tonnage were not penalized by LASSO and therefore included in every model. The regional fixed effects were similarly not penalized.

To populate the choice set of variables for LASSO, transforms of the relevant variables were generated: the natural logarithm and polynomials up to the fourth degree. Interaction terms were also tested, however the estimation became unstable with their inclusion, likely as a result of the irrepresentable condition.[5] Since there is no strong theoretical justification for the interaction of these variables, interactions were not included in the choice set for model selection. Lags of up to three years are included for variables such as weather conditions and bond yields that are available for the full span of the data. Lags from within the data were tested but are not included because this requires throwing out the first N observations for each vessel but in no cases were these lagged variables selected by the LASSO. In some cases the LASSO does not select any instruments, in which case the 2SLS estimate is not identified and no results are presented.

In light of the estimating equation, the variables for estimating the production function should include those that are related to the catchability coefficient and abundance, as these variables would directly impact catch. Therefore the variables considered for inclusion in this stage were the regulatory dummy, time trend, ocean and atmospheric conditions, and the applicable transformations and lags. The excluded variables considered for the instrument pool include fish prices, fuel prices, wage rates, and the number of vessels participating in each fishery as a measure of competition.

3.5.2 Comparison of Traditional and Lasso Methods

A summary of the relevant statistics for comparing the two methods is presented in Table 3.2, with a fuller view[6] of the results presented in Tables 3.3-3.7. The results strongly favor the use of LASSO for selecting variables in the context of fishery production. In the OLS models assuming exogenous effort the LASSO models show small improvements in the precision of the effort elasticity in eight cases, with only slight increases in the standard error for the other two cases. The 2SLS models show higher standard errors on the effort coefficient for seven of the

[5]Lasso performs poorly when there is significant collinearity in the variables included in the choice set.

[6]Due to the quantity and inconsistent inclusion of other covariates, in order to preserve readability only the results for the primary inputs to production are presented.

nine estimates, however one of the known confounds of weak instruments is estimated standard errors that are too small. The traditional instruments are universally weaker than the LASSO instruments, with the traditional instruments achieving the "rule of thumb" F-statistic of 10 or greater in only one model, while the LASSO instruments produce an F-statistic of 10 or greater in eight of the nine models. Therefore the seeming precision of the traditional 2SLS models is a misleading result of the weak instruments, with LASSO selecting stronger instruments and providing more accurate measures of the standard error. Finally, even though the LASSO is intended to reduce overfitting and improve out of sample fit, it produces an adjusted R^2 that is the same in two cases and improved in three cases. This is somewhat surprising, but in almost all cases the LASSO selects fewer variables than were included in the traditional models. Some of the variables that were dropped in the LASSO model were statistically significant in the traditional model, highlighting one of the major differences between the LASSO and step-wise regression; beyond just identifying correlated variables, it serves to identify variables with relatively strong predictive power.

Turning to the species by species results of Tables 3.3-3.7, it is apparent that the changes as a result of using LASSO are minimal in the case of OLS estimation, but for 2SLS there are marked differences in the estimates. Given the differences in the first stage F-statistics, it is not surprising that the point estimates are also quite different for the 2SLS models. The similarity between the OLS estimates is also unsurprising, since the difference between the LASSO and the traditional variable sets are "shifters" which may slightly change the catchability of a species (e.g. warmer waters cause fish to school further from shore), while most of the variation in annual catch is determined by the effort, abundance, and vessel characteristics which are included in both variable sets.

The difference-in-Sargan test for endogeneity, similar to the Hausman test but robust to heteroskedasticity, provides evidence that effort may be endogenously determined. The test statistic is less than 0.05 in at least one model for each species, For every species except market squid, the LASSO selects fewer control variables than might be included with traditional models, even selecting zero control variables in several instances. This shows that, on the one hand a simple and parsimonious model with reasonable results can be produced using data-driven variable selection, but on the other hand the inclusion of these extra control variables has had minimal impact on the results. There is no observable consistency across species of which ocean conditions are most predictive of catch, with different measures of ENSO, PNA, PDO, and NPGO all being selected for different species. Although the correlation between the various El Niño measures is high, the longer term cycles of PNA, PDO, and NPGO are less closely related, indicating that there are dangers to a "one model fits all" approach to estimating production functions for different species when we know that different oceanographic conditions affect the prosperity and the location of each species differently.

3.6 Concluding Remarks

The results clearly show that variable selection via LASSO can be a useful tool for the estimation of production in fisheries. The improvements are marginal in the case where effort is exogenous, with slight decreases in the standard errors and in some cases improvements in the adjusted R^2. When effort is endogenous, consistent with the claims of Belloni et al. (2012) there are dramatic improvements in the performance of 2SLS when LASSO is used to select the instruments. The first stage of two-stage least squares is fundamentally a prediction problem, which is precisely what LASSO is designed to address, and correspondingly it achieves higher first stage F-statistics.

While biologists are still debating about how different measures of oceanographic conditions impact different species, as can be seen in the literature regarding Pacific Sardine, it can be helpful to take an agnostic and data-driven approach to the inclusion of such variables. It is particularly important in multi-species fisheries to accommodate the differences in biology by

Table 3.2: Comparison of Traditional and LASSO methods

(a) Comparison of Standard Error for Elasticity of Effort

	Pacific Sardine		Market Squid		Northern Anchovy		Pacific Mackerel		Jack Mackerel	
	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO
OLS	0.043	0.042	0.040	0.038	0.067	0.066	0.087	0.082	0.081	0.052
FE OLS	0.032	0.031	0.025	0.026	0.060	0.066	0.043	0.042	0.058	0.057
IV	0.180	−.−	0.083	0.136	0.128	0.218	0.201	0.169	0.181	0.360
FE IV	0.120	0.085	0.060	0.073	0.109	0.140	0.148	0.171	0.149	0.182

(b) Comparison of Adjusted R^2 of OLS Models

	Pacific Sardine		Market Squid		Northern Anchovy		Pacific Mackerel		Jack Mackerel	
	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO
OLS	0.870	0.873	0.864	0.864	0.654	0.653	0.730	0.730	0.644	0.639
FE OLS	0.841	0.839	0.816	0.822	0.709	0.694	0.654	0.658	0.606	0.588

(c) Comparison of First Stage F-statistics of IV Models

	Pacific Sardine		Market Squid		Northern Anchovy		Pacific Mackerel		Jack Mackerel	
	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO	Trad.	LASSO
IV	2.21	−.−	8.89	29.44	7.95	21.98	3.03	23.92	6.98	9.31
FE IV	2.35	20.82	16.11	39.47	6.34	12.49	3.180	23.92	7.33	19.57

Table 3.3: Pacific Sardine

	OLS		Robust LASSO		2SLS		IV LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effort	1.27	1.16	1.28	1.18	0.85	0.77		1.08
	(0.043)	(0.032)	(0.042)	(0.031)	(0.180)	(0.120)		(0.085)
Abundance	−0.03	0.08	−0.06	−0.01	0.23	0.35		0.05
	(0.068)	(0.073)	(0.044)	(0.045)	(0.120)	(0.113)		(0.046)
Gross Tonnage	1.16		1.16		1.45			
	(0.125)		(0.122)		(0.176)			
Region Fixed Effects	X	X	X	X	X	X	X	X
Vessel Fixed Effects		X		X		X		X
Additional Controls	6	6	2[a]	1[b]	6	6	0	3[c]
Instrumental Variables					9	9	0	2[d]
Endogeneity p-value					0.022	0.031		0.058
Kleibergen-Paap Wald rk F-stat					2.31	2.35		20.82
Number of vessels	158	158	158	158	155	155	157	157
Number of observations	1,345	1,345	1,345	1,345	1,189	1,189	1,287	1,287

[a] Selected Controls: PNA (Dec-Feb) and PDO (Dec-Feb)
[b] Selected Controls: PDO (Dec-Feb)
[c] Selected Controls: PDO (lag 1)
[d] Selected Instruments: Price^2 of Dungeness Crab, Log Number of Sardine Vessels

Table 3.4: Market Squid

	OLS		Robust LASSO		2SLS		IV LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effort	1.35	1.22	1.36	1.20	1.47	1.52	1.29	1.27
	(0.040)	(0.025)	(0.038)	(0.026)	(0.083)	(0.060)	(0.136)	(0.073)
Abundance	0.12	0.18	0.17	0.13	0.12	0.14	0.09	−0.14
	(0.040)	(0.032)	(0.030)	(0.026)	(0.042)	(0.039)	(0.045)	(0.092)
Gross Tonnage	0.81		0.80		0.68		0.86	
	(0.063)		(0.064)		(0.099)		(0.134)	
Region Fixed Effects	X	X	X	X	X	X	X	X
Vessel Fixed Effects		X		X		X		X
Additional Controls	6	6	0	3[a]	6	6	7[b]	12[d]
Instrumental Variables					9	9	1[c]	1[c]
Endogeneity p-value					0.378	0.010	0.620	−. −[e]
Kleibergen-Paap Wald rk F-stat					8.89	16.11	29.44	39.47
Number of vessels	187	187	187	187	183	183	187	187
Number of observations	1,699	1,699	1,699	1,699	1,500	1,500	1,631	1,631

[a] Selected Controls: Log(Niño 1), Niño 1 (Dec-Feb), PDO (Lag 1)
[b] Selected Controls: Niño 1 (Annual), Niño 1 (Dec-Feb), ENSO Multivariate Index, ENSO Multivariate Index^3, PNA^4, PNA (Lag 2), and PDO (Dec-Feb).
[c] Selected Instruments: Log(Average price of Coastal Pelagic Species)
[d] Selected Controls: Nina 1 (Annual), Log(Niño 1), Niño 1 (Dec-Feb), ENSO Multivariate Index, ENSO Multivariate Index^3, PNA^4, PNA (Lag 2), PNA (Dec-Feb), Niño 3 (Dec-Feb), ENSO TS^3, PDOY^3, Abundance.
[e] Could not compute cluster-robust covariance matrix due to singleton dummy variable

Table 3.5: Northern Anchovy

	OLS		Robust LASSO		2SLS		IV LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effort	0.99	1.14	1.00	1.13	1.28	1.19	1.84	1.59
	(0.067)	(0.060)	(0.066)	(0.066)	(0.128)	(0.109)	(0.218)	(0.140)
Abundance	0.00	0.00	0.10	−0.03	−0.03	0.00	−0.04	−0.13
	(0.042)	(0.022)	(0.035)	(0.015)	(0.050)	(0.029)	(0.061)	(0.041)
Gross Tonnage	1.46		1.47		1.42		1.45	
	(0.195)		(0.186)		(0.214)		(0.311)	
Region Fixed Effects	X	X	X	X	X	X	X	X
Vessel Fixed Effects		X		X		X		X
Additional Controls	6	6	2[a]	0	6	6	2[b]	4[d]
Instrumental Variables					9	9	1[c]	2[e]
Endogeneity p-value					0.010	0.010	0.000	0.000
Kleibergen-Paap Wald rk F-stat					7.95	6.34	21.98	12.49
Number of vessels	134	134	134	134	129	129	129	129
Number of observations	830	830	830	830	734	734	750	750

[a] Selected Controls: Multivariate ENSO Index (Lag 1), PDO
[b] Selected Controls: Multivariate Enso Index (Lag 1), Scripps Pier SST (Lag 2)
[c] Selected Instruments: Price of Sardine[4]
[d] Selected Controls: Niño 1 (Dec-Feb), PNA (Dec-Feb), Niño 3 (Lag 1), Scripps Pier SST (Lag 2)
[e] Selected Instruments: Price of Sardine[3], Price of Sardine[4]

Table 3.6: Pacific Mackerel

	OLS		Robust LASSO		2SLS		IV LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effort	1.43	1.29	1.41	1.28	1.39	1.27	1.10	0.89
	(0.087)	(0.043)	(0.082)	(0.042)	(0.201)	(0.148)	(0.169)	(0.171)
Abundance	0.28	0.13	0.00	0.01	0.20	0.07	0.07	0.11
	(0.078)	(0.103)	(0.072)	(0.072)	(0.080)	(0.112)	(0.065)	(0.089)
Gross Tonnage	1.24		1.26		1.31		1.41	
	(0.139)		(0.132)		(0.169)		(0.141)	
Region Fixed Effects	X	X	X	X	X	X	X	X
Vessel Fixed Effects		X		X		X		X
Additional Controls	6	6	1[a]	1[a]	6	6	1[b]	2[d]
Instrumental Variables					9	9	2[c]	1[e]
Endogeneity p-value					0.931	0.557	0.278	0.015
Kleibergen-Paap Wald rk F-stat					3.03	3.18	23.92	42.37
Number of vessels	158	158	158	158	146	146	153	153
Number of observations	1,127	1,127	1,127	1,127	949	949	987	987

[a] Selected Controls: NPGO
[b] Selected Controls: PDO[4]
[c] Selected Instruments: Log(Number of Jack Mackerel Vessels), Log(Number of Market Squid Vessels)
[d] Selected Controls: Nina 3.4 (Lag 2), PDO[4]
[e] Selected Instruments: Log(Number of Market Squid Vessels)

Table 3.7: Jack Mackerel

	OLS		Robust LASSO		2SLS		IV LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effort	1.33	1.20	1.41	1.20	1.08	1.13	2.08	1.51
	(0.081)	(0.058)	(0.052)	(0.057)	(0.181)	(0.149)	(0.360)	(0.182)
Abundance	−0.07	−0.02	0.02	0.05	−0.07	−0.04	−0.02	−0.07
	(0.070)	(0.059)	(0.057)	(0.039)	(0.071)	(0.061)	(0.065)	(0.046)
Gross Tonnage	1.14		1.12		1.20	(3.805)	1.04	
	(0.166)		(0.167)		(0.169)		(0.199)	
Region Fixed Effects	X	X	X	X	X	X	X	X
Vessel Fixed Effects		X		X		X		X
Additional Controls	6	6	1[a]	1[b]	6	6	2[c]	2[e]
Instrumental Variables					9	9	2[d]	2[f]
Endogeneity p-value					0.465	0.942	0.000	−. −[g]
Kleibergen-Paap Wald rk F-stat					6.98	7.33	9.31	19.57
Number of vessels	114	114	114	114	111	111	112	112
Number of observations	760	760	760	760	661	661	674	674

[a] Selected Controls: PDO (Lag 3)
[b] Selected Controls: Log(Time Trend)
[c] Selected Controls: Scripps Pier SST (Lag 2), Log(Time Trend)
[d] Selected Instruments: Scripps Pier SST (Lag 2), Scripps Pier SST (Lag 3), Log(Time Trend)
[e] Selected Controls: Number of Jack Mackerel Vessels, Log(Number of Jack Mackerel Vessels)
[f] Selected Instruments: Price of Sardine, Log(Number of Jack Mackerel Vessels)
[g] Could not compute cluster-robust covariance matrix due to singleton dummy variable

including the variables which are most related to each species' production.

In addition to the benefits of LASSO in the estimation itself, there is an additional benefit in the form of transparency, as the researcher picks only the choice set of variables and not the variables themselves. This makes the estimation procedure less vulnerable to researcher interference such as p-hacking or targeting a particular point estimate. Unfortunately because the variable selection is data driven, the selected variables are specific to that data set and that research question, and should not be viewed as prescriptive variable choices for future research.

3.7 Acknowledgments

Chapter 3, in full, is an unpublished manuscript. The dissertation author is the sole author. I would like to acknowledge Dr. James Hilger who selected me to work on the Coastal SEES project of which this research is an offshoot, and whose curiosity regarding the application of big data techniques to fisheries was the spark for this paper. Thanks also to Professor Art Miller for his feedback on the relevant climate and oceanographic variables to include, Dr. Kristin Roll for comments and feedback, and to Dr. Emmanis Dorval for useful conversations regarding market squid biology and fishing.

References

Acemoglu, D., Johnson, S., Robinson, J. A. 2001. The Colonial Origins of Comparative Development: An Empirical Investigation. American Economic Review. 91 (5), 1369–1401.

Agnolucci, P., Barker, T., Ekins, P. 2004. Hysteresis and energy demand: the announcement effects and the effects of the uk climate change levy. In Applied Environmental Economics Conference, March, The Royal Society, UK, 1–12.

Agthe, D. E., Billings, R. B. 1987. Equity, Price Elasticity, and Household Income under Increasing Block Rates for Water. American Journal of Economics and Sociology. 46 (3), 273–286.

Ahrens, A., Hansen, C. B., Schaffer, M. E. 2018a. LASSOPACK: Stata module for lasso, square-root lasso, elastic net, ridge, adaptive lasso estimation and cross-validation. Statistical Software Components, Boston College Department of Economics.

Ahrens, A., Hansen, C. B., Schaffer, M. E. 2018b. PDSLASSO: Stata module for post-selection and post-regularization OLS or IV estimation and inference. Statistical Software Components, Boston College Department of Economics.

Ailawadi, K. L., Gedenk, K., Lutzky, C., Neslin, S. A. 2007. Decomposition of the sales impact of promotion-induced stockpiling. Journal of Marketing Research. 44 (3), 450–467.

Alexander, M. A., Bladé, I., Newman, M., Lanzante, J. R., Lau, N.-C., Scott, J. D. 2002. The atmospheric bridge: The influence of enso teleconnections on air–sea interaction over the global oceans. Journal of Climate. 15 (16), 2205–2231.

Andersen, P. K., Gill, R. D. 1982. Cox's regression model for counting processes: a large sample study. The annals of statistics. 1100–1120.

Athey, S., Imbens, G. W. 2006. Identification and inference in nonlinear difference-in-differences models. Econometrica. 74 (2), 431–497.

Baerenklau, K. A., Schwabe, K. A., Dinar, A. 2014. The residential water demand effect of increasing block rate water budgets. Land Economics. 90 (4), 683–699.

Belloni, A., Chernozhukov, V., Wang, L. 2011. Square-root lasso: pivotal recovery of sparse signals via conic programming. Biometrika. 98 (4), 791–806.

Belloni, A., Chen, D., Chernozhukov, V., Hansen, C. 2012. Sparse models and methods for optimal instruments with an application to eminent domain. Econometrica. 80 (6), 2369–2429.

Belloni, A., Chernozhukov, V. 2013. Least squares after model selection in high-dimensional sparse models. Bernoulli. 19 (2), 521–547.

Belloni, A., Chernozhukov, V., Fernández-Val, I., Hansen, C. 2017. Program evaluation and causal inference with high-dimensional data. Econometrica. 85 (1), 233–298.

Belloni, A., Chernozhukov, V., Hansen, C. 2014a. High-dimensional methods and inference on structural and treatment effects. Journal of Economic Perspectives. 28 (2), 29–50.

Belloni, A., Chernozhukov, V., Hansen, C. 2014b. Inference on treatment effects after selection among high-dimensional controls. The Review of Economic Studies. 81 (2), 608–650.

Belloni, A., Chernozhukov, V., Hansen, C., Kozbur, D. 2016. Inference in high-dimensional panel models with an application to gun control. Journal of Business & Economic Statistics. 34 (4), 590–605.

Bijwaard, G. E., Franses, P. H., Paap, R. 2006. Modeling purchases as repeated events. Journal of Business & Economic Statistics. 24 (4), 487–502.

Borenstein, S. 2009. To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing. Unpublished.

Brécard, D., Hlaimi, B., Lucas, S., Perraudeau, Y., Salladarré, F. 2009. Determinants of demand for green products: An application to eco-label demand for fish in europe. Ecological economics. 69 (1), 115–125.

Brown, P. 2005. Standardized Reference Evapotranspiration: A New Procedure for Estimating Reference Evapotranspiration in Arizona.Technical report, University of Arizona.

Brown, P., Kopec, D. 2000. Converting Reference Evapotranspiration into Turf Water Use.Technical report, University of Arizona.

Castro-Rodríguez, F., María Da-Rocha, J., Delicado, P. 2002. Desperately Seeking θ's: Estimating the Distribution of Consumers Under Increasing Block Rates. Journal of Regulatory Economics. 22, 29–58, 10.1023/A:1019951702523.

Chen, D. L., Yeh, S. 2010. The economic impacts of eminent domain.

Climate Prediction Center. 2019. Climate Prediction Center Oceanic Nino Index. https://catalog.data.gov/dataset/climate-prediction-center-cpcoceanic-nino-index, Accessed: 5-November-2019.

Daloonpate, A. 2002. Estimating Degree of Market Power and Price-Response Strategies in a Product-Differentiated Oligopoly: The Case of Canned Tuna Industry in a Local Market. Ph.D. dissertation, University of Tennessee.

Deaton, A., Muellbauer, J. 1980. An almost ideal demand system. The American economic review. 70 (3), 312–326.

Deyle, E. R., Fogarty, M., Hsieh, C.-h., Kaufman, L., MacCall, A. D., Munch, S. B., Perretti, C. T., Ye, H., Sugihara, G. 2013. Predicting climate effects on pacific sardine. Proceedings of the National Academy of Sciences. 110 (16), 6430–6435.

Di Lorenzo, E., Schneider, N., Cobb, K. M., Franks, P. J. S., Chhak, K., Miller, A. J., McWilliams, J. C., Bograd, S. J., Arango, H., Curchitser, E., Powell, T. M., Rivière, P. 2008. North pacific gyre oscillation links ocean climate and ecosystem change. Geophysical Research Letters. 35 (8).

Dorval, E., Crone, P. R., McDaniel, J. D. 2013. Variability of egg escapement, fishing mortality and spawning population in the market squid fishery in the california current ecosystem. Marine and Freshwater Research. 64 (1), 80–90.

Empen, J. J.-P. L., Weiss, C. 2011. Price Promotion and Brand Loyalty: Empirical Evidence for the German Ready-to-EatCerealMarket. In presentation at the Inra-Idei Seminar on Competition and Strategies in the Retailing Industry in Toulouse, on May, 16–17.

Felthoven, R. G., Paul, C. J. M., Torres, M. 2009. Measuring productivity and its components for fisheries: the case of the alaskan pollock fishery, 1994–2003. Natural Resource Modeling. 22 (1), 105–136, Uses North Pacific Index, Annual temperature, and Winter temperature. Uses polock and flatfish prices, gross tonnage, and an ice-cover index as instruments.

Ferraro, P. J., Price, M. K. 2013. Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. Review of Economics and Statistics. 95 (1), 64–73.

Glaser, L., Thompson, G. 1998. Demand for organic and conventional frozen vegetables. In Annual Meeting August, 8–11.

Glaser, L. K., Thompson, G. D. 2000. Demand for organic and conventional beverage milk. In Western Agricultural Economics Association Annual Meeting, June.

Gordinier, J. 2015. Seafood Restaurants Cast a Wider Net for Sustainable Fish. The New York Times.

Gordon, D. V. 2015. The endogeneity problem in applied fisheries econometrics: A critical review. Environmental and Resource Economics. 61 (1), 115–125.

Gordon, H. S. 1954. The economic theory of a common-property resource: The fishery. Journal of Political Economy. 62, 124–124.

Grubb, M. 2005. Asymmetrical price elasticities of energy demand. In Global warming and energy demand, 252–256, Routledge.

Gupta, S. 1991. Stochastic models of interpurchase time with time-dependent covariates. Journal of Marketing Research. 28 (1), 1–15.

Guttormsen, A. G., Roll, K. H. 2011. Technical efficiency in a heterogeneous fishery: the case of norwegian groundfish fisheries. Marine Resource Economics. 26 (4), 293–307.

Hallstein, E., Villas-Boas, S. B. 2013. Can household consumers save the wild fish? lessons from a sustainable seafood advisory. Journal of Environmental Economics and Management. 66 (1), 52–71.

Hannesson, R. 1983. Bioeconomic production function in fisheries: Theoretical and empirical analysis. Canadian Journal of Fisheries and Aquatic Sciences. 40, 968–982.

Helsen, K., Schmittlein, D. C. 1993. Analyzing duration times in marketing: Evidence for the effectiveness of hazard rate models. Marketing Science. 12 (4), 395–414.

Hendel, I., Lizzeri, A., Roketskiy, N. 2014. Nonlinear pricing of storable goods. American Economic Journal: Microeconomics. 6 (3), 1–34.

Hendel, I., Nevo, A. 2006. Measuring the Implications of Sales and Consumer Inventory Behavior. Econometrica. 74 (6), 1637–1673.

Herrick, S. 2000. Analysis of sardine markets.Technical report, National Marine Fisheries Service.

Hilger, J., Hallstein, E., Stevens, A. W., Villas-Boas, S. B. 2019. Measuring willingness to pay for environmental attributes in seafood. Environmental and Resource Economics. 73 (1), 307–332.

Hosken, D., Reiffen, D. 2001. Multiproduct retailers and the sale phenomenon. Agribusiness. 17 (1), 115–137.

Ito, K. 2010. Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. Job Market Paper.

Ito, K. 2014. Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. The American Economic Review. 104 (2), 537–563.

Jacob, L., Obozinski, G., Vert, J.-P. 2009. Group lasso with overlap and graph lasso. In Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09, 433–440, New York, NY, USA, Association for Computing Machinery.

Jaffry, S., Pickering, H., Ghulam, Y., Whitmarsh, D., Wattage, P. 2004. Consumer choices for quality and sustainability labelled seafood products in the uk. Food Policy. 29 (3), 215–228.

Jain, D. C., Vilcassim, N. J. 1991. Investigating household purchase timing decisions: A conditional hazard function approach. Marketing Science. 10 (1), 1–23.

Johnston, R. J., Roheim, C. A. 2006. A battle of taste and environmental convictions for ecolabeled seafood: A contingent ranking experiment. Journal of agricultural and resource economics. 283–300.

Johnston, R. J., Wessells, C. R., Donath, H., Asche, F. 2001. Measuring consumer preferences for ecolabeled seafood: An international comparison. Journal of Agricultural and Resource Economics. 26 (1), 20–39.

Kirkley, J., Squires, D., Strand, I. E. 1998. Characterizing Managerial Skill and Technical Efficiency in a Fishery. Journal of Productivity Analysis. 9 (2), 145–160.

Klaiber, H. A., Smith, V. K., Kaminsky, M., Strong, A. 2014. Measuring price elasticities for residential water demand with limited information. Land Economics. 90 (1), 100–113.

Liu, Y., Balachander, S. 2014. How long has it been since the last deal? consumer promotion timing expectations and promotional response. Quantitative Marketing and Economics. 12 (1), 85–126.

Machado, J. A. F., Silva, J. M. C. S. 2005. Quantiles for Counts. Journal of the American Statistical Association. 100 (472), 1226–1237.

Mansur, E. T., Olmstead, S. M. 2012. The value of scarce water: Measuring the inefficiency of municipal regulations. Journal of Urban Economics. 71 (3), 332–346.

Marian, L., Chrysochou, P., Krystallis, A., Thøgersen, J. 2014. The role of price as a product attribute in the organic food context: An exploration based on actual purchase data. Food Quality and Preference. 37, 52–60.

Maxwell, M. R., Jacobson, L. D., Conser, R. J. 2005. Eggs-per-recruit model for management of the california market squid (loligo opalescens) fishery. Canadian Journal of Fisheries and Aquatic Sciences. 62 (7), 1640–1650.

McClatchie, S., Goericke, R., Auad, G., Hill, K. 2010. Re-assessment of the stock–recruit and temperature–recruit relationships for pacific sardine (sardinops sagax). Canadian Journal of Fisheries and Aquatic Sciences. 67 (11), 1782–1790.

Melly, B., Santangelo, G. 2015. The changes-in-changes model with covariates. Unpublished.

Miranda, A. 2006. QCOUNT: Stata program to fit quantile regression models for count data. Statistical Software Components, Boston College Department of Economics.

Narasimhan, C., Neslin, S. A., Sen, S. K. 1996. Promotional elasticities and category characteristics. Journal of marketing. 60 (2), 17–30.

Nataraj, S., Hanemann, W. M. 2011. Does marginal price matter? a regression discontinuity approach to estimating water demand. Journal of Environmental Economics and Management. 61 (2), 198–212.

Neslin, S. A., Henderson, C., Quelch, J. 1985. Consumer promotions and the acceleration of product purchases. Marketing science. 4 (2), 147–165.

Nevo, A., Hatzitaskos, K. 2005. Why does the average price of tuna fall during Lent?.Technical report, National Bureau of Economic Research.

NOAA Fisheries. 2019a. NOAA Fisheries — Northern Anchovy. https://www.fisheries.noaa.gov/species/northern-anchovy, Accessed: 5-November-2019.

NOAA Fisheries. 2019b. NOAA Fisheries — Pacific Mackerel. https://www.fisheries.noaa.gov/species/pacific-mackerel, Accessed: 5-November-2019.

NOAA Fisheries. 2019c. NOAA Fisheries — Pacific Sardine. https://www.fisheries.noaa.gov/species/pacific-sardine, Accessed: 5-November-2019.

Nordin, J. A. 1976. A Proposed Modification of Taylor's Demand Analysis: Comment. The Bell Journal of Economics. 7 (2), 719–721.

Pacific Fisheries Information Network (PacFIN). www.psmfc.org, retrieval dated 12-August-2019.

Pacific Fishery Management Council. 2011. Status of the pacific coast coastal pelagic species fishery and recommended acceptable biological catches. stock assessment and fishery evaluation 2011.Technical report.

Pacific Fishery Management Council. 2019a. Coastal Pelagic Species: Background. https://www.pcouncil.org/coastal-pelagic-species/background-information/, Accessed: 5-November-2019.

Pacific Fishery Management Council. 2019b. Status of the pacific coast coastal pelagic species fishery and recommended acceptable biological catches. stock assessment and fishery evaluation for 2018.Technical report.

Pascoe, S., Coglan, L. 2000. Implications of differences in technical efficiency of fishing boats for capacity measurement and reduction. Marine Policy. 24 (4), 301–307.

Pascoe, S., Coglan, L. 2002. The contribution of unmeasurable inputs to fisheries production: an analysis of technical efficiency of fishing vessels in the english channel. American Journal of Agricultural Economics. 84 (3), 585–597.

Paul, C. J. M., Torres, M. d. O., Felthoven, R. G. 2009. Fishing revenue, productivity and product choice in the alaskan pollock fishery. Environmental and Resource Economics. 44 (4), 457.

Perretti, C. T., Sedarat, M. 2016. The influence of the el niño southern oscillation on paralarval market squid (doryteuthis opalescens). Fisheries oceanography. 25 (5), 491–499.

Pesendorfer, M. 2002. Retail sales: A study of pricing behavior in supermarkets*. The Journal of Business. 75 (1), 33–66.

Powell, D. 2016. Quantile Regression with Nonadditive Fixed Effects. Quantile Treatment Effects. Available at: http://works.bepress.com/david_powell/1/.

Renwick, M. E., Archibald, S. O. 1998. Demand side management policies for residential water use: who bears the conservation burden? Land economics. 343–359.

Roheim, C. A., Asche, F., Santos, J. I. 2011. The elusive price premium for ecolabelled products: evidence from seafood in the uk market. Journal of Agricultural Economics. 62 (3), 655–668.

Salladarré, F., Guillotreau, P., Perraudeau, Y., Monfort, M.-C. 2010. The demand for seafood eco-labels in france. Journal of agricultural & food industrial organization. 8 (1).

Schaefer, M. B. 1957. Some considerations of population dynamics and economics in relation to the management of the commercial marine fisheries. Journal of the Fisheries Board of Canada. 14 (5), 669–681.

Squires, D. 1987. Fishing effort: Its testing, specification, and internal structure in fisheries economics and management. Journal of Environmental Economics and Management. 14, 268–282.

Squires, D., Vestergaard, N. 2013. Technical Change and The Commons. The Review of Economics and Statistics. 95 (5), 1769–1787.

Squires, D., Vestergaard, N. 2018. Rethinking the commons problem: Technical change, knowledge spillovers, and social learning. Journal of Environmental Economics and Management. 91 (C), 1–25.

Sun, C.-H. J., Chiang, F.-S., Owens, M., Squires, D. 2017. Will american consumers pay more for eco-friendly labeled canned tuna? estimating us consumer demand for canned tuna varieties using scanner data. Marine Policy. 79, 62 – 69.

Taylor, L. D. 1975. The Demand for Electricity: A Survey. The Bell Journal of Economics. 6 (1), 74–110.

Teisl, M. F., Roe, B., Hicks, R. L. 2002. Can eco-labels tune a market? Evidence from dolphin-safe labeling. Journal of Environmental Economics and Management. 43 (3), 339–359.

Thompson, G., Glaser, L. 2001. National demand for organic and conventional baby food. In Selected Paper presented at Annual Meeting of the Western Agricultural Economics Association, Logan, Utah.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology). 58, 267–288.

Wang, E. Y. 2015. The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity. The RAND Journal of Economics. 46 (2), 409–441.

Wessells, C. R., Johnston, R. J., Donath, H. 1999. Assessing consumer preferences for ecolabeled seafood: The influence of species, certifier, and household attributes. American Journal of Agricultural Economics. 81 (5), 1084–1089.

Wichman, C. J. 2014. Perceived price in residential water demand: Evidence from a natural experiment. Journal of Economic Behavior & Organization. 107, 308–323.

Wolff, F.-C., Squires, D., Guillotreau, P. 2013. The Firm's Management in Production: Management, Firm, and Time Effects in an Indian Ocean Tuna Fishery. American Journal of Agricultural Economics. 95 (3), 547–567.

Yiridoe, E. K., Bonti-Ankomah, S., Martin, R. C. 2005. Comparison of consumer perceptions and preference toward organic versus conventionally produced foods: A review and update of the literature. Renewable Agriculture and Food Systems. 20 (4), 193–205.

Yoo, J., Simonit, S., Kinzig, A. P., Perrings, C. 2014. Estimating the price elasticity of residential water demand: the case of phoenix, arizona. Applied Economic Perspectives and Policy. 36 (2), 333–350.

Yuan, L., Liu, J., Ye, J. 2011. Efficient methods for overlapping group lasso. In Shawe-Taylor, J., Zemel, R. S., Bartlett, P. L., Pereira, F., Weinberger, K. Q. eds. Advances in Neural Information Processing Systems 24, 352–360, Curran Associates, Inc.

Zander, K., Feucht, Y. 2018. Consumers' willingness to pay for sustainable seafood made in europe. Journal of international food & agribusiness marketing. 30 (3), 251–275.

Zhao, P., Yu, B. 2006. On Model Selection Consistency of Lasso. Journal of Machine Learning Research. 7 (Nov), 2541–2563.

Zwolinski, J. P., Emmett, R. L., Demer, D. A. 2011. Predicting habitat to optimize sampling of pacific sardine (sardinops sagax). ICES Journal of Marine Science. 68 (5), 867–879.