Does Purchase Without Search Explain Counter Cyclic Pricing?

Avery Haviv
Rotman School of Management, University of Toronto

October 14, 2013

Abstract

Basic economic theory tells us to expect that an increase in demand should lead to an increase in price. However, previous studies have found the opposite trend in the prices of seasonal goods, such as canned soup. I propose an explanation of this phenomenon: consumers are more likely to purchase without search in low demand periods, reducing the gains of temporary price reductions, and decreasing estimated price sensitivity. Purchase without search is consistent with consumers using shopping lists to make their purchase decisions before observing prices. I test this explanation using a novel dynamic, structural inventory model where consumers make decisions on whether to search, which reveals price promotions, and which products to purchase given their search decision.

Estimating this model using previous methods is a computational challenge because of the expansion of the state space required to model seasonal preferences. To overcome this challenge, I develop a cyclic-successive approximation algorithm, which removes the computational burden of adding cyclic variables to the state space of a dynamic model.

I find that consumers purchase without search 71% of the time in winter. This causes price elasticities that are more than 60% larger in winter as they are in the summer. I find that the dominant cause of seasonal search is seasonal price variation, rather than seasonal consumption utility.

1 Introduction

Basic economic theory tells us to expect that an increase in demand should lead to an increase in price. However, previous studies have found the opposite trend in the prices of seasonal goods, such as canned soup, tuna, and ice cream, where prices have been observed to decrease in periods of high demand. This phenomenon has been termed "counter-cyclic pricing".

In this paper, I propose that counter-cyclic pricing is partially a reaction to seasonal changes in consumers' propensity to make purchase decisions without searching the category. For example, shoppers who use a shopping list will choose both the variety and quantity of a good to purchase before observing prices at the store. An increase in purchases without search will increase average prices because non-searching consumers will not react to discounts, and so holding a sale will simply decrease the price that these consumers pay. I posit that consumers make a higher proportion of their purchases without search in low demand periods for two reasons. First, the average purchase size is smaller, and so the expected savings resulting from finding a lower price are lower. Second, because the depth of discount is smaller in low demand periods, there is less reason to search for lower prices, which amplifies the first effect.

Several other explanations of this phenomenon have been put forward in the literature. Counter-cyclic pricing was initially observed by Warner and Barsky (1995) in retail stores and by Chevalier et al. (2003) in grocery stores.

Warner and Barsky (1995) propose that counter-cyclic pricing is caused by increased inter-store competition. Chevalier et al. (2003) find that these pricing trends are consistent with a loss-leader strategy, where the retailer lowers the price of goods to attract consumers to their store. In a re-analysis of their data, Nevo and Hatzitaskos (2006) find that underlying counter-cyclic pricing is a seasonal trend in the price sensitivity of consumers. They argue that prices are higher during low demand periods because consumers become less price sensitive. This paper complements their findings. Where Nevo and Hatzitaskos (2006) discover the trend in price sensitivity, this paper identifies an explanation of this trend.

Consumers strategically purchasing without observing prices has important implications beyond counter-cyclic pricing. Retailers can change the frequency of strategic consumer search by altering their price strategy. For example, strategic consumers will respond to a reduction in price variation by searching less frequently. Retailers benefit because strategic consumers would then make their purchase decisions without searching more often, making them less price sensitive, and allowing the retailer to charge higher prices. Consumers benefit because they avoid undergoing costly search. Studying counter-cyclic pricing allows me to demonstrate strategic consumer search, which practioners can account for when designing promotional strategies.

First, I find in my data that the concentrated soup industry exhibits counter-cyclic pricing. Second, I replicate the findings of Nevo and Hatzitaskos (2006) by showing that there is a significant seasonal trend in price sensitivity. Third, I find evidence that consumers sometimes purchase without search by looking at how price sensitivity is affected by factors which have been shown to impact search probability. Fourth, I find that the seasonal trend in price sensitivity is positively correlated with the seasonal trends in both overall volume sold and price variation. Both of these factors serve as proxies for search likelihood if consumers are strategically searching.

A dynamic consumer inventory model is necessary to study this phenomenon for two reasons: consumer stockpiling, and future expectations. First, many seasonal products, including canned soup, are storable. Consumers may take advantage of temporary price reductions by "stocking up" for future consumption. In this case, a static model would overestimate price sensitivity because it would misinterpret intertemporal substitution for an overall increase in demand. Previous studies have found that static models may overestimate price elasticities by as much as 30% Nevo and Hatzitaskos (2006). Second, rational search behaviour will depend on future expectations of price variation and consumption utility. Each of these expectations affect the expected purchase size, which in turn affects the benefits of finding a lower price. A dynamic consumer inventory model can structurally account for these factors.

Estimating this model using standard dynamic methods is a computational challenge because of the expansion of the state space required to accommodate seasonal shifts in consumption utility, price expectations, and search probabilities. The inclusion of these trends in the model necessitates that I solve for expected discounted payoffs separately for each seasonal period, which increases the size of the state space by a factor of 52, one for each week of the year. Increasing the size of the state space increases the burden of a value function iteration quadratically due to an increase in the number of transition probabilities that need to be calculated, and an increase in the number of states that must be updated. So, adding the seasonal period to the state space increases the computational burden of a value function iteration by a factor of 52^2.

To overcome this challenge, I develop the Cyclic Successive Approximation Algorithm (CSAA), which is an adjustment to the Successive Approximation Algorithm (SAA) of Rust (1987), that removes the computational burden of adding cyclic variables to the state space of a dynamic model. This is done by updating the states in the reverse order of the corresponding seasonal periods. I show that these algorithms converge to the true value function at the same rate, and in each iteration the results of the two algorithms will be identical for certain states. However, the cost of a CSAA iteration is invariant to the number of seasonal periods. In this problem, the technique reduces the theoretical burden of solving the value function by a factor of 52^2, and in practice reduces the computational time required by a factor of 52.

This paper fits a dynamic consumer inventory model, following Erdem et al. (2003), and Hendel and Nevo (2006). By assuming consumers do not observe prices in each period, I follow the price consideration model of Ching et al. (2009), and the dynamic implementation of this model presented in Seiler (2013). However, in those models, consumers must observe prices before making a purchase. In allowing consumers to selectively ignore price information, I follow the econometric framework of Mehta et al. (2003), and the descriptive evidence of Ray et al. (2012).

In my model, consumers respond to an increase in price variation by searching more frequently. Mela et al. (1998) show that, over time, consumers change their shopping strategies in response to high levels of price variation through stockpiling. Consumers can choose whether to search, which allows them to react to prices, or to purchase without search, in which case they are not price sensitive. In this respect, the estimated model can be considered a structural, dynamic implementation of Bucklin and Lattin (1991).

The rest of the paper is organized as follows: in Section 2, I describe the data set and report summary statistics; in Section 3, I present descriptive evidence that is consistent with the existence of counter-cyclic pricing and seasonal trends in purchase without search; in Section 4, I detail the dynamic, structural inventory model of the concentrated soup industry that allows consumers to purchase without search; in Section 5, I outline the estimation procedure; in Section 6, I introduce the cyclic successive approximation algorithm, which reduces the computational burden of estimating dynamic models with cyclic state variables, to make the computational burden manageable; in Section 7, I present the results of the estimation; and in Section 8, I conclude.

2 Data

This project used the panel data in the “IRI Marketing Data Set” (Bronnenberg, Kruger, and Mela 2008). Panel data on consumers in Eau Claire, Wisconsin and Pittsfield, Massachusetts is reported for 30 categories over the seven years between January 1st 2001 to December 31st 2006.

I focus on the purchases of concentrated soup to study counter-cyclic pricing trends for four reasons. First, the concentrated soup category clearly exhibits strong seasonality, with purchase volume rising dramatically in the winter. Second, the category exhibits counter-cyclic pricing trends. Third, in the data I analyze, concentrated soup is monopolistic, with Campbell’s having more than a 97% market share. This simplifies the pricing problem the retailer and manufacturer face. Fourth, concentrated soup is purchased almost exclusively in 10.75 oz. cans, which are typically consumed in one sitting. This limits consumer inventories to a discrete number of cans, which simplifies the construction of the dynamic model.

The dataset initially has 6,535 panelists. I focus the analysis on panelists in that make at least one concentrated soup purchases at a single store in Pittsfield, Massachusetts, which reduces the sample to 1,441 panelists. I focus on the purchases made at a single store to ensure that the price distribution is constructed accurately. Accurately estimating price distribution is important because consumers partially base their search decision on the price distribution. However, I use all purchase data when constructing consumer inventories to ensure that my sample limitation does not bias the estimated inventories.

The panel is unbalanced: participants participate for an average of 1.91 years. The data was collected in two ways: For 89% of panelists, purchases were recorded electronically when the panelist used a loyalty card at check out. In addition to using the loyalty card, 2% of panelists scanned their own purchases using a "key". 9% of panelists switched from using a key to using a loyalty card over the course of the sample.

Concentrated soup is sold in 29 varieties in the panel data. The market shares of each variety are presented in Table 1. To ensure each variety I analyze has an adequate sample size, I collect the 21 least popular varieties into a residual category, termed "Missing", which represents 4% of observed sales.

I identify the prices of each flavor in each week by looking at the cost and units purchased by panelists. The week-to-week price of each flavor changes in 24% of weeks. Flavors tend not to go on sale at the same time: in more than 50% of weeks, 3 or fewer flavors are discounted. The non-promoted price for soup only changes 4 times over the 7 years of observed data.

3 Model Free Evidence

In this section, I provide descriptive evidence which suggests that counter-cyclic pricing is caused by seasonality in purchase without search. This serves

Variety	Market Share	Average Price	Average Discount
Cream Of Mushroom	0.25	1.13	0.12
Cream Of Chicken	0.20	1.18	0.19
Chicken Noodle	0.13	1.11	0.16
Mega Noodle	0.10	1.16	0.10
Chicken Goldfish	0.09	1.15	0.11
Cream Of Celery	0.07	1.22	0.09
Missing	0.04	1.15	0.11
Tomato	0.04	0.99	0.44
Tomato Goldfish	0.02	1.08	0.11

Table 1: Flavor Summary Statistics

to motivate the construction of the dynamic model, while highlighting which trends in the data allow me to identify various features of the model.

3.1 Seasonality and Counter Cyclic Pricing

Counter-cyclic pricing denotes the simultaneous presence of three patterns in the category sales data: seasonal demand trends, seasonal pricing trends, and a negative correlation between these two trends. One might expect soup to be consumed more often during cold weather, or to help soothe a sore throat, both of which are more common in the winter months. This trend is typified by the sales data presented in Figure 1, where the demand for soup is highest during the winter and lowest during the summer. To test whether this trend is statistically significant, I regressed average weekly sales onto a 4-degree polynomial based on the time of year (Table 2, Column 1). The trend is found to be statistically significant, and explains 27.0% of the variation in soup sales, compared to the 6.7% that is explained exclusively by price fluctuations (Table 2, Column 2).

Controlling for price variation is required to accurately estimate underlying seasonal demand. Seasonal pricing trends may amplify or cause the observed seasonal demand trends because increased demand is a natural consequence of reduced prices. The seasonality in demand persists when I control for price fluctuations (Table 2, Column 3). Comparing Column 2 and Column 3 in table 2 highlights the dangers of omitting seasonality from the analysis of this category. The estimated price coefficient is almost halved when I include seasonal trends in the underlying demand for concentrated soup, which suggests that the omission of seasonal trends could lead to biased estimates of price elasticities.

Seasonal pricing trends in the category can be seen in Figure 2, which plots the average discount in each soup flavor over the course of the year. In contrast to quantity demanded, the average discount is lowest during the summer months and highest during the winter months.

3.2 Seasonality in Estimated Price Sensitivity

Nevo and Hatzitaskos (2006) found that counter-cyclic pricing is induced by a simultaneous trend in consumer price sensitivity. They observe a reduction in price sensitivity in low demand periods. They argue that counter-cyclic pricing is a retailer's reaction to this seasonal change in price sensitivity. That is, because consumers are less price sensitive in low demand periods, retailers charge higher prices.

Table 2: Descriptive Regression Results

	Average Weekly Units Sold		
	(1)	(2)	(3)
Week In Year - Linear	3.411^{**}		3.174^{**}
	(1.392)		(1.372)
Week In Year - Quadratic	-0.479^{***}		-0.448^{***}
	(0.105)		(0.104)
Week In Year - Cubic	0.017^{***}		0.016^{***}
	(0.003)		(0.003)
Week In Year - Quartic	-0.0002^{***}		-0.0002^{***}
	(0.00003)		(0.00003)
Average Price		-38.277^{***}	-23.856^{***}
		(7.546)	(6.759)
Constant	25.097^{***}	71.318^{***}	53.032^{***}
	(5.438)	(8.934)	(9.555)
R^2	0.270	0.067	0.295

I find evidence of this trend in the IRI data by using a logit model to predict consumers' flavor decisions over the course of the year. Hendel and Nevo (2006) show that if a shopper's consumption does not depend on the flavor purchased then the flavor decision in a dynamic inventory model simplifies to a logit model. I can use this model to identify flavor preferences and estimate price sensitivity over the course of the year.

Suppose that the utility gained for purchasing a product of flavor f_t is given by:

$$(\alpha_{s_t} p_{ft} + \eta_f + \varepsilon_{ft})$$

where α_{s_t} is the price sensitivity in seasonal period s_t, p_{vt} is the price of flavor v at time t, η_v is the flavor dummy, and ε_{vt} is an IID shock with a type-1 extreme value distribution. Then, assuming consumers are utility maximizing and comparing the log-market shares of each flavor in each period, we have

$$\log(s_{v_1 t}) - \log(s_{v_2 t}) = \alpha_{s_t}(p_{v_1 t} - p_{v_2 t}) + (\eta_{v_1} - \eta_{v_2}).$$

I approximate the seasonal price coefficient α_{s_t} with a 3 degree polynomial based on time of year. Consistent with Nevo and Hatzitaskos (2006), I find seasonal trends in price sensitivity are statistically significant (Figure 4)(Table 3, Column 2).

3.3 Purchase Without Search

Purchase without search cannot be directly identified because I do not observe whether shoppers search in any given period. However, there are two ways to check for purchase without search in the scanner data. First, Seiler (2013) finds that shoppers are more likely to search when they make large shopping trips. This is because searching the category takes time, and on average shoppers making large trips have more time to spend at the grocery store. If shoppers sometimes purchase without searching the category, and those who do cannot react to price promotions, then there would be a correlation between price sensitivity and the overall size of the shopping trip.

To test this, I classify a shopping trip as a "big trip" if it is larger than each consumers' median shopping trip across all other categories. By interacting prices with this dummy variable, I find that consumers are significantly more price sensitive on big shopping trips, suggesting that consumers are sometimes purchasing without checking prices (Table 3, Regression 3).

Second, consumers are more likely to purchase a large quantity of soup when they've found a good discount. If they found a good discount, then they must have checked prices. This leads to a correlation between search likelihood and quantity purchased. Since search likelihood is correlated with both effective price sensitivity, and quantity purchased, if consumers sometimes purchase without search there should be a relationship betweeen quantity purchased and effective price sensitivity. This correlation is observed in Table 3, Regression 4.

	Difference in Log Market Shares			
	(1)	(2)	(3)	(4)
Price Sensitivity - Constant	−1.024***	−1.219***	−0.600	−0.914***
	(0.060)	(0.264)	(0.390)	(0.240)
Price Sensitivity - Linear		0.082*	0.087**	0.120***
		(0.044)	(0.043)	(0.041)
Price Sensitivity - Quadratic		−0.003*	−0.005***	−0.006***
		(0.002)	(0.002)	(0.002)
Price Sensitivity - Cubic		0.00004*	0.0001***	0.0001***
		(0.00002)	(0.00002)	(0.00002)
Price * Big Shopping Trip			−0.664***	
			(0.304)	
Price * Multiple Cans				−0.893***
				(0.108)

*Flavor Dummies Not Shown

Table 3: Reduced Form Regresson Results

Ideally, I could demonstrate the importance of purchase without search by estimating a correlation between price sensitivity and search probability. I cannot do this directly because I do not observe whether a consumer searches in each period. However, if consumers are strategically searching, then they are more likely to observe prices when there are larger incentives to do so. Consumers would then be more likely to search when they have a high demand for soup, because the expected savings from finding a lower price are higher. Consumers would also be more likely to search when there is high price variation, because they want to take advantage of the larger sales. Hence, both underlying demand and price variation can serve as proxies for search probability if consumers are strategically searching.

I find that the estimated seasonality in price sensitivity (Table 3, Column 2) is highly correlated with the reduced form estimate of underlying demand from Table 2, Column 3 ($cor = -.796$, $p < 0.0001$), and with average weekly discount ($cor = .25$, $p < 0.0001$). If purchase volume and price variation are accurate proxies for search probability, then this suggests that there is a relationship between the seasonal trends in search probability and price sensitivity.

4 Model

I model purchases in the concentrated soup industry using a dynamic, structural, inventory model. The model is dynamic because shoppers make their decisions while being cognizant of future price expectations, consumption utilities, and the implications of their decisions on their future inventories.

In each week, consumers make decisions in three sequential stages: search, purchase, and consumption. In the search stage, consumers decide whether to search the category before making their purchase decision. If a shopper searches, they incur a search cost, but observe the prices of all varieties. If a consumer does not search, they instead make their purchase decision on the assumption that there are no price promotions. When making this decision, consumers take into account the distribution of prices, and the amount of soup they want to purchase.

In the purchase stage, consumers simultaneously decide on the variety and number of cans to purchase. The prices they use to make this decision depend on their search decision. When making this decision, consumers take into account the amount of soup in their inventory, and how much they want to consume soup in this seasonal period.

Finally, in the consumption stage, consumers choose the number of cans of soup to consume while taking into account the effect on their future inventory. Consumption utility varies by time of year because soup is a seasonal good. At the end of this stage, consumers incur a storage cost based on the number of cans of soup in their inventories. Consumers can only store up to i_{max} cans of soup, and will increase their consumption to ensure that their inventories are below i_{max}.

I outline the model in the following three steps. First, I define the utility that a consumer receives in each stage. Second, I define the expected discounted payoffs and Bellman equation for the problem, which are used to calculate the likelihood. Third, I formally define the choice that a consumer makes in each stage, solve for the probability of any particular choice, and calculate the expected discounted payoffs the consumer receives during the current substage, any remaining substages, and all subsequent periods.

4.1 Flow Utility

The utility a consumer receives in the search stage in period t is specified as

$$u_s(r_t; \eta_t) = \mathbb{1}(r_t = 1)(-\rho + \eta_t) \tag{1}$$

where r_t is an indicator that equals 1 if the consumer searches, ρ is the search cost, η_t has an IID type-1 extreme value distribution with standard deviation σ_η. If consumers choose to search, then they incur the search cost ρ.

The utility a consumer receives in the purchase stage depends on their search decision. If the consumer searches, then they observe prices and receive utility

$$u_p(q_t, f_t; r_t = 1, p_t, \varepsilon_t) = \alpha p_{f_t t} q_t + \eta_{f_t} + \varepsilon_{f_t q_t t}. \tag{2}$$

where q_t is the number of cans of soup purchased, f_t is the flavor of soup purchased, p_t is a vector of prices, ε_t is a vector of IID shocks to the purchase utility, $p_{f_t t}$ is the price of flavor f_t at time t, η_{f_t} is a flavor fixed effect, and $\varepsilon_{f_t q_t t}$ is an IID shock to the utility of purchasing q_t cans of flavor f_t that has a type-1 extreme value distribution. I assume that consumers only purchase one flavor in each period.

If the consumer does not search, they make their purchase decision on the assumption that there are no price discounts. In this case, the utility they receive is

$$u_p(q_t, f_t; r_t = 0, p_t, \varepsilon_t) = \alpha \widehat{p_{f_t t}} q_t + \eta_{f_t} + \varepsilon_{f_t q_t t}$$

where $\widehat{p_{f_t t}}$ is the undiscounted price of flavor f_t in period t. I assume that consumers only purchase one flavor in each period.

In each period, a consumer gains the following utility from consumption

$$u_c(c_t, q_t; v_t, s_t, i_t) = c_t \times k(s_t) + v_{c_t t} \tag{3}$$

where c_t is the quantity of soup consumed, v_t is a vector of IID shocks to consumption, s_t is the current seasonal period, i_t is the starting inventory in period t and $k(s_t)$ is the seasonally varying returns to consumption. The random shocks $v_{c_t t}$ ensures the model can accommodate varying levels of consumption, and allows for a simple computation of consumption probabilities. Consumers cannot consume more than $c_{\max}$ in any period. Seasonal variation in consumption is modelled through variations in $k(s_t)$ throughout the year.

Note that consumption utility only depends on the number of cans consumed, and not the flavor of those cans. Instead, flavor preferences are modelled in the purchase decision with a fixed effect. This assumption, initially in Hendel and Nevo (2006), will allow me to separate the flavor decision from dynamic considerations because the flavor decision now only affects the static utility in the purchase stage. This reduces the size of the state space because I only need to track the total number of cans of soup in inventory, rather than the number of cans of each flavor.

4.2 State Variables and Value Functions

The seasonal period is a cyclic state variable which updates deterministically as follows:

$$\begin{aligned} s_{t+1} &= s_t + 1 \text{ if } s_t < |S| \\ &= 1 \text{ if } s_t = S \end{aligned}$$

where $|S|$ is the total number of seasonal periods. This deterministic updating allows me to apply the cyclic successive approximation algorithm, developed in Section 6, which removes the computational burden of adding this variable to the state space.

Inventory levels are increased through purchase, and decreased through consumption:

$$i_{t+1} = i_t + x_t - c_t$$

Any combination of purchase and consumption decisions that would result in consumers having an inventory larger than $i_{\max}$ or smaller than 0 are forbidden.

Because consumers are forward looking, they consider how their decisions impact their expected discount future utility. Formally, consumers seek to maximize their expected discounted utility in each period. Let Ω_t be a vector of all of the transient state variables $\Omega_t = (\eta_t, p_t, \varepsilon_t, v_t)$, and let a_t be the vector of the actions a consumer can take $a_t = (r_t, f_t, q_t, c_t)$. I define the total utility that a consumer receives in period t as $u(a_t, s_t, i_t, \Omega_t)$. Note that consumers do not observe all the state variables simultaneously; they are revealed as the consumer progresses through the three stages. Consumers make their decisions to maximize their total expected discounted payoff

$$V(s_t, i_t, \Omega_t) = \max_{\{a_t, a_{t+1}, \ldots\}} \sum_{t=\tau}^{\infty} \delta^{\tau-t} E(u(a_\tau, s_\tau, i_\tau, \Omega_\tau) | s_t, i_t, \Omega_t) \tag{4}$$

where δ is the discount factor. This can be conveniently expressed in the following Bellman equation:

$$V(s_t, i_t, \Omega_t) = \max_{a_t} u(a_t, s_t, i_t, \Omega_t) + \delta E(V(s_{t+1}, i_{t+1}, \Omega_{t+1}) | s_t, i_t, \Omega_t) \tag{5}$$

I define the value function as the expected discounted payoff, given starting persistant states s_t and i_t

$$V(s_t, i_t) = E(\max_{a_t} u(a_\tau, s_\tau, i_\tau, \Omega_\tau) + \delta E(V(s_{t+1}, i_{t+1}, \Omega_{t+1}) | s_t, i_t, \Omega_t)) \tag{6}$$

This integrated Bellman equation will allow me to solve for the value function, which is required to estimate the model. I now formally state the consumer decisions and values in each stage.

4.3 Decision at the Consumption Stage

To solve the model, I work backwards through the stages in each period. To make the consumption decision, consumers take the purchase and search decisions as given, and then optimize the sum of their consumption utility, inventory costs, and expected future values:

$$\max_{c \in \{\max(0, i_{\max} - i_t - q_t), i_t + q_t)\}} c_t \times k(s_t) + v_{ct} + \delta E(V(s_{t+1}, i_t + q_t - c_t) | s_t, i_t).$$

Note that consumers must end the period with an inventory of at least $i_{\max}$. If after the purchase stage they have more than $i_{\max}$ cans of soup, then they will at least consume enough to bring their final inventory down to $i_{\max}$.

Because v_{ct} has an IID type-1 extreme value distribution, this leads to the following consumption probabilities:

$$P(c_t | x_t, s_t, i_t) = \frac{e^{c_t \times k(s_t) + \delta E(V(s_{t+1}, i_t + q_t - c_t) | s_t, i_t)}}{\sum_{c' = \max(0, i_{\max} - i - x)}^{\min(c_{\max}, i_t + q_t)} e^{c' \times k(s_t) + \delta E(V(s_{t+1}, i_t + q_t - c') | s_t, i_t)}}$$

To solve for the decision in the purchase stage, I need to combine values of the consumption stage, and the expected discounted payoffs in all future periods, which I define as $V_c(s_t, i_t, q_t)$:

$$V_c(s_t, i_t, q_t) = \log\left(\sum_{c'=\max(0, i_{\max}-i_t-q_t)}^{\min(c_{\max}, i_t+q_t)} e^{c'\times k(s_t)+\delta E(V(s_{t+1}, i_t+q_t-c')|s_t, i_t)}\right).$$

4.4 Decision at the Purchase Stage

During the purchase stage, consumers choose the quantity and flavor of concentrated soup to purchase by optimizing the sum of their purchase utility and their expected discounted utility in the search stage and in future periods. If the consumer searches, then they observe prices, and so their decision solves:

$$\arg\max_{q_t, f_t} \alpha p_{f_t t} q_t + \eta_{f_t} + \varepsilon_{f_t q_t t} + V_c(s_t, i_t, q_t). \tag{7}$$

If the consumer does not search, then they use the non-discounted price $\widehat{p_{fqt}}$:

$$\arg\max_{q_t, f_t} \alpha \widehat{p_{fqt}} q_t + \eta_{f_t} + \varepsilon_{f_t q_t t} + V_c(s_t, i_t, q_t). \tag{8}$$

The flavor decision only affects the purchase utility, and is independent of the payoffs in future periods. Calculating the probability of picking each flavor:

$$P(f_t|q_t, p_t, r_t = 1, s_t, i_t) = \frac{\exp(\alpha p_{f_t t} q_t + \eta_{f_t})}{\sum_{f' \in F} \exp(\alpha p_{f' t} q_t + \eta_{f'_t})}$$

$$P(f_t|q_t, p_t, r_t = 0, s_t, i_t) = \frac{\exp(\alpha \widehat{p_{fqt}} q_t + \eta_{f_t})}{\sum_{f' \in F} \exp(\alpha \widehat{p_{fqt}} q_t + \eta_{f'_t})}$$

where F is the full set of flavors. Integrating over the flavor decision reduces the decision in the purchase stage to a quantity decision. In the case where consumers search:

$$\arg\max_{q_t} \log\left(\sum_{f' \in F} e^{\alpha p_{f'_t} q_t + \eta_{f'_t}}\right) + V_c(s_t, i_t, q_t) + \varepsilon_{q_t t},$$

and in the case where consumers do not search:

$$\arg\max_{q_t} \log\left(\sum_{f' \in F} e^{\alpha \widehat{p_{fqt}} q_t + \eta_{f'_t}}\right) + V_c(s_t, i_t, q_t) + \varepsilon_{q_t t}$$

where $\varepsilon_{q_t t}$ has an IID type-1 extreme value distribution. Integrating over $\varepsilon_{q_t t}$, I can calculate the probability of the consumer choosing any purchase size as:

$$P(q_t | r_t = 1, p_t, s_t, i_t) = \frac{e^{\log\left(\sum_{f \in F} e^{\alpha p_{f_t} q + \eta_f}\right) + V_c(s_t, i_t, q_t)}}{\sum_{q'=0}^{i_{\max} + c_{\max} - i_t} e^{\log\left(\sum_{f \in F} e^{\alpha p_{f q' t} q' + \eta_f}\right) + V_c(s_t, i_t, q')}}.$$

In the case where the consumer searches, the decision is made according to the undiscounted price, $\widehat{p_{fqt}}$

$$P(q_t | r_t = 0, s_t, i_t) = \frac{e^{\log\left(\sum_{f \in F} e^{\alpha \widehat{p_{fqt}} q + \eta_f + \xi_q}\right) + V_c(s_t, i_t, q_t)}}{\sum_{q'=0}^{i_{\max} + c_{\max} - i_t} e^{\log\left(\sum_{f \in F} e^{\alpha \widehat{p_{fqt}} q' + \eta_f + \xi_{q'}}\right) + V_c(s_t, i_t, q')}}.$$

To solve for the decision in the search stage, I need to combine values of the purchase stage, consumption stage, and the expected discounted payoffs in all future periods, which I define as $V_p(s_t, i_t, q_t)$:

$$V_p(s_t, i_t, r_t = 1) = E_{p_t}\left(\log\left(\sum_{q'=0}^{i_{\max}+c_{\max}-i_t} e^{\log\left(\sum_{f\in F} e^{\alpha p_{f_t t} q' + \eta_f}\right) + V_c(s_t, i_t, q')}\right)\right)$$

$$V_p(s_t, i_t, r_t = 0) = \log\left(\sum_{q'=0}^{i_{\max}+c_{\max}-i_t} e^{\log\left(\sum_{f\in F} e^{\alpha \widehat{p_{fqt}} q' + \eta_f}\right) + V_c(s_t, i_t, q')}\right)$$

4.5 Decision at the Search State

When making the search decision, consumers compare the expected discounted value of searching with the expected discounted value of making their purchase decision without search. Because the search decision is made before prices are observed, the expectation is taken over prices to calculate the expected benefit of search. The search decision is then

$$\max_{r_t \in \{0,1\}} \mathbb{1}(r_t = 1)(V_p(s_t, i_t, r_t = 1) - \rho + \eta_t) + \mathbb{1}(r_t = 0)V_p(s_t, i_t, r_t = 0)$$

which leads to the following search probability:

$$P(r_t = 1|s, i_t) = \frac{e^{\frac{V_p(s_t, i_t, r_t=1)-\rho}{\sigma\eta}}}{e^{\frac{V_p(s_t, i_t, r_t=1)-\rho}{\sigma\eta}} + e^{\frac{V_p(s_t, i_t, r_t=0)}{\sigma\eta}}}.$$

This gives the overall expected value for the search stage, the purchase stage, and the consumption stage $V_s(s_t, i_t)$:

$$V(s_t, i_t) = \sigma_\eta \times E_p(\log(e^{\frac{V_p(s_t, i_t, r_t=1)-\rho}{\sigma_\eta}} + e^{\frac{V_p(s_t, i_t, r_t=0)}{\sigma_\eta}})).$$

After calculating the value function $V(s_t, i_t)$ using the Bellman equation 6, I use $V(s_t, i_t)$ to solve for $V_p(s_t, i_t, q_t)$, and $V_c(s_t, i_t, q_t)$. With these values in hand, I calculate the joint probability

$$P(r_t, f_t, q_t, c_t | s_t, i_t) = P(r_t | s_t, i_t) P(f_t | r_t, s_t, i_t) P(q_t | f_t, r_t, s_t, i_t) P(c_t | q_t, s_t, i_t).$$

Each of these terms can be calculated using the previous equations.

5 Estimation

5.1 Identification

I provide an informal discussion of the identification of the parameters of the model. The flavor dummies η_{f_t} are identified by comparing the sales of different flavors when they have the same price, as discussed in the next section. The seasonal consumption preferences $k(s_t)$ are identified by observing how the quantity of canned soup purchased varies throughout the year. The search parameters ρ and σ_η are identified by how trends in seasonal price sensitivity are correlated to changes in consumption and price variation. The overall price sensitivity α is identified by the average level of price sensitivity.

6 Cyclic Successive Approximation Algorithm

The addition of seasonal periods to the state space drastically increases the computational burden of solving for the value function, as the cost of a value function iteration grows quadratically with the size of the state space. The addition of 52 seasonal periods increases the overall burden by a factor of 2702. In real terms, this could bring the time taken to solve the value function from one minute to over a day and a half. During estimation, the value function must be solved many times, and so this additional computation cost might make this problem too burdensome to solve.

To overcome this cost, I develop the Cyclic Successive Approximation Algorithm (CSAA), which is a variant of the Successive Approximation Algorithm (SAA) of Rust (1987). The cyclic successive approximation algorithm eliminates the computational burden of adding cyclic variables to the state space when solving for the value function in any dynamic model. I show that these algorithms converge to the true value function at the same rate, and in each iteration the results of the two algorithms will be identical for certain states. Furthermore, the CSAA is simple to implement: states are updated in the same way as in the SAA, just in a different order.

The additional computational burden of evaluating value function iteration with a larger state space comes from two sources. Suppose the state space increases by a factor of N_s. First, the number of states one can transition to increases by a factor of N_s. For each destination state, a transition probability must be calculated. Second, the number of states that the value function must be solved for increases by a factor of N_s.

In the CSAA, I remove both of these increases in cost when adding a cyclic variable to the state space. First, I note that the transitions of cyclic state variables are deterministic, and so only a fraction of the transition probabilities need to be calculated. This is simple to show and has likely been incorporated in other implementations. Second, I show that one only needs to update a fraction of the state space in each iteration. The result is that the cost of a value function iteration does not grow with the state space.

The closest method to the CSAA is the cyclic-inversion algorithm of Paarsch and Rust (2009). They show that, when solving for the value function using policy function iterations, the matrix inversion can be decomposed so that the total computational cost of a policy function iteration grows linearly in the number of cyclical periods, rather than cubically. The CSAA has three advantages over the cyclic-inversion algorithm. First, the CSAA removes the burden of adding cyclic variables to the state space, while in the cyclic-inversion algorithm, additional seasonal state variables increase the computational burden linearly. Second, the CSAA applies to value function iterations rather than policy function iterations. Value function iterations have been found to be more efficient when the state space is large (Santos and Rust (2004)), which is when reducing the computational burden is most important. Third, the CSAA is arguably easier to implement as it is a small deviation from the standard methodology. On the other hand, because the cyclic-inversion algorithm uses policy function iterations, it will likely converge faster when discount rates approach 1.

In the SAA, a starting guess for the value function is chosen. Using the Bellman equation, value function iterations are used to update all states of the dynamic model. This process is repeated until the value function converges. In the CSAA, the process is the same, except that in each iteration only the states associated with a single seasonal period are updated. The seasonal period that is updated is the one that precedes the seasonal period that was updated in the previous iteration. That is, if there are $|S|$ seasonal periods, $S = \{1, \ldots, |S|\}$, and in the previous iteration the values in seasonal period s was updated, then, if $s > 1$, the current iteration updates the values in seasonal period $s - 1$. If $s = 1$, then seasonal period S is updated.

I formally define the CSAA algorithm below. Suppose there is a rational, forward-looking consumer making decisions based on their current state. Let S be a vector of seasonal state variables which are arranged in chronological order ranging from 1 to $|S|$. Let I be a vector of the remaining state variables. Let A be a vector of actions available to the agent in each period. Let $u(a, s, i)$ be the utility gained from taking action a_t in state (s, i). Let δ be the agents discount factor. Let $V(s, i)$ be the value function, which is the expected discounted payoff to the agent given that the agent starts in state

(s, i). I can use the value function to write the Bellman equation:

$$V(s_t, i_t) = E(\max_{a_t} u(a_t, \{s_t, i_t\}) + \delta \sum_{s' \in S, i' \in I} P(s_{t+1} = s', i_{t+1} = i' | a_t, s_t, i_t) V(s_{t+1}, i_{t+1})) \tag{9}$$

Integrating over the action chosen a_t,

$$V(s_t, i_t) = \sum_{a' \in A} P(a_t = a') E\left(u(a', \{s_t, i_t\}) + \delta \sum_{s' \in S, i' \in I} P(s_{t+1} = s', i_{t+1} = i | a', s_t, i_t) V(s', i') | a_t = a' \right).$$

Because the seasonal period updates deterministically, s_{t+1} is known, so the above expression can be simplified:

$$V(s_t, i_t) = \sum_{a' \in A} P(a_t = a') E\left(u(a', s_t, i_t) + \delta \sum_{i' \in I} P(i_{t+1} = i | a', s_t, i_t) V(s_{t+1}, i') | a_t = a' \right). \tag{10}$$

where $s_{t+1} = (s_t \bmod |S|) + 1$. This simplification reduces the number of $P(s_{t+1} = s', i_{t+1} = i | a', s_t, i_t)$, the state transitions probabilities, from $|S| \times |I|$, to $|I|$. The "mod" operation ensures that the seasonal period updates cyclically, and is defined as $s_t \bmod |S| = s_t - \left\lfloor \frac{s_t}{|S|} \right\rfloor \times |S|$. To clarify, this operator takes the following values:

$$\begin{aligned} (s_t \bmod |S|) + 1 &= s_t + 1 \text{ if } 0 < s_t < |S| \\ (s_t \bmod |S|) + 1 &= 1 \text{ if } s_t = |S| \end{aligned}$$

Let V be an $|S| \times |I|$ matrix of values, where the element in row s_t and column

i_t corresponds to the value of being in state $\{s_t, i_t\}$. Let $V_{s,i}$ represent row s and column i of matrix V. Let $T_\theta(V)$ be the standard Bellman operator on this matrix representation of the value functions, which outputs a V such that, for each row s and column i,

$$T(V)_{s_t,i_t} = \sum_{a' \in A} P(a_t = a')E\left(u(a', s_t, i_t) + \delta \sum_{i' \in I} P(i_{t+1} = i'|a', s_t, i_t)V_{s_{t+1},i'}|a_t = a'\right).$$

This Bellman operator applies the Bellman equation in equation 10 to the values V. This equation is a contraction mapping, and so the repeated application of equation 7 will yield a sequence of V that approaches the true values.

Let U^f be an operator on V defined as

$$\begin{aligned} U^f(V)_{s_t,i_t} &= \sum_{a' \in A} P(a_t = a')E\left(u(a', s_t, i_t) + \delta \sum_{i_{t+1} \in I} P(i_{t+1} = i|a', s_t, i_{t+1})V_{s_{t+1},i_{t+1}}|a_t = a'\right) \text{ if } s = f \\ &= V_{s,i}\text{otherwise} \end{aligned}$$

This operator is identical to $T(V)_{s,i}$ for the states where $s = f$. In other cases, the operator simply returns the input state. Note that, because U^f only operates on $\frac{1}{|S|}$ of the state space, each application of T is $|S|$ times more costly than an application of U^f.

The SAA algorithm is defined as

1. Make an initial guess at the value function V^0

2. Update the current guess V^k by applying the operator $V^{k+1}{=}T(V^k)$

3. Repeat step 2 until desired convergence

The CSAA algorithm is defined as

1. Make an initial guess at the value function $\tilde{V}^0$

2. Update the current guess $\tilde{V}^k$ by applying the operator $\tilde{V}^{k+1}{=}U^{(|S|-k \text{ mod } |S|)+1}(\tilde{V}^k)$

3. Repeat step 2 until desired convergence

Given the same starting guess, the guesses of the value function in each algorithm will be identical in certain states after k iterations of SAA and CSAA,. Specifically, the guesses of the values will be identical in the states most recently updated in the CSAA, which are states where the seasonal period is equal to $(S - k \text{ mod } S) + 1$.

Theorem 6.1. *Let V^k be the guess of the value function in the kth iteration of the SAA algorithm, and let $\tilde{V}^k$ be the guess of the value function in the kth iteration of the CSAA algorithm. Then, if $V^0_{0,i_t} = \tilde{V}^0_{0,i_t}$ for all $i_t \in I$ and $k \in \mathbb{N}$:*

$$V^k_{(|S|-k\ mod\ |S|)+1,i_t} = \tilde{V}^k_{(|S|-k mod\ |S|)+1,i_t}$$

Proof. By induction. Suppose $V^0_{0,i_t} = \tilde{V}^0_{0,i_t} \forall i_t \in I$. I first prove the induction hypothesis is true for $k = 1$, that is, that $\widetilde{V}^1_{s,i_t} = V^1_{s,i_t}fors = (|S| - 1\text{mod } |S|)+$

$1 = |S|$.

$$\begin{aligned}
\widetilde{V}^1_{|S|,i_t} &= U_\theta^{(|S|-1 \text{ mod } S)+1}(\tilde{V}^0)_{1,i_t} \\
&= \sum_{a' \in A} P(a_t = a')E\left(u(a',|S|,i_t) + \delta \sum_{i_{t+1}\in I} P(i_{t+1} = i|a',|S|,i_{t+1})\tilde{V}^0_{1,i_{t+1}}|a_t = a'\right) \\
&= \sum_{a' \in A} P(a_t = a')E\left(u(a',|S|,i_t) + \delta \sum_{i_{t+1}\in I} P(i_{t+1} = i|a',|S|,i_{t+1})V^0_{1,i_{t+1}}|a_t = a'\right) \\
&\quad \text{because the initial guesses are the same and so } \tilde{V}^0_{1,i_{t+1}} = V^0_{1,i_{t+1}} \\
&= T(V^0)_{|S|,i_t} \\
&= V^1_{|S|-1,i_t} \text{ as desired.}
\end{aligned}$$

Now suppose that the induction hypothesis is true for iteration k, that is, $\tilde{V}^k_{(|S|-k \text{ mod } |S|)+1,i_t} = V^k_{(|S|-k\text{mod } |S|)+1,i_t}$. I now show that this implies the induction hypothesis is true for $k+1$, that is, $\tilde{V}^{k+1}_{(|S|-k-1\text{mod } |S|)+1,i_t} = V^{k+1}_{(|S|-k-1 \text{ mod } |S|)+1,i_t}$. For simplicity, let $s_j = (|S| - k - 1 \text{ mod } |S|) + 1$, and $s_{j+1} = (|S| - k \text{ mod } |S|) + 1$.

$$
\begin{aligned}
\tilde{V}^{k+1}_{s_j,i} &= U^{s_j}_{\theta}(\tilde{V}^k)_{s_j,i} \\
&= \sum_{a' \in A} P(a_t = a')E\left(u(a', s_j, i_t) + \delta \sum_{i_{t+1} \in I} P(i_{t+1} = i | a', s_j, i_{t+1}) \tilde{V}^k_{s_{j+1}, i_{t+1}} | a_t = a' \right) \\
&\quad \text{because the operator updates seasonal period } s_j \\
&= \sum_{a' \in A} P(a_t = a')E\left(u(a', s_j, i_t) + \delta \sum_{i_{t+1} \in I} P(i_{t+1} = i | a', s_j, i_{t+1}) V^k_{s_{j+1}, i_{t+1}} | a_t = a' \right) \\
&\quad \text{by the induction hypothesis} \tilde{V}_{s_{j+1}, i_t} = V^k_{s_{j+1}, i_t} \\
&= T(V^k)_{s_j, i_t} \\
&= V^{k+1}_{s_{j+1}, i}
\end{aligned}
$$

As desired. □

Corollary 6.2. *Let $V^k_{s,i}$ be the guess of the value function for states $\{s, i\}$ in the kth iteration of SAA algorithm, and let $\tilde{V}^k_{s,i}$ be the guess of the value function for states $\{s, i\}$ in the kth iteration of CSAA algorithm. If $V^0_{1,i} = \tilde{V}^0_{1,i} \forall i_t \in \{0, \ldots, I\}$, then $V^{n|S|}_{1,i_t} = \tilde{V}^{n|S|}_{1,i_t} \forall i_t \in I, n \in \mathbb{N}$.*

Proof. This follows directly from theorem 1 with $k = n$. □

Corollary 6.3. *Given $V_{1,i}$ for all $i_t \in I$, we can compute $V_{s_t,i}$ for each s_t by applying CSAA iterations $S - 1$ times.*

Proof. If $V_{1,i} \forall i$ is known, then $V(|S|, i_t) \forall i$ can be calculated using equation

10:

$$\begin{aligned} V(|S|, i_t) &= \sum_{a' \in A} P(a_t = a') E\left(u(a', |S|, i_t) + \delta \sum_{i' \in I} P(i_{t+1} = i | a', 1, i') V(1, i') | a_t = a' \right) \\ &= U^{|S|}(V)_{|S|, i_t} \end{aligned}$$

which gives all of the $V(|S|, i_t)$ with one CSAA iteration. The remaining V can be calculated by repeating this process. □

By Corollary 2, SAA and CSAA have the same values in seasonal period 1 after $n|S|$ iterations for all $n \in \mathbb{N}$. Suppose there is some N such that after $N|S|$ SAA iterations the value function has converged. Because both algorithms implement the same Bellman equation, they have the same asymptotic convergence rate.

When adding a state variable that takes on $|S|$ values to the dynamic problem, the computational burden of updating a state increases by a factor of $|S|^2$ because the number of transition probabilities and the number of states to update each increase by a factor of $|S|$. The CSAA iteration eliminates both of these increases in computational burden, while arriving at the same result as the SAA algorithm.

By the simplification between equation 9 and equation 10, I reduced the number of transition of probabilities that need to be calculated by a factor of $|S|$. Furthermore, the CSAA only updates $|I|$ states in each iteration. In total, applying the SAA requires the calculation of $|S| \times |I|$ transition probabilities to update $|S| \times |I|$ states, resulting in a total burden of $|S|^2 \times$

$|I|^2$. The CSAA requires the calculation of $|I|$ transition probabilities to update $|I|$ states, resulting in a total computational burden of $|I|^2$, and so the computational burden does not grow with the number of cyclic periods.

This algorithm can apply to any cyclic variable. Conventional cyclic variables include the hours of the day, the months or seasons of the year, or the phases of the business cycle. In some cases, the state space can be transformed to make part of the state space cyclic. For example, consider a dynamic model of a firm that sells a product with quality q_t, and q_t can only increase or decrease by 1 unit in each year. In this case, I can restructure the state space around variables g_t and h_t such that:

$$\begin{aligned} g_t &= q_t \bmod 2 \\ h_t &= 2 \times \left\lfloor \frac{q_t}{2} \right\rfloor \end{aligned}$$

Then, $q_t = h_t + g_t$. In this case g is a cyclic variable that alternates between 0 and 1. The CSAA can then be applied, reducing the burden of solving the model by a factor of 4. Such a model with 3 firms would have the overall burden reduced by a factor of 64.

I test the effectiveness of this algorithm by solving for the value function with both the SAA and the CSAA in 100 randomly generated parameterizations in the presented dynamic model with 52 cyclic periods. I seed both algorithms with the same starting guess of 0 in all states. In both cases, I implement the savings associated with limiting the transition probabilities as this has likely been part of previous implementations. As such, I expect the CSAA to outperform the SAA by a factor of 52. I find that the CSAA, on average, outperforms the SAA by a factor of 43.34; this discrepency is likely due to a fixed cost in setting up each algorithm. Regardless, the CSAA provides a substantial savings, and allows cyclic variables to be added to the state space at no cost.

7 Results

The resulting parameter estimates are available in Table 7. The sign of each parameter makes intuitive sense: the price coefficient is negative, the search cost is positive, while consumption utility peaks in winter and is at its at lowest in summer.

Ssearch probability varies substatially over the course of the year. I estimate that consumers only search 71% of the time in summer (Figure 5). In dollar terms, the average search costs 39 cents.

This seasonality in search leads to seasonality in the price elasticity, which varies from -1.36 in the summer to -2.18 in the winter (Figure 6). These imply that the firm would set lower prices in the winter.

To find out whether the seasonal variation in search is caused by seasonality in consumption utility or price variation, I recalculated search probability while removing all seasonal variation in each of these terms (Figure 7). I find that the majority of the yearly variation in search costs is due to seasonal price expectations, though seasonal consumption utility plays an important role.

Finally, to demonstrate how search probabilities change with search incentives, I increased the promotional depth by 25% and recalculated search probabilities (Figure 8). As expected, I find that search probability increases with promotional depth.

	Dynamic Results
Consumption Spline 1	-0.0321
	(0.1202)
Consumption Spline 2	-0.4959^{***}
	(0.1751)
Consumption Spline 3	-0.2165^{*a*}
	(0.1343)
Consumption Constant	6.0435
	(5.3479)
Price Coefficient	-0.8803^{***}
	(0.1386)
Search Cost	0.3415^{***}
	(0.6967)
log(Search Variation)	-3.1595^{***}
	(0.6071)
log(Consumption Variation)	-2.8847^{*}
	(1.165)

Table 4: Dynamic Variables

Cream of Chicken	-0.1118^{***}
	(0.0421)
Chicken Noodle	0.1038^{***}
	(0.0549)
Mega Noodle	-0.7567^{***}
	(0.0558)
Chicken Goldfish	-0.3308^{***}
	(0.0540)
Cream of Celery	-0.7741^{***}
	(0.0562)
"Missing"	-0.5848^{***}
	(0.0600)
Tomato	-0.7623^{***}
	(0.1003)
Tomato Goldfish	-1.4226^{***}
	(0.0888)

Table 5: Flavor Dummy Variables

8 Conclusions

In summary, I find that seasonal purchase without search can explain counter cyclic pricing with an explanation that is plausible and fits the data well. I find that price changes are significantly correlated with seasonal trends in price elasticity. One limitation of this paper is that it does not directly compare the impact of purchase without search with the impact of other explanations of counter cyclic pricing, such as loss leader pricing, and consumer heterogeneity. Counter cyclic pricing could be working in tandem with these other effects, and future work might compare their magnitude.

I find that seasonal trends in consumption utility are important to modelling the concentrated soup industry. Other researchers might consider adding seasonal variables to their dynamic models because seasonal variation might be important to their problem, and seasonal variables can be added to the state space with no additional computational burden using the CSAA. Finally, I show that consumers may not observe prices before making their purchase decision. Informing consumers of price promotions might increase sales, but at the cost of increasing the effective price sensitivity.

References

Bucklin, R. and J. Lattin (1991). A two-state model of purchase incidence and brand choice. *Marketing Science 10*(1), 24–39.

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

Ching, A., T. Erdem, and M. Keane (2009). The price consideration model of brand choice. *Journal of Applied Econometrics 24*(3), 393–420.

Erdem, T., S. Imai, and M. Keane (2003). Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics 1*(1), 5–64.

Hendel, I. and A. Nevo (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica 74*(6), 1637–1673.

Mehta, N., S. Rajiv, and K. Srinivasan (2003). Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing science 22*(1), 58–84.

Mela, C., K. Jedidi, and D. Bowman (1998). The long-term impact of promotions on consumer stockpiling behavior. *Journal of Marketing Research 35*(2), 250–262.

Nevo, A. and K. Hatzitaskos (2006). Why does the average price paid fall during high demand periods?

Paarsch, H. J. and J. Rust (2009, March). Valuing programs with deterministic and stochastic cycles. *Journal of Economic Dynamics and Control 33*(3), 614–623.

Ray, S., C. Wood, and P. Messinger (2012). Multicomponent Systems Pricing: Rational Inattention and Downward Rigidities. *Journal of Marketing 76*(September), 1–17.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society 55*(5), 999–1033.

Santos, M. S. and J. Rust (2004, January). Convergence Properties of Policy Iteration. *SIAM Journal on Control and Optimization 42*(6), 2094–2115.

Seiler, S. (2013). The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics 11*(2), 155–203.

Warner, E. and R. Barsky (1995). The timing and magnitude of retail store markdowns: evidence from weekends and holidays. *The Quarterly Journal of Economics 110*(2), 321–352.

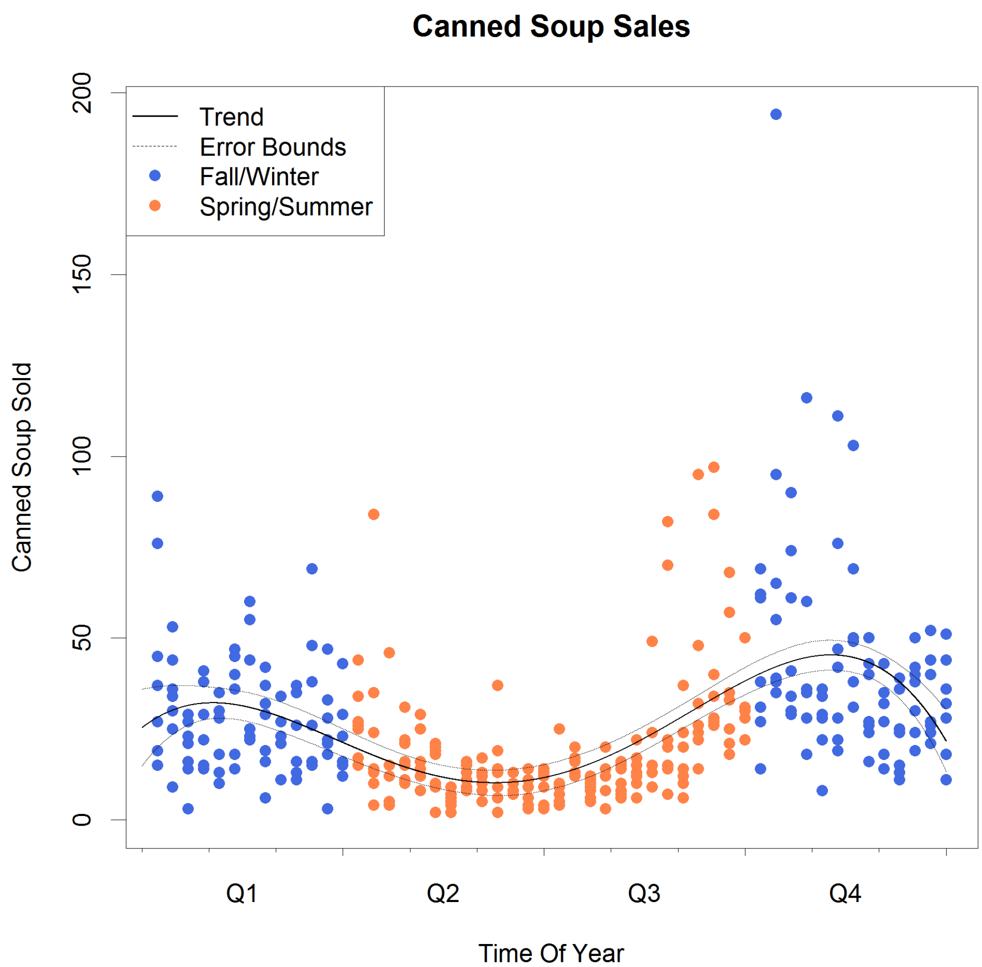

Figure 1: Seasonality in Demand

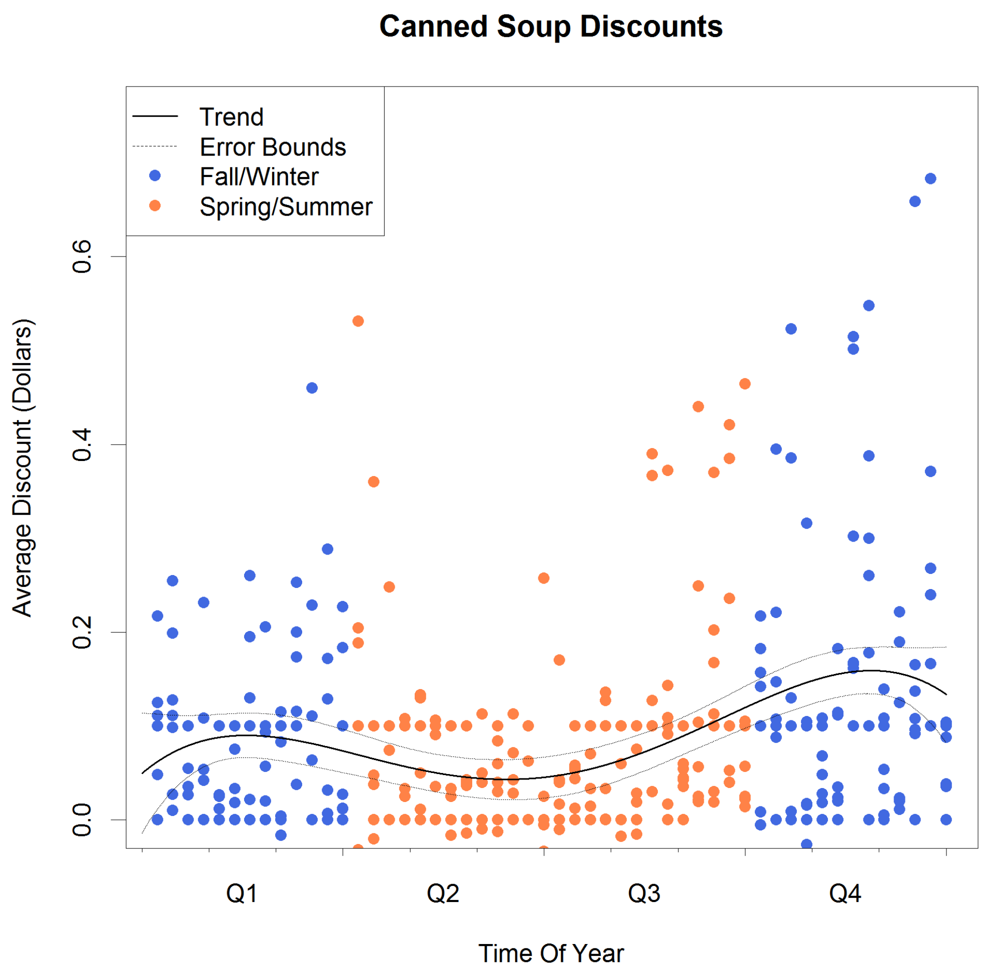

Figure 2: Seasonality in Pricing

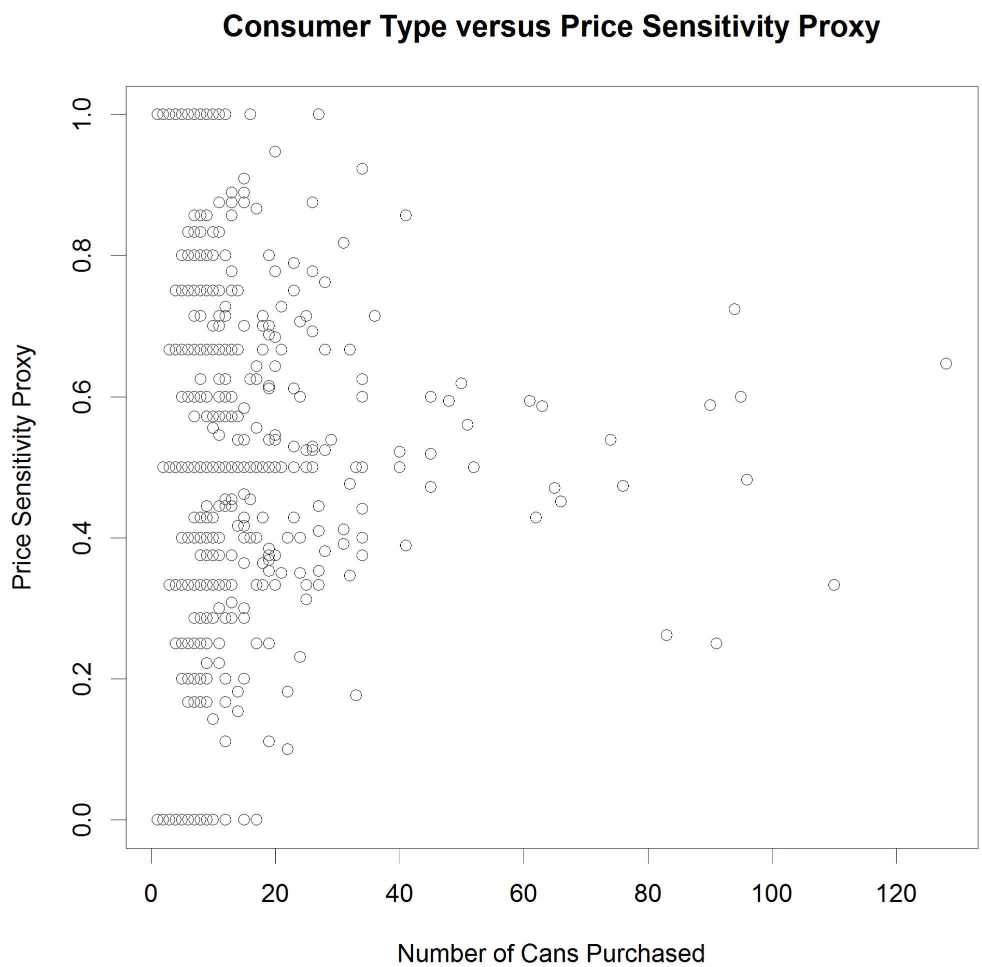

Figure 3: Price Sensitivity Proxy By Consumer Type

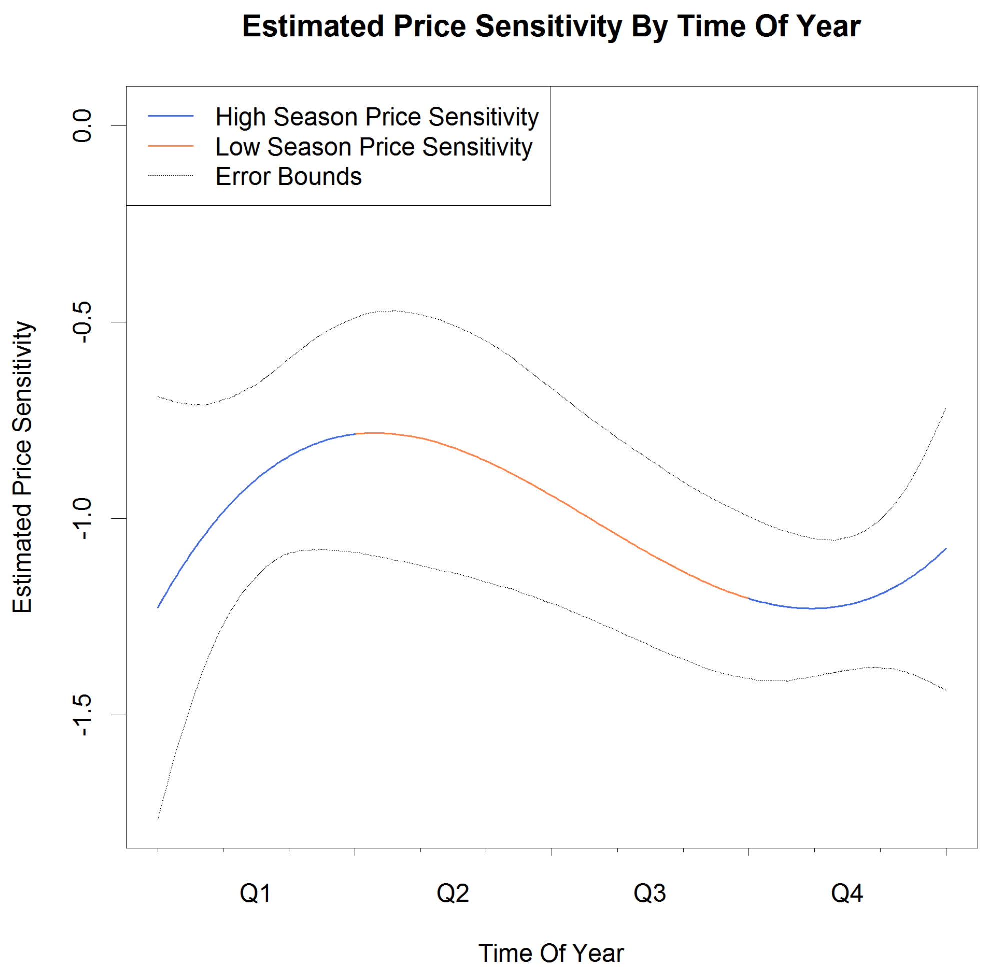

Figure 4: Price Sensitivity By Time Of Year

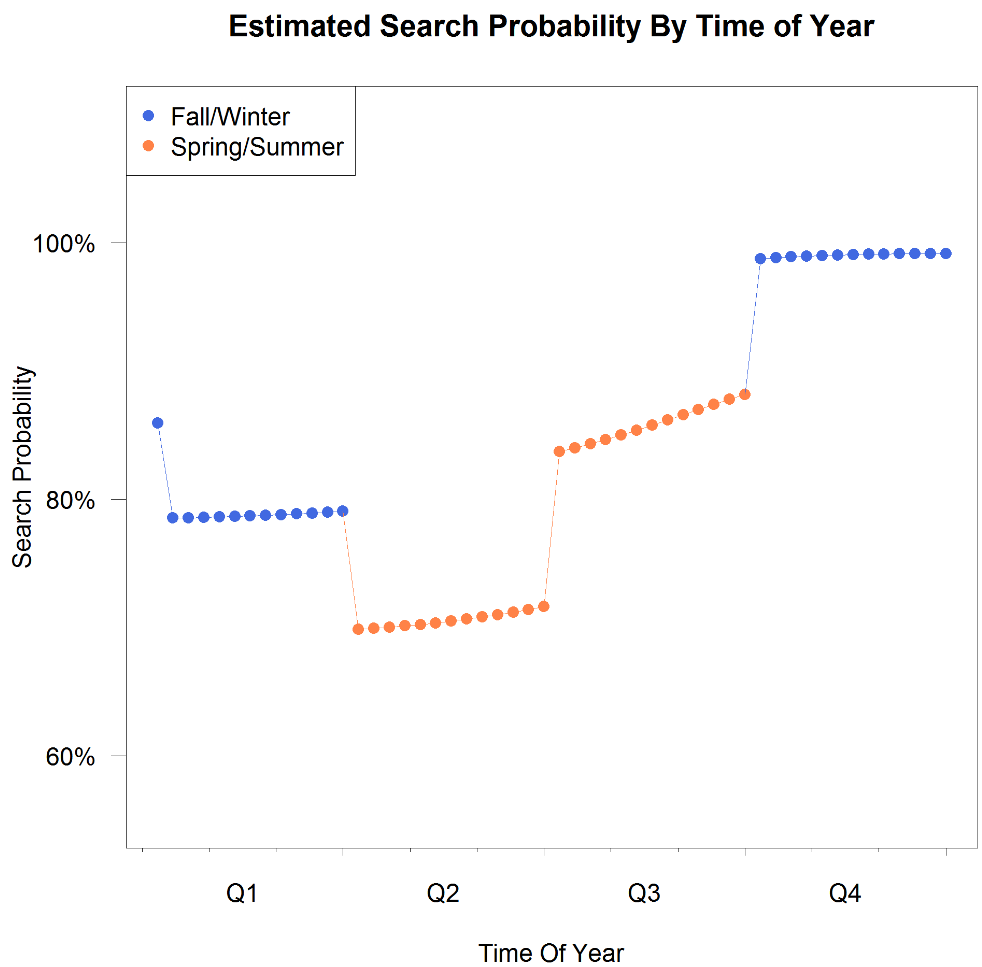

Figure 5: Estimated Search Probability

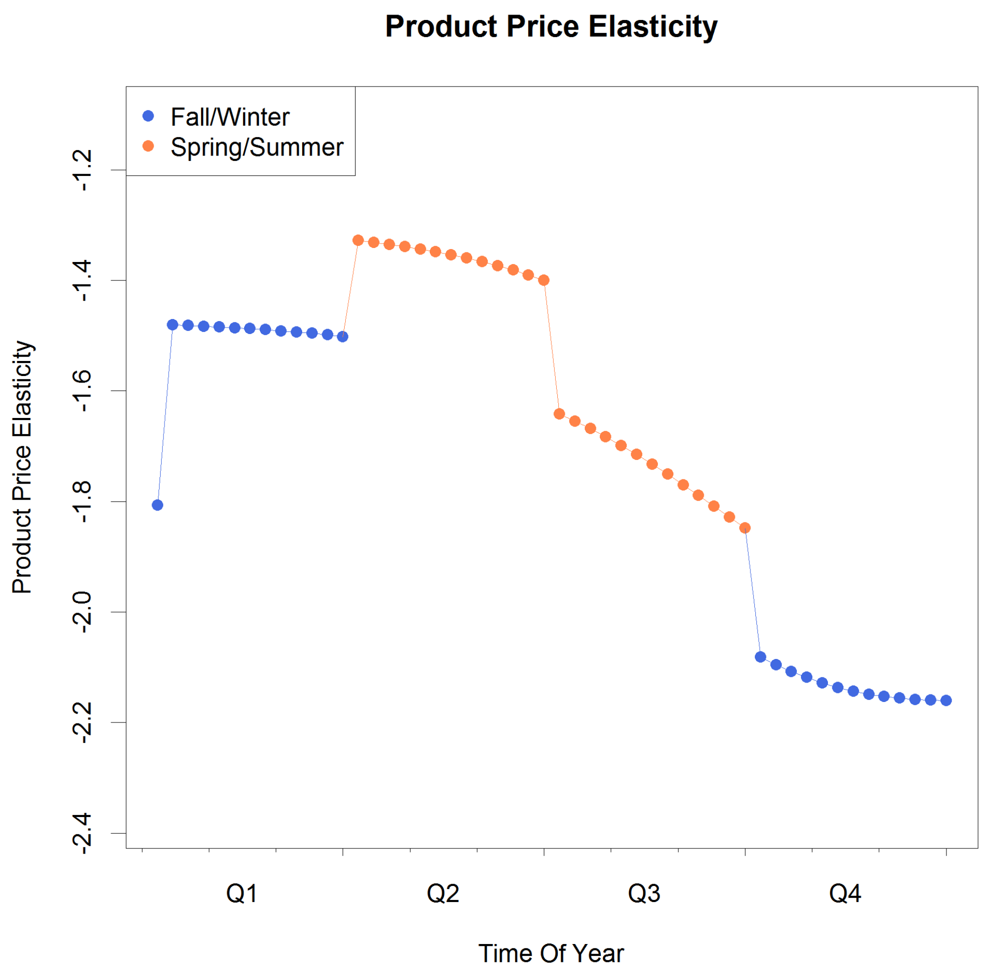

Figure 6: Estimated Price Elasticity

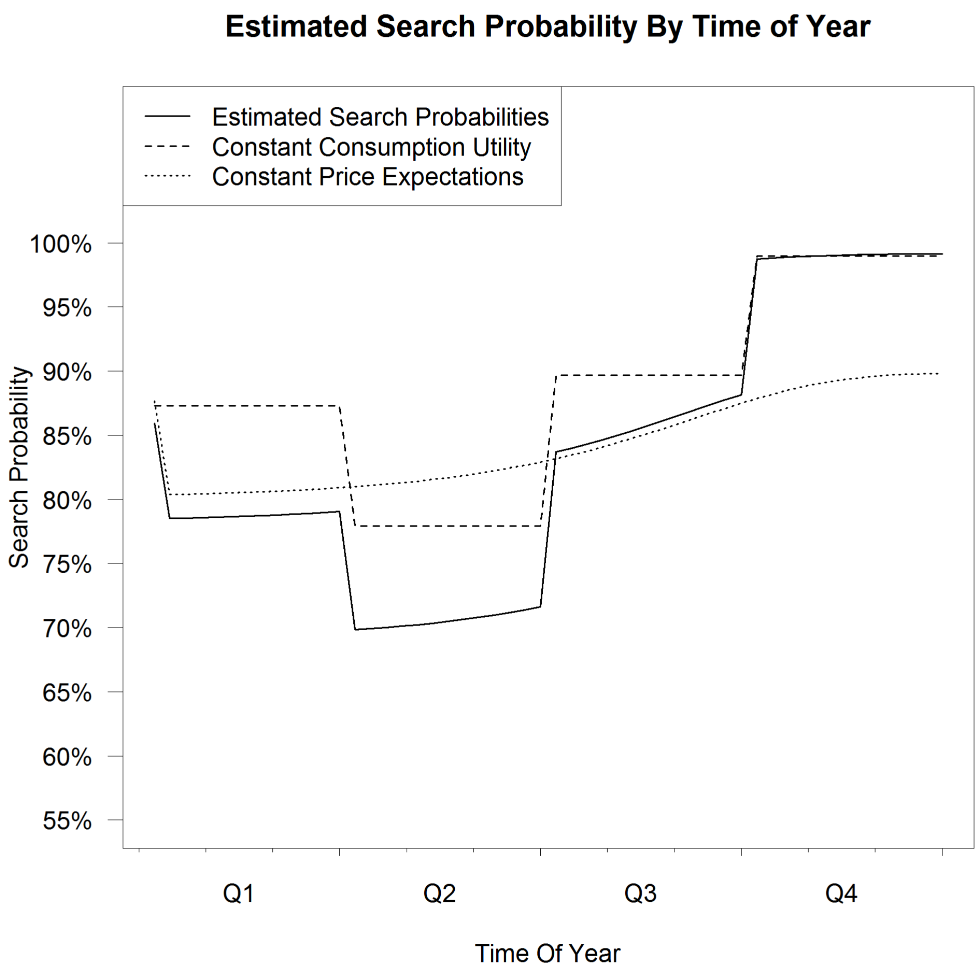

Figure 7: Counterfactual Search Probability - Determinants of Search

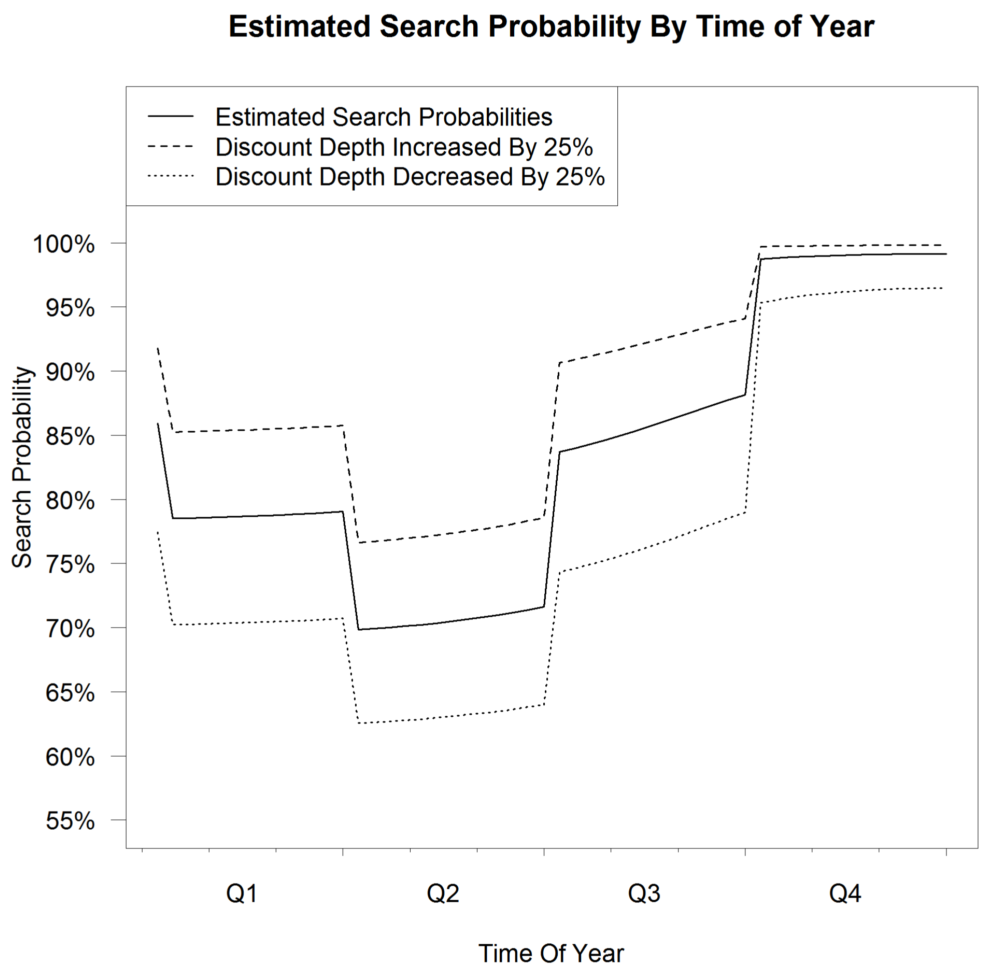

Figure 8: Counterfactual Search Probability - Promotional Depth