Displays, Sales, and In-Store Search in Retail Markets*

Matthew Gentry†

First version: October 2011

This version: September 2016

Abstract

This paper develops and estimates a structural model of differentiated-products demand in an environment with in-store displays and costly consumer price search. This investigation is motivated by two stylized facts on retail markets. First, the prevalence of retail "sales" means that relative prices in most product categories vary substantially from week to week. Second, quantities sold often respond at least as much to changes in relative locations (in-store displays) as to changes in relative prices. The model proposed here incorporates both effects: sale-induced price variation provides a reason to search, and displays convey information about prevailing prices. This model is then applied to store-level data on laundry detergent purchases, using short-run price fluctuation to recover preference parameters and short-run display fluctuations to recover search parameters. The resulting structural estimates suggest that information frictions have substantial effects on purchase outcomes, with mean search costs of between \$0.30 and \$1.50 across chains in the market. I further explore the potential relationship between consumer search and demand analysis, and find that accounting for displays and other promotions substantially lowers elasticity estimates.

*I thank my advisor, Tong Li, for his tremendous guidance and support on this and many other projects, and my committee members Ali Hortacsu, Yanqin Fan, and Eric Bond for their valuable contributions of time, energy, insight, and advice. I would not be where I am today without their support. Earlier versions of this paper have been presented at the 2011 Chicago Meeting of the Midwest Econometric Group and in seminars at the University of Alabama, the University of Tennessee, the University of Oklahoma, the Federal Reserve Bank of Kansas City, the University of Notre Dame, Baylor University, the London School of Economics, Ohio State University, and Bowdoin College; I thank participants at all presentations for a myriad of helpful comments. I have also benefited greatly from conversations with Stephan Seiler, Andrew Sweeting, Steve Berry, Matthew Shum, Marc Rysman, PJ Glandon, Caleb Stroup and others, whose insights I gratefully acknowledge. All remaining errors are my own.

†Department of Economics, London School of Economics. Email: m.l.gentry@lse.ac.uk.

1 Introduction

A large body of literature has established two clear and consistent empirical regularities in retail markets. First, supermarket prices in most product categories vary substantially from week to week, and most of this variation takes the form of "sales": temporary, quickly reversed reductions from a prevailing modal price.[1] Second, even controlling for sales and other advertising, in-store product displays (e.g. end-of-aisle features, in-store banners, and other measures designed to call attention to displayed products) have large effects on final purchase outcomes.[2] Stigler (1961)'s costly search paradigm provides a natural framework within which to analyze both effects: the week-to-week price churn induced by sales creates a natural economic motive for search, and the presence of search in turn helps to explain observed display effects. To date, however, structural work on supermarket search has been relatively limited.[3]

In this paper, I develop a structural model of retail demand in an environment with differentiated products, in-store product displays, and costly price search. I then apply this model to data on laundry detergent purchases, using temporary price reductions (sales) to recover preference parameters and temporary informational variations (in-store displays) to recover search parameters. The preference component of the model follows Berry, Levinsohn, and Pakes (1995) (henceforth BLP), but unlike BLP I assume that only prices on displayed products are observed *ex ante*. This structure is motivated by the large and significant effects of in-store displays on purchase outcomes, a fact not readily explainable in standard full-information demand models.

This work makes several core contributions to the literature on retail markets. First, I

[1]See, e.g., Pesendorfer (2002), Hendel and Nevo (2006), Griffith et al. (2009) and Chevalier and Kashyap (2011) for empirical descriptions of retailer pricing practices, and Varian (1980) for a theoretical exploration of potential links between consumer search and retail sales.

[2]See, e.g., Roberts and Lattin (1991), Andrews and Srinivasan (1995), Mehta et al. (2003), and van Nierop et al. (2010), among many others.

[3]Search models have a long history in empirical industrial organization; Sorensen (2000, 2001) to prescription drug stores, Hortacsu and Syverson (2004) to mutual funds, and Hong and Shum (2006) to online book markets are three representative examples. Historically, most such applications have focused on the case of identical products, which would exclude classic supermarket settings. More recent work has begun to reverse this trend; this literature is summarized below. There is also a long history of "consideration" models in marketing; examples include Roberts and Lattin (1991), Andrews and Srinivasan (1995), Mehta et al. (2003), and van Nierop et al. (2010), among many others. These studies draw on motivations similar to search, but their underlying methodology is typically quite different.

structurally analyze in-store search in retail markets, a subject about which relatively little is known. My estimates suggest that information frictions have substantial effects on purchase outcomes, a finding with potential implications for both the retail supply-side literature (which frequently references search as an explanation for sales) and the literature on consumer behavior. In particular, parameter medians suggest roughly 52 percent of consumers have positive search costs, with a mean search cost of roughly \$1.68 among this sub-population. Second, I explore the relationship between consumer search and demand analysis, finding that accounting for in-store displays and other promotions substantially lowers elasticity estimates. Finally, econometrically, I exploit a novel source of informational variation (displays) to estimate search effects *without* direct data on search. I thus contribute both to the literature on demand estimation in retail markets and to the emerging literature on search with differentiated products.[4]

This work is most closely related to three prior studies in industrial organization.[5] First, using panel data on laundry detergent purchases, Hendel and Nevo (2006) estimate a discrete-choice demand model incorporating dynamic responses to sales, finding that intertemporal substitution is an important avenue by which consumers respond to sales. My focus on search is obviously different from Hendel and Nevo's focus on dynamics, but the underlying sales motivation is similar. Second, Goeree (2008) develops a model purchase under limited information in differentiated-product markets, which she estimates using data on advertising in the personal computer industry. I similarly exploit advertising variation in an environment with limited information, but whereas Goeree bases estimation on a reduced-form pure consideration model, I directly incorporate endogenous price search. Finally, building on Hendel and Nevo (2006), Seiler (2013) develops a dynamic choice model with a preliminary market entry decision, finding that in-store search is an important factor explaining purchase incidence in

[4] As noted above, search applications have historically focused on the case of identical products. However, recent work has begun to reverse this trend. See, e.g., Moraga-Gonzalez and Wildenbeest (2008), *An Empirical Model of Search with Vertically Differentiated Products* (2011), Moraga-Gonzalez et al. (2010), Santos et al. (2012), and Seiler (2013) to mention just a few.

[5] Another closely related study in marketing is Mehta et al. (2003), which estimates a structural model of consideration and choice using data on laundry detergent and ketchup purchases. While I reference similar empirical patterns, my structural interpretation of these patterns is considerably different. Among other things, Mehta et al. (2003) focus primarily on learning and incorporate search via a simple reduced-form specification, while I formally characterize search using the full empirical price distribution.

detergent markets. My main innovation relative to these studies is to explicitly incorporate in-store displays as a source of identifying variation in a structural search model. Insofar as displays represent a direct proxy of information availability, my approach thus exploits a potentially powerful source of identifying variation on the role of search in differentiated product markets.

This paper is also closely related to a number of recent and contemporary papers on search and consideration set formation in retail and other markets. Pires (2012, 2014) develops and estimates a dynamic structural model of consumer search and consideration set formation; similar to the current study, he finds that displays and other retail promotional activities substantially reduce consumer search costs and therefore increase consumer access to information. Koulayev (2014) estimates a model of search behavior in online markets using variation in on-page displays to identify search frictions; this identifying strategy parallels that in the current paper, although in a different informational context. Lu (2016) estimates a demand model allowing for flexible consideration set formation, using moment inequalities applicable under any consideration set to construct identified sets for model primitives. Pinna and Seiler (2016) use data gathered from radio transmitters on shopping carts to analyze consumer shopping patterns, finding substantial reduced-form evidence of information frictions. Finally, Mojir et al. (2014) develop and estimate a structural model of spatiotemporal search in which consumers search for prices both across stores and over time. This paper contributes to this recent literature documenting information frictions in retail markets and exploring the role of promotional activities (advertising and displays) in alleviating these. It differs from current work, however, on at least two notable dimensions. First, the analysis here employs only store-level scanner data – thereby permitting much wider application than approaches which require panel data. Second, in contrast to most consideration-based models of consumer search in differentiated-products markets, I allow displayed price promotions to affect the probability of subsequent search. Displayed price promotions can thereby foreclose consideration of rival products – a potentially important business-stealing channel consistent with patterns I document in the data.

The rest of this paper is organized as follows. Section 2 summarizes my data and surveys price, display, and other promotional variation in laundry detergent markets. Section 3 presents descriptive evidence on promotional effects in the target market, and Section 4 describes the structural model I develop based on this evidence. Section 5 outlines details of my estimation procedure, and Section 6 presents key results. Finally, Section 7 concludes.

2 Data, industry, and market

Data comes from the Information Resources Incorporated (IRI) marketing dataset, which contains both UPC-level scanner data for 30 product categories in 47 geographic regions and household-level panel data for two selected markets for years 2002-2007. I here focus on the scanner sample, which includes data on revenue, units sold, temporary price promotions, displays, and other advertising activity at the store-UPC-week level for all stores and categories in the sample. Following standard practice, I divide weekly revenue by weekly to recover store-UPC-level price series.[6] Further description of this dataset is given in Bronnenberg et al. (2008).

While the structural features motivating my analysis are common to many retail markets, my empirical application in this paper focuses on the laundry detergent industry. Prior work suggests that information frictions matter in laundry detergent purchasing; Mehta et al. (2003) and Seiler (2013) directly explore aspects of search, while Hendel and Nevo (2006) find strong reduced-form display effects. The existing body of work on laundry detergents also provides a natural frame of reference within which to interpret my structural results.[7]

Finally, the structural model I develop below is primarily designed to capture store-level economic effects, and would be computationally infeasible to estimate on the entire IRI de-

[6] As typical in scanner datasets, this means that UPC-level prices are observed only in weeks with positive sales. My structural estimates focus on "important" products defined at the brand-size level, so missing UPC-level prices are not a major problem in practice. Where necessary, I fill in missing prices using the regular price series constructed below.

[7] The main potential caveat here is dynamics: Hendel and Nevo (2006) show that consumers respond to sales in laundry detergent markets by concentrating purchases in sale periods. Notably, however, Seiler (2013) finds that incorporating costly search substantially reduces estimates of dynamic response. Given this result and the vast computational cost of a full dynamic implementation, I here focus primarily on a static choice framework, with reduced-form proxies for dynamic effects included in estimation.

tergent sample. For the moment, therefore, I focus my descriptive regressions on the Atlanta market, and my primary structural estimates on the three largest detergent retail chains (IRI chain identifiers 78, 100, and 140) in this market. Taken together, these chains cover more than 90 percent of the total units of detergent sold in the Atlanta market within the IRI sample.[8] Selections in all cases are arbitrary with respect to the question under investigation, and quite typical of the broader IRI sample; I plan to continue expanding the estimation sample in future work. Except where noted otherwise, in what follows I refer to the Atlanta market simply as "the market."

2.1 The Market

Laundry detergents come in liquid and powdered forms, with liquid detergent accounting for roughly 66 percent of market-wide sales. I here focus on liquid detergents; this both simplifies computation and promotes comparability with prior studies. Table 1 summarizes prices, market shares, and marketing variables for the top 10 liquid detergent brands in the market. As this table illustrates, the market is highly concentrated, with the top 5 (10) brands accounting for roughly 80 (97) percent of volume sold. Sales, displays, and features all occur regularly, with sales occurring most frequently and displays least frequently.[9] Finally, sales induce substantial short-term variation in prices, with average discounts in the neighborhood of 12-15 percent.[10] For completeness, Table 6 in Appendix 1 presents corresponding summary statistics for stores in the estimation sample; as expected, these look similar to those for the overall Atlanta market.

[8] Note, however, that the IRI dataset does not contain all stores in the market; in particular, Wal-Mart is known to be excluded from the dataset.

[9] The IRI dataset also includes information on type of features and displays, but for current purposes I simply aggregate to "any feature" and "any display" indicators.

[10] As discussed in detail in Sections 2.2 below, the vast majority of price variation in this market is driven by "sales," temporary, quickly reversed reductions from a prevailing "regular" price. "Discount if sale" gives the average percentage discount (relative to regular price) in periods where a sale occurred.

Table 1: Liquid laundry detergent sales, Atlanta market

BRAND	Price / oz	Sale	Feat	Disp	Disc if sale	Share
ALL	0.0234	0.264	0.111	0.0645	0.149	0.126
ARM&HAMMER	0.0169	0.300	0.0599	0.0354	0.149	0.0969
CHEER	0.0347	0.146	0.0480	0.0312	0.162	0.0379
FAB	0.0202	0.329	0.0531	0.0338	0.209	0.0218
GAIN	0.0272	0.194	0.107	0.0670	0.167	0.104
PUREX	0.0162	0.276	0.0900	0.0926	0.166	0.186
SURF	0.0262	0.254	0.0890	0.0856	0.190	0.0171
TIDE	0.0333	0.223	0.138	0.0807	0.189	0.312
WISK	0.0287	0.312	0.157	0.0918	0.174	0.0490
XTRA	0.0102	0.211	0.0820	0.128	0.187	0.0192

Notes: *Sale*, *Feat*, and *Disp* are UPC-by-week indicators for sale, feature, and display promotions. *Disc if sale* represents average discount from "regular" price in weeks a sale occurs, where regular price series are constructed as in Section 2.2.2 below.

2.2 Prices and promotions

The interaction between prices, promotions, and search is at the heart of my structural analysis, so a thorough investigation of price and promotional variation in the data is essential in motivating the particular structure employed. Three patterns in particular will play a key role in my subsequent analysis: most price variation is driven by temporary "sales," most products have "regular prices" about which sales take place, and non-price promotions (displays and features) vary substantially apart from sales. Each of these patterns is explored in more detail below.

2.2.1 Price patterns

The large and growing literature on retail sales has shown that short-term, quickly reversed price reductions are a pervasive feature of most retail markets.[11] Surveying this literature, Hosken and Reiffen (2007) identify five key empirical regularities in supermarket pricing, of which the three most relevant for current purposes are (1) that there tends to be a large mode in the pricing distributions for all types of goods, (2) that most deviations from this mode are

[11] Relevant studies include Kehoe and Midrigan (2007), Chevalier and Kashyap (2011), Nakamura and Steinsson (2011), and Glandon (2011) on the macro side, and Griffith et al. (2009), Konieczny and Skrzypacz (2004), Pesendorfer (2002), and of course Hendel and Nevo (2006) on the micro side.

price reductions, and (3) that most such price reductions are temporary. Figure 1 presents one representive price history for the current market; Figures 8 and 9 in Appendix 1 present two additional examples.[12] All clearly illustrate the core pattern in question: a dominant regular price which changes infrequently, punctuated by frequent temporary sales.

Figure 1: Price history for Tide 100oz, IRI store 683960 (2002-2007)

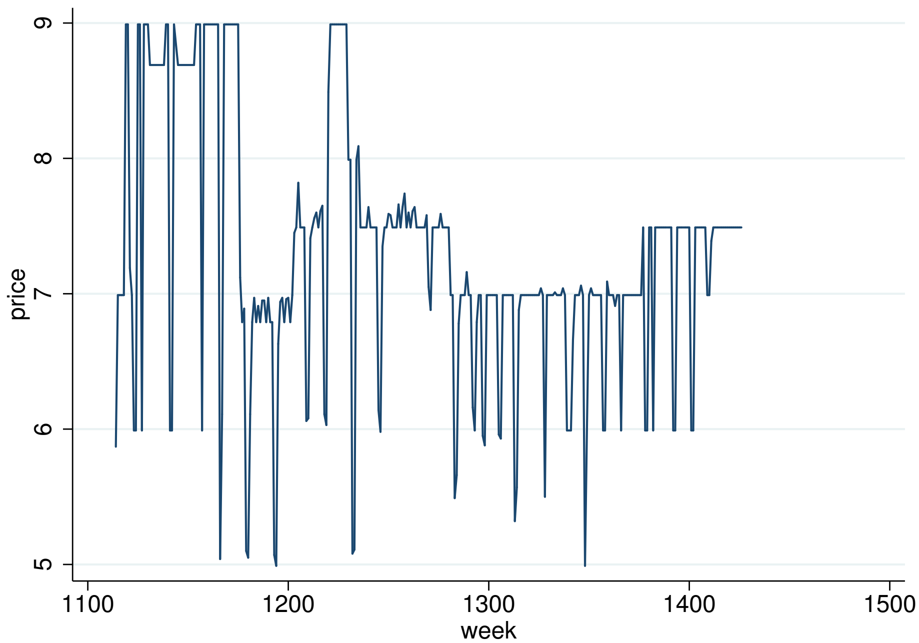

In turn, this consistent empirical pattern suggests at least two structural implications. First, sale-induced price churn motivates the hypothesis of costly price search: while consumers may be able to learn regular prices over time, they are likely to observe actual price realizations only after conscious effort. Second, insofar as average tastes should be relatively stable from week to week, short-run price variation induced by sales should help to identify price elasticities in a standard preference model.[13] Both insights play key roles in my subsequent structural

[12]The examples chosen are the best-selling UPCs for the three largest brands (Tide, Purex, and All) in the Atlanta market.

[13]Of course, price *distributions* will be endogenous to consumer tastes, but an examination of price histories suggests that week-to-week price *realizations* are not. This is particularly true of a market like laundry detergent, where the underlying fundamentals driving demand presumably shift only gradually. In some other markets, there is evidence that sale frequencies *increase* during periods of high demand (see, e.g., Chevalier et al. (2003)), but even this seems best interpreted as endogeneity of the price distribution rather than endogeneity of week-to-week price realizations. I will return to these issues in more detail below.

analysis.

2.2.2 Regular prices

In my structural analysis, regular prices will be employed primarily to characterize expected gains from search: consumers will be assumed to know regular prices from prior experience but only observe sale realizations following costly search. I construct regular price series using a modified rolling median algorithm: drop all periods listed as sales, calculate 9-week rolling price medians over remaining periods, and fill in promotional periods forward or backward based on least deviations from observed prices.[14] Figure 2 illustrates this procedure applied to the price series for Tide 100oz in Figure 1; a formal description of the algorithm and additional examples are given in Appendix 2.

Figure 2: Price vs Regprice for Tide 100oz, IRI store 683960 (2002-2007)

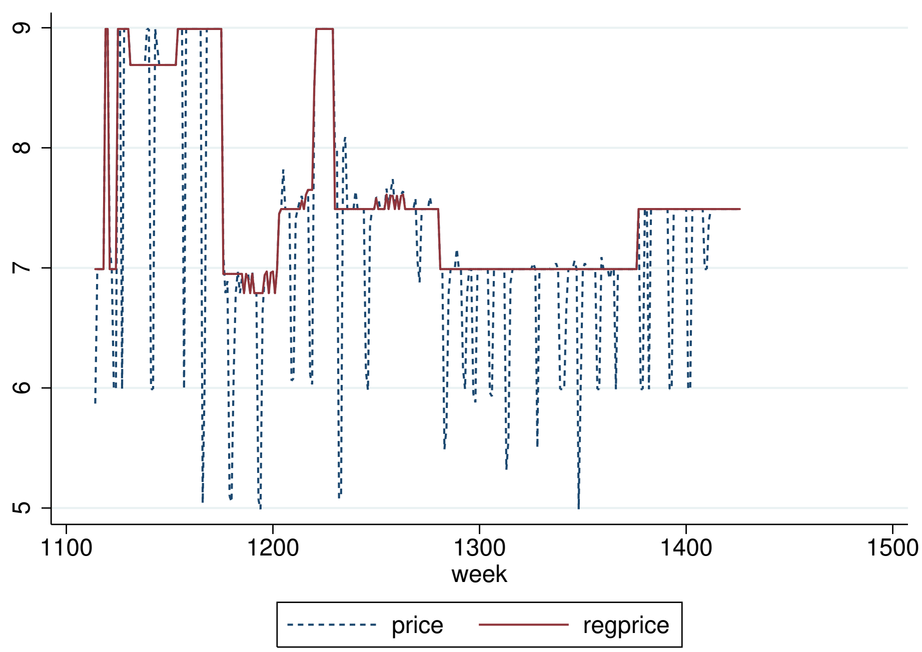

[14]Several other definitions have been proposed in the retail sales literature: for instance, Hendel and Nevo (2006) use a simple price mode, Eichenbaum et al. (2011) use a quarterly price mode, and Kehoe and Midrigan (2007) and Chevalier and Kashyap (2011) track regular prices as price changes not reversed within 5 weeks. My definition produces results similar to these approaches, but more consistently incorporates the concept of a regular price (by construction, periods with sales are not "regular") and more fully exploits the detailed promotional information available in the IRI data.

Table 2: UPC-level promotional variation, Atlanta market 2002-2007

	Conditional means			
Outcome	$sale = 1$	$disp = 1$	$feat = 1$	Overall
sale	1	0.579	0.896	0.237
disp	0.172	1	0.258	0.0704
feat	0.398	0.387	1	0.105

Notes: Cells represent frequencies of row outcomes conditional on column outcome; variables at store-week-UPC level.

2.2.3 Sales, displays, and features

My identification strategy hinges on exploiting week-to-week price variation to recover preference effects and week-to-week display variation to recover search effects. Strong correlation in sale and display realizations could pose problems for this approach, so Table 2 explores the nature of promotional variation in the data. In particular, note that less than 60 percent of products on display are also on sale, and only 17 percent percent of products on sale are also on display. This level of distinct variation should be more than sufficient to support the identification strategy pursued here.

3 Descriptive regressions: price and promotion effects

A vast body of evidence from economics and marketing has established that both in-store displays and external advertising have important effects on consumer behavior.[15] Table 3 highlights a number of descriptive regressions which strongly suggest that this pattern extends to the current sample. In these regressions, *sale*, *disp*, and *feat* are the IRI-generated promotion indicators summarized above, *discount* is percentage difference between current price and regular price, *discxdisp* and *discxfeat* are the corresponding interaction terms, and the dependent variable *qnorm* is a normalized store-UPC-level sales measure (percentage by which quantity sold this week exceeds average weekly quantity sold). Table 7 in Appendix 1 reports related regressions using several alternative outcome measures and covariate specifications; qualitative results are similar in all cases.

[15] See for instance Hauser and Wernerfelt (1990). Roberts and Lattin (1991), Andrews and Srinivasan (1995), Mehta et al. (2003), and van Nierop et al. (2010), among many others.

Table 3: UPC-level promotion effects, Atlanta detergent 2002-2007

VARIABLES	(1) Prices only	(2) Promo dummies	(3) Only interacts	(4) All channels
price	-0.0104***	-0.00952***	-0.00987***	-0.00912***
	(0.000341)	(0.000322)	(0.000355)	(0.000328)
discount	-3.263***	-2.230***	-1.205***	-1.144***
	(0.0513)	(0.0415)	(0.0355)	(0.0423)
sale	0.145***	-0.0299***		0.0528***
	(0.00823)	(0.00584)		(0.00621)
feat		0.734***		0.323***
		(0.00685)		(0.0155)
disp		0.563***		0.304***
		(0.00735)		(0.0151)
discxfeat			-3.461***	-2.286***
			(0.0633)	(0.0951)
discxdisp			-3.066***	-2.204***
			(0.107)	(0.129)
Constant	0.113	-0.159	-0.0649	-0.146
			(553.0)	
Observations	528,801	528,801	528,801	528,801
R-squared	0.202	0.282	0.292	0.303

Notes: Dependent variable is *qnorm*, percent by which store-UPC-week quantity sold exceeds average store-UPC weekly quantity sold. *sale*, *disp*, and $feat$ are store-UPC-week promo indicators, $discount \equiv (price - regprice)/regprice$, and $discxdisp$ and $discxfeat$ are corresponding interaction terms. Robust standard errors in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Several key patterns in Table 3 suggest the presence of search effects. First, as expected, display and feature effects are large and significant in every regression where they appear. Second, the magnitude of the coefficient on *discount* falls consistently as promotional covariates are added: nearly two-thirds from Column (1) to Column (4). Finally, displays and features strongly increase the quantity effects of price reductions. This can be seen most clearly from Column (4), which suggests that displayed and featured price reductions have quantity effects roughly three times larger than those not advertised. This last finding in particular strongly supports the idea that promotional activities convey price information, which in turn motivates the structural model I develop below.[16]

More broadly, Table 3 also highlights the importance of properly accounting for promo-

[16]Of course, more work would be needed to claim a definite causal connection. Combined with prior marketing work on displays, advertising, and consideration, however, the patterns reported here are highly suggestive.

tional effects in demand analysis. Since *discount* and *qnorm* are expressed as percentage deviations, coefficients on *discount* will roughly parallel own-price elasticities.[17] Column (1) of Table 3 thus roughly corresponds to the expected output of a demand model with no promotional effects, and Column (2) to that of a demand model incorporating displays and features as preference shifters. Comparison with Column (4) suggests that either partial approach will substantially overstate own-price elasticities: relative to the model with all interaction effects, the model with no promotion controls overestimates the coefficient on *discount* by a factor of three and the model with only promotion dummies overestimates the coefficient on *discount* by a factor of two. This finding suggests the need to explore the structural relationship between price elasticities and promotional effects, a question to which I return in more detail below.

4 Structural search-plus-choice demand model

I consider choice among $\mathcal{J} = \{1, ..., J\}$ differentiated products in a single retail market. Each product j is characterized by a vector of attributes x_j (here assumed time-invariant for simplicity), and is marketed via a time-varying two-part process: a weekly price p_{jt} and a weekly display indicator d_{jt}.[18] Consumers are assumed to know the equilibrium price distribution $F_{pt}(\cdot)$ for each period t, but may not observe price realizations $\mathbf{p}_t$ *ex ante*. Purchasing decisions therefore involve aspects of price search. The no-purchase option is modeled via an outside good 0, which has utility normalized to 0 and is always displayed.

Within this environment, I model consumer i's choice problem as a two-stage process. First, upon entering the store, consumer i costlessly observes prices for all products in the *display*

[17] The equivalence is rough because the percentages used to construct *pgap* and *qnorm* do not exactly correspond to those used to define elasticities, and because the specification here ignores rival prices. Price endogeneity is not a major concern since variation in *pgap* primarily reflects week-to-week changes in sales realizations, and as noted in Section 2.2 these realizations are almost certainly not responding to week-to-week shifts in demand.

[18] In practice, this process will presumably reflect an equilibrium outcome of competition among profit-maximizing retailers, but for purposes of demand estimation I take it as given. A substantial existing literature explores this supply-side equilibrium more formally; see, e.g., Salop and Stiglitz (1977), Butters (1977), Varian (1980), Burdett and Judd (1983), Rob (1985), Pesendorfer (2002), and Glandon (2011), to name just a few. Notably, several of these studies cite costly consumer search as a leading explanation for observed retail practices, and in future work I hope to use the model described here to explore this connection more fully.

set $\mathcal{D}_t = \{j|d_{jt} = 1\}$. Next, in Stage 1, consumer i chooses whether to search remaining prices at effort cost $c_i \sim F_c(\cdot)$. Finally, in Stage 2, consumer i chooses the utility-maximizing alternative j from the set of products searched. This search-plus-choice process is summarized more formally in Assumptions 1-3 below.

Assumption 1. *All consumers know the prevailing price distribution $F_{pt}(\cdot)$, but week-to-week price realizations are random from the perspective of at least some consumers.*

Assumption 2. *For each consumer i, the Stage 1 search decision is binary: either search only the display set $\mathcal{D}$ (free) or search all products $\mathcal{J}$ at cost $c_i \geq 0$.*

Assumption 3. *Consumer i's choice set is the set of products searched; i.e. consumer i can purchase product j only if consumer i first searches product j.*

Assumption 1 is standard in the search literature, but also has a natural economic motivation here: general knowledge of the price distribution can arise from repeated shopping experience, but price churn induced by sales implies that price realizations are likely to differ substantially from visit to visit.[19] Assumption 2 is more restrictive, but some additional structure seems essential in this context: the leading fully endogenous alternatives (sequential or simultaneous search) would be computationally infeasible and would likely derive identification only from functional-form assumptions.[20] The particular structure chosen is motivated by a "supermarket story" in which consumers face a choice between buying from an end-of-aisle display or searching the entire aisle. While obviously not trivial, this represents a reasonable simplification of the economic fundamentals involved. Finally, Assumption 3 is standard in the search literature. I thus seek to distill the interactions between search, displays, and choice noted above into an econometrically tractable structural choice model.

As usual in the discrete-choice demand literature, consumers are assumed to make Stage 2 purchasing decisions to maximize net utility. For current purposes, I specialize preferences

[19] A slight caveat here: the model I estimate will permit a fraction of consumers with zero search costs. This is primarily intended to account for the fact that some people enjoy shopping, but would also be consistent with a model where some hyper-rational consumers perfectly predict retail behavior.

[20] As noted above, there is a rapidly growing literature on estimation of search models in environments with differentiated products. To my knowledge, this literature has not yet established identification results for the context explored here.

using a linear structure similar to BLP (1995; 2004):

Assumption 4. *Conditional on facing choice set $C \in \{\mathcal{D}, \mathcal{J}\}$, consumers have BLP-style linear preferences over alternatives with heterogeneous tastes for product attributes x:*

$$\begin{aligned} v_{ijt}^{C} &= \mu_{ijt} - \alpha_i p_{jt} + \epsilon_{ijt}^{C}, \\ \mu_{ijt} &\equiv \beta_i x_{jt} + \xi_{jt}; \end{aligned}$$

where parameter α_i measures i's (potentially idiosyncractic) marginal utility of income, random coeffiencients β_i describe i's idiosyncratic preferences over product attributes, ξ_{jt} represents average chain-level residual taste for product j at time t, and errors ϵ_{ijt}^{C} capture instantaneous idiosyncratic choice-specific consideration shocks.

This specification takes unobserved types $\omega_i \equiv (\alpha_i, \beta_i, c_i)$ to describe consumers' fundamental preferences over alternatives in the market, reflecting factors such as brand loyalty which are likely to be persistent across stages of the shopping process. Consumers will select into search on the basis of these fundamental preferences, thereby allowing both the set of products on display and the relative magnitude of displayed discounts to endogenously affect the composition of consumers continuing to search. Meanwhile, I interpret idiosyncratic choice-specific errors $\left\{\epsilon_{ijt}^{\mathcal{D}}\right\}_{j\in\mathcal{D}}$ and $\left\{\epsilon_{ijt}^{\mathcal{J}}\right\}_{j\in\mathcal{J}}$ as reflecting instantaneous preference shocks arising from consideration of a given choice set, with these drawn identically and independently across consideration rounds.[21] I formalize this structure as follows:

Assumption 5. *Consumer types $\omega_i \equiv (\alpha_i, \beta_i, c_i)$ are drawn from a joint distribution function $F_\omega(\cdot|\zeta_{it})$ that is continuous, time-invariant, and depends (at most) on a vector of individual- or market-level observables ζ_{it}. Meanwhile, idiosyncratic choice-specific errors $\left\{\epsilon_{ijt}^{\mathcal{D}}\right\}_{j\in\mathcal{D}}$,*

[21] The assumption that idiosyncratic shocks are drawn independently across consideration rounds follows Seiler (2013) among others. In prior versions of this paper, I also explored a model where the idiosyncratic choice-specific errors were fixed across consideration rounds. While greatly increasing computational complexity, this alternative approach led to very similar qualitative findings with little if any increase in explanatory power. I therefore focus on independent draws as a natural simplification of the model.

$\left\{\epsilon^{\mathcal{J}}_{ijt}\right\}_{j\in\mathcal{J}}$ *are drawn identically and independently from a Type I Extreme Value distribution:*

$$\textit{For each } C \in \{\mathcal{D}, \mathcal{J}\}, \ \epsilon^{C}_{ijt} \sim F_{\epsilon}(\epsilon) = \exp(-e^{-\epsilon}) \ \textit{ for all } i, j, t.$$

In practice, some consumers may enjoy shopping for its own sake, in which case the net "effort cost" of search c_i may be negative. For completeness, in estimation I also permit a subpopulation of "shoppers": a mass λ of consumers who exogenously search. The standard full-information discrete-choice demand model is formally nested within this structure by setting $\lambda = 1$. In practice, however, I find this subpopulation to be negligible, suggesting that information frictions play an important role in shaping final purchase outcomes.

4.1 Consumer choice problem

Under assumptions 1-3, consumer i's Stage 1 search problem reduces to a binary decision: buy a product from the display set $\mathcal{D} \cup \{0\}$, or search all available products $\mathcal{J}$ at effort cost c_i. Consumer i will thus choose to search when expected utility gain from search exceeds c_i, and will purchase from the display set otherwise.

Search stage Suppose that consumer i with private type $\omega_i = (\alpha_i, \beta_i, c_i)$ elects to search all products, and consider the post-search purchase problem facing i at price realization $\mathbf{p}$. Suppressing for the moment dependence of parameters on t, let $\delta_{ij}(p_j) \equiv \mu_{ijt} - \alpha_i p_j$ be the fundamental component of indirect utility i derives from purchasing good j at price p_j. Then normalizing the utility of the outside good to zero, the probability that i chooses good j at price realization p is

$$P^{\mathcal{J}}(j|\mathbf{p};\omega_i) = \frac{\exp(\delta_{ij}(p_j))}{1+\sum_{k=1}^{J}\exp(\delta_{ik}(p_k))}.$$

Furthermore, the expected net utility i derives from this purchase choice is

$$V_i^S(p) \equiv \log\left(\sum_{j=1}^{J}\exp(\delta_{ij}(p_j))\right) + \gamma,$$

where γ is Euler's constant. This describes i's expected payoff from continuing to search over unknown prices when the actual (unknown) price realization is $\mathbf{p}$.

Display stage Now consider the continuation problem consumer i faces in the display stage. At this stage, i may either purchase from the set of products currently displayed, search over remaining prices and products at effort cost c_i, or purchase the outside good and exit the market. Suppressing subscript t as above, purchasing a product $j \in \mathcal{D}$ yields utility

$$v_{ij}^{\mathcal{D}} = \delta_{ij}(p_j) + \epsilon_{ij}^{\mathcal{D}}.$$

Meanwhile, in choosing whether to search, i must form expectations over prices for products not currently on display. Define $\mathbf{d} = j : j \in \mathcal{D}$ and partition the price vector $\mathbf{p}$ into *displayed prices* $\mathbf{p}^D$ and *searched prices* $\mathbf{p}^S$ (so that $\mathbf{p} \equiv (\mathbf{p}^S, \mathbf{p}^D)$). Observing displays then conveys information of $(\mathbf{p}^D, \mathbf{d})$, which may both be relevant for forecasting the unknown prices $\mathbf{p}^S$. Let $EV_i^S(\mathbf{p}^D, \mathbf{d})$ denote the expectation of i's net search-stage utility $V_i(p)$ with respect to unknown prices P^S conditional on information conveyed by observing displays:

$$EV_i^S(\mathbf{p}^D, \mathbf{d}; \omega_i) = \int V_i^S(P^S; \mathbf{p}^D)\, F_p(dP^S | \mathbf{p}^D, \mathbf{d}).$$

Given continuation payoffs $EV_i^S(\mathbf{p}^D, \mathbf{d}; \omega_i)$ and effort cost c_i, I assume that carrying on to search remaining products yields expected net continuation utility

$$u_i^s(\mathbf{p}^D, \mathbf{d}) = EV_i^S(\mathbf{p}^D, \mathbf{d}) - c_i + \epsilon_i^S,$$

where ϵ_i^S is an instantaneous choice shock associated with search assumed to be drawn identically and independently from other choice shocks from a Type I EV distribution.[22] Consumer i's display-stage problem thus reduces to a utility-maximizing choice over a set of logit-style alternatives. Conditional on choosing to purchase in the display stage, the probability that i

[22]The choice shock ϵ_i^S serves primarily to smooth and simplify predicted choice probabilities; its presence is theoretically inessential.

chooses good $j \in \mathcal{D}$ is thus

$$P^{\mathcal{D}}(j|\mathbf{p}^D, \mathbf{d}; \omega_i) = \frac{\exp(\delta_{ij}(p_j))}{1 + \sum_{k \in \mathcal{D}} \exp(\delta_{ik}(p_k))},$$

where clearly $Pr(j|p, \mathcal{D}) = 0$ for all j if $\mathcal{D} = \emptyset$. Meanwhile, the associated probability that i continues to search remaining products given type ω_i and display information $\Omega_{jt}^{\mathcal{D}}$ is

$$P^S(\mathbf{p}^D, \mathbf{d}; \omega_i) = \frac{\exp(EV_i^S(\mathbf{p}^D, \mathbf{d}; \omega_i) - c_i)}{1 + \exp(EV_i^S(\mathbf{p}^D, \mathbf{d}) - c_i) + \sum_{k \in \mathcal{D}} \exp(\delta_{ik}(p_k))}.$$

Note that for a given realization of consumer parameters ω_i we can approximate the expectation $EV_i^S(\mathbf{p}^D, \mathbf{d})$ using draws from the empirical distribution of prices observed at display set $\Omega_t^{\mathcal{D}}$:

$$EV_{it}^S(\mathbf{p}^D, \mathbf{d}) \approx \frac{1}{R} \sum_{r=1}^{R} V_{it}^S(\mathbf{p}_r^S; \mathbf{p}^D),$$

where $\mathbf{p}_r^S \sim F_p(P^S|\mathbf{p}^D, \mathbf{d})$ and displayed prices p^D are held fixed across resampling iterations.

4.2 Predicted market shares

From above, we have derived the three key elements of choice problem for a consumer of type ω_i; the probability that good $j \in \mathcal{J}$ is purchased conditional on search $(P^{\mathcal{J}}(j|p_t; \omega_i))$, the probability that good $j \in \mathcal{D}$ is purchased conditional on not searching $(P^{\mathcal{D}}(j|\Omega_t^{\mathcal{D}}; \omega_i))$, and the probability of search given information communicated by displays $(P^S(\Omega_t^{\mathcal{D}}; \omega_i))$. It only remains to aggregate these into the overall market shares predicted by the model. Toward this end, note that the probability that a consumer of type ω_i purchases a good $j \in \mathcal{D}$ at marketing realization $(\mathbf{p}, \mathbf{d})$ is:

$$\begin{aligned} P(j|\mathbf{p}, \mathbf{d}; \omega_i) &= P^S(\mathbf{p}^D, \mathbf{d}; \omega_i) \cdot P^{\mathcal{J}}(j|\mathbf{p}; \omega_i) + \left[1 - P^S(\mathbf{p}^D, \mathbf{d}; \omega_i)\right] \cdot P^{\mathcal{D}}(j|\mathbf{p}^D, \mathbf{d}; \omega_i) \\ &= P^{\mathcal{D}}(j|\mathbf{p}^D, \mathbf{d}; \omega_i) + \left(P^{\mathcal{J}}(j|\mathbf{p}; \omega_i) - P^{\mathcal{D}}(j|\mathbf{p}, \mathbf{d}; \omega_i)\right) P^S(\mathbf{p}, \mathbf{d}; \omega_i). \end{aligned}$$

Meanwhile, the probability that a consumer of type ω_i purchases good $j \in \mathcal{J} \backslash \mathcal{D}_t$ at prices p_t is

$$P(j|\mathbf{p}, \mathbf{d}) = P^S(\mathbf{p}^D, \mathbf{d}; \omega_i) \cdot P^{\mathcal{J}}(j|\mathbf{p}; \omega_i).$$

Integrating across types ω_i for each good $j \in \mathcal{J}$, the model thus predicts an unconditional market share $\sigma_j(\mathbf{p}, \mathbf{d}; \zeta)$:

$$\begin{aligned} \sigma_j(\mathbf{p}, \mathbf{d}; \zeta) \quad = \quad & d_j \int_{\Omega} \left[1 - P^S(\mathbf{p}^D, \mathbf{d}; \omega_i)\right] \cdot P^{\mathcal{D}}(j|\mathbf{p}^D, \mathbf{d}; \omega_i)\, dF_{\omega}(\omega; \zeta) \\ & + \int P^S(\mathbf{p}^D, \mathbf{d}; \omega_i) \cdot P^{\mathcal{J}}(j|\mathbf{p}; \omega_i)\, dF_{\omega}(\omega; \zeta). \end{aligned} \tag{1}$$

Finally, stacking predicted market shares $\sigma_j(\mathbf{p}, \mathbf{d}; \zeta)$ across products j gives the overall $(J + 1) \times 1$ predicted search-plus-choice demand system $\sigma(\mathbf{p}, \mathbf{d})$ corresponding to marketing realizations $(\mathbf{p}, \mathbf{d})$. This system in turn provides the basis for my estimation algorithm.

5 Econometrics: implementation and estimation

As usual, estimation will involve matching predicted market shares $\sigma(\mathbf{p}, \mathbf{d}; \xi_t)$ to empirical choice probabilities in an estimation sample. I implement this match using simulated maximum likelihood under the following identification hypothesis:

Assumption 6. *Price, display, and feature* distributions *are endogenous to consumer tastes, but week-to-week variation in price, display, and feature* realizations *is not.*

This assumption is motivated by the empirical patterns noted in Section 2, and is consistent with empirical surveys of retail price-setting behavior.[23] It should be plausible in almost any market exhibiting a clear regular-sale-regular price sequence, but is particularly so for products like detergent where underlying demand fundamentals are likely to change only slowly over time. Short-run changes in promotional realizations will then provide informative variation

[23] Two observations from the sales literature seem particularly relevant here. First, Kehoe and Midrigan (2007) suggest that retailers typically update pricing strategies on a quarterly system, which would imply that week-to-week demand shifts are not driving short-term price variation. Second, Hosken and Reiffen (2007) note that sales are typically not driven by retailer margins, which would undercut typical price instruments at the store-week level.

on underlying preference fundamentals.[24]

5.1 Implementation

In practice, I estimate separate model specifications for each "large chain" in the market, where the set of "large chains" is defined by first sorting chains in the market by units of detergent sold in the sample period, then choosing the first n of these such that taken together "large chains" represent at least 90 percent of the total units of detergent sold. In the IRI subsample for the Atlanta market, this leaves me with three large retailers (IRI chain identifiers 78, 100, and 140) for the sample period January 1 2003 - December 31, 2005. Consistent with the model outlined above, I take the unit of observation to be the store-week period. Given that I seek to analyze behavior within a series of small markets (store-week observations), I assume that the product-specific component of preferences is constant over time within a given chain: in other words, for each chain c in the market, I assume $\xi_{jct} = \xi_{jc}$ for all t in the sample.[25] For the moment, I adopt the following additional parametric structure on the distributions of taste-related parameters.[26]

Assumption 7 (Parametrics)**.** *The fundamental type distributions $F_\omega(\cdot)$ and $F_c(\cdot)$ are specialized as follows:*

[24] More formally, estimation under Assumption 6 turns on a preference orthogonality restriction: short-run changes in promotional realizations do not reflect short-run changes in preferences. This restriction could be implemented in at least three ways. First, one could invert market shares $\sigma(\mathbf{p}, \mathbf{d}; \xi_t)$ to recover product-market-level errors ξ_{jt}, and proceed based on assumed orthagonality of week-to-week *changes* in tastes and preferences (e.g. $(\Delta p_{jt}, \Delta d_{kt}) \perp \Delta\xi_{lt} \forall j, k, l$). Second, one could impose structural restrictions on the evolution of unobserved preferences ξ_t; for instance, assume a Markov chain or higher-order autoregressive process. Finally, one could simply assume that preferences ξ_t are constant over the period in question. I pursue the latter here, but plan to explore the other approaches in future work.

[25] While one could in principle incorporate time-varying ξ_t, I here assume fixed ξ for several reasons. First, since I pursue estimation at the store-week level, the large-market justification for standard BLP demand inversion may not hold (in particular, zero shares are common). Second, calculating market shares $\sigma(\mathbf{p}, \mathbf{d}; \xi_t)$ is computationally costly, and assuming fixed ξ significantly reduces the number of such calculations required. Finally, other work in similar contexts (e.g. Hendel and Nevo (2006)) also assume fixed ξ, so this assumption helps to maintain comparability.

[26] An interesting question I do not address here is whether the search-plus-demand model derived above is nonparametrically identified. Berry and Haile (2014, 2010) have established that the standard BLP-style discrete choice structure is nonparametrically identified under relatively weak index restrictions on the underlying preference model. The presence of search and display effects renders the corresponding analysis more complicated for the model considered here. My conjecture at present is that the full search-plus-demand model may not be point-identified, but that estimation based on general functional forms could in principle proceed a long way toward flexible implementation.

1. *Idiosyncratic attribute-level tastes are joint normally distributed and IID across consumers:* $\beta_i \sim N(\bar{\beta}_x, \Sigma_x)$ *for all* i.

2. *Search costs are distributed across the population as follows:*

   - *Fraction* λ *of consumers are shoppers: have zero search cost and always search.*
   - *The remaining fraction* $(1-\lambda)$ *of consumers have positive search costs drawn from a log-normal distribution:* $\log(c_i) \sim N(\mu_c, \sigma_c^2)$.

3. *Price sensitivity is constant across consumers:* $\alpha_i \equiv \alpha$ *for all* i.

The assumption that $\beta_i \sim N(\bar{\beta}_x, \Sigma_x)$ is standard in the literature; see, e.g., BLP (1995) and many subsequent studies. The assumption that a fraction of consumers always search is common in the sales literature (see, e.g., Pesendorfer (2002) and Chevalier and Kashyap (2011)) and is particularly attractive in my case since it nests the standard full-information demand model. Assuming a constant price sensitivity α is somewhat more restrictive, but follows Hendel and Nevo (2006) and simplifies computation. For notational compactness, I index free preference parameters via the vector $\theta \equiv \{\alpha, \bar{\beta}_x, \Sigma_x, \mu_c, \sigma_c^2, \lambda, \xi\}$.

Since I pursue estimation at the store level rather than the city or regional level, standard large-sample equivalence beteween market shares and empirical purchase frequencies may not apply. I therefore model purchase outcomes in each week as a size-N_t draw from a multinomial distribution characterized by probability vector $\sigma(\cdot;\theta)$, where N_t is a proxy for number of consumers "in the market" in week t (to be formalized below). My final store-level log-likelihood function is thus

$$\ln L(\theta) = \sum_{t=1}^{T} \left[\ln \frac{N_t!}{q_{ot}! \cdots q_{Jt}!} + \sum_{j=0}^{J} q_{jt} \ln(\sigma_j(\mathcal{M}_t;\theta)) \right]. \tag{2}$$

where $\mathcal{M}_t = \{\mathbf{p}_t, \mathbf{d}_t, \mathbf{f}_t\}$ are marketing outcomes in week t, $\mathbf{q}_t$ is the vector of units sold in week t, and $\sigma_j(\mathcal{M}_t;\theta)$ are predicted purchase probabilities at realization $\mathcal{M}_t$ and parameters θ.[27]

[27] One potential problem with maximum likelihood this context is that I must ultimately simulate the likeli-

5.1.1 Constructing N_t

Unfortunately, actual store-level market size is unobserved in practice: by nature, scanner data contains information on quantities sold, but not on number of consumers visiting a store. Consequently, I consider two distinct approaches to estimation, which together span the set of plausible measures of N_t. The first of these, labeled *narrow* N_t, takes market size to be the total units of detergent sold (both liquid and solid) in the target store in week t. This approach probably understates the true number of potential customers in the marketplace, but may be a plausible approximation for the informational reach of displays.[28] The second, labeled *broad* N_t, combines estimates derived from IRI-household-level panel surveys with scanner data on total purchases in the target store to derive a prediction for market size in week t. This approach almost certainly overstates the number of consumers actively shopping for detergent, but likely represents a reasonable proxy for total number of consumers visiting the store. Encouragingly, both proxies yield qualitatively similar estimates, which suggests that results are not overly sensitive to definition of N_t. The rest of this subsection gives details on the construction of my second (broad) proxy.

As noted in Section 2, the IRI dataset contains two primary types of data: store-level scanner data for 47 geographic markets, and BehaviorScan household-level panel data for two selected markets (Eau Claire, Wisconsin and Pittsfield, Massachusetts). The BehaviorScan survey includes information on both household-level purchases and locations purchased, so I can construct a dataset relating the total number of BehaviorScan households visiting each store in each week to total purchases (across all product categories) by BehaviorScan households in that store for that week. I then reverse this relationship via the following predictive regression model:

$$trips_{st} = \beta_1 \cdot qcat_{1,st} + \beta_2 \cdot qcat_{2,st} + \cdots + \beta_{30} \cdot qcat_{30,st} + e_{st}, \tag{3}$$

hood function, and simulated maximum likelihood estimates are guaranteed to be consistent only as simulation size approaches infinity. I recognize this as a potential problem, and plan to explore estimation based on method of moments in future work.

[28] If anything, this approach should understate total display effects, since it ignores the possibility that displays increase effective market size.

Figure 3: Predicted vs actual visits, IRI store 652159

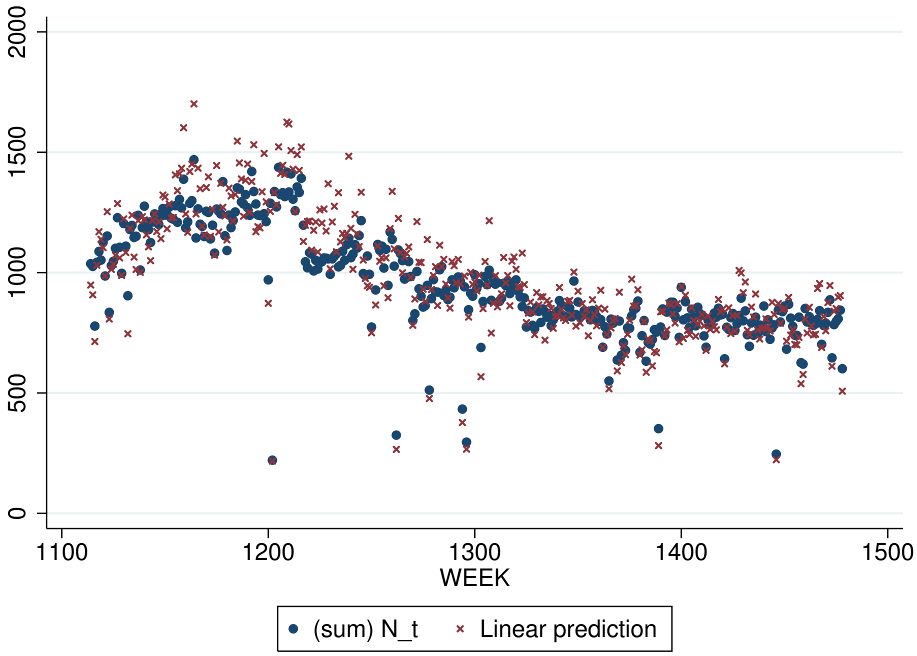

where $trips_{st}$ is the total number of BehaviorScan households visiting store s in week t, and $qcat_{k,st}$ is the total number of units purchased (across all UPCs) in IRI product category k by these households at store s in week t. I then use the estimated coefficients $\hat{\beta}$ plus data on total purchases in each target store to generate a proxy for number of store visits.[29] Figure 12 in Appendix 1 illustrates the predicted N_t resulting from this procedure.

Encouragingly, the simple predictive regression (3) fits the BehaviorScan data very well: $R^2 > 0.98$. This fact is illustrated in Figure 3, which plots predicted versus actual visits to IRI store 652159, the store with the most observed trips in the BehaviorScan sample. Since household purchases are likely to be fairly predictable in aggregate, this finding is not particularly surprising, but it does build confidence that the N_t proxy thus constructed provides a reasonable estimate for trips to the target store.

[29]Obviously, N_t is an integer, so one could also use a count data model. In this case, however, N_t is large (more than 5000 in all periods), rounding is unlikely to be important. Further, a simple regression model fits the data very well. Thus a more complex approach would seem unnecessary.

5.1.2 Product aggregation

To simplify computation of market shares, I aggregate products to the brand-size level in analysis.[30] This is both because of the very large number of UPC-level products in the market (more than 400 in the Atlanta sample), and because many of these products are likely to be very close substitutes in practice (e.g. same size and brand but different scents). I aggregate prices and displays within groups as follows. First, to best reflect final prices paid, I define the weekly price for each group as the weekly quantity-weighted average of UPC-level prices for products in the group.[31] Figure 4 illustrates this procedure for the Tide 100oz group (15 UPCs); Figures 10 and 11 in Appendix 1 give corresponding plots for Purex 100oz and All 100oz groups. As evident from these figures, prices within brand-size groups track very closely across virtually all products in the market, with the sales-weighted average closely following the minimum. This suggests that aggregation of prices to the brand-size level involves minimal loss of precision. Second, consistent with my fundamental hypothesis that promotions convey price information, I take the *maximum* display or feature status among UPCs in a product group as the aggregate display or feature status for each week. This approach is corroborated by the close correlation in prices noted above, and should if anything bias my findings away from the effects of interest.[32]

Table 6 in Appendix 1 presents market-level summary statistics for the brand-size product aggregates defined above. By construction, displays and features occur somewhat more frequently in the aggregated sample than in the UPC-level data, but summary variables otherwise look very similar. For completeness, Table 8 in Appendix 1 replicates the descriptive regressions in Section 3 on brand-size product aggregates; coefficient magnitudes change somewhat, but all qualitative patterns noted in Section 3 extend. The brand-size aggregates defined above thus seem to preserve the essential economic features of the underlying UPC-level sample.

[30] This is again similar to Hendel and Nevo (2006), though my focus on promotions leads to a slightly different aggregation procedure.

[31] Building on the intuition that products within brand-size groups are close subsitutes, I also explored taking the minimum UPC-level price within each group as the group price. Since the two price series track very closely, the two approaches are essentially equivalent in practice.

[32] In particular, UPC-level displays can never lead to a larger proportional increase in group-level quantity sold than in UPC-level quantity sold. Hence if displays matter only at the UPC level, my aggregation procedure will tend to dilute display effects.

Figure 4: Alternative price aggregates for Tide 100oz, IRI store 683960

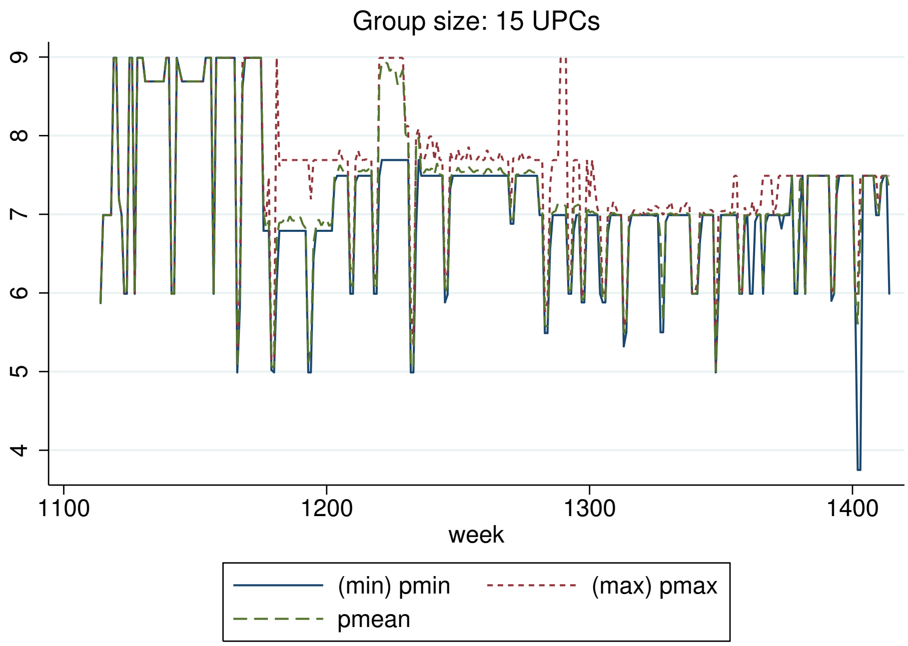

5.1.3 Sample refinement

Aggregating to the brand-size level reduces the dimensionality of the choice set considerably, but estimation using the full remaining sample is still computationally infeasible for most chains. Consequently, for each target chain, I refine the estimation sample as follows. First, to minimize the impact of assuming constant market-level preferences ξ, I focus on one-year subsets of the overall sample period. Baseline estimates are based on the midpoint years in the sample (2003-2006); robustness checks on other years yield similar results. Second, since many brand-size pairs are almost never purchased, I restrict attention to "important" brands, defined as the chain's private label plus the top ten brands by share of total liquid detergent units sold in the target chain within the sample period. This typically yields a set of between 15 and 25 "important" brand-sizes for each chain considered; due to high market concentration, these "important" products typically account for well over than 90 percent of total chain liquid detergent purchases. For each target chain, the baseline estimation sample is thus a set of "important" brand-size products for the calendar year 2004-2005.[33]

5.1.4 Incorporating features

My discussion thus far has focused on displays, which have a clear structural interpretation in the demand context I consider. The descriptive analysis in Section 3 suggests that feature advertisements also have large and significant effects on purchase outcomes, but their structural interpretation in a model of in-store choice is somewhat less obvious. Consequently, my baseline specification simply assumes that features shift the distribution of tastes in the market. This approach is based on the hypothesis that consumers with high tastes for a particular product may be more likely to enter the market in weeks that product is featured. I have also explored estimation under the assumption that features are equivalent to displays, an approach motivated by the possibility that features convey information to at least some consumers. However, preliminary estimates suggest that the baseline approach provides a substantially better fit, so my subsequent discussion concentrates on the baseline case.

5.2 Estimation

Computation of the market demand system $\sigma(\cdot)$ involves two interrelated complications: purchase outcomes must be simulated conditional on endogenous search, and simulation of search probabilities $\pi_s(\cdot)$ must account for the fact that consumers will select into search based on preferences for nondisplayed goods. To overcome these challenges, I implement estimation based on simulated maximum likelihood, where simulation is used to approximate two dimensions of uncertainty: over the characteristics of individuals in the market ($\omega_{i}{}_{i=1}^{I}$ drawn from the distribution of types implied by given parameters), and over the price uncertainty faced by each individual ($\{\mathbf{p}_r^S\}_{r=1}^R$ resampled from an approximation to $F_p(P^S|\mathbf{p}^d, \mathbf{d})$). Since in practice most price variation is induced by week-to-week sale discounts, I implement the price

[33]This final sample is obviously a relatively small subset of the IRI dataset. As noted above, this is primarily due to computation costs, which increase roughly linearly in the number of brand-store-weeks considered. Further, by design, my selection rules at every level are orthogonal to my questions of interest: first alphabetical IRI market, largest stores in this market, and sample midpoint in this store. I have also explored estimation based on other markets, chains, and years, with qualitatively similar results.

resampling component of this simulation based on the following specialization of Assumption 1:

Assumption 8. *Consumers know regular prices* $\mathbf{r}_t$ *but must search over sale realizations* $\Delta\mathbf{p}_t \equiv \mathbf{p}_t - \mathbf{r}_t$*. Further, the distribution of sale realizations* $F_{\Delta p}(\cdot)$ *prevailing in each chain is stable over the sample period.*

For simplicity, the resampling algorithm I employ here further assumes that sale realizations Δp_j are independent across products given realizations of $\mathbf{d}$. In particular, maintaining this assumption, I implement price resampling as follows. First, consistent with Assumption 8 and the economic motivation of the search model, I construct price discounts $\{\Delta\mathbf{p}_t\}_{t=1}^{T}$ for all store-week observations in each chain-level sample. Second, for each store-week observation $t \in \{1, ..., T\}$, I resample a set of price differences $\left\{\Delta p_r^S\right\}_{r=1}^{R}$ based on the following matching algorithm: first resample (without replacement) from periods s in which $\mathbf{d}_s = \mathbf{d}_t$, then resample (without replacement) from remaining periods s weighted by the number of elements in $\mathbf{d}_t$ which are also in $\mathbf{d}_s$, then resample (with replacement) from all periods randomly. Finally, conver this differenced sample $\left\{\Delta p_r^S\right\}_{r=1}^{R}$ to an unconditional price sample $\left\{p_r^S\right\}_{r=1}^{R}$ by setting $\mathbf{p}_r^S = \mathbf{r}_t^S + \Delta p_r^S$. Under the auxiliary assumption $\Delta p_j \perp \Delta p_k | \mathbf{d}$ noted above, as $T \to \infty$ this in turn will converge to a sample from $F_p(P^S|\mathbf{p}^d, \mathbf{d})$. This auxiliary assumption is likely to provide a reasonable approximation to consumer expectations, and significantly simplifies the resampling algorithm. It does, however, involve some loss of generality, and could be relaxed in future work.[34]

6 Results

This section reports results from applying three structural demand specifications to data on large chains in the Atlanta market. My main focus in this section is of course the search-plus-demand model developed above, where displays convey information as in Section 4 and

[34] In particular, studies of retail sales typically find that sale realizations are somewhat negatively correlated across competing products. One natural extension would be to model sale realizations using a copula structure. This would slightly complicate price resampling, but simulation could otherwise proceed as above.

features enter as preference shifters. For elasticity and model fit comparisons, I also estimate two standard full-information demand models: a naive specification ignoring promotions altogether, and a more sophisticated model *a la* Hendel and Nevo (2006) where displays and features enter as preference shifters. In all specifications, product attributes X_j include dummies for all brands, a bulk dummy for sizes above 125oz, and dummies for "premium" and "budget" brands, defined as those with average volume-weighted regular prices above \$1.00 and below \$0.60 per 16 ounces respectively.[35] Consumers have random preferences $\beta_i \sim N(0, \Sigma_x)$ over these attributes, where in my baseline specification I assume Σ_x diagonal as in BLP. As a reduced-form proxy for potential dynamic effects, all specifications also include 4 lags of units sold in the utility function.[36] Simulations in all cases are based on 300 preference draws and search simulations are based on 100 price draws.[37]

6.1 Structural parameters

Table 4 reports core structural parameters resulting from application of the baseline search specification to six large stores in the Atlanta market. Results in Table 4 are based on the narrow N_t proxy (total category purchases) defined in Section 5.1.1; corresponding estimates for my broad N_t proxy (total store visits) are given in Table **??** in Appendix 1. As above, α is price sensitivity, λ is the fraction of shoppers in the market, and γ is the mean of the exponential search cost distribution among non-shoppers. The parameters lag_1-lag_4 denote utility effects of lagged total purchases, where all lagged-purchase variables have been centered and rescaled in terms of store-level standard deviations. Finally, $ftaste$ coefficients give the estimated effect of a feature promotion represented as a utility shifter; as noted above, this representation is motivated by the fact that consumers with high preferences for a particular product may be more likely to enter the market when that product is featured. 95 percent

[35] Descriptive analysis suggests that volume-weighted regular prices are tri-modal, with modes at approximately \$0.50, \$0.70, and \$1.10 respectively. Products within a price class are likely to be perceived as closer substitutes than products of different price classes, motivating the specification given.

[36] This approach is similar to that of Pesendorfer (2002), whose reduced-form model incorporates dynamic effects via a lagged time-since-sale variable. I use lagged units sold rather than lagged time-since-sale to better reflect Hendel and Nevo (2006)'s insight that current product stocks should affect utility of additional purchase.

[37] Other simulation sizes give similar results; I settled on (300, 100) as a reasonable balance of accuracy and computation speed.

Table 4: Core structural parameters, large Atlanta chains (narrow N_t)

	Chain 78	Chain 100	Chain 140
α	0.387	0.444	0.449
	(0.385, 0.389)	(0.441, 0.446)	(0.446, 0.452)
μ_c	-1.83	-2.2	-0.382
	(-2.41, -1.25)	(-3.24, -1.17)	(-0.52, -0.244)
σ_c	1.4e-14	0.559	5.78e-08
	(-0.309, 0.309)	(0.103, 1.02)	(-0.142, 0.142)
λ	5.89e-12	0.000699	0.000654
	(-0.0707, 0.0707)	(-0.0867, 0.0881)	(-0.0582, 0.0595)
lag_1	0.0219	0.0783	0.0681
	(0.0197, 0.0241)	(0.0767, 0.0799)	(0.0656, 0.0706)
lag_2	0.0198	0.0756	0.0177
	(0.0175, 0.0221)	(0.074, 0.0772)	(0.0153, 0.0202)
lag_3	0.012	0.0681	0.0193
	(0.00986, 0.0141)	(0.0665, 0.0696)	(0.0165, 0.0221)
lag_4	0.0249	0.0515	0.0154
	(0.0234, 0.0263)	(0.0501, 0.0528)	(0.0132, 0.0176)
$ftaste$	0.392	0.762	0.54
	(0.389, 0.395)	(0.758, 0.766)	(0.534, 0.545)
Mean c_i	0.388	0.29	1.53

maximum likelihood confidence intervals are given in parentheses.[38]

Parameter values in Table 4 suggest the presence of substantial informational effects, though estimated search patterns differ somewhat across stores. The estimated population of shoppers (λ) is negligible in all specifications, suggesting that the baseline search-plus-purchase model provides a better fit to observed purchase behavior for virtually all consumers. In two of three chains (78 and 140), estimated search cost distributions are concentrated tightly around a point, while estimates of σ_c for the remaining chain (100) suggest somewhat more heterogeneity in search costs across consumers. Recalling that $\log(c_i) \sim N(\mu_c, \sigma_c^2)$ and standardizing estimated distributions of search costs by their corresponding price sensitivity parameters, the estimates in Table 4 suggest that mean search costs are approximately \$0.39 in Chain 78, \$0.29 in Chain 100, and \$1.53 in Chain 140 respectively.[39] Taken together, these estimates

[38] These standard errors are not corrected for potential simulation error. As the number of simulation draws is large, however, bias due to simulation is likely to be small.

[39] While the magnitudes of these estimates seem natural, my point estimates are somewhat lower than those reported by Seiler (2013), who reports mean costs of roughly \$3.00-\$4.50 for a consumer entering the market. My lower estimate is natural given that I explore the search intensity following market entry, whereas Seiler focus on the market entry decision itself.

Figure 5: Implied structural search cost distributions, narrow N_t

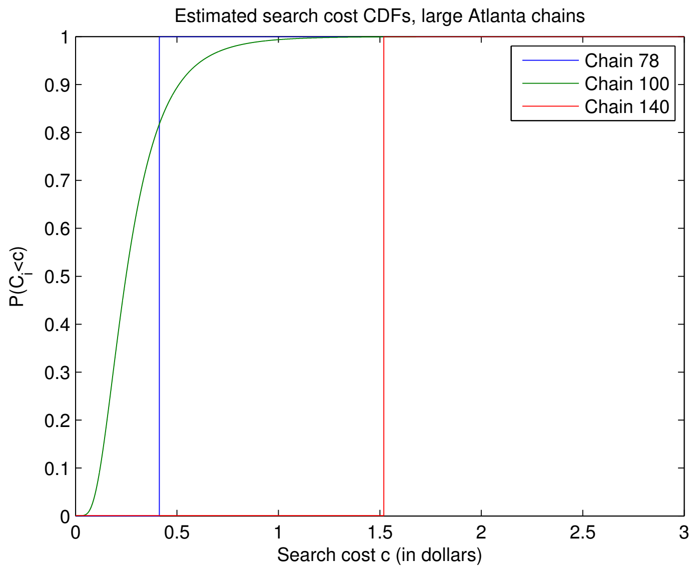

suggest the presence of substantial search effects, with the magnitude of these effects potentially heterogeneous across chains in the population.

Figure 5 (in text) and Tables 11 and 12 (in Appendix 1) give some additional interpretation of the estimated structural parameters. Figure 5 graphically represents the search cost distributions implied by the chain-level parameters in Table 4, where the shape of each distribution above zero is determined by the interaction between the location parameter μ_c and the scale parameter σ_c, and values on the horizontal axis have been normalized by α to express costs in dollar equivalents. This figure clearly indicates the patterns noted above: for Chains 78 and 140 the estimated search cost distributions are concentrated around their means, whereas for Chain 100 the estimated search cost distribution involves substantial dispersion. Meanwhile, Tables 11, and 12 summarize elasticities derived from the structural estimates: Table 11 presents own-price elasticities for all stores in the estimation sample, and Table 12 gives a full set of cross-price elasticities for one representative chain (100), where elasticities in each case are calculated relative to a no-promotion baseline period.

6.2 Model validation

To assess the performance of the structural search-plus-demand model, I compare empirical purchase frequencies with corresponding market shares predicted by the model. Figure 16-13 in Appendix 1 illustrate two examples of this validation exercise, plotting actual versus predicted market shares for the four most purchased brand-size pairs in the four Chain 78 stores with the most units sold in the sample period. On balance, these examples suggest the search-plus-demand model fits the data reasonably well; predicted shares typically closely match both average shares and promotion-induced spikes, though some store-level heterogeneity is evident in the figures. Thus while the underlying model is obviously stylized, it seems to provide a satisfactory match to the key patterns of interest in the data.

6.3 Elasticity comparisions

One key objective of this study is to explore how accounting for interactions between limited consumer information and in-store displays might influence elasticities derived from a structural demand model. To address this question, I estimate two full-information demand models for each of the six stores in my estimation sample: the first a naive specification ignoring promotions altogether, and the second incorporating promotions via preference dummies *a la* Hendel and Nevo (2006). I then compare structural price elasticities implied by these models with those based on the full search-plus-demand model above.

Table 5 presents one example of this cross-specification comparison; in this case, implied own-price elasticities for all specifications estimated for Chain 78. Consistent with the reduced-form patterns noted in Section 3, this table strongly confirms that failing to account for promotions can substantially bias demand analysis: the naive full-information model yields much larger elasticities than either comparison specification. In contrast to the reduced-form results, however, there is virtually no difference between own-price elasticities from the search model and own-price elasticities from the model with in which displays enter as preference shifters. While not formally reported here, very similar patterns appear in estimates of cross-price elasticities in Chain 78 and own- and cross-price elasticities in other chains estimated.

Table 5: Selected own-price elasticities, Chain 78

	Full Info		
	Naive	Promo FX	Search
AH100	2.12	1.37	1.37
ALL100	5.59	3.59	3.6
ALL200	8.14	5.24	5.25
FAB100	2.55	1.64	1.65
GAIN100	3.35	2.09	2.08
GAIN200	6.62	4.25	4.25
PUREX100	1.98	1.29	1.3
PUREX200	3.84	2.51	2.52
SURF100	3.54	2.28	2.28
TIDE100	8.28	5.13	5.16
TIDE200	13	8.23	8.27
WISK100	3.34	2.16	2.16

Notes: Elasticities simulated relative to baseline market with no displays, features, or promotions.

Thus while preference dummies cannot shed light on search behavior *per se*, the analysis here suggests that they may be practically sufficient for many questions not directly involving search.

6.4 Displays, in-store search, and consumer behavior

I next apply my structural search-plus-demand model to explore three counterfactual questions of interest in informational economics. First, how do limited attention, price promotions, and in-store displays interact in shaping consumer purchase behavior? Second, how do displays affect market cannibalization? I.e. what fraction of the market share increase due to a displayed price reduction comes from substitution within a brand versus substitution between brands? Third, how large are the costs of limited information – in terms of both foregone purchase opportunities and missed price promotions – to consumers in the Atlanta detergent market? Finally, interpreted as a form of informational advertising, how important are in-store displays in alleviating these information costs? Building on the structural search-plus-demand estimates obtained above, I propose answers to each of these in turn.

I first explore how limited attention, displays, and price promotions interact in shaping consumer behavior. More precisely, how much of the effect of a displayed price reduction is due to the price, how much to the display, and how much to changes in search behavior induced by the interaction of both? To answer this question, I decompose the overall effect of a displayed price reduction into three channels: an effect on search, an effect on purchase given search, and an effect on purchase given no search. Figure 6 plots the results of this exercise for two representative products (Tide 50oz and Tide 100oz in Chain 78), where counterfactuals are relative to a baseline market with no other displays or sales. This figure in turn suggests three interesting patterns. First, as expected, displays induce both a level shift in quantity sold and an increased quantity response to price reductions. This can be seen by comparing the overall purchase probability $\Pr(buy\, j)$ with the corresponding no-display baseline $\Pr(buy\ j\ if\ no\ D)$: the intercept difference reflects a level shift, and the slope difference reflects an increase in price responsiveness. Second, endogenous search tends to amplify the purchase effects of a displayed price reduction: search probability decreases as the sale discount increases, thereby shifting more consumers toward the displayed product. Finally, the choice of which product within a brand to display substantially influences consumers' willingness to search remaining products, and thereby the effectiveness of displays in shaping purchase outcomes. In particular, displaying a relatively popular size within a brand (Tide 50oz) generates almost twice the search response as displaying a relatively less popular brand (Tide 100oz), even in the absence of any price reduction. All three effects are consistent with the stylized patterns noted in Section 3, and confirm the importance of information frictions in analysis of retail competition and consumer choice.

Given that the choice of product to display can substantially affect both purchase incidence and search behavior, I next turn to explore the market cannibalization effect of displays: specifically, is a display mainly inducing substitution from other products in the same brand, or from consumers who otherwise would have purchased another product? As a first take on this question in the Atlanta detergent market, I next decompose the market share gain induced by displaying a price reduction – that is, the change in j's market share induced by switching from a nondisplayed sale to a displayed sale – into three components: the fraction of consumers

Figure 6: Counterfactual simulation: Search vs purchase effects

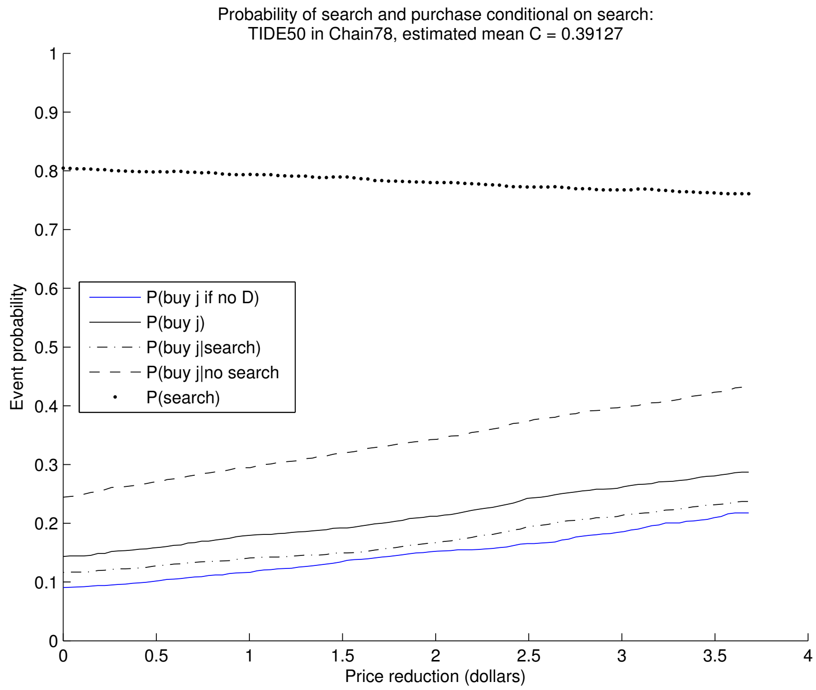

(a) Search and purchase probabilities, brand-leading product displayed

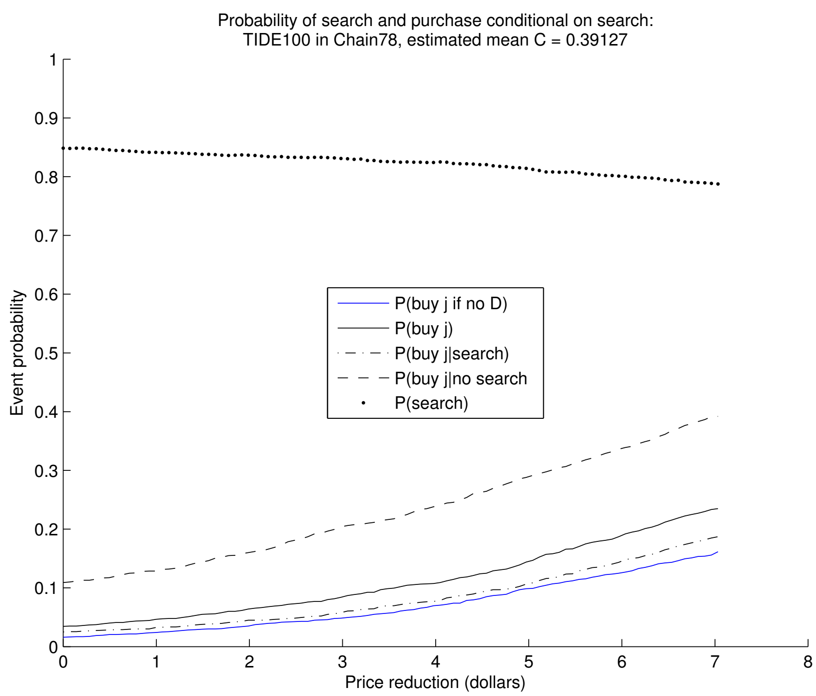

(b) Search and purchase probabilities, non-leading product displayed switching from other products of the same brand as j, the fraction of consumers switching to j from other brands, and the fraction of consumers switching to j from the outside good. Figure 7 presents results of this procedure for the two products analyzed above in Chain 78: Panel (a) decomposes market share changes induced by display of Tide 50oz (the Tide brand leader in Chain 78), and Panel (b) decomposes changes induced by display of Tide 100oz (the second most popular Tide product in Chain 78). Taken together, these figures highlight a very interesting pattern: whereas substitution from other Tide products is relatively similar for Tide 50oz and Tide 100oz (about 0.5 percent of consumers switching after display in both cases), substitution from both other brands and from the outside good is much larger for Tide 50oz than for Tide 100oz. In other words, from a strategic perspective, display of more popular products can substantially amplify market share gains from cross-brand predation relative to market share losses in other products of the same brand. This in turn helps to rationalize the stylized empirical fact that more popular products are more likely to be displayed.

Finally, I turn to a pair of questions of longstanding interest in information economics: how large are the potential welfare costs of limited information in the Atlanta detergent market, and how effective are in-store displays – interpreted as a form of informational advertising – in alleviating these? To answer these questions, I consider the following counterfactual exercise. Using my baseline structural estimates, I compute predicted consumer welfare under three alternative informational scenarios: the *empirical baseline*, in which consumer preferences and search costs are as estimated above and in-store pricing, promotional, and display activities are as observed in the data, the *no-display counterfactual*, in which consumer primitives and store pricing policies are held fixed but stores are counterfactually assumed never to use in-store displays to promote detergent, and the *full-information counterfactual*, in which consumer brand preferences and store marketing activities are held fixed but all consumers are assumed to be shoppers (observe all prices at zero cost) in the sense defined above.[40] This exercise is obviously only partial in the sense they consider neither equilibrium price responses (e.g.,

[40] To avoid biasing this comparison by the fact that consumers redraw idiosyncratic utility shocks upon search, I assume that consumers in the full information model draw separate idiosyncratic utility shocks for displayed and nondisplayed versions of the same product. This ensures that consumers maximize over the same number of idiosyncratic utility draws under both counterfactual alternatives, thereby permitting meaningful comparison of corresponding utility levels.

Figure 7: Decomposing the sources of display-related gains

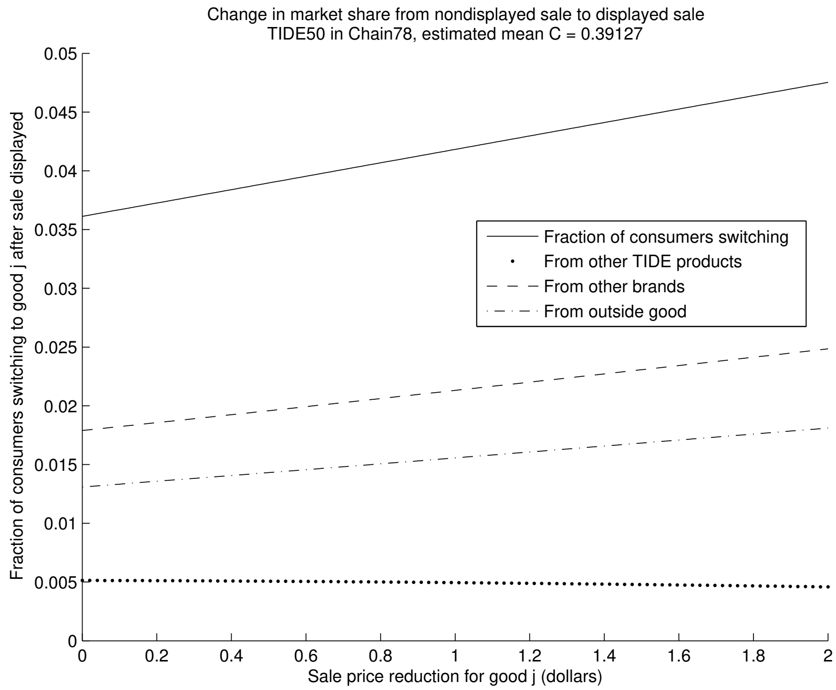

(a) Brand-leading product displayed

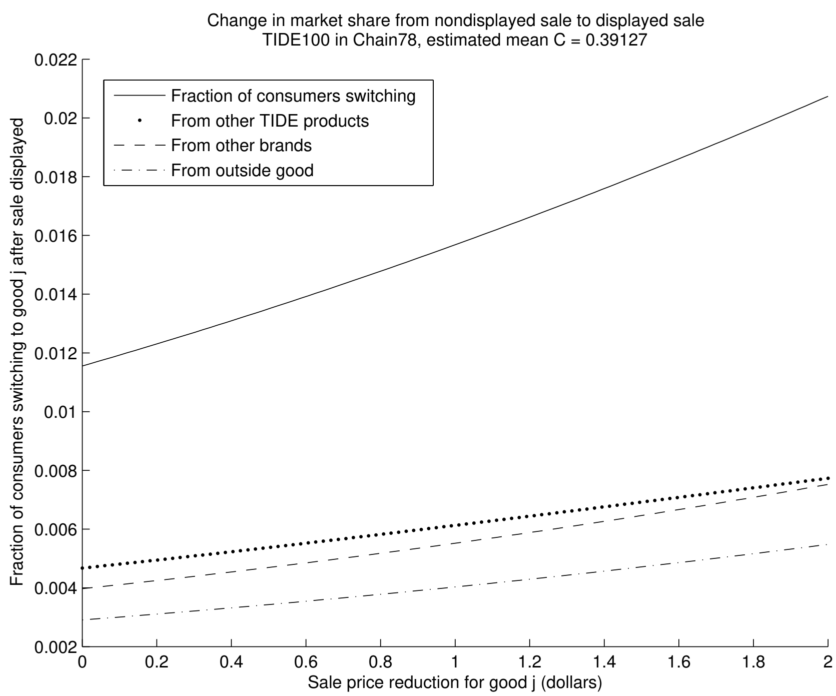

(b) Non-leading product displayed

detergent firms could substitute toward price promotions in the absence of displays) or potential gains from alternative uses of floor space (e.g. replacing end-aisle detergent displays with end-aisle displays for products in other categories). Bearing these caveats in mind, however, it is hoped that the underlying counterfactual will be informative.

The main results from this counterfactual exercise are as follows. First, information frictions generate substantial losses in this market: moving from the full-information counterfactual to the empirical baseline leads to an in-category welfare loss of roughly 5.4 percent. However, these welfare losses – although nontrivial – are substantially ameliorated by in-store displays: welfare in the no-display counterfactual is more than 7.5 percent below that in the full-information counterfactual, more than 2 percentage points lower than the empirical baseline. In other words, informative advertising (in the form of in-store displays) leads to roughly 30 percent smaller welfare losses than would arise absent such advertising. These findings underscore both the importance of information frictions in this market and the potentially important role of informative advertising in alleviating such frictions.

7 Conclusion

Motivated by several prominent features of retail markets, this paper develops a structural model of consumer choice in an environment with informative in-store displays and costly price search. This model is then applied to obtain structural search-plus-choice estimates for a representative target market, using short-run sale-induced price variation to recover preference parameters and short-run display-induced informational variation to recover search parameters. Consistent with preliminary reduced-form analysis, structural parameter estimates imply substantial search effects, with estimated search costs between approximately \$0.30 and approximately \$1.50 across chains. Displays and promotions are also shown to have important effects on estimated elasticities, though the results presented here suggest these can be accounted for in large part using preference dummies. Finally, I use my structural results to explore potential interaction between search and purchase decisions and to simulate counterfactual effects of price, feature, and display promotions.

While the results presented here are suggestive, there is still room to refine several dimensions of the underlying structural model. First, while I here aggregate products to the brand-size level, IRI scanner data in fact provides display and feature information at the (much narrower) UPC level. Reduced-form evidence suggests that aggregation masks substitution within a brand-size induced by displays, and estimation exploiting this additional information (for instance, via a nested logit structure) could permit more precise decompositions of the search and purchase channels explored above. Finally, while one objective this paper was to develop a model estimable using only scanner data, one could also integrate an application to IRI household-level panel data. This should both allow direct observation of no-purchase outcomes and permit a considerably richer household-level preference model.

Dynamics represent another fruitful avenue for future exploration. Notably, while Hendel and Nevo (2006) find that intertemporal substitution is an important channel by which consumers respond to price discounts, Seiler (2013) obtains much smaller intertemporal purchase responses in a model with costly search. This finding suggests that accounting for information frictions reduces the scope for potential dynamic effects, which in turn supports the static approach considered here.[41] Nevertheless, formal modeling of potential dynamic effects would be valuable both in terms of robustness and to deepen our understanding of the relationship between search and other possible channels for promotional effect.

Finally, my work thus far raises broader questions about the economics of search in retail markets. The most obvious of these is also the simplest: what do consumers actually know? Motivated by empirical patterns in retail markets, I extend the classical full-information demand paradigm to a simple model of rational search. This approach is distinct from standard marketing models of consideration, which typically take more of a reduced-form behavioral approach.[42] The practical interplay between "search" in the Stigler (1961) sense and "consideration" in the marketing sense is an interesting topic worth further study. The structural search model explored here may also have implications for understanding the supply side of retail markets: sales are frequently interpreted as a means for retailers to discriminate between

[41] I thank Stephan Seiler for clarifying my interpretation on this point.

[42] See, e.g., Goeree (2008) for an example of a pure consideration model in the economics literature; Santos et al. (2012) directly compare several leading search models in the context of online book markets.

consumers with different willingness to search, and displays may well serve a similar function.

References

Andrews, R.L. and TC Srinivasan, "Studying Consideration Effects in Empirical Choice Models Using Scanner Panel Data," *Journal of Marketing Research*, 1995, *32* (1), 30–41.

An Empirical Model of Search with Vertically Differentiated Products

An Empirical Model of Search with Vertically Differentiated Products*, RAND Journal of Economics, 2011, 42 (4), 729 – 757.*

Berry, Stephen T. and Philip A. Haile*, "Identification in Differentiated Products Markets Using Market Level Data," Econometrica, 2014, 82 (5), 1749 – 1797.*

Berry, Steven, James Levinsohn, and Ariel Pakes*, "Automobile Prices in Market Equilibrium," Econometrica, July 1995, 63 (4), 841–890.*

_ , _ , and _ *, "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," The Journal of Political Economy, February 2004, 112 (1), 68–105.*

Berry, Steven T. and Philip A. Haile*, "Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers," 2010, NBER Working Paper Series.*

Bronnenberg, B.J., M.W. Kruger, and C.F. Mela*, "The IRI Marketing Data Set," Marketing Science, 2008, 27 (4), 745–748.*

Burdett, Kenneth and Kenneth L. Judd*, "Equilibrium Price Dispersion," Econometrica, July 1983, 51 (4), 955–969.*

Butters, Gerard R.*, "Equilibrium Distributions of Sales and Advertising Prices," The Review of Economic Studies, October 1977, 44 (3), 465–491.*

Chevalier, J.A., A.K. Kashyap, and P.E. Rossi, *"Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," The American Economic Review, 2003, 93, 15–37.*

Chevalier, J.A. and A.K. Kashyap, *"Best Prices," 2011, NBER Working Paper Series.*

Eichenbaum, M., N. Jaimovich, and S. Rebelo, *"Reference Prices and Nominal Rigidities," American Economic Review, 2011, 101 (1), 234–262.*

Glandon, Phllip J., *"The Effect of Sales on Average Price Paid," 2011, Working Paper, Vanderbilt University.*

Goeree, Michelle Sovinsky, *"Limited Information and Advertising in the U.S. Personal Computer Industry," Econometrica, 2008, 76 (5), 1017–1074.*

Griffith, R., E. Leibtag, A. Leicester, and A. Nevo, *"Consumer Shopping Behavior: How Much Do Consumers Save?," The Journal of Economic Perspectives, 2009, 23 (2), 99–120.*

Hauser, John R. and Birger Wernerfelt, *"An Evaluation Cost Model of Consideration Sets," The Journal of Consumer Research, March 1990, 16 (4), 393–408.*

Hendel, Igal and Aviv Nevo, *"Measuring the Implications of Sales and Consumer Inventory Behavior," Econometrica, November 2006, 74 (6), 1637–1673.*

Hong, Han and Matthew Shum, *"Using Price Distributions to Estimate Search Costs," The RAND Journal of Economics, July 2006, 37 (2), 257–275.*

Hortacsu, Ali and Chad Syverson, *"Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds*," Quarterly Journal of Economics, Feb 2004, 119 (2), 403–456.*

Hosken, D. and D. Reiffen, *"Pricing Behavior of Multiproduct Retailers," BE Journal of Theoretical Economics, 2007, 7 (1).*

Kehoe, P. and V. Midrigan, *"Sales and the Real Effects of Monetary Policy," 2007, Working Paper.*

Konieczny, J.D. and A. Skrzypacz, *"Search, Costly Price Adjustment and the Frequency of Price Changes–Theory and Evidence," 2004, Working Paper.*

Koulayev, Sergey, *"Search for Differentiated Products: Identification and Estimation," RAND Journal of Economics, 2014, 45 (3), 553–575.*

Lu, Zhentong, *"Estimating Multinomial Choice Models with Unobserved Consideration Sets," Working paper, Shanghai University of Finance and Economics, 2016.*

Mehta, Nitin, Surendra Rajiv, and Kannan Srinivasan, *"Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation," Marketing Science, January 2003, 22 (1), 58–84.*

Mojir, N., K. Sudhir, and A. Khwaja, *"Spatiotemporal Search," 2014, Working Paper, Yale University.*

Moraga-Gonzalez, J.L., Z. Sandor, and M.R. Wildenbeest, *"On the Identification of the Costs of Simultaneous Search," 2010.*

Moraga-Gonzalez, Jose Luis and Matthijs R. Wildenbeest, *"Maximum Likelihood Estimation of Search Costs," European Economic Review, Jul 2008, 52 (5), 820–848.*

Nakamura, Emi and Jon Steinsson, *"Five Facts about Prices: A Reevaluation of Menu Cost Models*," Quarterly Journal of Economics, February 2011, 123 (4), 1415–1464.*

Pesendorfer, Martin, *"Retail Sales: A Study of Pricing Behavior in Supermarkets," The Journal of Business, January 2002, 75 (1), 33–66.*

Pinna, Fabio and Stephan Seiler, *"Estimating Search Benefits from Path-Tracking Data: Measurement and Determinants," Marketing Science, 2016, Conditionally accepted.*

Pires, Tiago, *"Consideration Sets in Storable Goods Markets," 2012, Working Paper, University of North Carolina.*

— , *"Measuring the Effects of Search Costs on Equilibrium Prices and Profits," 2014, Working Paper, University of North Carolina.*

Rob, Rafael, *"Equilibrium Price Distributions," The Review of Economic Studies, July 1985, 52 (3), 487–504.*

Roberts, John H. and James M. Lattin, *"Development and Testing of a Model of Consideration Set Composition," Journal of Marketing Research, November 1991, 28 (4), 429–440.*

Salop, Steven and Joseph Stiglitz, *"Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," The Review of Economic Studies, October 1977, 44 (3), 493–510.*

Santos, B.I., A. Hortacsu, and M.R. Wildenbeest, *"Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior," American Economic Review, 2012, 102 (6), 2955–2980.*

Seiler, S., *"The Impact of Search Costs on Consumer Behavior: a Dynamic Approach," Quantitative Marketing and Economics, 2013, 11 (2), 155–203.*

Sorensen, Alan T., *"Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," Journal of Political Economy, 2000, 108 (4), 833–850.*

— , *"An Empirical Model of Heterogeneous Consumer Search for Retail Prescription Drugs," October 2001, NBER Working Paper Series (8548).*

Stigler, George J., *"The Economics of Information," The Journal of Political Economy, June 1961, 69 (3), 213–225.*

van Nierop, E., B. Bronnenberg, R. Paap, M. Wedel, and P.H. Franses, *"Retrieving Unobserved Consideration Sets from Household Panel Data," Journal of Marketing Research, 2010, 47 (1), 63—74.*

Varian, Hal R., *"A Model of Sales," The American Economic Review, 1980, 70 (4), 651–659.*

Appendix 1: Supplemental tables and figures

Table 6: Liquid laundry detergent sales, estimation stores

VARIABLES	Price / oz	Sale	Feat	Disp	Discount	Share
ALL	0.0238	0.279	0.111	0.0773	0.114	0.143
ARMHAMMER	0.0165	0.290	0.0455	0.0436	0.126	0.124
CHEER	0.0343	0.190	0.0661	0.0335	0.141	0.0578
ERA	0.0252	0.212	0.127	0.0763	0.128	0.0115
FAB	0.0213	0.351	0.00771	0.0129	0.159	0.0120
GAIN	0.0272	0.235	0.146	0.0600	0.190	0.0884
PUREX	0.0166	0.272	0.0788	0.0586	0.137	0.150
SURF	0.0267	0.242	0.103	0.113	0.186	0.0217
TIDE	0.0335	0.252	0.158	0.0820	0.187	0.295
WISK	0.0287	0.350	0.163	0.113	0.149	0.0685

Notes: *Sale*, *Feat*, and *Disp* are UPC-by-week indicators for sale, feature, and display promotions. *Disc if sale* represents average discount from "regular" price in weeks a sale occurs, where regular price series are constructed as in Section 2.2.2.

Figure 8: Price history for Purex 100oz, IRI store 683960 (2002-2008)

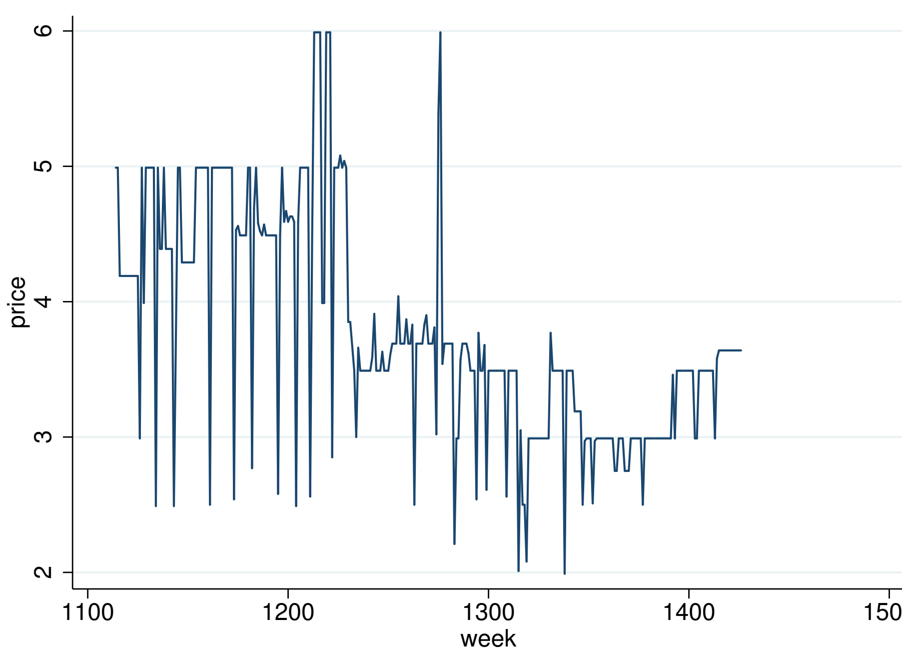

Figure 9: Price history for All 100oz, IRI store 683960 (2002-2008)

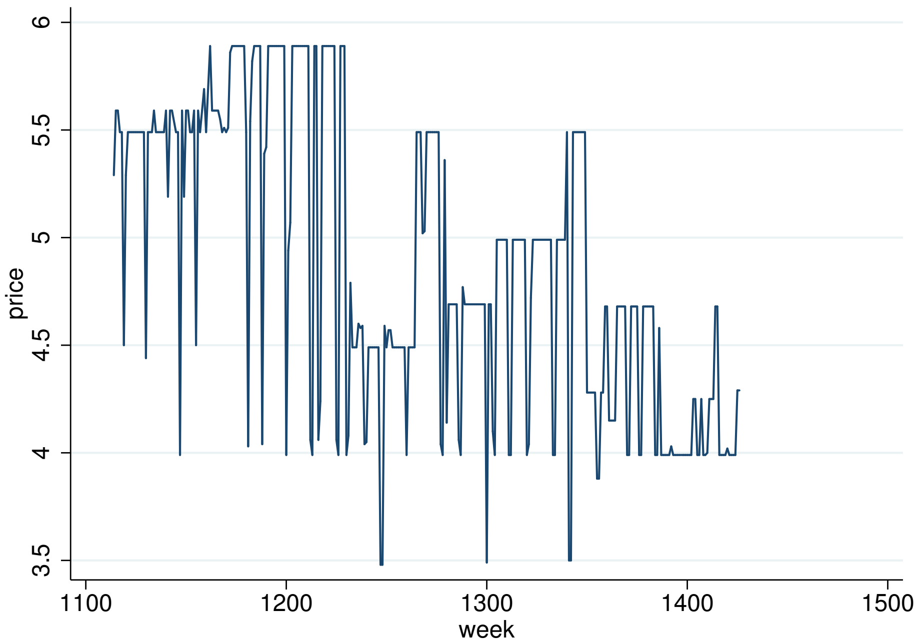

Figure 10: Alternative price aggregates for Purex 100oz, IRI store 683960

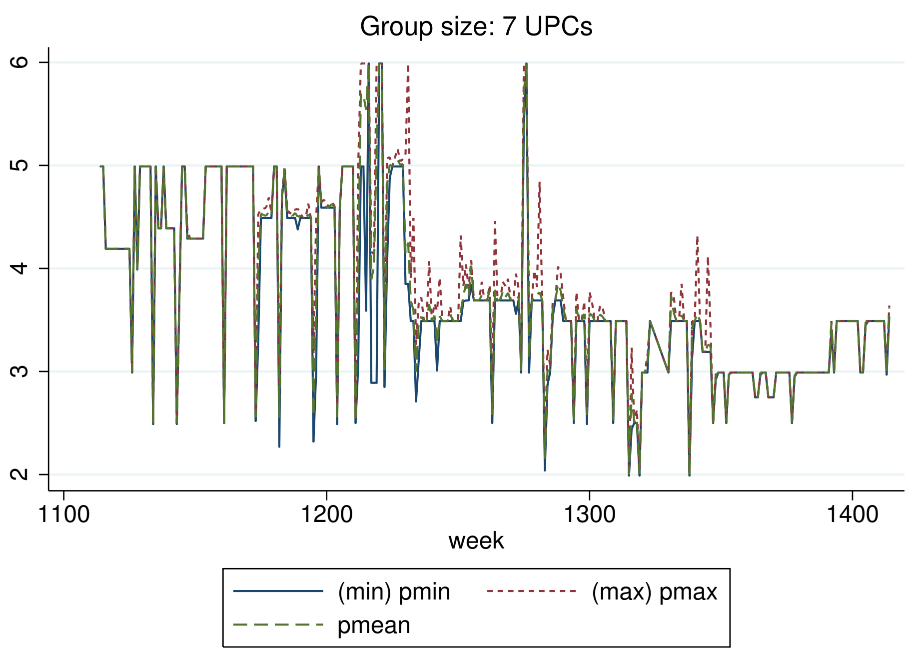

Figure 11: Alternative price aggregates for All 100oz, IRI store 683960

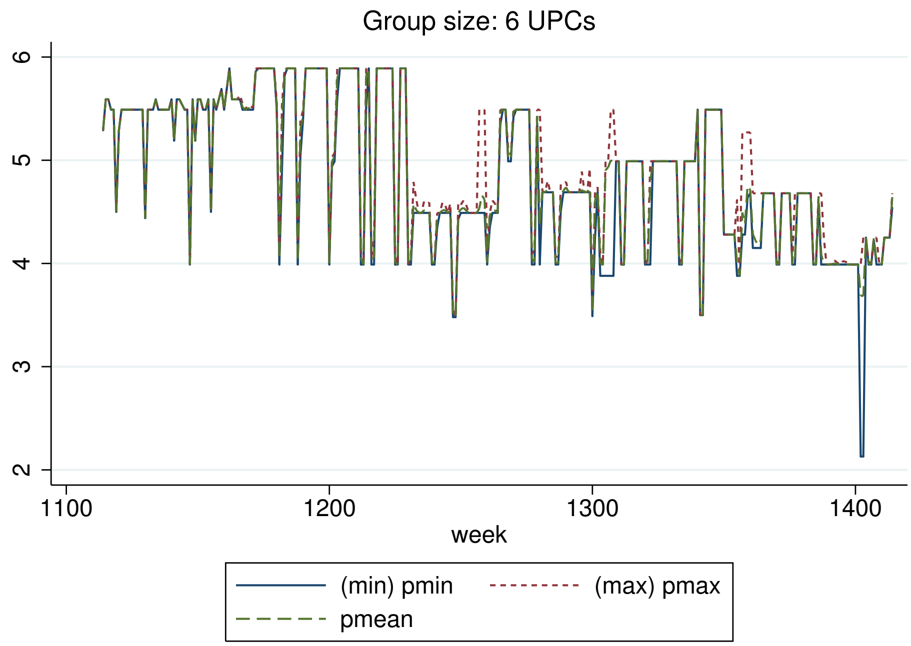

Figure 12: Predicted N_t for IRI store 683960

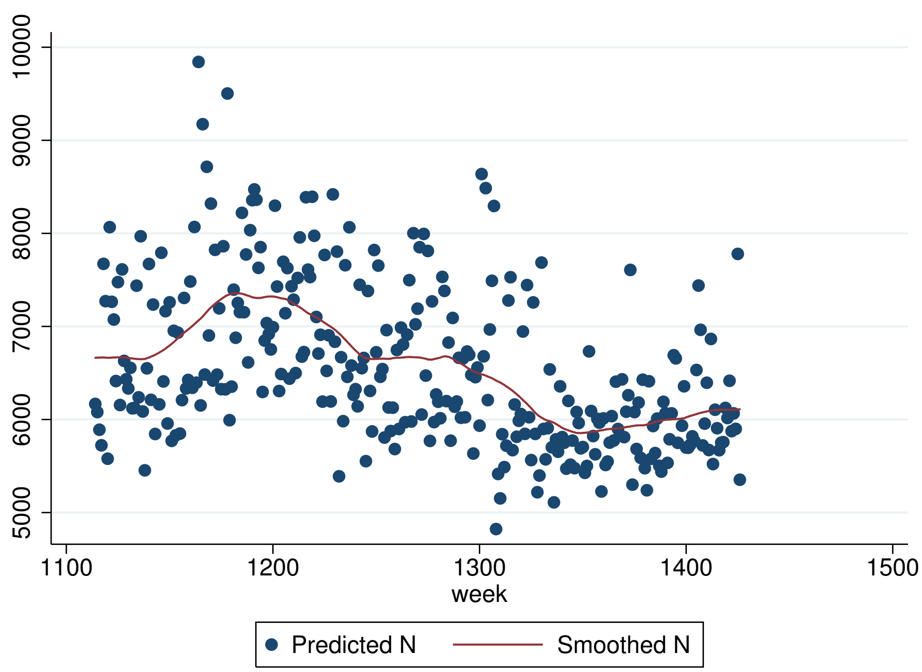

Table 7: Promotional effects in Atlanta detergent markets, alternative specifications

	Other outcomes					Regprice spec			
VARS	qnorm	wkshare	units	lnunits	VARS	Prices only	Dummies	Interacts	All channels
price	-0.00912***	-0.00148***	-0.453***	-0.0729***	regprice	-0.00883***	-0.00931***	-0.00922***	-0.00894***
	(0.000328)	(9.63e-06)	(0.00236)	(0.000249)		(0.000347)	(0.000334)	(0.000330)	(0.000330)
pgap	-1.144***	-0.0181***	-4.839***	-0.509***	pgap	-3.326***	-2.288***	-1.268***	-1.199***
	(0.0423)	(0.00144)	(0.331)	(0.0189)		(0.0517)	(0.0419)	(0.0351)	(0.0425)
sale	0.0528***	-0.000866***	0.0155	0.0600***	sale	0.147***	-0.0287***		0.0540***
	(0.00621)	(0.000217)	(0.0463)	(0.00319)		(0.00839)	(0.00594)		(0.00633)
feat	0.323***	-0.00304***	0.742***	0.245***	feat		0.736***		0.324***
	(0.0155)	(0.000672)	(0.114)	(0.00682)			(0.00687)		(0.0155)
disp	0.304***	0.0109***	3.054***	0.388***	disp		0.563***		0.304***
	(0.0151)	(0.000528)	(0.108)	(0.00613)			(0.00735)		(0.0153)
pgapxfeat	-2.286***	-0.122***	-15.41***	-0.774***	pgapxfeat			-3.471***	-2.291***
	(0.0951)	(0.00391)	(0.750)	(0.0362)				(0.0639)	(0.0957)
pgapxdisp	-2.204***	-0.0893***	-18.33***	-0.586***	pgapxdisp			-3.063***	-2.199***
	(0.129)	(0.00463)	(1.009)	(0.0432)				(0.108)	(0.130)
Constant	-0.146	0.0137	2.555	1.759	Constant	0.0871	-0.164	-0.0768	-0.150
		(6.383)		(393.6)		(406.4)	(304.3)		
R-squared	0.303	0.441	0.309	0.391	R-squared	0.201	0.282	0.292	0.303

Robust standard errors in parentheses

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

Table 8: Brand-size level promotion effects, Atlanta market (2002-2007)

VARIABLES	Prices only	Promo dummies	Only interacts	All channels
regprice	-0.00144**	-0.00596***	-0.00474***	-0.00584***
	(0.000731)	(0.000710)	(0.000714)	(0.000708)
pgap	7.030***	5.262***	3.349***	3.282***
	(0.325)	(0.289)	(0.146)	(0.282)
sale	0.132**	-0.186***		0.0619
	(0.0640)	(0.0501)		(0.0496)
feat		1.331***		0.542***
		(0.0348)		(0.0711)
disp		0.857***		0.677***
		(0.0230)		(0.0414)
pgapxfeat			5.966***	4.076***
			(0.297)	(0.438)
pgapxdisp			3.667***	1.699***
			(0.375)	(0.410)
Constant	-0.186***	-0.198***	-0.154***	-0.187***
	(0.00626)	(0.00602)	(0.00674)	(0.00603)
Observations	278,491	278,491	278,491	278,491
R-squared	0.150	0.196	0.193	0.203

Notes: Products aggregated to brand-size level. Dependent variable is *qnorm*, pct by which store-product-week quantity sold exceeds average weekly quantity sold. *sale*, *disp*, and *feat* are store-product-week promo indicators, *discount* is pct price below regprice, and *discxdisp* and *discxfeat* are interaction terms. Robust standard errors in parentheses.

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

Table 9: Brand-size level promotion effects by store, whole sample (2002-2007)

	Store-level regressions						
VARIABLES	653369	250094	263568	266100	243785	683960	Market
regprice	-0.00258	-0.00556*	-0.00869	-0.00826**	-0.00459*	-0.00454	-0.00584***
	(0.00316)	(0.00319)	(0.00539)	(0.00344)	(0.00252)	(0.00295)	(0.000708)
discount	3.204***	3.368***	3.702***	3.329***	3.300***	3.663***	3.282***
	(0.404)	(0.520)	(0.725)	(0.445)	(0.454)	(0.585)	(0.282)
sale	-0.0835	0.0195	-0.0609	0.0888	0.0724	-0.0239	0.0619
	(0.0649)	(0.112)	(0.0999)	(0.0881)	(0.0673)	(0.0907)	(0.0496)
feat	0.216*	0.183	0.323	-0.0347	-0.0412	0.149	0.542***
	(0.115)	(0.127)	(0.229)	(0.137)	(0.101)	(0.109)	(0.0711)
disp	0.0197	0.0421	-0.00160	0.339***	0.796***	0.344***	0.677***
	(0.0507)	(0.0591)	(0.181)	(0.0724)	(0.185)	(0.0884)	(0.0414)
discxfeat	1.823**	1.504*	0.764	2.683***	1.421**	2.688***	4.076***
	(0.816)	(0.883)	(2.003)	(0.872)	(0.657)	(0.865)	(0.438)
discxdisp	4.411***	3.454***	6.959***	2.834***	2.124**	1.481*	1.699***
	(0.798)	(0.700)	(2.058)	(0.799)	(0.975)	(0.777)	(0.410)
Constant	-0.140***	-0.128***	-0.127***	-0.112***	-0.104***	-0.134***	-0.187***
	(0.0330)	(0.0325)	(0.0449)	(0.0330)	(0.0288)	(0.0293)	(0.00603)
Observations	8,191	7,774	7,926	8,197	8,489	8,270	278,491
R-squared	0.163	0.167	0.203	0.177	0.111	0.193	0.203

Notes: Products aggregated to brand-size level. Dependent variable is *qnorm*, percent by which store-product-week quantity sold exceeds average weekly quantity sold. *sale*, *disp*, and *feat* are store-product-week promo indicators, *discount* is percent*price* below *regprice*, and *discxdisp* and *discxfeat* are interactions. Robust SE in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Table 11: Selected structural own-price elasticities, all chains (narrow N_t)

Product / Chain	Chain 78	Chain 100	Chain 140
AH100	1.37	1.7	2.16
ALL100	3.6	4.04	2.16
ALL200	5.25	4.94	4.08
CHEER80	2.06	2.14	2.62
FAB100	1.65	1.84	2.05
GAIN100	2.08	4.76	5.09
GAIN200	4.25	1.65	1.55
PUREX100	1.3	1.5	1.51
PUREX200	2.52	3.28	2.85
SURF100	2.28	2.17	2.65
TIDE100	5.16	5.67	6.07
TIDE200	8.27	8.56	8.85
WISK100	2.16	2.41	2.76

Table 10: Brand-size level promotion effects by store, estimation sample (2004-2005)

	Store-level regressions						
VARIABLES	653369	250094	263568	266100	243785	683960	Market
regprice	0.00419	-0.00869**	-0.0113**	0.00401	-0.00639	0.00393	-0.00659***
	(0.00532)	(0.00383)	(0.00482)	(0.00501)	(0.00460)	(0.00514)	(0.00115)
discount	2.963***	1.627**	2.511***	3.297***	3.463***	3.495***	3.337***
	(0.777)	(0.663)	(0.617)	(0.743)	(0.736)	(1.106)	(0.249)
sale	0.279*	0.417**	0.432***	0.300	0.0125	0.179	0.0505
	(0.169)	(0.211)	(0.131)	(0.183)	(0.144)	(0.198)	(0.0524)
feat	-0.588	-0.945**	-0.412	-1.131**	-0.580*	-0.887**	0.374**
	(0.412)	(0.406)	(0.343)	(0.496)	(0.329)	(0.389)	(0.174)
disp	-0.146	0.0804	-0.101	0.0222	0.368*	0.258**	0.419***
	(0.106)	(0.0671)	(0.130)	(0.152)	(0.191)	(0.131)	(0.0356)
discxfeat	3.587	6.722***	3.734**	6.429**	4.776***	7.041***	3.949***
	(2.896)	(2.343)	(1.841)	(2.636)	(1.625)	(2.237)	(0.786)
discxdisp	4.426**	2.942**	4.587***	2.825**	1.086	-0.482	2.628***
	(2.009)	(1.271)	(1.272)	(1.383)	(1.117)	(1.229)	(0.459)
Constant	-0.161***	-0.170***	-0.152***	-0.178***	-0.170***	-0.147***	-0.221***
	(0.0508)	(0.0397)	(0.0496)	(0.0509)	(0.0411)	(0.0541)	(0.00949)
Observations	1,305	1,321	1,349	1,351	1,447	1,334	51,806
R-squared	0.239	0.305	0.302	0.247	0.277	0.231	0.249

Notes: Products aggregated to brand-size level. Dependent variable is *qnorm*, percent by which store-product-week quantity sold exceeds average weekly quantity sold. *sale*, *disp*, and *feat* are store-product-week promo indicators, *discount* is percent *price* below *regprice*, and *discxdisp* and *discxfeat* are interactions. Robust SE in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 12: Selected cross-price elasticities, Chain 78 (narrow N_t)

	OG	AH100	ALL100	ALL200	FAB100	GAIN100	GAIN200	PUREX100	PUREX200	SURF1
AH100	-0.0908	**1.7**	-0.0856	-0.0817	-0.0907	-0.0803	-0.0817	-0.0914	-0.089	-0.0909
ALL100	-0.0401	-0.0402	**4.04**	-0.0619	-0.0385	-0.0357	-0.0348	-0.0403	-0.0375	-0.0399
ALL200	-0.0322	-0.0321	-0.0519	**4.94**	-0.031	-0.0288	-0.046	-0.0322	-0.0509	-0.0324
FAB100	-0.0549	-0.0546	-0.0494	-0.0474	**1.84**	-0.0496	-0.0486	-0.0548	-0.0531	-0.0541
GAIN100	-0.0474	-0.0482	-0.0456	-0.044	-0.0495	**4.76**	-0.108	-0.0476	-0.046	-0.0483
GAIN200	-0.0103	-0.0106	-0.00965	-0.0152	-0.0105	-0.0234	**1.65**	-0.0104	-0.016	-0.0104
PUREX100	-0.115	-0.115	-0.108	-0.103	-0.114	-0.0999	-0.1	**1.5**	-0.113	-0.115
PUREX200	-0.0494	-0.0496	-0.0444	-0.0721	-0.0491	-0.0427	-0.0683	-0.0499	**3.28**	-0.0494
SURF100	-0.0564	-0.056	-0.0523	-0.0507	-0.0554	-0.0497	-0.0493	-0.0563	-0.0547	**2.17**
TIDE100	-0.0884	-0.0872	-0.0922	-0.103	-0.0883	-0.0952	-0.0977	-0.0882	-0.0941	-0.0891
TIDE200	-0.0449	-0.0441	-0.047	-0.1	-0.0448	-0.0443	-0.068	-0.0443	-0.0778	-0.0442
WISK100	-0.0383	-0.0361	-0.0409	-0.0338	-0.0396	-0.0425	-0.0356	-0.0383	-0.0326	-0.0375

Note: Elasticities represent percentage changes in market share due to one percent *decrease* in price of product on horizontal axi with no displays, features, or sales. Elasticities reported for standard sizes (100oz and 200oz) of major brands in Chain 78.

Table 13: Cross-specification comparison, chain 78 (narrow N_t)

Parameter	Naive	Promo FX	Search
α	0.596	0.387	0.387
μ_c	–	–	-1.83
σ_c^2	–	–	1.4e-14
γ	–	–	0.652
lag_1	0.00309	0.0205	0.0219
lag_2	0.0129	0.0187	0.0198
lag_3	0.00734	0.0112	0.012
lag_4	0.0249	0.0255	0.0249
$ftaste$		0.385	0.392
$dtaste$		0.502	–
Objective	-169428.9	-153471.9	-153594.3

Figure 13: Model validation: actual vs predicted market shares, Store 645502

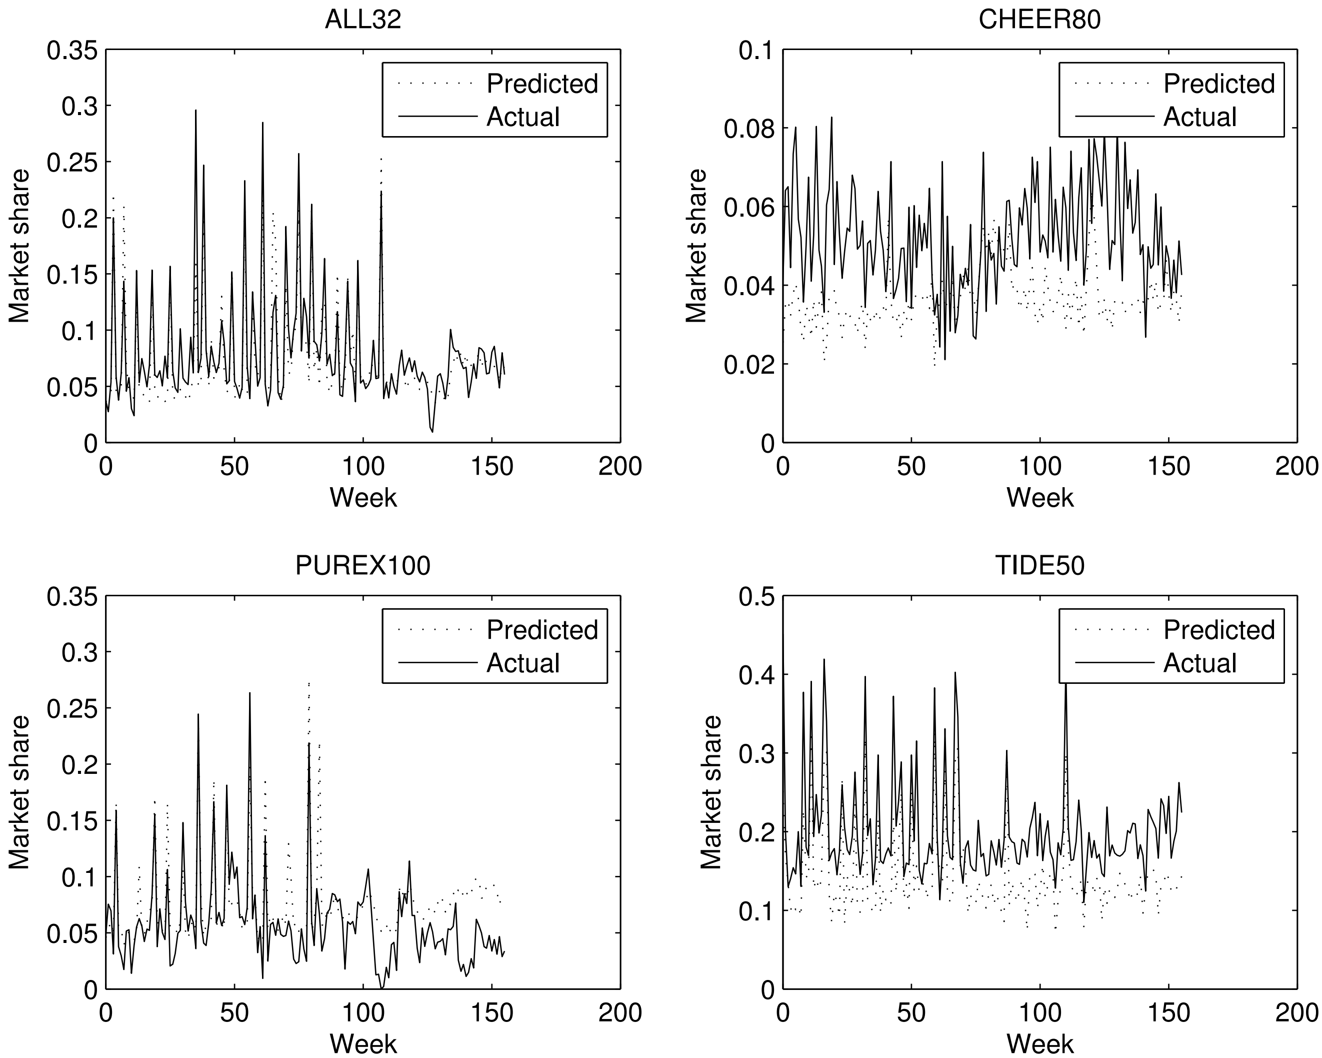

Figure 14: Model validation: actual vs predicted market shares, Store 647840

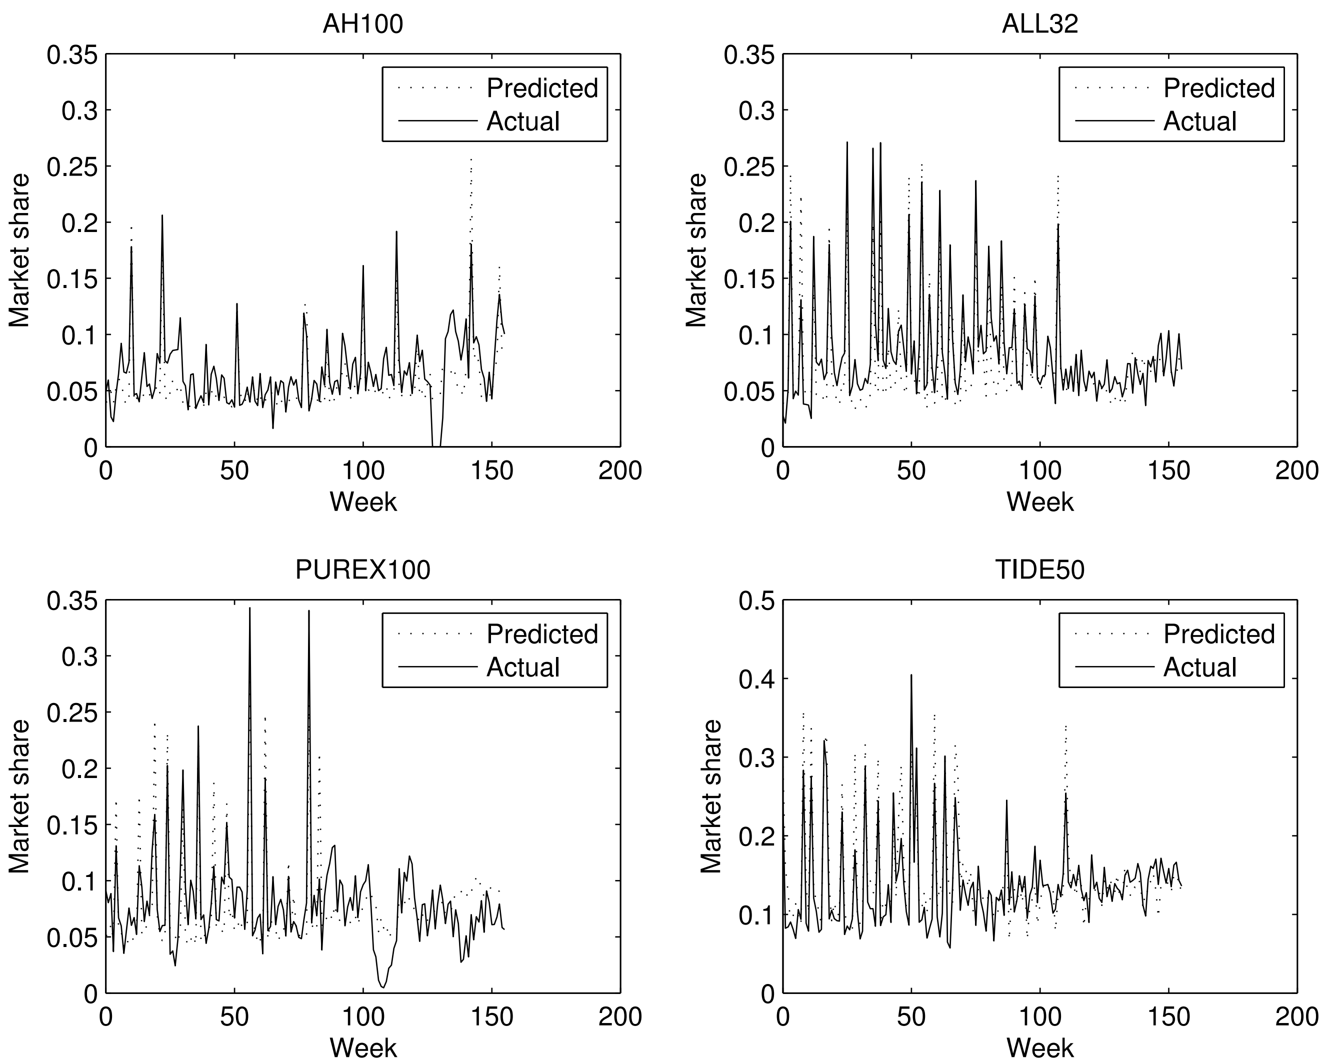

Figure 15: Model validation: actual vs predicted market shares, Store 648768

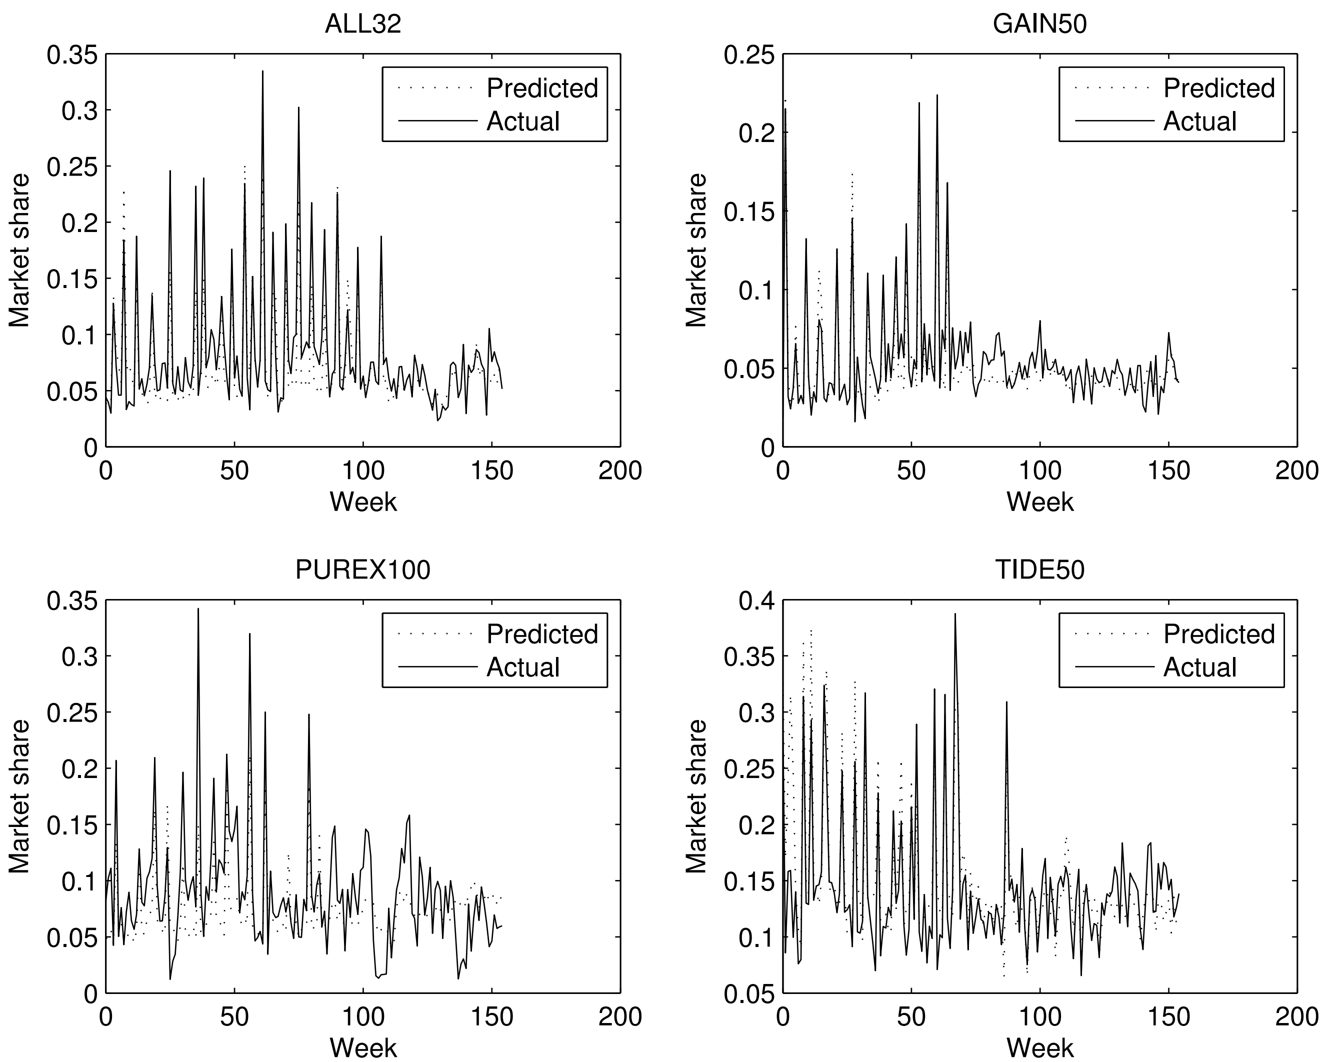

Figure 16: Actual vs predicted market shares, store 654982

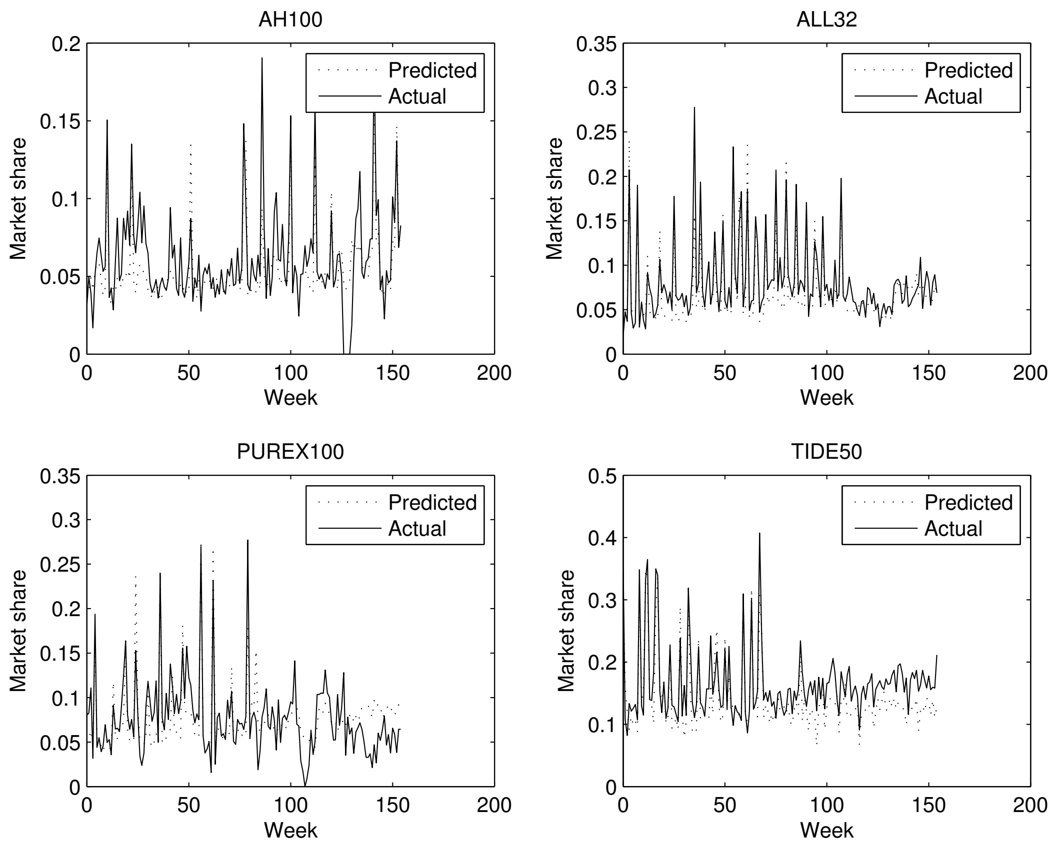

Appendix 2: Regular prices

Regular price series used in structural estimation are defined as follows:

Algorithm (Constructing regular prices)**.**

1. *Drop all periods listed as sales.*

2. *For each remaining period, calculate forward-looking, backward-looking, and centered 9-week rolling price medians.*

3. *If current price equals either forward or backward rolling median, take this value as the regular price; otherwise use the centered rolling median.*[43]

4. *For promotional periods, fill in regular prices from the regular price immediately preceding or the regular price immediately following based on least deviation from price observed.*
   \

5. *Fill in missing values rolling forward.*

Figures 2, 17, and 18 illustrate the product of this algorithm on three representative price series. On balance, the algorithm performs well: it successfully isolates persistent local modes in the price distribution, and thereby permits distinction between secular price shifts and short-run price variation due to sales. Hence I take the resulting regular price series as a basis for estimating sale-induced gains from search.[44]

[43]The comparison with forward and backward rolling prices permits better identification of the regular price in periods near a shift in the regular price.

[44]In preliminary work, I also explored a regular price filter based on rolling modes. Experiments suggest the rolling median filter is more robust to potential noise in the price data. However, qualitative results were vere similar in both cases.

Figure 17: Price vs Regprice for Purex 100oz, IRI store 683960 (2002-2008)

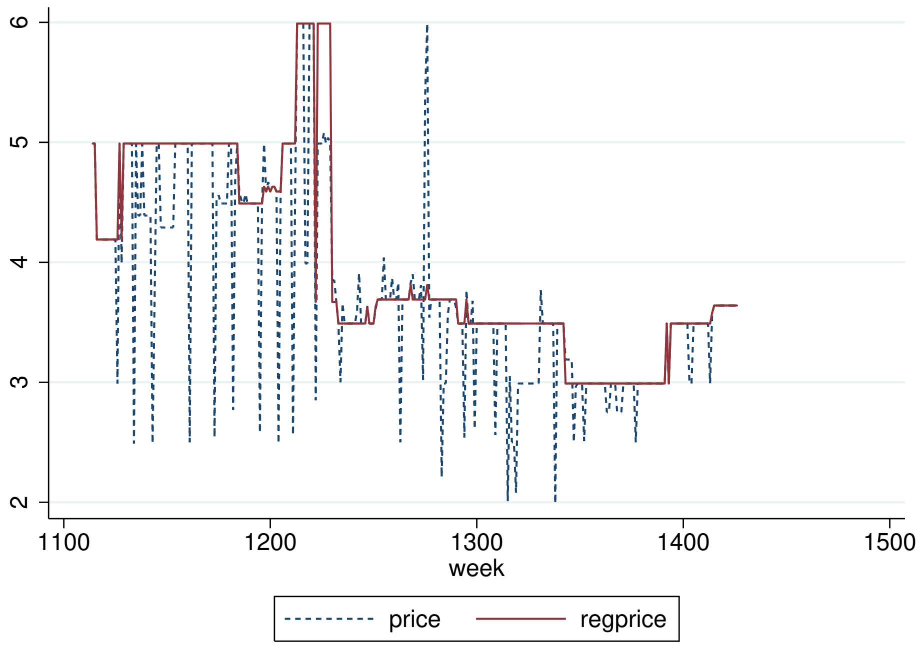

Figure 18: Price vs Regprice for All 100oz, IRI store 683960 (2002-2008)