ESSAYS ON THE STRUCTURAL ANALYSIS OF SELECTION AND SEARCH

By

Matthew Loren Gentry

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

August, 2012

Nashville, TN

Approved:

Professor Tong Li

Professor Yanqin Fan

Professor Eric Bond

Professor Ali Hortacsu

To God be the glory,

great things He has done.

ACKNOWLEDGMENTS

I thank my advisor, Professor Tong Li, for his tremendous guidance, encouragement and support on this and many other projects. His tireless concern on my behalf has been a model of all that a mentor could convey. To the members of my committee, Professors Yanqin Fan, Eric Bond, and Ali Hortacsu, I also extend profound thanks for their invaluable contributions of time, effort, insight, and advice. I would not be where I am today without their support. My research has benefited greatly from conversations with Martin Pesendorfer, Pasquale Schiraldi, Stephan Seiler, Andrew Sweeting, Steve Berry, PJ Glandon, Roger Crawford, Caleb Stroup and others too numerous to mention. Financial support from Professor Tong Li, the National Science Foundation (SES-0922109), and a Vanderbilt University Graduate Fellowship is gratefully acknowledged.

To my parents Dave and Carylene, I say thank you for being the best role models one could hope for. Their love and encouragement are constants on which I can always depend, and their wise counsel has been a profound instrument for good in my life. To my brother Jake, I say thanks for being the best friend, supporter, and sibling I could imagine. To Elissa, I extend profound thanks for her incredible patience and mental, emotional, and spiritual support over the past year, which has been an encouragement beyond anything I could ask or imagine. Finally, I thank my Lord and Savior Jesus Christ for His faithfulness in bringing me thus far, and for His grace that gives assurance for the future.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

AUCTIONS WITH SELECTIVE ENTRY

Entry is a quantitatively and qualitatively important aspect of many real-world auction processes, but theoretical analysis of auctions with entry has to date been limited to a few notable but restrictive special cases. Two paradigmatic examples in the literature are Samuelson [1985] (henceforth S), who proposes a simultaneous entry model in which potential bidders know their valuations *ex ante* but must incur a fixed cost to submit bids, and Levin and Smith [1994] (henceforth LS), who consider simultaneous entry under the alternative assumption that bidders learn their valuations after incurring the fixed cost. A common theme in this literature is that different assumptions on entry can produce very different practical and policy conclusions. For example, under the LS model a revenue-maximizing seller will set a zero reserve price and maximize social welfare, whereas in the S model revenue maximization requires a binding, socially inefficient reserve price. Hence while the existing literature contains many important insights on auctions with entry, it permits few overarching theoretical and policy conclusions.

This paper seeks to generalize existing work on auctions with entry using a framework we call the *Affiliated Signal (AS)* model. First suggested by Ye [2007], the AS model assumes that potential bidders receive imperfect signals of their valuations prior to entry, make simultaneous entry decisions based on these signals, then learn their valuations and submit bids. This structure imposes minimal *a priori* restrictions on pre-entry information, requiring only that signals and values be *affiliated* in the sense of Milgrom and Weber [1982] (so that higher signals are "good news"). It also includes both the S and LS models as polar cases: the former when signals and values are perfectly correlated, and the latter when signals and values are independent. The AS model thus represents an ideal basis for a general theoretical analysis of auctions with entry.

Building on the general AS entry model, the paper makes the following specific contributions. We consider IPV auction environments with AS entry, focusing on a class of mechanisms we call *RS auctions* (after Riley and Samuelson [1981]).[1] For this class of auctions, we establish the following four results. First, we characterize equilibrium entry and bidding behavior induced by any auction in the class considered. Second, building on this result, we establish a generalized revenue equivalence theorem applicable to auctions with general AS entry. Third, we characterize efficiency in auctions with AS entry, and show that the seller's optimal auction is inefficient in general; that is, that Levin and Smith [1994]'s finding that revenue maximization implies efficiency applies only to the polar case of LS entry. Finally, we explore revenue-maximizing reservation prices and entry fees directly, and establish that these will be positive in the "regular case" where potential bidders prefer lower costs of entry. These findings have potentially important implications for welfare and policy analysis in auctions with entry, and to our knowledge none have been established at the level of generality we consider.

This study is related to a large and growing literature on the theory and empirics of auctions with entry. The empirical branch of this literature includes studies establishing the importance of entry in a wide range of applications: Bajari and Hortacsu [2003] in online auctions, Hendricks et al. [2003] in outer continental shelf "wildcat" auctions, Li and Zheng [2009] and Krasnokutskaya and Seim [2009] in highway construction procurement auctions, and Li and Zheng [2012], Li and Zhang [2010a], Athey et al. [2011] and others in timber auctions, to mention just a few. The theoretical literature also contains a number of notable contributions not yet mentioned. McAfee

[1] Roughly, mechanisms such that only the high bidder has a positive probability of award, and the probability of award depends only on the highest bid.

and McMillan [1987] explore a model of sequential entry where entry is interpreted as value discovery; this value-discovery paradigm has been adopted by much of the subsequent entry literature, though these studies typically consider simultaneous rather than sequential entry. More recent work by Lu [2008, 2009a,b, 2010] and Moreno and Wooders [2011] extends the basic LS entry model to incorporate heterogeneous entry costs; Lu characterizes equilibrium, efficiency, and optimal auction design in this extended model, while Moreno and Wooders note that the presence of private entry costs overturns core efficiency results in Levin and Smith [1994]. Finally, Marmer et al. [2007], Gentry and Li [2012], and Roberts and Sweeting [2010b,a] explore specification testing, nonparametric identification, and an empirical application of the general AS model respectively. This study provides a theoretical complement to this recent application-oriented work.

The rest of the paper is organized as follows. Section 1 outlines the structure of the AS model, and Section 2 characterizes symmetric equilibrium entry and bidding behavior under any RS auction rules. Section 3 establishes revenue equivalence in the class of auctions considered. Section 4 establishes that the seller's optimal auction will in general be inefficient, and Section 5 explores revenue-maximizing policy directly. Finally, Section 6 concludes.

I.1 The AS model

We consider allocation of a single indivisible good among N potential bidders via a two-stage auction mechanism M, where bidders have independent private values for the good being sold. Timing of the auction game is as follows. First, in Stage 1, each potential bidder i observes a private signal s_i of her (unknown) private value v_i, and all potential bidders simultaneously choose whether to enter the auction. Each entering bidder must pay an entry cost c; this may be interpreted as the net of opportunity, learning, and bid preparation costs. Then, in Stage 2, the n bidders who chose to enter in Stage 1 learn their true values v_i and submit bids for the object being sold. Finally, auction outcomes are determined according to the rules of the mechanism M, which are common knowledge to all participants. Consistent with institutional features common to many official procurement lettings, we assume that bidders observe the number of potential bidders N prior to entry, but do not observe the number of entrants n prior to bidding.[2]

We frame our analysis in terms of a general class of mechanisms we call *RS auctions* (after the work of Riley and Samuelson [1981]):

Definition I.1. A *RS auction* is any bidding mechanism having the following properties:

1. Mechanism rules are anonymous.
2. If award is made, it is to the bidder submitting the highest bid.
3. The probability of award depends only on the highest bid.
4. For any distribution of rival values, there exists a unique symmetric bidding equilibrium such that bids submitted are strictly increasing in bidder values.

The class of RS auctions includes all four standard auctions (first-price, Vickery, English ascending, and Dutch), plus many less common auction types. It therefore represents a natural focal point for our current investigation.

We formalize the remaining assumptions of the AS entry model as follows.

Assumption 1. *The seller and all potential bidders are risk-neutral.*

Assumption 2. *All bidders are* ex ante *symmetric, and draw value-signal pairs* (V, S) *independently from a continuous joint distribution* $f(v, s)$ *satisfying the following properties:*

[2] Allowing bidders to observe n prior to bidding would slightly change the details of the derivation, but would not substantially alter any of our core results.

(i) *The marginal density of second-stage values ($f(V)$) has positive support on a bounded interval* $[\underline{v}, \bar{v}]$.

(ii) *WLOG, we normalize first-stage signals S to have a uniform marginal distribution on* $[0, 1]$*:* $S_i \sim U[0, 1]$.

(iii) *For each bidder i, the random variablesV_i and S_i* are affiliated *in the sense of Milgrom and Weber [1982].*

Assumption 3. *Information structure:*

(i) *Each bidder i observes own signal s_i prior to entry, but does not learn own value v_i until after entry.*

(ii) *The number of potential bidders N is known to all participants, but the number of entrants n is not revealed until the auction concludes.*

Assumption 4. *The second-stage auction mechanism M and all other model fundamentals are common knowledge.*

Assumption 4 (common knowledge) is entirely standard, Assumption 1 (risk neutrality) is strong but typical in the auction literature, and the signal normalization 2(ii) is feasible since any monotone transformation of a signal preserves information. The affiliation assumption 2(iii) formalizes the sense in which a higher signal is "good news," but otherwise imposes minimal restrictions on the nature of selection; in particular, it nests both the S model (perfect dependence) and the LS model (independence) as polar cases. Leaving n unobserved prior to bidding differs slightly from the corresponding assumption in Levin and Smith [1994], but is motivated by institutional features typical of many real-world auctions and in any event is entirely incidental to our core results.[3]

[3]For instance, in sealed-bid procurement auctions, the auctioneer typically does not disclose information on bids received until announcing auction results. The form of the equilibrium profit function would be slightly different if n were observed prior to bidding, but substantive features of the auction system would be otherwise unchanged.

Finally, as usual, we frame our theoretical analysis in terms of direct mechanisms; by the Revelation Principle, any mechanism has an equivalent direct mechanism, so this is without loss of generality (see Krishna [2009]). Let the *award rule* $\alpha_M(y)$ denote the probability RS auction M results in a sale when the highest reported value among entrants is y; by definition, RS auctions can award only to the highest bidder, so the performance of any RS auction M can be fully characterized by its award rule $\alpha_M(\cdot)$. For current purposes, we add three further regularity conditions on the mechanism M:

Assumption 5. *The second-stage mechanism M is an RS auction with a direct equivalent such that:*

(i) *The award rule $\alpha_M(y)$ is weakly increasing in the maximum entrant value y.*

(ii) *A low-type bidder (entrant with value $\underline{v}$) weakly prefers less Stage 2 competition.*

(iii) *WLOG, the mechanism M is specified such that when $N-1$ potential rivals report values $\mathbf{z}_{-i} \in [0,\bar{v}]^{N-1}$, the payment of a potential bidder reporting value $z_i \leq \underline{v}$ takes the form*

$$p(z_i; \mathbf{z}_{-i}) = \mathbf{1}[n=1]\left\{\alpha(z_i)z_i - \int_0^{z_i} \alpha(y)dy\right\} + \rho(\mathbf{z}_{-i}), \qquad \text{(I.1)}$$

where $\rho(\mathbf{z}_{-i})$ is a symmetric function and reports $\mathbf{z}_{-i}$ incorporate the possibility of non-entry.

Conditions (i) and (ii) are standard and satisfied by almost all mechanisms used in practice. Condition (iii) may look restrictive at first glance, but is actually without loss of generality since it applies only to *out-of-equilibrium* reports $z_i \leq \underline{v}$. This latter fact is important in ensuring the generality of our results, so we state it formally as a lemma:

Lemma I.1. *Any RS auction satisfying Assumptions 1-5(ii) is payoff- and performance-equivalent to some RS auction also satisfying Assumption 5(iii).*

Finally, note that Assumption 5 intuitively nests all leading special cases of interest. We illustrate this fact via two simple examples:

Example I.1. Consider a first-price auction with a reserve price $r \leq \underline{v}$. This structure can be nested under Assumption 5 by setting $\alpha(y) = \mathbf{1}[y \geq r]$ and $\rho(\mathbf{z}_{-i}) = 0$ for all $\mathbf{z}_{-i}$.

Example I.2. Consider a second-price auction with a secret reserve price $r \sim F_r(\cdot)$ and an entry fee $e > 0$. This structure can be nested under Assumption 5 by setting $\alpha(y) = F_r(y)$ and $\rho(\mathbf{z}_{-i}) = 0$ for all $\mathbf{z}_{-i}$.

I.2 Equilibrium

Since bidders are *ex ante* symmetric by hypothesis, we focus on the class of symmetric, subgame perfect, pure strategy Bayesian Nash equilibria. An equilibrium of this form involves two components: a Stage 1 *entry set* S^* such that "enter if $s_i \in S^*$" is optimal given equilibrium continuation play, and a Stage 2 *bidding function* $\beta(v; S)$ (monotonic in its first argument) such that an entrant with value v_i optimally submits bid $\beta(v_i; S)$ in response to $N-1$ rivals who enter according to S and bid according to $\beta(\cdot; S)$. We restrict attention to RS auctions, so by definition $\beta(\cdot; S)$ exists and is unique for any nonempty entry set S.

Stage 2: Bidding equilibrium

Temporarily suppose that the Stage 1 entry set S can be characterized by a symmetric *entry threshold* $\bar{s} \in [0, 1]$ such that bidder i chooses to enter if and only if $s_i \geq \bar{s}$; we subsequently establish that any Stage 1 equilibrium must take this form. Then the (selected) distribution of values among entrants at $\bar{s}$ is given by

$$F^*(v;\bar{s}) \equiv \frac{1}{1-\bar{s}} \int_{\bar{s}}^{1} F(v|t)\,dt. \tag{I.2}$$

By hypothesis, the equilibrium bid function $\beta(\cdot;\bar{s})$ is monotonic in its first argument. An entrant with value v will thus win against potential rival j in one of two events: either j does not enter, or j enters but draws a value less than v. Let $F_w^*(v;\bar{s})$ denote the joint probability of these events:

$$F_w^*(v;\bar{s}) = \bar{s} + (1-\bar{s}) \cdot F^*(v;\bar{s}).$$

We will frequently reference stochastic ordering of the distributions $F^*(v;s)$ and $F_w^*(v;s)$ in s, so we formally establish these orderings via the following lemma.

Lemma I.2. *For any $(v;s)$ and any $s' \geq s$, the ex-post value distribution $F^*(v;\cdot)$ satisfies*

$$F^*(v;s) \geq F^*(v;s')$$

and the ex-post win probability distribution $F_w^*(v;\cdot)$ satisfies

$$F_w^*(v;s) \leq F_w^*(v;s').$$

In other words, $F^*(\cdot;s')$ first-order stochastically dominates $F^*(\cdot;s)$, and $F_w^*(\cdot;s)$ first-order stochastically dominates $F_w^*(\cdot;s')$.

The form of the equilibrium bidding function $\beta(\cdot;\bar{s})$ will obviously depend on the rules of the mechanism M. However, via standard arguments in mechanism design, we can characterize expected Stage 2 *profit* in any RS auction as follows.

Proposition I.1. *In any symmetric Stage 2 equilibrium of any RS mechanism, the expected Stage 2 profit of an entrant with value v facing $N-1$ potential rivals who enter according to threshold $\bar{s}$ is given by*

$$\pi(v;\bar{s}) = \int_0^v \alpha_M(y) \cdot F_w^*(y;\bar{s})^{N-1} dy - E[\rho|N,\bar{s}], \tag{I.3}$$

where the low-type payment function $\rho(\cdot)$ is defined as in Assumption 5 and

$$E[\rho|N,\bar{s}] \equiv \int \cdots \int \rho(z_1,...,z_{N-1}) dF_w^*(z_1;\bar{s}) \cdots dF_w^*(z_{N-1};\bar{s}). \tag{I.4}$$

Stage 1: Entry equilibrium

Given this symmetric Stage 2 equilibrium profit function $\pi(v;\bar{s})$, we next characterize the symmetric Stage 1 entry threshold $\bar{s}$. Toward this end, consider the Stage 1 decision faced by potential bidder i with signal s_i facing $N-1$ potential rivals who enter according to $\bar{s}$. Bidder i's *ex ante* expected Stage 2 profit is given by

$$\begin{aligned} \Pi(s_i;\bar{s}) &= E_v[\pi(v;\bar{s})|s_i] \\ &= \int_0^{\bar{v}} f(v|s_i) \int_0^v \alpha(y) \cdot F_w^*(y;\bar{s})^{N-1} dy - E[\rho|N,\bar{s}] \\ &= \int_0^{\bar{v}} \alpha(y) \cdot [1 - F(y|s_i)] \cdot F_w^*(y;\bar{s})^{N-1} dy - E[\rho|N,\bar{s}], \end{aligned}$$

where the second line follows from Proposition I.1 and the third follows from integration by parts. The key properties of this ex ante profit function are stated in the following lemma.

Lemma I.3. Ex ante *expected Stage 2 profit for a bidder with Stage 1 signal s_i facing $N-1$ rivals entering according to $\bar{s}$ under RS auction M is*

$$\Pi(s_i;\bar{s}) = \int_0^{\bar{v}} \alpha(y) \cdot [1 - F(y|s_i)] \cdot F_w^*(y;\bar{s})^{N-1} dy - E[\rho|N,\bar{s}]. \tag{I.5}$$

This function is weakly increasing in s_i for all $(\bar{s}, N)$, strictly decreasing in $\bar{s}$ for all (s_i, N), and strictly decreasing in N for all s_i and any $\bar{s} < 1$.

Bidder i will choose to enter whenever expected net profit from entry is positive, i.e. whenever

$$\Pi(s_i; \bar{s}) \geq c. \tag{I.6}$$

This fact in turn implies a breakeven condition which must hold at any candidate *interior* threshold $\bar{s} \in (0, 1)$:

$$\Pi(\bar{s}; \bar{s}) \equiv c$$

that is, a bidder drawing signal $S_i = \bar{s}$ must be indifferent to entry when facing $N-1$ potential rivals who also enter according to $\bar{s}$.

Finally, note that that any Stage 1 equilibrium can be represented in threshold form. To see this, first generalize the expected Stage 2 profit function in Equation I.5 to an arbitrary entry set S:

$$\Pi(s_i; S) = \int_0^{\bar{v}} \alpha(y) \cdot [1 - F(y|s_i)] \cdot F_w^*(y; S)^{N-1} dy - E[\rho|N, S].$$

Affiliation implies that $F(v|s)$ is decreasing in s for all v, so $\Pi(s_i; S)$ will be (weakly) increasing in its first argument. Hence, if S is an equilibrium, we must have $\Pi(s, S) \geq \Pi(\min(S); S) \geq c$ for every $s \geq \min(S)$. If either inequality is strict, it will be optimal for a bidder with signal s to enter, so we must have $s \in S^*$. Otherwise, we can replace $\min(S)$ with s' to obtain a new entry set S' which is payoff- and performance-equivalent to S, and iterating this argument will eventually produce an equivalent threshold set.

We combine these arguments to establish Proposition I.2, which formally characterizes Stage 1 equilibrium.

Proposition I.2. *A symmetric entry equilibrium in the AS model is characterized by a signal threshold $\bar{s}$ such that only bidders with $s_i \geq \bar{s}$ choose to enter. This signal threshold is uniquely determined as follows.*

- *If* $\Pi(0;0) > c$, *then* $\bar{s} = 0$ *and all potential bidders always enter.*
- *If* $\Pi(1;1) < c$, *then* $\bar{s} = 0$ *and no potential bidder ever enters.*
- *Otherwise, the signal threshold* $\bar{s}$ *satisfies the breakeven condition*

$$\Pi(\bar{s};\bar{s}) \equiv c, \tag{I.7}$$

where $\Pi(s_i;\bar{s})$ *is defined as in Lemma I.3.*

Taken together, Propositions I.1 and I.2 characterize the unique symmetric Bayesian Nash equilibrium of the AS model under any Stage 2 RS auction.

I.3 Revenue equivalence

By definition, *ex ante* expected seller revenue in mechanism M under entry structure (s,N) is the difference between total social welfare and surplus accruing to bidders at (s,N):

$$R_M(s;N) = W_M^*(s;N) - N\Pi_M^*(s;N),$$

where $W_M^*(s;N)$ is expected social welfare generated by mechanism M under entry structure (s,N), and $\Pi_M^*(s;N)$ is expected *ex ante* equilibrium profit for any given potential bidder at this equilibrium.

First consider social welfare $W_M^*(s;N)$. Let $Y_{k:N}$ be the kth highest realized value among N potential entrants, where bidder i's *realized value* is defined as v_i if bidder i enters and 0 otherwise. By definition, RS auction M awards either to the high bidder or not at all, with the probability of award conditional on $Y_{1:N}$ determined by the award rule $\alpha_M(\cdot)$. Mechanism M thus generates expected social welfare

$$\begin{aligned} W_M^*(s,N) &= E\left[Y_{1:N}\alpha_M(Y_{1:N}) + v_0[1-\alpha_M(Y_{1:N})]|s,N\right] - N(1-s)c \\ &= \int \left[y\alpha_M(y) + v_0[1-\alpha_M(y)]\right] dG_{1:N}^*(y;s,N) - N(1-s)c, \end{aligned}$$

where $G_{1:N}^*(y;s,N) \equiv F_w^*(y;s)^N$ denotes the distribution of $Y_{1:N}$ given (s,N). Note that this function depends on M only through the award rule $\alpha_M(\cdot)$.

Now consider equilibrium profit $\Pi_M^*(\bar{s},N)$. By Proposition I.1, equilibrium Stage 2 profit for an entrant with value v_i is given by

$$\pi_M(v;s,N) = \int_0^v \alpha_M(y)\cdot F_w^*(y;s)dy - \rho_M,$$

where for ease of exposition we assume $\rho(\cdot) = \rho_M$.[4] $\Pi_M^*(s,N)$ is the expectation of this function with respect to the distribution $F_w^*(v;s)$ less expected entry costs at (s,N); that is, the *ex ante* expected profit of an arbitrary potential entrant:

$$\Pi_M^*(s,N) = \int_0^{\bar{v}}\int_0^v \alpha_M(y)\cdot F_w^*(y;s)dy\, dF_w^*(v;s) - (1-s)(\rho_M + c).$$

Again, note that this function depends on $\Pi_M^*(s,N)$ only through $\alpha_M(\cdot)$ and ρ_M.

Finally, by Proposition I.2, any two mechanisms M_1 and M_2 having the same award rule and low-type payoff must induce the same entry behavior: $\rho_1 = \rho_2$ and $\alpha_1(y) = \alpha_2(y)$ for all y implies $\bar{s}_1 = \bar{s}_2$ in equilibrium. Revenue equivalence of M_1 and M_2 then follows immediately from the definitions of $W_M^*(\cdot)$ and $\Pi_M^*(\cdot)$ above. But $\rho_1 = \rho_2$ and $\alpha_1(y) = \alpha_2(y)$ are exactly the conditions of the standard fixed-n equivalence theorem. This in turn establishes our core equivalence result:

[4] The argument can easily be extended to general $\rho(\cdot)$.

Proposition I.3 (Revenue Equivalence)**.** *Suppose RS mechanisms M_1 and M_2 are revenue-equivalent for fixed n. Then M_1 and M_2 are also revenue-equivalent under endogenous and selective AS entry.*

We thus extend the classic equivalence results of Riley and Samuelson [1981] and Levin and Smith [1994] to the case of endogenous and selective AS entry.

I.4 Revenue maximization and efficiency

In this section we study the relationship between revenue maximization and social efficiency in the class of RS auctions with AS entry. This investigation is motivated by a key result of Levin and Smith [1994]: when bidders enter without selection, a revenue-maximizing seller will also maximize social welfare. We show that this conclusion applies only in the polar LS case: in the broader AS model, the seller will generally prefer an inefficient mechanism. For current purposes, we focus on two policy instruments: a public reserve price r and an entry fee e. For simplicity, we also normalize $v_0 = 0$.

First consider social welfare at an arbitrary entry threshold s. Specializing the arguments in the last section, we obtain

$$W_M^*(s) = \int_r^{\bar{v}} y dG_{1:N}^*(y; s, N) - N(1-s)c. \tag{I.8}$$

A binding reserve price $r > \underline{v}$ is obviously inefficient, and a nonbinding reserve price $r \in (0, \underline{v}]$ can always be offset by an appropriate entry fee. Hence we can set $r = 0$ WLOG, and characterize the efficient auction via only choice of e.

Rearranging Equation I.8 via integration by parts and substituting using the definition of $G_{1:N}^*(y; s, N)$ produces the following equivalent representation for social welfare at threshold s:

$$W_M^*(s) = \bar{v} - \int_0^{\bar{v}} F_w^*(y; s)^N dy - N(1-s)c. \tag{I.9}$$

It can be shown that this function is concave in s. Social welfare is thus maximized at a threshold $\hat{s}$ satisfying the necessary and sufficient first-order condition

$$-N \int_0^{\bar{v}} F_w^*(y; \hat{s})^{N-1}[1 - F(y|\hat{s})] + Nc \equiv 0. \tag{I.10}$$

Let $\bar{s}_e$ be the entry threshold corresponding to entry fee e. By Proposition I.2, $\bar{s}_e$ satisfies

$$\int_0^{\bar{v}} F_w^*(y; \bar{s}_e)^{N-1}[1 - F(y|\bar{s}_e)] \equiv c + e.$$

Setting $e = 0$ thus ensures $\int_0^{\bar{v}} F_w^*(y; \bar{s}_0)^{N-1}[1 - F(y|\bar{s}_0)] \equiv c$, which by Equation I.10 implies $\bar{s}_0 = \hat{s}$. This observation translates into the following characterization of the socially optimal mechanism:

Proposition I.4 (Efficiency). *Social welfare is maximized when the seller sets no reserve price or entry fee: setting $\hat{r} = 0$ and $\hat{e} = 0$ produces socially optimal entry.*

Now let $m = (e, r)$ be the seller's policy choice, and consider the seller's revenue maximization problem. By definition, seller revenue is the difference between social surplus and expected profits among potential bidders, which with slight abuse of notation we write as follows:

$$R^*(m) = W^*(m) - \Pi^*(m).$$

By construction, the seller's optimal policy m^* satisfies

$$\frac{\partial R^*(m^*)}{\partial m} \equiv 0.$$

Meanwhile, the social optimum $\hat{m}$ satisfies

$$\frac{\partial S^*(\hat{m})}{\partial m} = \frac{\partial R^*(\hat{m})}{\partial m} + \frac{\partial \Pi^*(\hat{m})}{\partial m} \equiv 0,$$

which in turn implies

$$\frac{\partial R^*(\hat{m})}{\partial m} = -\frac{\partial \Pi^*(\hat{m})}{\partial m}. \tag{I.11}$$

Since bidders earn positive expected profits when entry is selective, the RHS of Equation (I.11) will not be zero in general. Hence the seller's optimal policy need not correspond to the social optimum. We state this result formally as a lemma:

Lemma I.4. *In general, a revenue-maximizing seller will not maximize social welfare.*

As noted above, this result contrasts with the corresponding finding in Levin and Smith [1994]. Intuitively, when potential bidders have no private information, entry will involve mixed strategies and potential bidders will compete away all profits. Hence $\Pi^*(m)$ will be identically zero for all m, so $R^*(m) \equiv W^*(m)$ and a revenue-maximizing seller will maximize total surplus. In contrast, when entry is selective, $\Pi^*(m) > 0$ in general, and a revenue-maximizing seller will induce distortion to capture part of this additional surplus. The welfare results of Levin and Smith [1994] thus depend crucially on the particular informational assumptions of the LS model, and in general do not apply outside that polar case.

I.5 Revenue-maximizing policy

The last section proceeded via negation, first characterizing the socially efficient auction, then noting that in general a revenue-maximizing seller will not choose this auction. This section proceeds more positively, first characterizing revenue-maximizing choices of the seller's policy variables e and r, then deriving conditions under which these will be positive.

Seller's optimal entry fee

Set $r = 0$, and consider second-price auctions without loss of generality. Proposition I.2 implies a one-to-one correspondence between the entry fee e and the equilibrium entry threshold s, so we can derive optimal policy in terms of either. We thus consider maximization of expected revenue corresponding to threshold s:

$$R^*(s) = \int_0^{\bar{v}} y \, dG^*_{2:N}(y; s) + N(1-s) \cdot e(s),$$

where $e(s) \equiv \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy - c$ is the entry fee required to produce threshold s as an equilibrium. Integrating by parts, substituting for $e(s)$, and rearranging then produces the following expression for seller revenue:

$$R^*(s) = \left\{ \bar{v} - \int_0^{\bar{v}} F^*_w(y; s)^N dy - N(1-s)c \right\} - \left\{ N(1-s) \int_0^{\bar{v}} [F(y|s) - F^*(y; s)] F^*_w(y; s)^{N-1} dy \right\},$$

where the first term gives total surplus (see Equation (I.9)) and the second term gives total profit among potential bidders. In the "regular case" where potential bidders strictly prefer lower total entry costs, the second term will be decreasing in s. Combining these observations with a standard FOC yields the following characterization of the seller's optimal entry fee e^*:

Proposition I.5 (Optimal entry fee)**.** *Suppose the seller's only policy variable is e. Then the seller will choose e^* to induce an equilibrium threshold s^* satisfying $R^*_s(s^*) = 0$, where:*

$$\begin{aligned} R^*_s(s) \;\; &\equiv \;\; c - \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy \\ &- \;\; (1-s) \int_0^{\bar{v}} F_s(y|s) \cdot F^*_w(y; s)^{N-1} dy \\ &- \;\; (1-s) \int_0^{\bar{v}} [F(y|s) - F^*(y; s)] \cdot (N-1) F^*_w(y; s)^{N-2} [1 - F(y|s)] dy \end{aligned} \quad \text{(I.12)}$$

The first line of $R^*_s(s)$ is negative, zero, or positive as $s \gtreqless \hat{s}$, the second line is always *positive, and the third line is always negative. Hence $R_s^*(s)$ may be either positive or negative at $\hat{s}$, but will not be zero in general.*

If potential bidders strictly prefer lower total cost of entry, the sum of the last two lines will be positive and a revenue-maximizing seller will choose $e^* > 0$.

Finally, for completeness, we specialize Proposition I.5 to the S and LS polar cases.

Corollary I.1 (Optimal entry fee, LS case)**.** *Suppose LS entry obtains (bidders have no prior information on own values). Then the seller optimally choosese$^* = 0$ and induces efficient entry.*

Proof. Under LS entry, bidders cannot select into participation, so $F(y|s) \equiv F^*(y;s) \equiv F(y)$ and $F_s(y|s) \equiv 0$. Condition (I.12) then reduces to

$$0 \equiv \int_0^{\bar{v}} [1 - F(y)] F_w^*(y; s^*)^{N-1} dy - c,$$

which we know from above is uniquely satisfied at the social optimum $e^* = 0$. □

Corollary I.2 (Optimal entry fee, S case)**.** *Suppose S entry obtains (bidders know values exactly prior to entry). Then the seller optimally induces an entry threshold s^* such that $R_y^*(F_y^{-1}(s^*)) = 0$, where*

$$R_y^*(y) \equiv N[1 - F(y)]F(y)^{N-1} - N[yF(y)^{N-1} - c]f(y).$$

Further, $R_y^*(\hat{y}) > 0$ at the social optimum $\hat{y}$, so the seller always prefers less than efficient entry.

Proof. Follows by setting $F^*(y;\bar{s}) = [F(y) - \bar{s}]/[1 - \bar{s}]$, $F(y|\bar{s}) = 1[y \geq F_y^{-1}(\bar{s})]$, $F_s(y|\bar{s})$ equal to the Dirac delta function $\delta(y - F_y^{-1}(\bar{s}))$, and $\bar{s} \equiv F(\bar{y})$ in Proposition I.5. A direct proof is given in the appendix. □

Seller's optimal reserve price

Now set $e = 0$ and consider the seller's choice of r in a second-price auction. Seller revenue corresponding to reserve price $r \geq 0$ at entry threshold s is given by:

$$R(r, s) = \bar{v} - \int_r^{\bar{v}} G^*_{2:N}(y; s)dy - rF^*_w(r; s)^{N-1}.$$

As above, we frame the seller's problem in terms of choosing an entry threshold s to maximize $R^*(s) \equiv R(r(s), s)$, where $r(s)$ is the reserve price inducing equilibrium entry s. In this case, $R(r, s)$ is non-separable and $r(s)$ is defined implicitly, so while a formal characterization of the solution is feasible, it is not particularly helpful. We thus focus instead on conditions under which the seller will prefer a positive reserve price. Toward this end, we take appropriate partial derivatives, evaluate at the social optimum $\hat{s}$, and simplify to obtain

$$\begin{aligned} R^*_s(\hat{s}) \quad \propto \quad & -\int_0^{\bar{v}} F_s(y|\hat{s}) \cdot F^*_w(y; \hat{s})^{N-1} dy \\ & -(N-1) \int_0^{\bar{v}} [F(y|\hat{s}) - F^*(y; \hat{s})] F^*_w(y; \hat{s})^{N-2} [1 - F(y|\hat{s})] dy. \end{aligned}$$

The RHS of this expression is identical to the last two terms of Equation (I.12) in Proposition I.5. Gains from setting a positive reserve price are therefore feasible under exactly the same condition as gains from a positive entry fee; namely, in the "regular case" when potential bidders strictly prefer lower total entry costs.

Proposition I.6 (Seller's optimal reserve price)**.** *Suppose the seller's only policy variable is r. Then the seller will set a positive reserve price if*

$$-\int_0^{\bar{v}} F_s(y|\hat{s}) \cdot F^*_w(y; \hat{s})^{N-1} dy > (N-1) \int_0^{\bar{v}} [F(y|\hat{s}) - F^*(y; \hat{s})] F^*_w(y; \hat{s})^{N-2} [1 - F(y|s)] dy.$$

If potential bidders prefer lower total entry costs, this condition will be satisfied and the seller will optimally set $r^* > 0$.

Finally, we specialize Proposition I.6 to the polar cases of S and LS entry.

Corollary I.3 (Optimal reserve price, LS case)**.** *Suppose LS entry obtains (bidders have no prior information on own values). Then the seller optimally chooses* $r^* = 0$ *and induces socially optimal entry.*

Proof. Levin and Smith [1994] derive this result from Proposition I.4 plus the fact that the seller captures all social surplus. It also follows immediately from Proposition I.6, since $F_s(y|s) \equiv 0$ and $F(y|s) \equiv F^*(y; s)$ in the LS case. □

Corollary I.4 (Optimal reserve price, S case)**.** *Suppose S entry obtains (bidders know values exactly prior to entry). Then the seller's optimal entry fee* r^* *induces entry threshold* $s^* \equiv F_y(y^*)$*, where* y^* *satisfies the first-order condition*

$$N[1 - F(y^*)]F(y^*)^{N-1} - r^* F(y^*)^{N-1} f(y^*).$$

Further, $r^* > 0$ *and the seller induces less than optimal entry.*

Proof. $r^* > 0$ follows from Proposition I.6 plus Corollary I.2. The first-order condition is derived in the Appendix. □

I.6 Conclusion

This study proposes a general analysis of auctions with entry based on a framework we call the Affiliated-Signal (AS) model. From a theoretical perspective, this framework has several major advantages: it relaxes existing restrictions on pre-entry information, permits endogenous and arbitrarily selective entry, and nests several leading special cases in the literature. The AS model thus represents an ideal basis for a general theoretical approach to auctions with entry.

Within this general AS entry framework, we focus on the broad class of mechanisms considered by Riley and Samuelson [1981]: roughly, auctions such that award (if made) is only to the high bidder. For this class of auctions, we establish the following four results. First, we characterize equilibrium entry and bidding behavior in the AS model under general RS auction rules. Second, we extend the classic revenue equivalence results of Riley and Samuelson [1981] and Levin and Smith [1994] to auctions with endogenous and arbitrarily selective entry. Third, we characterize efficient entry in the AS model, and show that a revenue-maximizing seller will in general induce suboptimal entry decisions. Finally, we explore revenue-maximizing policy in the AS model, and derive conditions under which the seller will prefer positive reservation prices and entry fees. While some of these results were available for special cases of our model, to our knowledge none have been established at the level of generality we consider.

Our findings on efficiency in particular contrast sharply with those of Levin and Smith [1994], and this contrast is worth discussing further in the context of the literature. Levin and Smith [1994] study entry under two key assumptions: that all potential bidders face the same entry cost, and that potential bidders have no specific information on own values prior to entry. Under these assumptions, entry is in mixed strategies, the seller captures all auction gains, and the revenue-maximizing auction is thus by definition efficient. Recent work by Moreno and Wooders [2011] has established that the first assumption is pivotal: if potential bidders have private entry costs, the revenue-maximizing auction will in general be inefficient. This study establishes that the second is also pivotal. Taken together with the work of Moreno and Wooders [2011], our findings thus suggest that the Levin and Smith result is best considered a “corner case”: coincidence between revenue maximization and efficiency is not a general consequence of entry, but rather arises only very special assumptions. This observation in turn has implications for both policy design and welfare analysis.

CHAPTER II

IDENTIFICATION IN AUCTIONS WITH SELECTIVE ENTRY

Endogenous participation clearly matters in real-world auction markets. Studies of auctions in a variety of economic contexts routinely find that large fractions of eligible bidders elect not to submit bids. For example, Hendricks et al. [2003] report an overall participation rate of less than 25 percent in US Minerals Management Service "wildcat auctions" held from 1954-1970. Li and Zheng [2009] find that only about 28 percent of planholders in Texas Department of Transportation mowing contracts actually submit bids. Similar results have been reported for timber auctions (Athey et al. [2011], Li and Zhang [2010a,b]), in online auction markets (Bajari and Hortacsu [2003]), and in other procurement settings (Krasnokutskaya and Seim [2009]). Such endogenous participation can overturn core predictions of classical auction theory: for instance, Levin and Smith [1994] show that the possibility of entry can lead to a zero optimal reserve price, and Li and Zheng [2009] show that it can cause a seller to prefer *less* potential competition. Hence properly accounting for entry is practically important in applied research.

The *Affiliated-Signal (AS) model* described in Chapter I represents an ideal theoretical basis for the structural analysis of auctions with entry: it nests both the *S model* (after Samuelson [1985]) and the *LS model* (after Levin and Smith [1994]) as polar cases, and permits a wide variety of policy dynamics. This model has recently begun to receive attention in empirical applications; Marmer et al. [2007] (henceforth MSX) propose nonparametric tests of the AS, S, and LS models, and Roberts and Sweeting [2010a,b] apply a parametric variant of the model to data on timber auctions. Unfortunately, the identification properties of the AS model are not well established, and in many applications the model will not be point-identified.[1] Consequently, empirical work on the AS has to date been limited and reliant on strong parametric assumptions.

This chapter explores identification in the general AS model, building on the partial identification paradigm pioneered by Manski.[2] In particular, we establish three core identification results for auctions with AS entry. First, using exogenous variation in entry behavior (induced by variation in, e.g., potential competition or entry costs), we derive identified bounds on fundamentals under endogenous and arbitrarily selective entry in a general class of auction mechanisms considered in Riley and Samuelson [1981]. Second, we translate these bounds on fundamentals into bounds on seller revenue corresponding to a wide range of counterfactual mechanisms (again accounting for endogenous and selective entry). Finally, we state conditions under which all bounds collapse to exact identification. To our knowledge, these are the first such results in the identification literature.

Within the existing literature on partial identification in auctions, our work is most similar in spirit to that of Haile and Tamer [2003], who derive bounds on fundamentals in ascending auctions under weak behavioral assumptions, then translate their results into bounds on the seller's optimal reserve price. However, we focus on a very different problem (endogenous entry, not considered in Haile and Tamer) and relax a different set of assumptions (those governing the nature of selection). Another related paper is Tang [2011], who provides bounds for counterfactual revenue

[1]Nonparametric identification of auction models focuses on using observables such as bids to identify model primitives such as the distribution of values. This literature was established by Guerre et al. [2000], who address nonparametric identification and estimation in first-price IPV auctions, and has focused primarily on auction models without entry. See, e.g., Li et al. [2002] for the affiliated private value (APV) model, Li et al. [2000] for the conditionally independent private information model, Krasnokutskaya [2009] for an asymmetric auction with unobserved auction heterogeneity, Hortacsu [2002] for treasury bond auctions, and Athey and Haile [2002] for other standard auction models/formats. Athey and Haile [2005] provide a comprehensive survey of the literature.

[2]See Manski [2003] for a summary of the early partial-identification literature; recent additions to the literature include Manski and Tamer [2002], Magnac and Maurin [2008], Molinari [2008], Fan and Park [2009], and Tamer [2003], to name only a few.

in affiliated values (AV) auction settings. Again, however, our work is set in a different context (auctions with selective entry) and focuses on a much different set of problems. We thus contribute both to the literature on partial identification and to the literature on the econometrics of auctions with entry.[3]

The plan of this chapter is as follows. Sections II.1 and II.2 explore identification of the AS entry model for the general class of RS auctions and present our core partial-identification results. Section II.3 translates these core partial-identification results into bounds on seller revenue corresponding to a wide range of counterfactual RS auctions. Finally, Section II.4 concludes. Detailed proofs and a numerical example are included as appendices. The model and equilibrium are as defined in Chapter I.

II.1 The identification problem

Identification involves recovery of model fundamentals $F(v, s)$ and c. We consider this problem based on a large sample of auctions from some AS process $\mathcal{L}$, where for each auction ℓ the following variables are observed: number of potential bidders N_ℓ, number of actual bidders n_ℓ, and a vector of submitted bids $\mathbf{b}_\ell$. In applications, N_l is typically proxied by variables such as number of planholders (e.g. Li and Zheng [2009]) or number of bidders in related auctions (e.g. Roberts and Sweeting [2010a]) and n_ℓ is taken to be the number of bids submitted. Optionally, the econometrician may also observe a vector of cost shifters z_ℓ, which are assumed to be excluded from $F(v, s)$.[4] For current purposes, we impose three further restrictions on the DGP $\mathcal{L}$.

First, our core contribution in this paper is to derive bounds on AS model fundamentals from statistical objects already established as identified by prior work.

[3]Though much of the early structural auction literature focused on the no-entry case, auctions with entry have also received substantial attention in recent years. See, e.g., Athey et al. [2011], Bajari and Hortacsu [2003], Li [2005], Li and Zhang [2010a,b], Li and Zheng [2012, 2009], Hendricks et al. [2003], Krasnokutskaya and Seim [2009], MSX, and Roberts and Sweeting [2010b,a] among others.

[4]As usual, it is trivial to extend all results to the case of nonexcludable covariates x_ℓ.

We therefore focus on RS auctions which are *Stage 2 identified:*

Definition (Stage-2 identified)**.** RS mechanism M is *Stage-2 identified* if, for any marginal distribution $F^*(v)$ and any $n > 1$, a sample of observed bids $\mathbf{b}_\ell$ generated by n bidders competing (without entry) under M based on draws from $F^*(v)$ would permit consistent nonparametric estimation of $F^*(v)$.

Assumption 6 (Stage 2 identification)**.** *Process $\mathcal{L}$ involves an auction mechanism M that is Stage 2 identified.*

The class of Stage-2 identified mechanisms includes all standard auctions (first-price, English, Vickery, and Dutch), so this focus is not particularly restrictive. We refer readers to Athey and Haile [2005] for further details.[5]

Second, we assume that the distribution $F(v, s)$ and the cost function $c(z)$ are invariant across auctions in $\mathcal{L}$:

Assumption 7. *For each auction ℓ generated by process $\mathcal{L}$, $F_\ell(v, s) = F(v, s)$ and $c_\ell = c(z_\ell)$.*

As usual, our results generalize immediately to the case of observable auction-level heterogeneity x_ℓ: simply condition all statistical objects on x_ℓ in the arguments below. We return to the case of unobservable auction-level heterogeneity u_ℓ as an extension.

Finally, to obtain meaningful restrictions on the joint distribution $F(v, s)$, we require at least some observable variation attributable to signals s. In the context of the AS model, this translates into requiring variation in factors affecting the equilibrium entry threshold $\bar{s}$: namely, potential competition N_ℓ and instruments z_ℓ. We therefore introduce a key exclusion restriction:

[5] An interesting question we do not address here is how to apply our method when the Stage 2 distribution is only partially identified (e.g. the model of Haile and Tamer [2003]). In principle, this should be a relatively straightforward extension, but would complicate notation and discussion.

Assumption 8 (Exogenous entry variation)**.** *Process $\mathcal{L}$ involves either exogenous variation in N_ℓ for fixed z, or exogenous variation in z_ℓ for fixed N, or both.*

Exogenous variation in N_ℓ has been used in several prior studies for testing purposes: for instance, Haile et al. [2003] use such variation to construct a test for common values, and MSX use it to construct tests for competing entry specifications. Exogenous variation in z_ℓ directly extends a long tradition of instrumental variables in econometrics.[6] Both induce variation in the equilibrium threshold $\bar{s}$, which we in turn exploit as a source of identifying information on model fundamentals.

Directly identified objects

For each $(z, N) \in \mathcal{L}$, a large sample from process $\mathcal{L}$ will directly identify two statistical objects. First, given an equilibrium threshold $\bar{s}$, the probability that any particular bidder enters is simply $1 - \bar{s}$. We can thus identify the equilibrium threshold $\bar{s}(z, N)$ corresponding to (z, N) directly from observed entry decisions:

$$\bar{s}(z, N) \equiv 1 - E[n_\ell | z, N].$$

Second, by hypothesis, the mechanism M is Stage-2 identified. For each $(z, N) \in \mathcal{L}$, we can therefore recover the value distribution $F^*(\cdot|z, N)$ corresponding to bids submitted at (z, N). By Proposition I.2, this will represent the *ex post* distribution of values among entrants at (z, N):

$$F^*(v; \bar{s}(z, N)) \equiv F^*(\cdot|z, N),$$

where as above $F^*(v; \bar{s}) \equiv F(v|S_i \geq \bar{s})$.

[6] Recent work using instrumental variables to address identification of nonparametric models includes Chesher [2005] for nonparametric identification of models with discrete endogenous variables and Berry and Haile [2010a] for nonparametric identification of multinomial choice demand models, to name only a few.

Let $\mathcal{S}(\mathcal{L})$ denote the set of equilibrium thresholds induced by process $\mathcal{L}$:

$$\mathcal{S}(\mathcal{L}) \equiv \{s \in [0,1] | s = \bar{s}(z, N) \text{ for some } (z, N) \in \mathcal{L}\}.$$

The objects identified by process $\mathcal{L}$ are then $\mathcal{S}(\mathcal{L})$ itself and the *ex post* distribution $F^*(v; \bar{s})$ for each $\bar{s} \in \mathcal{S}(\mathcal{L})$. For notational compactness, we state all subsequent identification results in terms of the identified set $S(\mathcal{L})$.

Identified objects versus model fundamentals

Identification in the AS model requires recovery of the joint distribution $F(v, s)$, but the AS process $\mathcal{L}$ directly identifies only distributions of values conditional on entry: $F^*(v; \bar{s}) \equiv F(v|S_i \geq \bar{s})$ for each $\bar{s} \in \mathcal{S}(\mathcal{L})$. The relationship between these two distributions can be expressed as follows:

$$F(v, s) = F^*(v; 0) - (1 - s)F^*(v; s). \tag{II.1}$$

This fact in turn suggests the AS-model equivalent of a full-support condition: if the identified set $\mathcal{S}(\mathcal{L})$ spans the unit interval, then $F(v, s)$ will be fully identified. Unfortunately, this condition is likely to fail in many applications of interest. In particular, if no instrument z_ℓ is available, then all variation in $\bar{s}$ will be driven by variation in N_ℓ and $\mathcal{S}(\mathcal{L})$ will be a finite set. In such cases $F(v, s)$ will not be fully identified, so nonparametric analysis must fall back on identified bounds.

II.2 Identified bounds on fundamentals

This section establishes our core identification results, building on the insight that exogenous variation in $\bar{s}$ generates practically useful information on model fundamentals. The precision of this information depends on the nature of variation in $\bar{s}$: the AS model will be *partially identified* in DGPs where $\bar{s}$ takes a discrete set of values in equilibrium, but may be *point identified* when equilibrium $\bar{s}$ takes a continuum of values. We therefore frame results in this section as a map from the identified set $\mathcal{S}$ to natural identified bounds on model fundamentals. As above, the objects of interest are the joint distribution $F(v, s)$ and the entry cost c.

To establish the main results in this section, we will need some additional notation. Define *nearest-neighbor* functions $t^+(s)$ and $t^-(s)$ as follows:

$$t^+(s) = \begin{cases} \inf\{t \in \mathcal{S} | t > s\} & \text{if } \max\{\mathcal{S}\} > s; \\ 1 & \text{otherwise.} \end{cases}$$

$$t^-(s) = \begin{cases} \sup\{t \in \mathcal{S} | t < s\} & \text{if } \min\{\mathcal{S}\} < s \\ 0 & \text{otherwise.} \end{cases}$$

The intuition behind these functions is simple: for any $t \in [0, 1]$, return the nearest upper and lower neighbors of t in the identified set $\mathcal{S}$. The implementation captures three important special cases: return uninformative bounds if t is outside the scope of $\mathcal{S}$, return t if $t \in \text{int}(\mathcal{S})$, and return nearest elements in $\mathcal{S}$ not equal to t otherwise.

Bounds on distributions: $F(v|s)$ and $F(v, s)$

As a first step toward obtaining bounds on $F(v, s)$ and $c(\cdot)$, we derive bounds on the conditional distribution $F(v|s)$. Toward this end, rewrite the *ex post* distribution $F^*(v; \bar{s})$ in terms of $F(v|s)$ as follows:

$$F^*(v; \bar{s}) \equiv \frac{1}{1 - \bar{s}} \int_{\bar{s}}^1 F(v|t)dt.$$

This identity in turn implies

$$F(v|\bar{s}) = -\frac{\partial}{\partial \bar{s}}[(1 - \bar{s})F^*(v; \bar{s})].$$

Both $\bar{s}$ and $F^*(v;\bar{s})$ are identified for $\bar{s} \in \mathcal{S}$. Consequently, if $\bar{s} \in \text{int}(\mathcal{S})$, we can identify the RHS derivative exactly. Otherwise, we can approximate via a small finite difference $\Delta\bar{s}$:

$$F(v|\bar{s}) \approx -\frac{\Delta[(1-\bar{s})F^*(v;\bar{s})]}{\Delta\bar{s}}.$$

In practice, the smallest differences we can recover will be $\Delta\bar{s} = [t^+(\bar{s}) - \bar{s}]$ and $\Delta\bar{s} = [\bar{s} - t^-(\bar{s})]$. These differences yield natural identified approximations to $F(v|\bar{s})$, which by affiliation can also be shown to bound $F(v|\bar{s})$. We thus obtain the following lemma:

Lemma II.1. *Choose any $\bar{s} \in \mathcal{S}$, and define $\check{F}^+(v|\bar{s})$ and $\check{F}^-(v|\bar{s})$ as follows:*

$$\check{F}^+(v|\bar{s}) = \begin{cases} \lim_{t \uparrow t^-(\bar{s})} \left\{ \frac{(1-t)F^*(v;t) - (1-\bar{s})F^*(v;\bar{s})}{\bar{s}-t} \right\} & \text{if } t^-(\bar{s}) \in \mathcal{S}; \\ 1 & \text{otherwise.} \end{cases}$$

$$\check{F}^-(v|\bar{s}) = \begin{cases} \lim_{t \downarrow t^+(\bar{s})} \left\{ \frac{(1-\bar{s})F^*(v;\bar{s}) - (1-t)F^*(v;t)}{t-\bar{s}} \right\} & \text{if } t^+(\bar{s}) \in \mathcal{S}; \\ 0 & \text{otherwise.} \end{cases}$$

Then $\check{F}^+(v|\bar{s})$ and $\check{F}^-(v|\bar{s})$ are identified for all v, represent distributions over $[\underline{v}, \bar{v}]$, and bound the conditional distribution $F(v|\bar{s})$:

$$\check{F}^+(v|\bar{s}) \geq F(v|\bar{s}) \geq \check{F}^-(v|\bar{s}) \, \forall v,$$

with equality whenever $\bar{s} \in \text{int}(\mathcal{S})$.

We can thus obtain identified bounds on $F(v|\bar{s})$ for any $\bar{s} \in \mathcal{S}$. The next proposition extends this result into bounds on $F(v|s)$ for any $s \in [0,1]$:

Proposition II.1 (Bounds on $F(v|s)$). *For any $s \in [0,1]$, define $F^+(v|s)$ and $F^-(v|s)$ as follows:*

$$F^+(v|s) = \begin{cases} \check{F}^+(v|s) & \textit{if } s \in \mathcal{S}; \\ \check{F}^+[v|t^-(s)] & \textit{if } s \notin \mathcal{S}. \end{cases}$$

$$F^-(v|s) = \begin{cases} \check{F}^-(v|s) & \text{if } s \in \mathcal{S}; \\ \check{F}^-[v|t^+(s)] & \text{if } s \notin \mathcal{S}. \end{cases}$$

Then $F^+(v|s)$ and $F^-(v|s)$ are identified, represent distributions over $[\underline{v}, \bar{v}]$, and bound $F(v|s)$:

$$F^+(v|s) \geq F(v|s) \geq F^-(v|s),$$

with equality whenever $s \in int(\mathcal{S})$.

This proposition naturally extends the logic of Lemma II.1: given any $s \in [0,1]$, find the nearest identified neighbors $\bar{s} \in \mathcal{S}$. Lemma II.1 gives identified bounds on $F(v|\cdot)$ at these neighbors, and by affiliation these bounds will also apply to $F(v|s)$. Local variation in $\bar{s}$ permits approximations become exact, and exact identification follows. For clarity, we restate this latter fact as a corollary:

Corollary II.1. *$F(v|\bar{s})$ is exactly identified for any $\bar{s} \in int(\mathcal{S}(\mathcal{L}))$.*

Finally, identified bounds on the conditional distribution $F(v|s)$ immediately imply identified bounds on the joint distribution $F(v,s)$:

Corollary II.2. *Define $F^+(v,s)$ and $F^-(v,s)$ as follows:*

$$F^+(v,s) = \int_0^s F^+(v|t)dt$$

$$F^-(v,s) = \int_0^s F^-(v|t)dt.$$

Then $F^+(v,s)$ and $F^-(v,s)$ are identified and $F^+(v,s) \geq F(v,s) \geq F^-(v,s)$.[7]

Bounds on $c(z)$

Now consider identification of $c(z)$. By Proposition I.2, at any (z, N) with nontrivial entry, the equilibrium entry threshold $\bar{s}_N(z)$ must satisfy a breakeven condition of the form:

$$\begin{aligned} c(z) &\equiv E_v\left[\pi(v; \bar{s}_N(z), N) | S_i = \bar{s}_N(z)\right] \\ &= \int_0^{\bar{v}} \alpha(y) \cdot [1 - F(y|\bar{s}_N(z))] \cdot F_w^*(y; \bar{s}_N(z))^{N-1} dy - \rho. \end{aligned} \tag{II.2}$$

ρ and $\alpha(\cdot)$ are known from the Stage 2 auction, $F_w^*(y;\bar{s}) \equiv \bar{s} + (1-\bar{s})F^*(y;\bar{s})$ is identified directly for $\bar{s} \in \mathcal{S}$, and the RHS integral is decreasing in $F(y|\bar{s})$. Identified bounds on $F(y|s)$ thus immediately imply identified bounds on $c(z)$:

Proposition II.2. *For any $(N, z) \in \mathcal{L}$, define $c_N^+(z)$ and $c_N^-(z)$ as follows:*

$$c_N^+(z) = \pi_0(\bar{s}_N(z), N) + \int_{\underline{\mathrm{v}}}^{\bar{v}} \alpha(y) \cdot [1 - F^-(y|\bar{s}_N(z))] \cdot F_w^*(y; \bar{s}_N(z))^{N-1} dy.$$

$$c_N^-(z) = \pi_0(\bar{s}_N(z), N) + \int_{\underline{\mathrm{v}}}^{\bar{v}} \alpha(y) \cdot [1 - F^+(y|\bar{s}_N(z))] \cdot F_w^*(y; \bar{s}_N(z))^{N-1} dy.$$

Then $c_N^+(z)$ and $c_N^-(z)$ are identified and $c_N^+(z) \geq c(z) \geq c_N^-(z)$, with equality if $\bar{s}_N(z) \in int(\mathcal{S}(\mathcal{L}))$.

As usual, we can pool across N to obtain sharper bounds on $c(z)$. Let $c^+(z)$

[7] If desired, we could extend this proposition to incorporate the Frechet-Hoeffding bounds (see Nelsen [1999]):

$$F_v(v) + s - 1 \leq F(v,s) \leq \min\{F_v(v), s\}.$$

In practice, $F_v(v)$ may be unknown, but can be bounded as follows:

$$F^*(v; \min \mathcal{S}) \leq F_v(v) \leq \min \mathcal{S} + (1 - \min \mathcal{S})F^*(v; \min \mathcal{S}).$$

Combining these results yields alternative bounds on $F(v,s)$, which in some cases might be tighter than those above.

and $c^-(z)$ denote these intersection bounds:

$$c^+(z) = \min_{N \in \mathcal{L}} c_N^+(z)$$
$$c^-(z) = \max_{N \in \mathcal{L}} c_N^-(z).$$

Finally, by Corollary II.1, $F(v|\bar{s})$ is exactly identified if $\bar{s} \in \text{int}(\mathcal{S}(\mathcal{L}))$. Equation (II.2) then implies an analogous condition for exact identification of entry costs:

Corollary II.3. *For any z such that $\bar{s}_N(z) \in int(\mathcal{S}(\mathcal{L}))$ for some N, $c(z)$ is exactly identified.*

Full identification

Corollaries II.1, II.2, and II.3 convey a common message: the AS model is exactly identified (almost) everywhere if we observe data generated at (almost) every $\bar{s} \in [0, 1]$. The conditions under which this will occur will depend on the nature of the underlying fundamentals, but will require an excluded instrument z that induces sufficient variation in entry cost $c(z)$. The next proposition formalizes this intuition.

Proposition II.3. *Suppose the econometrician observes an instrument z satisfying Assumptions 7 and 8 above, which has positive support on a set $Z \subset R^k$. Then the following statements hold:*

1. *If $z \in int(Z)$, then $c(z)$ is identified and $F(v|\bar{s})$ is locally identified at each $\bar{s} \in \mathcal{S}(\mathcal{L}(z))$.*

2. *If $[0, \bar{v}] \subset c(Z)$, then $F(v|s)$, $F(v, s)$, and $c(\cdot)$ are fully identified.*

The sufficient Condition 2 can probably be relaxed somewhat in many applications: as noted above, the fundamental property needed for full identification is $\text{cl}(\mathcal{S}(\mathcal{L})) = [0, 1]$. Condition 2 merely ensures that this property will hold absent further restrictions on N and $F(v, s)$.

Sharp bounds

Full identification in the general AS model depends on strong conditions, which can easily fail in practice. In such cases, nonparametric analysis must fall back on bounds like those in subsections II.2 and II.2. These bounds represent natural, intuitive, and directly estimable approximations to the fundamentals of interest, but may not exhaust all variation in the data. This subsection formally characterizes sharp bounds in the AS model.

Define a *candidate model* corresponding to process $\mathcal{L}$ as follows:

Definition II.1 (Candidate model)**.** A *candidate model* for process $\mathcal{L}$ is any pair $\{\tilde{F}(\cdot|\cdot), \tilde{c}(\cdot)\}$ satisfying the following conditions for all $z \in \mathcal{L}$:

1. **Distribution:** for all $s \in [0,1]$, $\tilde{F}(\cdot|s)$ defines a distribution over $[\underline{v}, \bar{v}]$.

2. **Selection:** $\tilde{F}(\cdot|\cdot)$ implies the set of distributions identified by sub-process $\mathcal{L}(z)$*:* $(1-s)F^*(v;s) = \int_s^1 F(v|t)dt$ *for all* v*, for all* $s \in \mathcal{S}(z)$*.*

3. **Affiliation:** $\tilde{F}(\cdot|\cdot)$ implies a joint distribution $\tilde{F}(v,s) \equiv \int_0^s \tilde{F}(v|t)dt$ satisfying affiliation.

4. **Entry:** $\Pi^*(s, N_s(z); \tilde{F}) \equiv \tilde{c}(z)$ for all $s \in \mathcal{S}(z)$, where $N_s(z)$ denotes the competition level N corresponding to s under $\mathcal{L}(z)$ and

$$\Pi^*(s, N, \tilde{F}) \equiv \int_V \pi(v; s, N) d\tilde{F}(v|s). \qquad \text{(II.3)}$$

As defined above, $F^+(v|s)$ and $F^-(v|s)$ directly exploit only the distribution and selection conditions of Definition II.1.[8] Hence there could exist a (v, s) pair such

[8]More precisely, we incorporate a slightly weaker "differences" version of the selection condition (2):

$$(1-s)F^*(v;s) - (1-s')F^*(v;s') = \int_s^{s'} F^+(v|t)dt$$

for $s' > s$, with F^- substituted for F^+ when $s' < s$.

that (say) *no* candidate $\tilde{F}(\cdot|\cdot)$ attaining the upper bound $F^+(v|s)$ at (v,s) could simultaneously satisfy the entry condition (4) for all other $s \in \mathcal{S}(z)$. Consequently, the bounds in sections II.2 and II.2 would not be sharp.

Extending this intuition yields a characterization of the sharp identified set. Let $\mathcal{M}(\mathcal{L})$ be the set of all candidate models $\{\tilde{F}(\cdot|\cdot), \tilde{c}(\cdot)\}$ for process $\mathcal{L}$. By assumption, the true model $\{F(v|s), c(\cdot)\}$ satisfies Definition II.1, so $\mathcal{M}(\mathcal{L})$ is nonempty. Then sharp bounds will be given by the upper and lower envelopes of $\mathcal{M}(\mathcal{L})$:

Lemma II.2. *Let $\mathcal{M}(\mathcal{L})$ be the set of all candidate models at $\mathcal{L}$. Then $\mathcal{M}(\mathcal{L})$ is the sharp identified set at $\mathcal{L}$, and implies sharp bounds $\{\tilde{c}^+, \tilde{c}^-\}$ and $\{\tilde{F}^+, \tilde{F}^-\}$ as follows:*

$$\begin{aligned}
\tilde{c}^+(z) &= \sup\{\tilde{c}(z) \ s.t. \ \tilde{c}(\cdot) \in \mathcal{M}(\mathcal{L})\} \\
\tilde{c}^-(z) &= \inf\{\tilde{c}(z) \ s.t. \ \tilde{c}(\cdot) \in \mathcal{M}(\mathcal{L})\}
\end{aligned}$$

and for each (v,s)

$$\begin{aligned}
\tilde{F}^+(v|s) &= \sup\{\tilde{F}(v|s) \ s.t. \ \tilde{F}(v|s) \in \mathcal{M}(\mathcal{L})\} \\
\tilde{F}^-(v|s) &= \inf\{\tilde{F}(v|s) \ s.t. \ \tilde{F}(v|s) \in \mathcal{M}(\mathcal{L})\}.
\end{aligned}$$

Implementation of Lemma II.2 could perhaps be attempted using sieve or other functional approximation methods, but would likely be challenging in practice. Hence we focus primarily on the (more interesting) constructive bounds in II.2 and II.2, but note that in principle sharper bounds might exist.[9]

[9] Our discussion in this respect parallels that of Ciliberto and Tamer [2009b] on sharpness in oligopoly entry models. In particular, at any $t \in \mathcal{S}$, our bounds on $F(v|t)$ exploits information generated by the nearest-neighbors $s^+(t)$ and $s^-(t)$, but not that potentially generated by more distant entry equilibria. Meanwhile, Ciliberto and Tamer [2009b] exploit zero-one bounds on probability magnitudes, but not the condition that probabilities must sum to one. In both cases the condition omitted is a cross-equation restriction which, though potentially informative, would be very difficult to implement.

II.3 Bounds on counterfactual revenue

This section considers counterfactual analysis based on the AS model. In particular, we start from the set of entry thresholds $S(\mathcal{L})$ identified by process $\mathcal{L}$, and seek to characterize expected seller revenue R_M corresponding to counterfactual RS auction M. This problem is complicated by the generality of the AS model: in the presence of endogenous and selective entry, M will affect seller revenue directly, through the Stage 1 entry threshold $\bar{s}$, and through the selected Stage 2 distribution $F^*(\cdot;\bar{s})$. We derive revenue bounds accounting for all three effects in the presence of partial identification induced by AS entry, where as usual we assume the no-sale outcome yields value $v_0 \leq \underline{v}$ to the seller.

Our first step toward this end is to characterize seller revenue $R_M(\bar{s}, N)$ generated by mechanism M under *arbitrary* (out-of-equilibrium) threshold $\bar{s}$:

Lemma II.3. *Under Assumptions 1-5, expected seller revenue under RS auction M at competition structure $(\bar{s}, N)$ is given by*

$$R_M(\bar{s};N) = \int_{v_0}^{\bar{v}} \{\alpha_M(y)[y - \lambda_M(y;\bar{s},N)] + [1-\alpha_M(y)]v_0\}\, dG^*_{1:N}(y;\bar{s}) + N(1-\bar{s})\rho_M \tag{II.4}$$

where

$$\lambda_M(v;s,N) \equiv \begin{cases} 0 & \text{if } \alpha_M(v) = 0; \\ \int_{v_0}^{v} \frac{\alpha_M(t)}{\alpha_M(v)} \cdot \frac{F^*_w(t;s)^{N-1}}{F^*_w(v;s)^{N-1}} dt & \text{otherwise.} \end{cases}$$

Further, considered as a function of $\bar{s}$, $R_M(\bar{s};N)$ satisfies the following properties:

1. *$R_M(\bar{s};N)$ is decreasing in $\bar{s}$ for all N.*

2. *$R_M(\bar{s};N)$ is identified for any $\bar{s} \in \mathcal{S}(\mathcal{L})$.*

We now consider counterfactual entry. When entry is endogenous, the equilibrium threshold $\bar{s}_M$ will depend on M, and when $F(v|s)$ and $c(\cdot)$ are not identified, this dependence will not be determined exactly. However, we can use the bounds on fundamentals derived above to *bound* the relationship between $\bar{s}_M$ and M:

Lemma II.4. *Choose any (z,N) and let $c^+(z)$ and $c^-(z)$ be identified bounds on $c(z)$, $F^+(v|\bar{s})$ and $F^-(v|\bar{s})$ be identified bounds on $F(v|s)$, and $\bar{s}_M(z,N)$ be the (unknown) equilibrium entry threshold induced by counterfactual mechanism M. Define $s_M^+(z,N)$ and $s_M^-(z,N)$ as follows:*

$$s_M^+(z,N) = \begin{cases} \inf\{s \in \mathcal{S} | \Pi_M(s,N;F^+) > c^+(z)\} & \textit{if } \exists \textit{such } s; \\ 1 & \textit{otherwise} \end{cases}$$

$$s_M^-(z,N) = \begin{cases} \sup\{s \in \mathcal{S} | \Pi_M(s,N;F^-) < c^-(z)\} & \textit{if } \exists \textit{such } s; \\ 0 & \textit{otherwise} \end{cases}$$

where $\Pi_M(\cdot;\tilde{F})$ denotes expected profit of an entrant with signal s given conditional distribution $\tilde{F}$ and mechanism M:

$$\Pi_M(s,N;\tilde{F}) = \int_{v_0}^{\bar{v}} [1 - \tilde{F}(y|s)]\alpha_M(y) F_w^*(y;s)^{N-1} dy - \rho_M.$$

Then $s_M^+(z,N)$ and $s_M^-(z,N)$ are identified and $s_M^+(z,N) \geq \bar{s}_M(z,N) \geq s_M^-(z,N)$, with equality if $s_M(z,N) \in int(\mathcal{S}(\mathcal{L}))$.

Finally, combining results in Lemmas II.3 and II.4 produces identified bounds on expected seller revenue $R_M(z,N)$ corresponding to counterfactual mechanism M:

Proposition II.4. *Choose any (z,N), define $s_M^+(z,N)$ and $s_M^-(z,N)$ as in Lemma II.4, and let $R_M(z,N)$ be (unknown) expected revenue under $\alpha(\cdot)$ at (z,N). Define*

$R^+_M(z,N)$ and $R^-_M(z,N)$ as follows:

$$R^-_M(z,N) = \begin{cases} R_M(s^+_M(z,N);N) & \textit{if } s^+_M(z,N) \in \mathcal{S}(\mathcal{L}) \\ 0 & \textit{otherwise} \end{cases}$$

$$R^+_M(z,N) = \begin{cases} R_M(s^-_M(z,N);N) & \textit{if } s^-_M(z,N) \in \mathcal{S}(\mathcal{L}) \\ \check{R}_M(0;N) & \textit{otherwise}, \end{cases}$$

where

$$\check{R}_\alpha(0;N) = \int_{v_0}^{\bar{v}} \left\{ \alpha(y)[y - \int_{v_0}^{y} \frac{\alpha(t)}{\alpha(y)} \cdot \frac{F^*(t;\min\mathcal{S})^{N-1}}{F^*(y;\min\mathcal{S})^{N-1}}]dt + [1-\alpha(y)]v_0 \right\} dG^*_{1:N}(y;\min\mathcal{S}) + N\rho_M$$

is a semi-informative upper bound applicable when $s^-_M(z,N) \equiv 0$.[10]

Then $R^+_M(z,N)$ and $R^-_M(z,N)$ are identified and $R^+_M(z,N) \geq R_M(z,N) \geq R^-_M(z,N)$, with equality if $s_M(z,N) \in int(\mathcal{S}(\mathcal{L}))$.

The intuition behind this result is straightforward: Lemma II.3 establishes that conditional revenue is decreasing in $\bar{s}$ for any M, Lemma II.4 establishes bounds s^+_M and s^-_M on the counterfactual entry threshold $\bar{s}_M$, and we know $R_M(\bar{s},N)$ is identified for any $\bar{s} \in \mathcal{S}(\mathcal{L})$. Plugging entry bounds s^+_M and s^-_M into $R_M(\bar{s},N)$ then yields bounds on expected revenue accounting for endogenous and selective entry.

Revenue bounds applicable to the special case of a public reserve price can easily be obtained by setting $\alpha_M(y) \equiv \mathbf{1}[y \geq r]$ in Proposition II.4. In some applications, however, researchers may wish to characterize not just expected seller revenue but also the seller's *optimal reserve price* (ORP). Consequently, following Haile and Tamer [2003], we translate the revenue bounds above into bounds on the seller's ORP:

Corollary II.4 (Bounds on Seller's ORP)**.** *Let $R^+(r;z,N)$ and $R^-(r;z,N)$ be defined*

[10] In particular, $\check{R}_\alpha(0;N)$ is the revenue that would result if all potential bidders always enter but draw values from distribution $F^*(\cdot;\min\mathcal{S})$. Since $\min\mathcal{S} \geq 0$, we know $F^*(v;\min\mathcal{S}) \leq F^*(v;0)$, so $\check{R}_\alpha(0;N) \geq R_\alpha(0;N) \geq R_\alpha(s_\alpha;N)$. Further, since $\min\mathcal{S} \in \mathcal{S}$, $\check{R}(0;N)$ is identified.

as in Proposition II.4, $R^*_- \equiv \sup_r R^-(r; z, N)$ *be the maximum value of the lower bound, and* $r^*(z, N)$ *be the seller's optimal reserve price at* (z, N)*. Define* $r^*_+(z, N)$ *and* $r^*_-(z, N)$ *as follows:*

$$\begin{aligned} r^*_+(z, N) &= \sup\{r | R^+(r; z, N) \geq R^*_-\} \\ r^*_-(z, N) &= \inf\{r | R^+(r; z, N) \geq R^*_-\}. \end{aligned}$$

Then r^*_+ *and* r^*_- *are identified and* $r^*_+(z, N) \geq r^*(z, N) \geq r^*_-(z, N)$*, with equality if* $s_{r^*}(z, N) \in int(\mathcal{S}(\mathcal{L}))$*.*

Thus, to summarize: we obtain identified bounds on expected revenue applicable to a wide range of counterfactual mechanisms in the general class of RS auctions accounting for endogenous and arbitrarily selective entry. To our knowledge, these are the first such results reported in the literature. Further, in the special case of a public reserve price, these revenue bounds can be translated into bounds on the seller's optimal reserve price following Haile and Tamer [2003]. We thus establish that the general AS model can support a rich variety of counterfactual and policy analyses under relatively weak assumptions on the nature of entry and selection.

II.4 Conclusion

In this paper, we explore a general approach to identification in auctions with entry based on a framework we call the AS model. In the process, we make three core contributions to the related literature. First, we derive nonparametric bounds on AS model fundamentals applicable to a general class of auctions with endogenous and arbitrarily selective entry. Second, we translate these bounds on fundamentals into bounds on expected revenue corresponding to a wide range of counterfactual award rules, again accounting for endogenous and selective entry. Finally, we outline conditions under which all bounds collapse to exact identification. To our knowledge, these are the first formal (partial) identification results applicable to the general AS model, and represent the most general treatment of identification in auctions with entry to date.

While our discussion thus far has focused on the case of symmetric bidders, our underlying logic extends readily to environments with asymmetry. In particular, suppose process $\mathcal{L}$ involves a set $\mathcal{T}$ of potential bidder types, where a bidder of type $\tau_i \in \mathcal{T}$ draws from affiliated joint distribution $F_{\tau_i}(v, s|\cdot)$ and has entry cost $c_{\tau_i}(\cdot)$.[11] Let $\tau \in \mathcal{T}^N$ be a vector of bidder types, and $\bar{s}(\tau)$ a corresponding vector of type-symmetric Stage 1 equilibrium entry thresholds. A repeated sample from equilibrium $\bar{s}(\tau)$ will permit identification of $\bar{s}_{\tau_i}(\tau)$ and $F^*_{\tau_i}(\cdot;\bar{s}_{\tau_i}(\tau))$ for each type τ_i, and pooling across other type vectors containing τ_i will generate a set of type-specific identified thresholds $\mathcal{S}_{\tau_i}(\mathcal{L})$. Identification of $F_{\tau_i}(\cdot)$ and $c_{\tau_i}(\cdot)$ can then proceed as above. The key additional complication in the asymmetric case is the potential presence of multiple Stage 1 equilibria, and the corresponding additional restriction is that the equilibrium played is either known or depends deterministically on τ. Given this restriction, however, asymmetry may substantially improve identification: in particular, a continuous typeset $\mathcal{T}$ may induce a continuous identified set $\mathcal{S}_{\tau_i}(\mathcal{L})$, and thus permit exact identification based only on variation in potential competition.

A second potentially important extension of our results is to environments with unobserved heterogeneity. As noted in our discussion of Assumption 7, this extension requires two additional restrictions. First, following Krasnokutskaya [2009], we assume auction-level unobserved heterogeneity u enters the value distribution either additively or multiplicatively: that is, $v_i = u\epsilon_i$ or $v_i = u + \epsilon_i$, where ϵ_i is IID across bidders.[12] Second, we assume the realization of u is revealed to bidders only after Stage 1 entry is complete. This second assumption seems plausible in applications where c is

[11] These types can be either discrete or continuous; we require only that continuous types affect model fundamentals continuously.

[12] Krasnokutskaya [2009] assumes u enters multiplicatively, but her argument can easily be adapted to the additive case as well.

interpreted as a cost of learning, and permits us to separate heterogeneity induced by the unobserved latent variable u from that induced by unobserved signals s_i.[13] The identification argument can then proceed as follows. First, following Subsection II.1, we know we can identify the set of equilibrium entry thresholds $\mathcal{S}(\mathcal{L})$ corresponding to process $\mathcal{L}$. For each $\bar{s} \in \mathcal{S}(\mathcal{L})$, the observed bid distribution will depend on two components: the ex-post selected distribution $F^*(\epsilon_i; \bar{s})$ and the distribution of unobserved heterogeneity $F_u(u)$. Second, following Krasnokutskaya [2009], each of these components can be recovered using deconvolution methods on an appropriate sample of observed bids. It follows that the ex-post distribution $F^*(\epsilon_i; \bar{s})$ is identified for any $\bar{s} \in \mathcal{S}(\mathcal{L})$. Finally, we can apply results in Section II.2 to the identified ex post distribution $F^*(\epsilon_i; \bar{s})$ to obtain bounds on remaining model fundamentals, and partial identification of the overall model follows immediately.

[13] For instance, in highway construction auctions, many features likely to generate project-level heterogeneity (e.g. exact location, specifics of work to be performed, soil type, etc) are observable only to bidders obtaining detailed project plans (the outcome typically taken to represent entry). Hence the assumption that u is revealed following entry seems reasonable. Without this separating restriction, identification would require us to successfully disentangle distribution effects directly attributable to heterogeneity (u) from those attributable to changes in endogenous selection (including those induced by shifts in u). In a fully nonparametric model with endogenous and arbitrarily selective entry, this is likely to prove impossible in practice.

Numerical example

Sections II.2 and II.3 develop identified bounds for model fundamentals and seller revenue in auctions with arbitrarily selective entry. In this appendix, we explore a simple numeric example designed to illustrate what these theoretical identified bounds might look like in practice. Consistent with our emphasis in the rest of the paper, this example focuses on *identification*, not estimation: the figures that follow illustrate the bounds that would obtain in an infinite auction sample. Nevertheless, this simple exercise should help to indicate what kind of information could in principle be recovered using the methods developed above.

Details of this example are as follows. We model the joint distribution $F(v, s)$ using a Gaussian copula $C_\rho(F_v, F_s)$, where the marginal distribution $F_v(\cdot) \sim N(\mu = 100, \sigma = 10)$ and as above we normalize $F_s(\cdot) \sim U[0, 1]$. The correlation parameter ρ measures the degree of affiliation between s and v, with $\rho = 0$ generating the no-information LS case and $\rho \to 1$ approaching the full-information S case. In what follows, we take $\rho = 0.75$ unless noted otherwise; other values of ρ yield quite similar results. Entry involves cost $c = 2$, and we assume potential competition varies exogenously on the set $N = \{4, 5, ..., 16\}$. These parameter values are chosen to be qualitatively similar to existing findings in the literature.[14]

Given this parametric specification, it is straightforward to calculate the set of equilibrium entry thresholds $\mathcal{S} = \{\bar{s}_4, ..., \bar{s}_{16}\}$ satisfying the the breakeven condition (I.7). From subsection II.1, we know these thresholds (and the corresponding value distributions $F^*(v; \bar{s})$ are the objects identified by a standard $(N, n, \mathbf{b})$ sample. We can then use results in Sections II.2 and II.3 to obtain identified bounds on quantities of interest as follows.

First, Proposition II.1 implies identified bounds on $F(v|s)$ for any (v, s). Figures

[14] See, e.g., Roberts and Sweeting [2010a, 2010b] and Li and Zheng [2009] for examples.

II.1 and II.2 illustrate these bounds (across s) for two values of v.

Second, based on Proposition II.2, we can translate identified bounds on $F(v|s)$ into identified bounds on c. Applying this proposition to our numeric example and pooling results across N yields identified bounds $c^+ = 2.026$ and $c^- = 1.971$, where (as noted above) true $c = 2$.

Third, following Section II.3, we can translate these bounds on $F(v|s)$ and c into identified bounds on counterfactual seller revenue. Results in Section II.3 are framed in terms of an arbitrary award rule $\alpha(\cdot)$, but for simplicity we here restrict attention to the special case of a counterfactual public reserve price r. As in Section II.3, the first step in this process is to obtain identified bounds on the counterfactual entry threshold $\bar{s}_r$ corresponding to each candidate reserve price r. These bounds are illustrated for $N = 6$ and $N = 9$ in Figures II.3 and II.4 below.

Finally, using Proposition II.4, we can translate bounds on $\bar{s}_r$ to bounds on counterfactual revenue R_r at any (N, r). These bounds are illustrated for $N = 6$ and $N = 9$ in Figures II.5 and II.6 below. In both cases, the seller's value is assumed to be $v_0 = 60$.

If so desired, we could also adapt the argument of Haile and Tamer [2003] to translate bound on counterfactual revenue to bounds on the optimal reserve price r as in Corollary II.4. In this example, the implied bounds on optimal r would be rather wide: an upper bound of roughly 90 when $N = 6$, and of roughly 100 when $N = 9$ (with uninformative lower bound $v_0 = 60$ in each case). However, in both cases the identified bounds on counterfactual *revenue* are surprisingly tight.

Figure II.1: IDed bounds on $F(v|s)$, $v = 95$

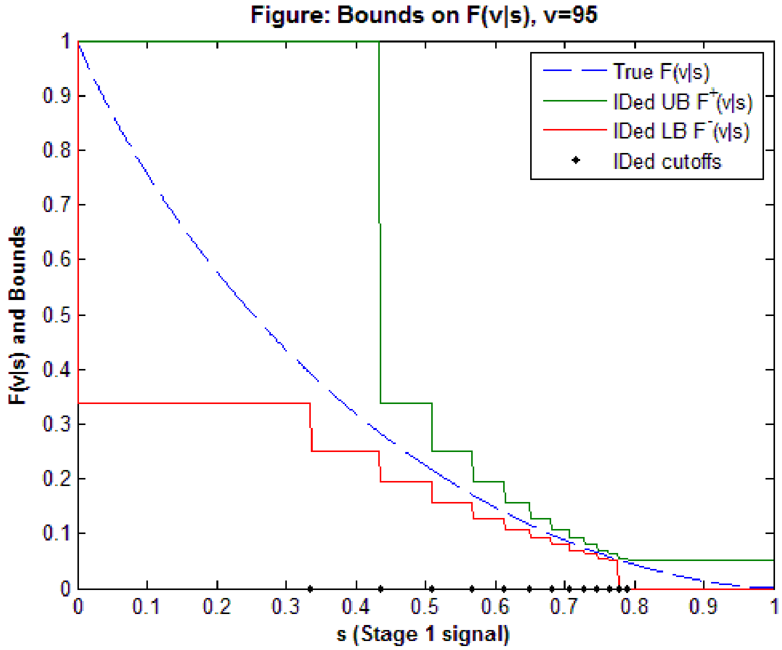

Figure II.2: IDed bounds on $F(v|s)$, $v = 105$

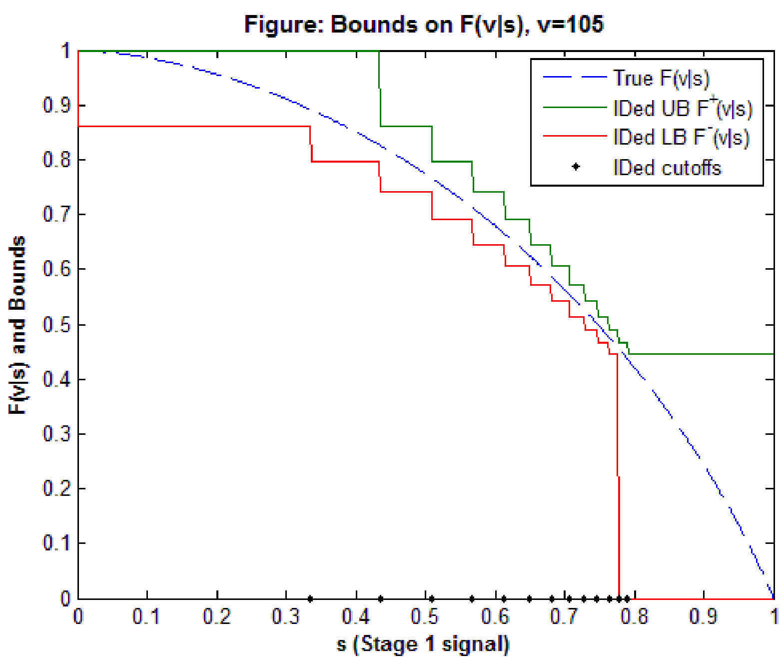

Figure II.3: IDed bounds on $\bar{s}_r$ at $N = 6$, $\rho = 0.75$

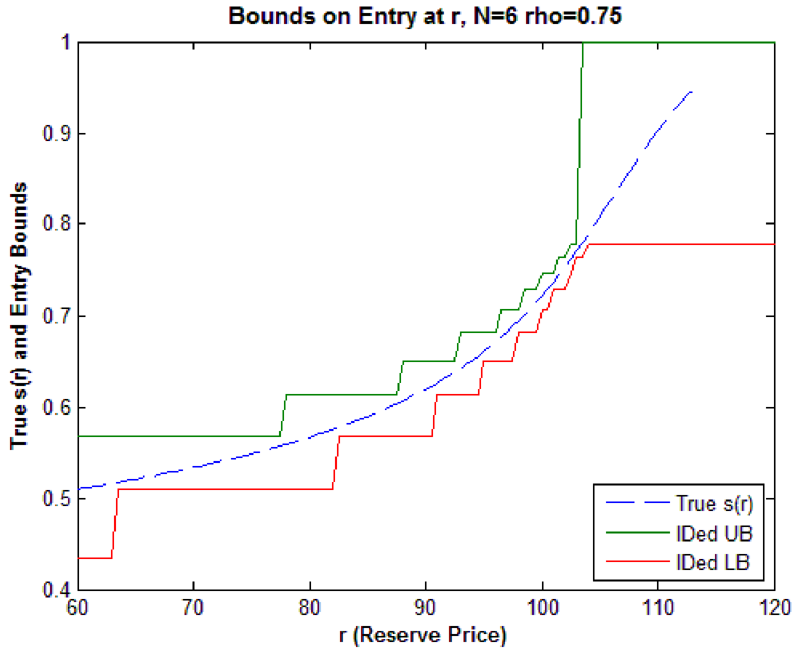

Figure II.4: IDed bounds on $\bar{s}_r$ at $N = 9$, $\rho = 0.75$

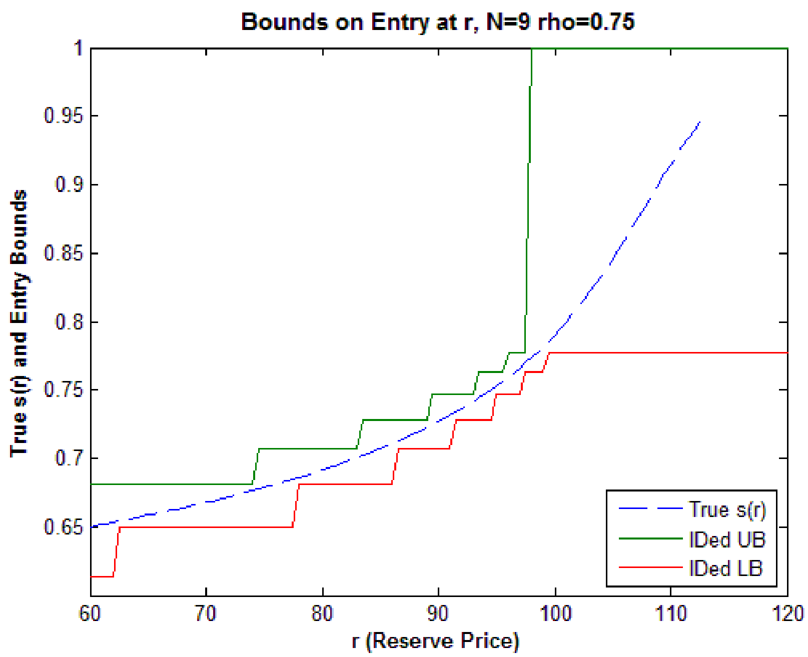

Figure II.5: IDed bounds on CF revenue at $N = 6$, $\rho = 0.75$

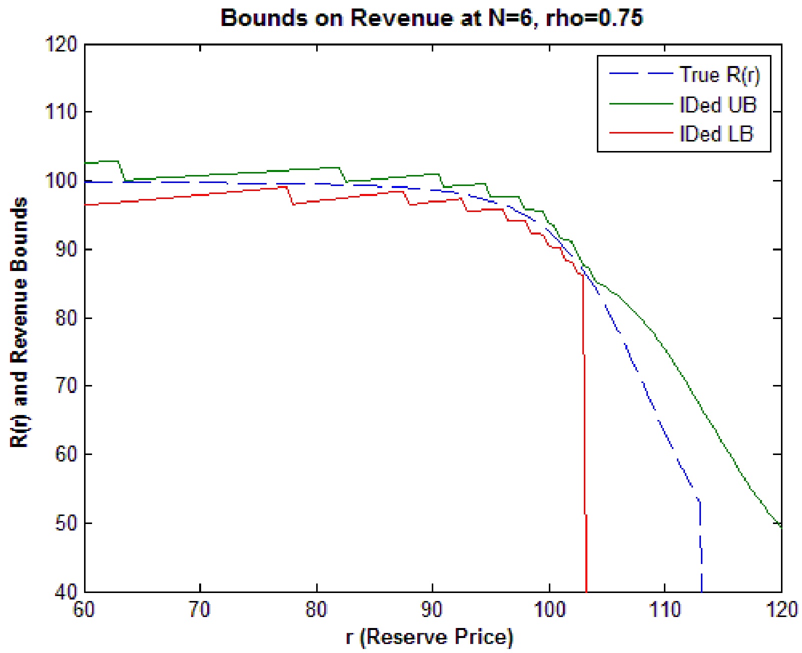

Figure II.6: IDed bounds on CF revenue at $N = 9$, $\rho = 0.75$

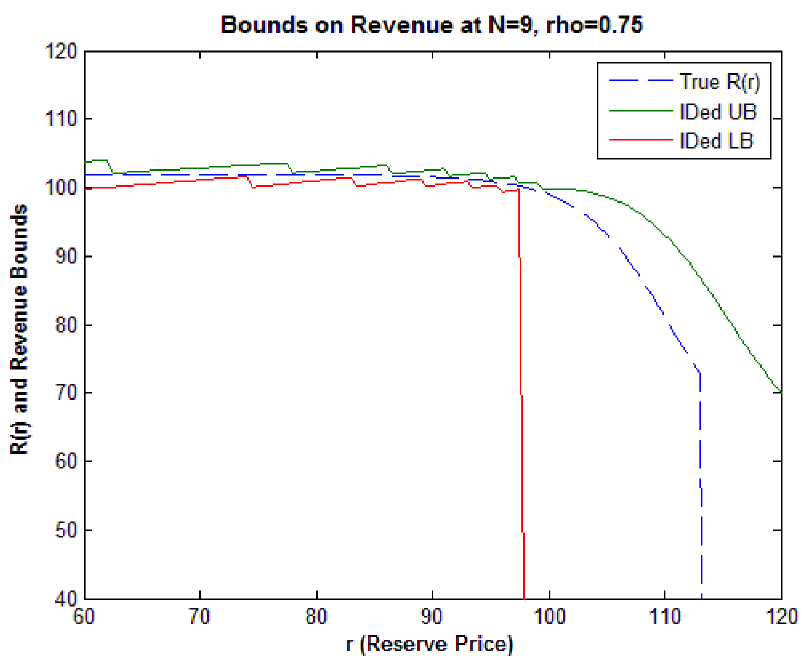

CHAPTER III

DISPLAYS, SALES, AND IN-STORE SEARCH IN RETAIL MARKETS

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 vari-

[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.

ations (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 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 existing studies in industrial organization.[5] First, using panel data on laundry detergent purchases, Hendel and Nevo [2006]

[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], Wildenbeest [2009], Moraga-Gonzalez et al. [2010], Santos et al. [2010], and Seiler [2011] to mention just a few.

[5]Another somewhat 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.

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 [2011] develops a dynamic choice model with a preliminary market entry decision, finding that in-store search is an important factor explaining purchase incidence in detergent markets. My main innovation 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 variation which to my knowledge has not been previously explored in the search literature.

The rest of this paper is organized as follows. Section III.1 summarizes my data and surveys price, display, and other promotional variation in laundry detergent markets. Section III.2 presents descriptive evidence on promotional effects in the target market, and Section III.3 describes the structural model I develop based on this evidence. Section III.4 outlines details of my estimation procedure, and Section III.5 presents key results. Finally, Section III.6 concludes.

III.1 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 [2011] 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 detergent sample. My descriptive regressions therefore focus on the Atlanta market, and my structural estimates on six of the largest detergent retailers (IRI identifiers 243785, 250094, 263568, 653369, and 683960) in this market. 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. In what follows, I refer to the Atlanta market simply as "the market."

[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 [2011] 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.

Table III.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 III.1 below.

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 III.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.[8] Finally, sales induce substantial short-term variation in prices, with average discounts in the neighborhood of 12-15 percent.[9] For completeness, Table B.1 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] 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.

[9] As discussed in detail in Sections III.1 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.

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.

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.[10] 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 price reductions, and (3) that most such price reductions are temporary. Figure III.1 presents one representive price history for the current market; Figures B.1 and B.2 in Appendix 1 present two additional examples.[11] All clearly illustrate the core pattern in question: a dominant regular price which changes infrequently, punctuated by frequent temporary sales.

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 av-

[10]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.

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

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

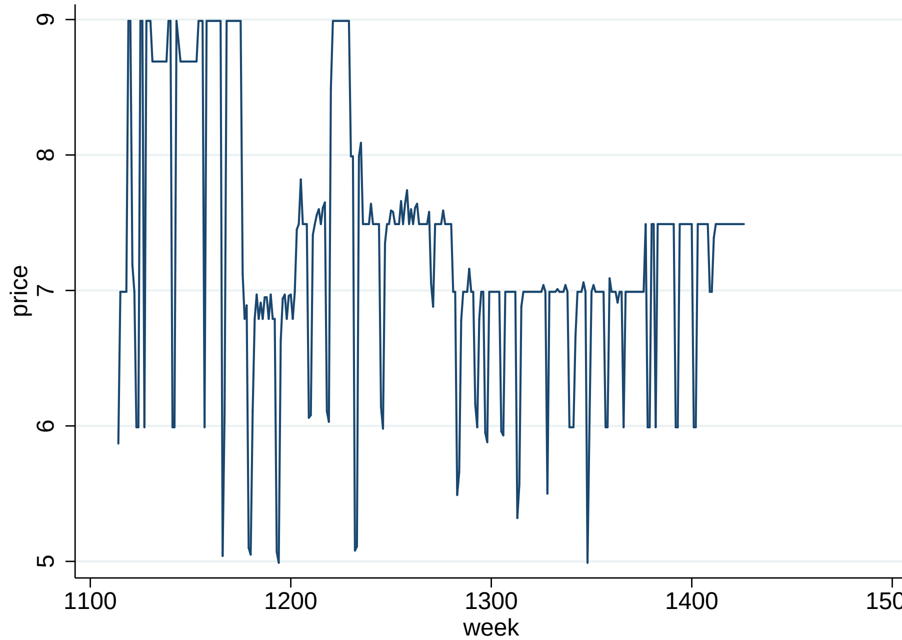

erage tastes should be relatively stable from week to week, short-run price variation induced by sales should help to identify price elasticities in a standard preference model.[12] Both insights play key roles in my subsequent structural analysis.

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

[12] 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.

prices.[13] Figure III.2 illustrates this procedure applied to the price series for Tide 100oz in Figure III.1; a formal description of the algorithm and additional examples are given in Appendix 2.

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

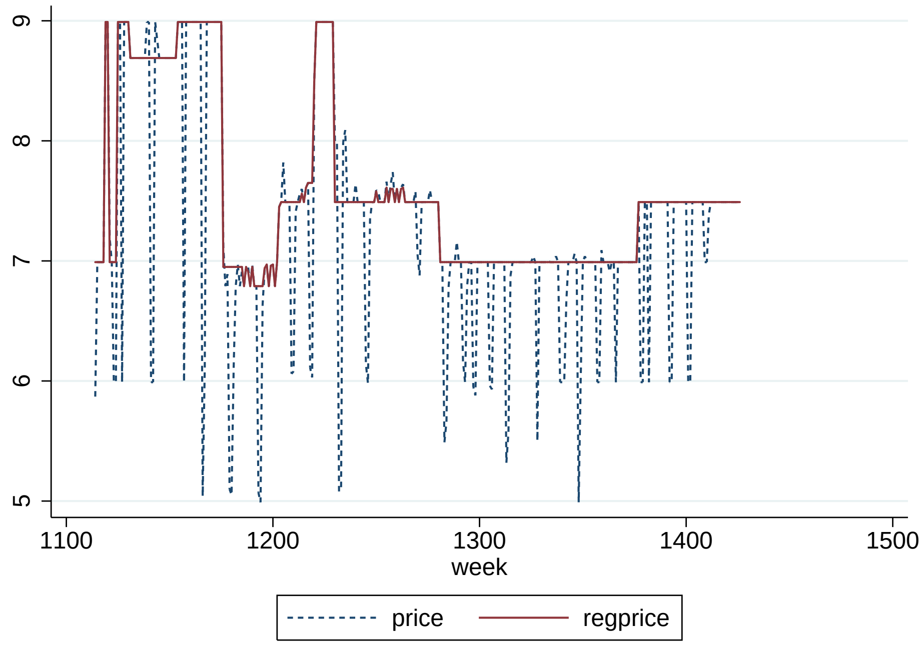

Sales, displays, and features

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 III.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

[13]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. [2008] 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 III.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.

should be more than sufficient to support the identification strategy pursued here.

III.2 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.[14] Table III.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 B.2 in Appendix 1 reports related regressions using several alternative outcome measures and covariate specifications; qualitative results are similar in all cases.

Several key patterns in Table III.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

[14] 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 III.3: UPC-level promotion effects, Atlanta detergent 2002-2007

	(1)	(2)	(3)	(4)
VARIABLES	Prices only	Promo dummies	Only interacts	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

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.[15]

More broadly, Table III.3 also highlights the importance of properly accounting for promotional effects in demand analysis. Since *discount* and *qnorm* are expressed as percentage deviations, coefficients on *discount* will roughly parallel own-price elas-

[15]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.

ticities.[16] Column (1) of Table III.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.

III.3 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, and is marketed via a time-varying two-part process: a weekly price p_{jt} and a weekly display indicator d_{jt}.[17] 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

[16] 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 III.1 these realizations are almost certainly not responding to week-to-week shifts in demand.

[17] 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.

process. First, upon entering the store, consumer i costlessly observes prices for all products in the *display 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 formally summarized in Assumptions 9-11 below.

Assumption 9. *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 10. *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 11. *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 9 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.[18] Assumption 10 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.[19] 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

[18]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.

[19]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.

simplification of the economic fundamentals involved. Finally, Assumption 11 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 using a linear structure similar to BLP [1995, 2004]:

Assumption 12. *Consumers have BLP-style linear preferences with heterogeneous tastes for product attributes x:*

$$v_{ijt} = -\alpha_i p_{jt} + \beta_i x_{jt} + \xi_{jt} + \epsilon_{ijt},$$

where (α_i, β_i) are idiosyncratic preference coefficients, ξ_{jt} is average residual taste for product j at time t, ϵ_{ijt} is consumer i's idiosyncratic residual taste for product j, and consumer types $\omega_i \equiv (\alpha_i, \beta_i, \epsilon_i)$ are drawn from a conditional 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}.

Finally, I assume search costs are drawn independent of tastes (but potentially conditional on observables) from a distribution with nonnegative support:

Assumption 13. *Search costs c_i are drawn conditionally independent of tastes ω_i from a distribution function $F_c(\cdot|\zeta_{it})$ with positive support on $\mathcal{C} \subseteq \mathbb{R}^+$ that depends (at most) on the individual- or market-level observable vector ζ_{it}.*

In practice, some consumers may enjoy shopping for its own sake, in which case the cost distribution $F_c(\cdot|\zeta_{it})$ will have a mass point at $c = 0$ and the corresponding subpopulation of consumers will always search. All specifications I estimate below permit such a mass point. The standard full-information discrete-choice demand model is formally nested by a cost distribution $F_c(\cdot)$ with all mass at 0.

Optimal search

Under assumptions 9-11, 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.

Let preferences $v_{ijt} = -\alpha_i p_{jt} + \beta_i x_{jt} + \xi_{jt} + \epsilon_{ijt}$ be as above, and let $u_{ijt} \equiv \beta_i x_{jt} + \xi_{jt} + \epsilon_{ijt}$ denote consumer i's *direct* utility from consuming good j. Then consumer i's total utility from purchasing good j at market prices $\mathbf{p}$ is given by

$$v_{ij} = u_{ij} - \alpha_i p_j,$$

with u_{ij} being the component of total utility that is always known ex ante.

Now consider consumer i's search problem. Without loss of generality, I partition the price vector $\mathbf{p}$ into *displayed prices* $\mathbf{p}_d$ and *searched prices* $\mathbf{p}_s$, so that $\mathbf{p} \equiv (\mathbf{p_d}, \mathbf{p}_s)$. By assumption, $\mathbf{p}_d$ is observed upon entering the market, but from consumer i's perspective $\mathbf{p}_s$ is a random vector realized only following costly search. Conditional on choosing to purchase from the display set, consumer i obtains realized value

$$\bar{v}_{id}(\mathbf{p}_d) = \max\{u_{ik} - \alpha_i p_k | k \in \mathcal{D}\}, \tag{III.1}$$

and this quantity is known with certainty ex ante.[20] Meanwhile, conditional on choosing to search, consumer i observes the full price vector $\mathbf{p}$, chooses the good yielding the highest net utility, and obtains realized value

$$\bar{v}_{is}(\mathbf{p}) = \max\{u_{ik} - \alpha_i p_k | k \in \mathcal{J}\}.$$

[20] Again, bear in mind that the outside good 0 is included in the display set by definition, though to streamline presentation Equation III.1 does not explicitly indicate this.

Consumer i's expected gain from search can thus be written

$$\begin{aligned} g_{is}(\mathbf{p}_d) &= E_{p_s}[\bar{v}_{is}(\mathbf{p}_d, \mathbf{p}_s) - \bar{v}_{id}(\mathbf{p}_d)] \\ &\equiv E_{p_s}[\max\{\mathbf{u}_i - \alpha_i \cdot (\mathbf{p}_d, \mathbf{p}_s)\}] - \bar{v}_{id}(\mathbf{p}_d) \end{aligned} \tag{III.2}$$

It will be optimal for consumer i to search whenever the expected net utility gain is positive; i.e. whenever

$$g_{is}(\mathbf{p}_d) \geq c_i.$$

Under assumption 13 above, c_i is drawn independently of other preference parameters. Consequently, the probability that consumer i will choose to search facing display prices $\mathbf{p}_d$ is simply

$$\pi_{is}(\mathbf{p}_d) = F[g_{is}(\mathbf{p}_d)], \tag{III.3}$$

where $g_{is}(\mathbf{p}_d)$ has the known form defined in equation (III.2).

Predicted market shares

Now consider consumer i with preference type $\omega_i = \{\alpha_i, \beta_i, \epsilon_i\}$ facing marketing realizations $(\mathbf{p}, \mathbf{d})$, and define a collection of binary *maximand indicator functions* $\iota_j(S; \cdot)$ over sets S and products $j \in \mathcal{J}$ as follows:[21]

$$\iota_j(S; \mathbf{p}, \omega_i) \equiv \mathbf{1}[u_{ij} - \alpha_i p_j \geq u_{ik} - \alpha_i p_k \forall k \in S]. \tag{III.4}$$

The probability that consumer i purchases product j can then be described as follows.

If consumer i with preferences ω_i chooses to search, all prices are observed, the choice set is $\mathcal{J}$, and i chooses good j if

$$\iota_j(\mathcal{J}; \mathbf{p}, \omega_i) = 1.$$

[21] That is, a set of functions such that $\iota_j(S; \mathbf{p}, \omega_i) \equiv 1$ when product j is the maximand of set S given tastes ω_i at prices $\mathbf{p}$ and $\iota_j(S; \mathbf{p}, \omega_i) \equiv 0$ otherwise.

Alternatively, if consumer i chooses not to search, only displayed prices $\mathbf{p}_d$ are observed, the choice set is $\mathcal{D}$, and i chooses good j if $j \in \mathcal{D}$ and

$$\iota_j(\mathcal{D}; \mathbf{p}, \omega_i) = 1.$$

Equation (III.3) implies that consumer i will search with probability

$$\pi_s(\mathbf{p}, \mathbf{d}; \omega_i) = F_c[g_s(\mathbf{p}_d; \omega_i)].$$

Conditional on type ω_i, the probability that consumer i chooses product j is thus

$$\sigma_j(\mathbf{p}, \mathbf{d}; \omega_i) = d_j \cdot [1 - \pi_s(\mathbf{p}, \mathbf{d}; \omega_i)] \cdot \iota_j(\mathcal{D}; \mathbf{p}, \omega_i) + \pi_s(\mathbf{p}, \mathbf{d}; \omega_i) \cdot \iota_j(\mathcal{J}; \mathbf{p}, \omega_i) \tag{III.5}$$

Integrating this function across types ω_i gives the unconditional probability of choosing product j:

$$\begin{aligned} \sigma_j(\mathbf{p}, \mathbf{d}; \zeta) = \; & d_j \int_\Omega [1 - \pi_s(\mathbf{p}, \mathbf{d}; \omega_i)] \cdot \iota_j(\mathcal{D}; \mathbf{p}, \omega_i) \, dF_\omega(\omega; \zeta) \\ & + \int_\Omega \pi_s(\mathbf{p}, \mathbf{d}; \omega_i) \cdot \iota_j(\mathcal{J}; \mathbf{p}, \omega_i) \, dF_\omega(\omega; \zeta), \end{aligned} \tag{III.6}$$

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 $(\mathbf{p}, \mathbf{d})$. Display effects in this system correspond to the first line in Equation III.6; additional technical properties are described in a supplemental appendix.

III.4 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 14. *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 III.1, and is consistent with empirical surveys of retail price-setting behavior.[22] 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 on underlying preference fundamentals.[23]

Implementation

As above, consumer i's utility function takes the linear form $v_{ijt} = -\alpha_i p_{jt} + \beta_i x_{jt} + \xi_{jt} + \epsilon_{ijt}$. For current purposes, I assume product-level unobserved preferences ξ_{jt} are stable across the estimation sample: $\xi_{jt} = \xi_j$ for all t.[24] I adopt the following

[22]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.

[23]More formally, estimation under Assumption 14 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.

[24]While I hope to incorporate time-varying ξ_t in future work, 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. Third, I focus on a relatively short sample period (one year), so assuming fixed ξ seems unlikely to substantially bias structural results. Finally, Hendel and Nevo [2006] also assume fixed ξ, so this assumption helps to maintain comparability.

additional parametric structure on the distributions of taste-related parameters.[25]

Assumption 15 (Parametrics)**.** *The fundamental type distributions $F_\omega(\cdot)$ and $F_c(\cdot)$ are specialized as follows:*

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. *Idiosyncratic taste residuals are IID standard normal across consumers, products and periods: $\epsilon_{ijt} \sim N(0,1)$ for all i,j,t.*

3. *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 exponential distribution: $c_i \sim 1 - \exp(-\gamma c_i)$.*

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

$\beta_i \sim N(\bar{\beta}_x, \Sigma_x)$ follows BLP exactly, and $\epsilon_{ijt} \sim N(0,1)$ differs from BLP only in substituting a normal distribution for a logit.[26] 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

[25]An interesting question I do not address here is whether the search-plus-demand model derived above is nonparametrically identified. Berry and Haile [2010b,c] 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.

[26]Normal errors permit computation of market shares using a modified GHK algorithm, a procedure formally described in Appendix 3 below.

greatly simplifies computation.[27] For notational compactness, I index free preference parameters via the vector $\theta \equiv \{\alpha, \bar{\beta}_x, \Sigma_x, \lambda, \gamma, \xi\}$.

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

$$\ln L(\theta) = \sum_{t=1}^{T} \left[\ln \frac{N_t!}{q_{ot}! \cdots q_{Jt}!} + \sum_{j=0}^{J} q_{jt} \ln(\sigma_j(\mathcal{M}_t;\theta)) \right]. \tag{III.7}$$

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.[28]

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

[27] In my context, letting α_i to vary across consumers would require recalculating predicted market shares for every possible value of α_i. If the set of possible values is small, this should be feasible, but it would increase computation time linearly in the number of types. I am also exploring alternative approaches to simulation, which might permit more significant variation in α_i at lower computational cost.

[28] One potential problem with maximum likelihood this context is that I must ultimately simulate the likelihood 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.

for the informational reach of displays.[29] 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 III.1, 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 model:

$$trips_{st} = \beta_1 \cdot qcat_{1,st} + \beta_2 \cdot qcat_{2,st} + \cdots + \beta_{30} \cdot qcat_{30,st} + e_{st}, \tag{III.8}$$

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.[30] Figure B.5 in Appendix 1 illustrates the predicted N_t resulting from this procedure.

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

[30] 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.

Figure III.3: Predicted vs actual visits, IRI store 652159

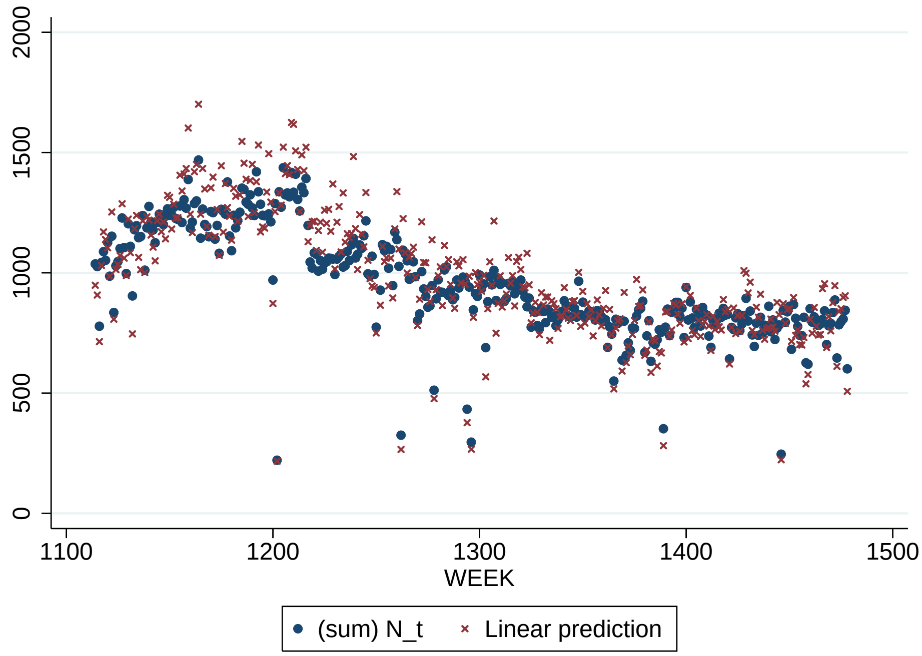

Encouragingly, the simple predictive regression (III.8) fits the BehaviorScan data very well: $R^2 > 0.98$. This fact is illustrated in Figure III.3, which plots predicted versus actual visits to IRI store 652159, the store with the most observed trips in the BehaviorScan sample. Since household purchases 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.

Product aggregation

To simplify computation of market shares, I aggregate products to the brand-size level in analysis.[31] This is both because of the very large number of UPC-level products in the market (217 in store 683960, and more than 400 in the Atlanta sample), and because many of these products are likely to be very close substitutes in practice (e.g.

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

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.[32] Figure III.4 illustrates this procedure for the Tide 100oz group (15 UPCs); Figures B.3 and B.4 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, with the sales-weighted average closely following the minimum. 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.[33]

Table B.1 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 B.3 in Appendix 1 replicates the descriptive regressions in Section III.2 on brand-size product aggregates; coefficient magnitudes change somewhat, but all qualitative patterns noted in Section III.2 extend. The brand-size aggregates defined above thus seem to preserve the essential economic features of the underlying UPC-level sample.

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

[32]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.

[33]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 III.4: Alternative price aggregates for Tide 100oz, IRI store 683960

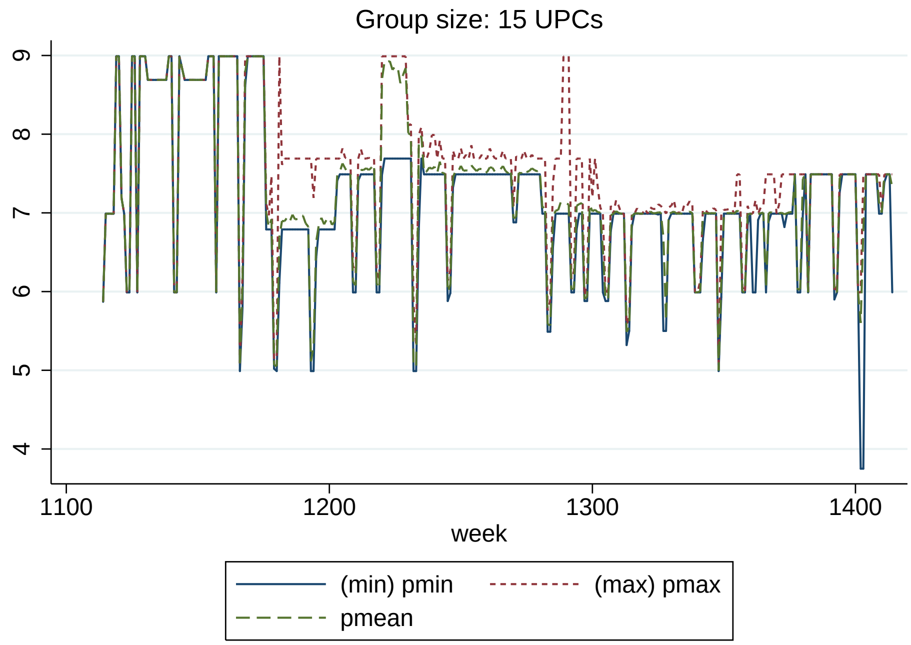

infeasible for most stores. Consequently, for each target store, I refine the estimation sample as follows. First, since many brand-size pairs are almost never purchased, I restrict attention to "important" brand-sizes, defined as those having weekly market shares (by purchase frequency among liquid detergents) of at least 1.5 percent on average in weeks sold by the target store.[34] The number of "important" brand-sizes varies between 9 and 16 for the stores considered; due to high market concentration, these "important" products typically account for more than 90 percent of store detergent purchases. Second, 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 year in the sample (2004-2005); robustness checks on other years yield similar results. For each target store, the baseline estimation sample is thus a set of "important" brand-size products for the calendar year 2004-2005.[35]

[34]Note that many products appear in only portions of the sample, so 1.5 percent market share on average in weeks sold is distinct from (and weaker than) 1.5 percent average market share over the whole sample.

[35]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-

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 III.2 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.

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. Further, robust implementation of the maximum likelihood approach above requires nonzero purchase outcomes to have nonzero purchase probabilities, a result not guaranteed under simple frequency methods. To address these concerns, I employ a modified Geweke, Hajivassiliou, and Keane (GHK) probit simulator, which reweights GHK draws by search probabilities simulated via price resampling to produce a consistent, always-positive,

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.

smooth-in-parameters simulator of $\sigma(\cdot)$.[36] Since in practice most price variation is induced by week-to-week sale discounts, I implement price resampling based on the following specialization of Assumption 9:

Assumption 16. *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 $F_{\Delta p}(\cdot)$ is stable over the sample period.*

For simplicity, the resampling algorithm I employ below further assumes that sale realizations Δp_j are independent across products. This additional restriction is not essential, however, and will be relaxed in future work.[37] Additional details on simulation are given in Appendix 3.

While this simulator for $\sigma(\cdot)$ renders maximum likelihood estimation tractable, the underlying objective function III.7 may not be well-behaved. Consequently, I implement estimation using an optimization approach pioneered by Chernozhukov and Hong [2003]: classical estimation based on Markov Chain Monte Carlo (MCMC) simulation. In my context, the main idea of this approach is to interpret the likelihood function III.7 as a Bayesian posterior distribution derived from a flat prior. I can then use the Metropolis-Hastings MCMC algorithm to generate a sample of draws from this distribution. For my purposes, this approach has three key advantages. First, by construction, MCMC sampling yields a sample representative of the full target distribution, so estimation based on this algorithm will eventually find the global maximum. Second, by definition, sample draws will be concentrated in areas with high likelihood, so the best objective value obtained in an extended sample should be close to the true maximum. Finally, if desired, one can interpret results as pure

[36] See, e.g., Geweke et al. [1994] for a description of the GHK algorithm; details of the extended simulator are described in Appendix 3.

[37] 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.

Bayesian estimates, which yields confidence bounds on parameters of interest. Thus the Chernozhukov and Hong [2003] approach is ideal for the maximum likelihood problem I consider.

III.5 Results

This section reports results from applying three structural demand specifications to data on six Atlanta stores using variants of the procedure above. My main focus in this section is of course the search-plus-demand model developed above, where displays convey information as in Section III.3 and 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 and a bulk dummy for sizes above 125oz. Consumers have random preferences $\beta_i \sim N(0, \Sigma_x)$ over these attributes, where as in BLP I assume Σ_x diagonal. As a reduced-form proxy for potential dynamic effects, all specifications also include 4 lags of units sold in the utility function.[38] GHK simulations in all cases are based on 200 preference draws and search simulations are based on 200 price draws.[39]

Structural parameters

Table III.4 reports core structural parameters resulting from application of the baseline search specification to six large stores in the Atlanta market via Chernozhukov-Hong MCMC estimation. Results in Table III.4 are based on the narrow N_t proxy (total category purchases) defined in Section III.4; corresponding estimates for my

[38]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.

[39]I experimented with other simulation sizes, but settled on (200, 200) as a reasonable balance of accuracy and computation speed.

broad N_t proxy (total store visits) are given in Table B.6 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 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 Bayesian confidence intervals based on the last 10,000 steps in each MCMC sequence are given in parentheses.[40]

Parameter values in Table III.4 suggest the presence of substantial informational effects, though estimated search patterns differ somewhat across stores. The cross-store medians of the core search-related parameters α, λ, and γ, are 0.428, 0.477, and 0.72 respectively, which taken together would imply that roughly $100 \cdot (1 - \lambda) =$ 53.3 percent of consumers have positive search costs, with a mean dollar search cost of approximately $(\gamma/\alpha) = \$1.68$ among this sub-population.[41] The first four cases are relatively heterogeneous, with estimated parameters similar to the cross-store median, but the last two cases involve estimated γ's much larger than the median. These latter two estimates would be too large to plausibly represent search costs, but can be interpreted as consistent with simple models of exogenous consideration appearing frequently; where, for instance, consumers make shopping decisions prior to entering the store, and hence do not respond actively to displayed prices.[42] In practice, however, the last two cases may be outliers driven by a short estimation sample.

[40] Burn-in for all MCMC simulations involved a combination of directed search and MCMC simulation, with length varying by algorithm used. Results reported reflect stationary series in the underlying objective.

[41] This cost estimate is consistent with although slightly lower than that reported by Seiler [2011], who reports mean costs of roughly \$3.00-\$4.50. My lower estimate is natural given that I explore the search intensity following market entry, whereas Seiler focus on the market entry decision itself.

[42] See, e.g., Pesendorfer [2002], Chevalier and Kashyap [2011], and Glandon [2011] for three representative examples of exogenous consideration models.

Table III.4: Core structural parameters, all stores (narrow N_t)

Parameter	653369	250094	263568	266100	243785	683960	Median
α	0.372	0.313	0.457	0.444	0.444	0.411	0.4275
	(0.37, 0.417)	(0.292, 0.315)	(0.458, 0.527)	(0.442, 0.485)	(0.436, 0.49)	(0.402, 0.438)	
λ	0.321	0.599	0.272	0.354	0.826	0.859	0.4765
	(0.139, 0.319)	(0.502, 0.622)	(0.112, 0.271)	(0.165, 0.352)	(0.792, 0.828)	(0.834, 0.861)	
γ	0.225	0.705	0.735	0.652	9.024	$>$30	0.72
	(0.181, 0.259)	(0.507, 0.719)	(0.554, 0.74)	(0.461, 0.641)	(5.443, 9.021)	–	
lag_1	-0.014	0.005	-0.02	-0.028	-0.015	-0.024	-0.0175
	(-0.033, -0.005)	(-0.01, 0.017)	(-0.041, -0.002)	(-0.047, -0.01)	(-0.034, 0.002)	(-0.041, -0.011)	
lag_2	-0.018	0.003	-0.028	-0.01	0.001	-0.003	-0.0065
	(-0.032, -0.006)	(-0.011, 0.014)	(-0.045, -0.012)	(-0.026, 0.009)	(-0.016, 0.018)	(-0.018, 0.012)	
lag_3	-0.008	-0.012	0.01	0.008	-0.003	-0.007	-0.005
	(-0.02, 0.009)	(-0.024, 0.001)	(-0.008, 0.031)	(-0.011, 0.028)	(-0.024, 0.013)	(-0.022, 0.008)	
lag_4	-0.016	0.005	0.001	0.003	-0.011	-0.015	-0.005
	(-0.03, -0.003)	(-0.005, 0.019)	(-0.02, 0.015)	(-0.016, 0.02)	(-0.03, 0.006)	(-0.031, -0.001)	
$ftaste$	0.283	0.421	0.312	0.226	0.281	0.341	0.2975
	(0.234, 0.304)	(0.389, 0.443)	(0.269, 0.369)	(0.198, 0.281)	(0.234, 0.305)	(0.302, 0.372)	
Objective	-1952.1	-2106.0	-2146.4	-2123.7	-1888.0	-2271.1	-2114.844763

Figure III.5: Implied structural search cost distributions, narrow N_t

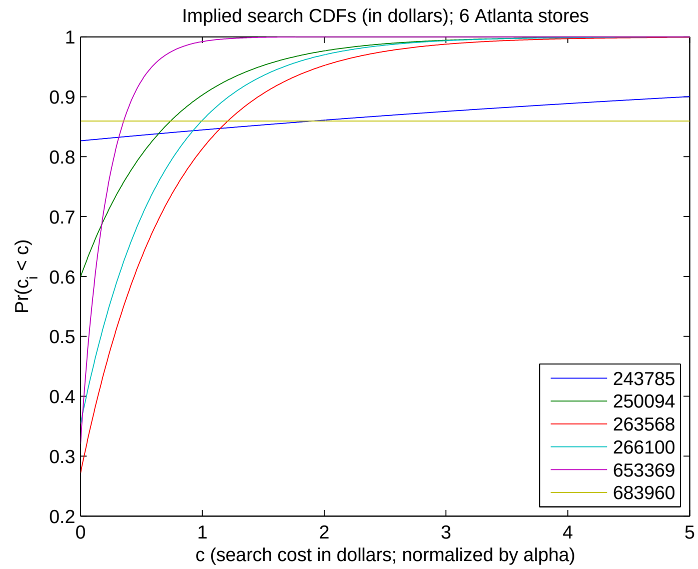

Intuitively, identification of search-related parameters turns on two key patterns in the data: level increases in quantity sold due to displays help to pin down the intercept parameter λ, and increases in price responsiveness for displayed products help to pin down the shape parameter γ. Both patterns hold robustly at the brand-size level for the entire market and for all six store-level subsamples over the entire 2002-2007 sample, but the second pattern obtains *only* for the first four store-level subsamples over the 2004-2005 estimation period (see Tables B.4 and B.5 in Appendix 1). It is therefore not particularly surprisingly that the two exceptions also produce irregular estimates of γ. Thus the results in Table III.4 both confirm the structural importance of display effects and motivate ongoing refinement of the estimation algorithm.[43]

Figure III.5 (in text) and Tables B.7 and B.8 (in Appendix 1) give some additional interpretation of the estimated structural parameters. Figure III.5 graphically represents the search cost distributions implied by the store-level parameters in Table III.4, where the vertical intercept of each distribution gives the fraction of shoppers

[43] In particular, I am currently implementing estimation based on longer sample periods and incorporating pooling across stores. This in turn should minimize the potential impact of small-sample irregularities in estimation.

λ, the shape of each distribution above zero is determined by the exponential parameter γ, and values on the horizontal axis have been normalized by α to express costs in dollar equivalents. Meanwhile, Tables B.7, and B.8 summarize elasticities derived from the structural estimates: Table B.7 presents own-price elasticities for all stores in the estimation sample, and Table B.8 gives a full set of cross-price elasticities for one representative store (266100), where elasticities in each case are calculated relative to a no-promotion baseline period.

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 III.5 and Figure III.6 summarize results of this procedure. Table III.5 presents one example of this cross-specification comparison; in this case, implied own-price elasticities for all specifications estimated for Store 266100. Consistent with the reduced-form patterns noted in Section III.2, this table suggests that promotions matter in demand analysis: the naive full-information model yields much larger elasticities than either comparison specification, with the full search model typically (though not always) producing the smallest estimates. Figure III.6 extends this comparison to the full estimation sample via a histogram plot, where observations correspond to store-product pairs and values on the horizontal axis represent the ratio of each full-information store-product elasticity to its full-search counterpart. This

Table III.5: Implied own-price elasticities, store 266100 (no-promo baseline)

Product	Naive	Full Info Promo FX	Search
ALL100	4.67	2.25	**2.15**
SURF100	5.57	**2.34**	3.34
AH125	2.33	1.51	**1.38**
CHEER80	2.44	1.66	**1.47**
ERA100	3.23	2.62	**1.40**
FAB100	3.28	**2.77**	3.01
GAIN100	4.34	2.87	**2.23**
PUREX100	3.58	2.13	**1.81**
PUREX200	5.54	**3.41**	3.42
TIDE100	2.53	2.06	**1.53**
TIDE200	7.86	**5.27**	5.82
WISK100	5.52	4.12	**2.93**

Notes: elasticities based on no-promo baseline period; row minimums in **bold**.

full-sample analysis strongly confirms the conclusions of Table III.5: incorporating promotions lead to large reductions in estimated elasticities, with naive elasticites biased upward approximately 64.5 percent relative to the full search model. In contrast to preliminary reduced-form estimates, however, there is relatively little difference between estimates based on preference and those fully incorporating search; only a 3.9 percent reduction on average. Thus while preference dummies may not elimate all sources of potential bias, Figure III.6 suggests they may be practically sufficient for many questions not directly involving search – a useful result in its own right, given the additional challenges of full search implementation.

Validation and comparisions

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 III.7 provides one example of this model validation exercise, plotting actual versus predicted market shares for the four most purchase brand-sizes in store 266100. Figure B.6 in Appendix 1 provides a corresponding plot for store 683960.

Figure III.6: Full-info demand elasticities relative to search-plus-demand equivalents

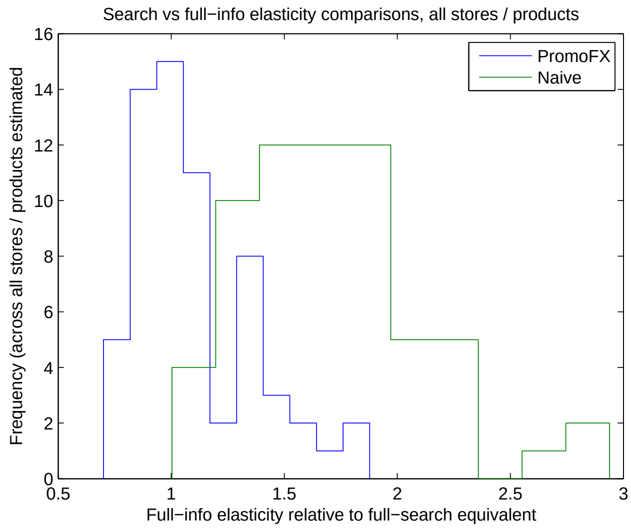

On balance, these examples suggest the search-plus-demand model fits the data well: predicted shares closely match both average shares and promotion-induced spikes, though with a slight tendency to underpredict promotional effects. Thus while the underlying model is obviously stylized, it seems to match important patterns in the data well.

Finally, I directly compare the full search-plus-demand model with my two full-information alternatives. Table B.9 in Appendix 1 presents one representative store-level parameter comparison; not surprisingly, the naive full-information model consistently yields larger price sensitivity estimates, but otherwise few clear patterns exist. Meanwhile, Table III.6 presents maximized likelihood values for each store and specification estimated. This comparison suggests that the full search-plus-demand model achieves the best in-sample fit of the three models, attaining the maximal likelihood value for all six stores considered (with the naive model typically worst by a substantial margin). This finding is both encouraging and consistent with the intuition in Section III.2: modeling displays as preference shifters will naturally generate increases in quantity sold, but not necessarily increases in price responsiveness. Such increased price responsiveness is an important empirical consequence of displays,

Figure III.7: Actual vs predicted market shares, store 266100

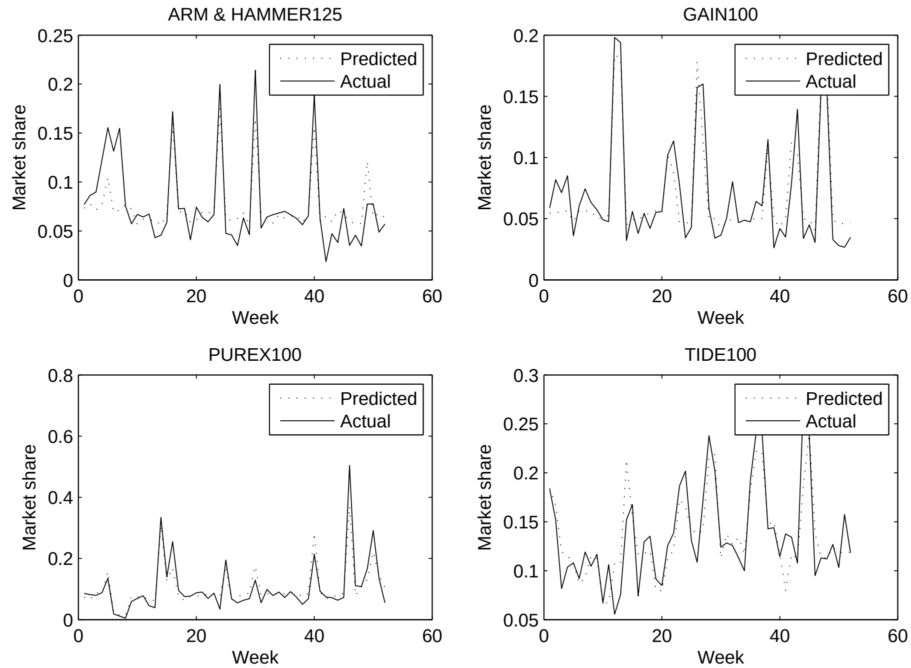

Table III.6: Model fit by store and specification (narrow N_t)

Spec / Store	250094	263568	266100	653369	243785	683960
Naive	-2468.4	-2505.9	-2310.8	-2104.9	-2030.6	-2526.3
Promo FX	-2161.7	-2188.8	-2149.2	-1968.4	-1949.2	-2303.8
Full search	-2141.3	-2132	-2113.8	-1942	-1895.4	-2270.5

Notes: Cells give best attained likelihood values; likelihood calculations involve fitting model-specific choice probabilities to store-specific estimation sample.

which in turn suggests that the full search-plus demand model may be capturing a channel of consumer response missed by standard full-information demand models.

Counterfactuals

The potential interaction between search and displays raises an interesting set of counterfactual questions: 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 explore this issue, I use the structural search-plus-demand estimates obtained above to decompose the overall effect of a displayed

Figure III.8: Counterfactual simulation: Search vs purchase effects

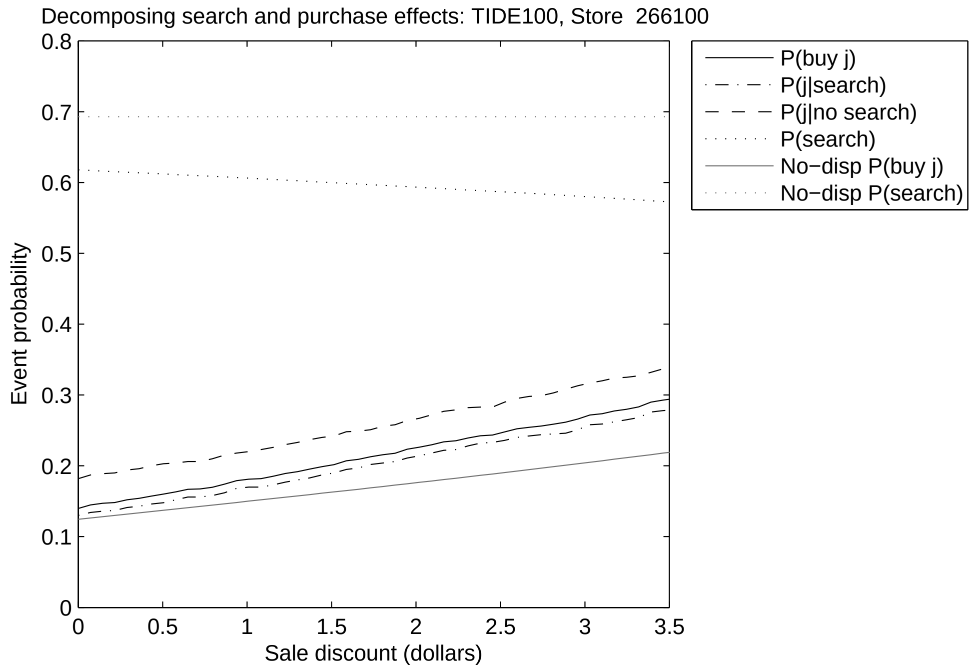

price reduction into three channels: an effect on search, an effect on purchase given search, and an effect on purchase given no search. Figure III.8 plots the results of this exercise for one representative product (Tide 100oz in Store 266100), where counterfactuals are relative to a baseline market with no other displays or sales. This figure in turn suggests two 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: 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. Both effects are consistent with the reduced-form patterns noted in Section III.3, and tend to confirm the potential structural importance of displays.

Finally, Table III.7 presents a counterfactual exercise illustrating interactions between various forms of promotional effects. Starting from a mid-sample baseline period (Week 27) in Store 27, I simulate predicted market shares under observed marketing realizations and four alternative hypotheses: no promotions, no displays, no features, and no sale discounts. This particular example involves a variety of interesting features: one displayed product (Tide 100oz), three featured products (Era 100oz, Gain 100oz, and Wisk 100oz), and three substantially discounted products (Fab 100oz, Gain 100oz, and Wisk 100oz), with Tide 100oz only displayed, Era 100oz featured with only a minimal discount (4 percent), and Fab 100oz only discounted. Simulation results illustrate both the effect of displays specifically and the presence of substantial cross-effects across promotions and products. Comparing columns (2) and (4) illustrates the potential importance of displays: the predicted purchase incidence of Tide 100oz increases 2.8 points by virtue of display alone (from 11.37 percent to 14.17 percent), of which about 1.3 points represent purchases diverted from other products. Comparisons across the remaining columns highlight potential effects and cross-effects of other forms of promotions. For instance, note that the combination of feature and sale nearly doubles the predicted share of Gain 100oz, and that shifting from "no displays" to "no promotions" increases the predicted market share of Tide 100oz substantially (from 11.37 to 12.52 percent). Since Tide 100oz is displayed but not on sale or featured, the latter change is coming entirely through cross-effects from other promotions.

III.6 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

Table III.7: Actual and counterfactual market shares, Store 266100, Week 27

			Counterfactuals				Promos		
	TRUE	Predicted	No promos	No Display	No Sale	No Feat	Disp	Feat	Disc
Outside	0.3125	0.3522	0.4231	0.3675	0.3887	0.3725	–	–	–
ALL100	0.0475	0.0476	0.0560	0.0482	0.0541	0.0499	0	0	0.000
ALL SURF100	0.0050	0.0110	0.0139	0.0111	0.0133	0.0118	0	0	0.000
ARM & HAMMER125	0.0350	0.0626	0.0732	0.0632	0.0709	0.0654	0	0	0.000
CHEER80	0.0425	0.0472	0.0546	0.0479	0.0527	0.0491	0	0	0.000
ERA100	0.0100	0.0120	0.0107	0.0121	0.0125	0.0104	0	1	0.039
FAB100	0.0225	0.0225	0.0127	0.0228	0.0122	0.0238	0	0	0.285
GAIN100	0.1600	0.1129	0.0540	0.1161	0.0684	0.0921	0	1	0.299
PUREX100	0.0550	0.0743	0.0933	0.0751	0.0894	0.0792	0	0	0.000
PUREX200	0.0050	0.0170	0.0200	0.0172	0.0194	0.0178	0	0	0.000
TIDE100	0.1725	0.1417	0.1252	0.1137	0.1478	0.1442	1	0	-0.005
TIDE200	0.0350	0.0267	0.0310	0.0295	0.0279	0.0272	0	0	0.000
WISK100	0.0975	0.0721	0.0321	0.0755	0.0426	0.0566	0	1	0.273
XTRA128	0.0017	0.0015	0.0017	0.0015	0.0017	0.0016	0	0	0

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 roughly 52 percent of consumers having positive search costs and a mean search cost of roughly \$1.68 among this sub-population. 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 substantial, there is still room to refine several dimensions of the underlying structural model. First, to improve precision of parameter estimates, I am currently implementing estimation based on longer samples and incorporating pooling across stores; possibilities in both dimensions are limited only by computing resources available. Second, to better capture potential unobserved heterogeneity, I am considering estimation permitting multiple consumer types. This extension will also involve some additional computational cost, but otherwise should be straightforward. Third, the IRI dataset also classifies displays and features by type, and in future work I hope to better employ this more nuanced data. Finally, while one objective this paper was to develop a model estimable using only scanner data, I ultimately intend to pursue 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 [2011] 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.[44] Nevertheless, a more thorough exploration 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. Consequently, I am currently investigating other possible simulation algorithms, which may permit full dynamic estimation using the IRI household panel sample in future work.[45]

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.[46] 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 consumers with different willingness to search, and displays may well serve a similar function. This is a fascinating question to which I hope to return in more detail in the future.

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

[45] Panel estimation would also permit exploration of dynamic *consideration* effects: evidence from marketing suggests that consumers are more likely to consider products they have previously purchased. This would provide additional motivation for manufacturers to encourage displays, and is an avenue I hope to explore in future work.

[46] See, e.g., Goeree [2008] for an example of a pure consideration model in the economics literature; Santos et al. [2010] directly compare several leading search models in the context of online book markets.

BIBLIOGRAPHY

R.L. Andrews and TC Srinivasan. Studying Consideration Effects in Empirical Choice Models Using Scanner Panel Data. *Journal of Marketing Research*, 32(1):30–41, 1995. ISSN 0022-2437.

S. Athey and P. Haile. Nonparametric approaches to auctions. *Handbook of Econometrics*, Vol. 6, 2005.

Susan Athey and Philip A. Haile. Identification of standard auction models. *Econometrica*, 70(6):2107–2140, 2002. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2002.00435.x. URL `http://dx.doi.org/10.1111/j.1468-0262.2002.00435.x`.

Susan Athey, Jonathan Levin, and Enrique Seira. Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics*, 126(1):207 –257, February 2011. URL `http://qje.oxfordjournals.org/content/126/1/207.abstract`.

Patrick Bajari and Ali Hortacsu. The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *The RAND Journal of Economics*, 34(2):329–355, 2003. ISSN 07416261. URL `http://www.jstor.org/stable/1593721`.

S. Berry and P. Haile. Nonparametric identification of multinomial choice demand models with heterogeneous consumers. Working Paper, Yale University, 2010a.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, July 1995. ISSN 00129682. URL `http://www.jstor.org/stable/2171802`.

Steven Berry, James Levinsohn, and Ariel Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. *The Journal of Political Economy*, 112(1):68–105, February 2004. ISSN 00223808. URL `http://www.jstor.org/stable/10.1086/379939`.

Steven T. Berry and Philip A. Haile. Identification in differentiated products markets using market level data. NBER Working Paper Series, 2010b.

Steven T. Berry and Philip A. Haile. Nonparametric identification of multinomial choice demand models with heterogeneous consumers. NBER Working Paper Series, 2010c.

B.J. Bronnenberg, M.W. Kruger, and C.F. Mela. The IRI Marketing Data Set. *Marketing Science*, 27(4):745–748, 2008. ISSN 0732-2399.

Kenneth Burdett and Kenneth L. Judd. Equilibrium price dispersion. *Econometrica*, 51(4):955–969, July 1983. ISSN 00129682. URL `http://www.jstor.org/stable/1912045`.

Gerard R. Butters. Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies*, 44(3):465–491, October 1977. ISSN 00346527. URL `http://www.jstor.org/stable/2296902`.

Victor Chernozhukov and Han Hong. An mcmc approach to classical estimation. *Journal of Econometrics*, 115:293–346, 2003.

A. Chesher. Nonparametric identification under discrete variation. *Econometrica*, 73: 1525–1550, 2005.

J.A. Chevalier and A.K. Kashyap. Best prices. NBER Working Paper Series, 2011.

J.A. Chevalier, 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*, 93:15–37, 2003.

Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828, 2009b. ISSN 1468-0262. doi: 10.3982/ECTA5368.

M. Eichenbaum, N. Jaimovich, and S. Rebelo. Reference prices and nominal rigidities. Working Paper, 2008.

Y. Fan and S. Park. Partial identification of the distribution of treatment effects and its confidence sets. *Advances in Econometrics*, 25:3–70, 2009.

M. Gentry and T. Li. A revenue equivalence result for auctions with selective entry. Working Paper, Vanderbilt University, 2012.

John Geweke, Michael Keane, and David Runkle. Alternative computational approaches to inference in the multinomial probit model. *Review of Economics and Statistics*, 76(4):609–632, November 1994.

Phllip J. Glandon. The effect of sales on average price paid. Working Paper, Vanderbilt University, 2011.

Michelle Sovinsky Goeree. Limited information and advertising in the U.S. personal computer industry. *Econometrica*, 76(5):1017–1074, 2008. ISSN 1468-0262. URL `http://dx.doi.org/10.3982/ECTA4158`.

R. Griffith, E. Leibtag, A. Leicester, and A. Nevo. Consumer shopping behavior: How much do consumers save? *The Journal of Economic Perspectives*, 23(2):99–120, 2009. URL `http://www.ingentaconnect.com/content/aea/jep/2009/00000023/00000002/art00006`.

E. Guerre, I. Perrigne, and Q. Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.

P. Haile, H. Hong, and M. Shum. Nonparametric tests for common values in first-price sealed bid auctions. Working Paper, 2003.

Philip A. Haile and Elie Tamer. Inference with an incomplete model of english auctions. *The Journal of Political Economy*, 111(1):1–51, February 2003. ISSN 00223808. URL `http://www.jstor.org/stable/10.1086/344801`.

John R. Hauser and Birger Wernerfelt. An evaluation cost model of consideration sets. *The Journal of Consumer Research*, 16(4):393–408, March 1990. ISSN 00935301. URL `http://www.jstor.org/stable/2489451`.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, November 2006. ISSN 00129682. URL `http://www.jstor.org/stable/4123086`.

K. Hendricks, J. Pinske, and RH Porter. Empirical implications of equilibrium bidding in first price, symmetric, common value auctions. *Review of Economic Studies*, 70 (1), 2003.

Han Hong and Matthew Shum. Using price distributions to estimate search costs. *The RAND Journal of Economics*, 37(2):257–275, July 2006. ISSN 07416261. URL `http://www.jstor.org/stable/25046241`.

A. Hortacsu. Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. Working Paper, University of Chicago, 2002.

Ali Hortacsu 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*, 119(2):403–456, Feb 2004. ISSN 0033-5533. doi: doi: 10.1162/0033553041382184. URL http://dx.doi.org/10.1162/0033553041382184.

D. Hosken and D. Reiffen. Pricing behavior of multiproduct retailers. *BE Journal of Theoretical Economics*, 7(1), 2007.

P. Kehoe and V. Midrigan. Sales and the real effects of monetary policy. Working Paper, 2007.

J.D. Konieczny and A. Skrzypacz. Search, costly price adjustment and the frequency of price Changes–Theory and evidence. Working Paper, 2004.

E. Krasnokutskaya. Identification and estimation in highway procurement auctions under unobserved auction heterogeneity. *Review of Economic Studies*, forthcoming, 2009.

E. Krasnokutskaya and K. Seim. Bid preference programs and participation in highway procurement auctions. *American Economic Review*, forthcoming, 2009.

V. Krishna. *Auction theory*. Academic press, 2009.

Dan Levin and James L. Smith. Equilibrium in auctions with entry. *The American Economic Review*, 84(3):585–599, June 1994. ISSN 00028282. URL http://www.jstor.org/stable/2118069.

T. Li and B. Zhang. Affiliation and entry in First-Price auctions with heterogeneous bidders. Working Paper, Vanderbilt University, 2010a.

T. Li, I. Perrigne, and Q. Vuong. Conditionally independent private information in ocs wildcat auctions. *Journal of Econometrics*, 98(1):129–162, 2000.

Tong Li. Econometrics of first-price auctions with entry and binding reservation prices. *Journal of Econometrics*, 126(1):173 – 200,

2005. ISSN 0304-4076. doi: DOI: 10.1016/j.jeconom.2004.06.002. URL http://www.sciencedirect.com/science/article/pii/S0304407604001344.

Tong Li and Bingyu Zhang. Testing for affiliation in first-price auctions using entry behavior. *International Economic Review*, 51(3):837–850, 2010b. ISSN 1468-2354. URL http://dx.doi.org/10.1111/j.1468-2354.2010.00603.x.

Tong Li and X. Zheng. Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions. *Journal of Econometrics*, 168:29–46, 2012.

Tong Li and Xiaoyong Zheng. Entry and competition effects in First-Price auctions: Theory and evidence from procurement auctions. *Review of Economic Studies*, 76(4):1397–1429, 2009. ISSN 1467-937X. URL http://dx.doi.org/10.1111/j.1467-937X.2009.00558.x.

Tong Li, Isabelle Perrigne, and Quang Vuong. Structural estimation of the affiliated private value auction model. *The RAND Journal of Economics*, 33(2):171–193, 2002. ISSN 07416261. URL http://www.jstor.org/stable/3087429.

J. Lu. Optimal entry in auctions with valuation discovery costs. *Applied Economics Research Bulletin*, 2(1):22–30, 2008.

J. Lu. Auction design with opportunity cost. *Economic Theory*, 38(1):73–103, 2009a.

J. Lu. Entry coordination and auction design with private costs of information acquisition. *Economic Inquiry*, 48(2):274–289, 2010.

Jingfeng Lu. Why a simple Second-Price auction induces efficient endogenous entry. *Theory and Decision*, 66(2):181–198, February 2009b. ISSN 0040-5833. URL http://dx.doi.org/10.1007/s11238-007-9092-9.

Thierry Magnac and Eric Maurin. Partial identification in monotone binary models: Discrete regressors and interval data. *Review of Economic Studies*, 75(3): 835–864, 2008. ISSN 1467-937X. doi: 10.1111/j.1467-937X.2008.00490.x. URL http://dx.doi.org/10.1111/j.1467-937X.2008.00490.x.

C.F. Manski. *Partial Identification of Probability Distributions*. Berlin: Springer-Verlag, 2003.

Charles F. Manski and Elie Tamer. Inference on regressions with interval data on a regressor or outcome. *Econometrica*, 70(2):519–546, 2002. ISSN 1468-0262. doi: 10.1111/1468-0262.00294. URL http://dx.doi.org/10.1111/1468-0262.00294.

V. Marmer, A. Shneyerov, and P. Xu. What model for entry in first-price auctions? a nonparametric approach. Working Paper, University of British Columbia, 2007.

R.P. McAfee and J. McMillan. Auctions with entry. *Economics Letters*, 23(4):343–347, 1987.

Nitin Mehta, Surendra Rajiv, and Kannan Srinivasan. Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science*, 22(1):58–84, January 2003. ISSN 07322399. URL http://www.jstor.org/stable/4129724.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982. ISSN 00129682. URL http://www.jstor.org/stable/1911865.

P.R. Milgrom. *Putting auction theory to work*. Cambridge Univ Pr, 2004.

Francesca Molinari. Partial identification of probability distributions with misclassified data. *Journal of Econometrics*, 144(1):81–117, May 2008.

ISSN 0304-4076. doi: doi: DOI: 10.1016/j.jeconom.2007.12.003. URL http://www.sciencedirect.com/science/article/B6VC0-4RHP9DG-1/2/de271405b5f73b9d1

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

Jose Luis Moraga-Gonzalez and Matthijs R. Wildenbeest. Maximum likelihood estimation of search costs. *European Economic Review*, 52(5):820–848, Jul 2008. ISSN 0014-2921. doi: doi: DOI: 10.1016/j.euroecorev.2007.06.025. URL http://www.sciencedirect.com/science/article/B6V64-4P59XD2-1/2/6ed7b31d2d60aa590

D. Moreno and J. Wooders. Auctions with heterogeneous entry costs. *The RAND Journal of Economics*, 42(2):313–336, 2011.

Emi Nakamura and Jon Steinsson. Five facts about prices: A reevaluation of menu cost models*. *Quarterly Journal of Economics*, 123(4):1415–1464, February 2011. ISSN 0033-5533. doi: doi: 10.1162/qjec.2008.123.4.1415. URL http://dx.doi.org/10.1162/qjec.2008.123.4.1415.

Roger B. Nelsen. *An Introduction to Copulas*. Lecture Notes in Statistics. Springer-Verlag New York, 1999.

Martin Pesendorfer. Retail sales: A study of pricing behavior in supermarkets. *The Journal of Business*, 75(1):33–66, January 2002. ISSN 00219398. URL http://www.jstor.org/stable/10.1086/323504.

J. Riley and W. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.

Rafael Rob. Equilibrium price distributions. *The Review of Economic Studies*, 52(3):487–504, July 1985. ISSN 00346527. URL http://www.jstor.org/stable/2297666.

J. Roberts and A. Sweeting. Competition versus auction design. Working Paper, Duke University, 2010a.

John H. Roberts and James M. Lattin. Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4):429–440, November 1991. ISSN 00222437. URL http://www.jstor.org/stable/3172783.

J.W. Roberts and A. Sweeting. Entry and selection in auctions. Working Paper, Duke University, 2010b.

Steven Salop and Joseph Stiglitz. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, 44(3):493–510, October 1977. ISSN 00346527. URL http://www.jstor.org/stable/2296903.

William F. Samuelson. Competitive bidding with entry costs. *Economics Letters*, 17(1-2):53–57, 1985. ISSN 0165-1765. doi: doi: DOI: 10.1016/0165-1765(85)90126-0. URL http://www.sciencedirect.com/science/article/B6V84-45BC766-3W/2/e99cd7047d041bb9

B.I. Santos, A. Hortacsu, and M.R. Wildenbeest. Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior. Working Paper, 2010.

S. Seiler. The impact of search costs on consumer behavior: a dynamic approach. Working Paper, 2011.

Alan T. Sorensen. Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy*, 108(4):833–850, 2000.

Alan T. Sorensen. An empirical model of heterogeneous consumer search for retail prescription drugs. NBER Working Paper Series(8548), October 2001. URL http://www.nber.org/papers/w8548.

George J. Stigler. The economics of information. *The Journal of Political Economy*, 69(3):213–225, June 1961. ISSN 00223808. URL http://www.jstor.org/stable/1829263.

Elie Tamer. Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies*, 70(1):147 –165, January 2003. URL http://restud.oxfordjournals.org/content/70/1/147.abstract.

X. Tang. Bounds on revenue distributions in counterfactual auctions with reserve prices. *RAND Journal of Economics*, 42(1):175–203, 2011.

E. van Nierop, B. Bronnenberg, R. Paap, M. Wedel, and P.H. Franses. Retrieving unobserved consideration sets from household panel data. *Journal of Marketing Research*, 47(1):63—74, 2010.

Hal R. Varian. A model of sales. *The American Economic Review*, 70(4):651–659, 1980. ISSN 00028282. URL http://www.jstor.org/stable/1803562.

M.R. Wildenbeest. An empirical model of search with vertically differentiated products. Working Paper, 2009.

Lixin Ye. Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior*, 58(1):181–207, jan 2007. ISSN 0899-8256. doi: doi: DOI: 10.1016/j.geb.2005.12.004. URL http://www.sciencedirect.com/science/article/B6WFW-4J5T600-2/2/24ca547972cef1692

APPENDIX A

TECHNICAL DETAILS

A.1 Proofs from Chapter 1

Proof of Proposition I.1 (following Krishna [2009]). WLOG, restrict attention to direct mechanisms with equilibria such that participants truthfully report types. In particular, let M be an arbitrary direct mechanism involving allocation rule $Q(\mathbf{v}; E)$ and payment rule $P(\mathbf{v}; E)$, where $\mathbf{v}$ is a vector of (realized) bidder values and $E \equiv (\bar{s}, N)$ is an entry structure, and let $q(v_i; E) \equiv \int_{V_{-i}} Q(v_i, \mathbf{v}_{-i}, E) f(\mathbf{v}_{-i}|E) d\mathbf{v}_{-i}$ and $p(v_i; E) \equiv \int_{V_{-i}} P(v_i, \mathbf{v}_{-i}, E) f(\mathbf{v}_{-i}|E) d\mathbf{v}_{-i}$ be the corresponding (expected) allocation and payment functions facing bidder i.

Now consider an arbitrary bidder with value $v_i \in V$, and note that by construction the mechanism $\alpha(\cdot)$ permits bidders to report any signal z in $Z = [0, \bar{v}]$. Thus for truth-telling to be an equilibrium, we must have

$$\pi(v_i; E) \equiv q(v_i; E) \cdot v_i - p(v_i; E) \geq \max_{z \in Z} \{q(z; E) \cdot v_i - p(z; E)\}.$$

It follows that $\pi(v_i; E)$ is the maximum of a family of affine functions, which in turn implies that $\pi(\cdot; E)$ is a convex function on V.

By the Integral Form Envelope Theorem (see Milgrom [2004]), this restriction in turn implies that any incentive-compatible direct mechanism must yield equilibrium bidder profit $\pi(\cdot; E)$ of the form

$$\pi(v; E) = \pi_0(E) + \int_{\underline{v}}^{v} q(y; E)\, dy,$$

where $\pi_0(E)$ is the (mechanism-determined) profit of the lowest entering bidder.

Now consider RS auctions specifically. By Definition I.1, the probability of allocation to an entering bidder with value y is

$$\begin{aligned} q(y; E) &= \alpha(y) \cdot \Pr(W_j \leq y \, \forall j) \\ &= \alpha(y) \cdot \prod_{j \neq i} \Pr(W_j \leq y) \\ &= \alpha(y) \cdot F_w^*(y; \bar{s})^{N-1}, \end{aligned}$$

where $W_j \equiv \mathbf{1}[s_j \geq \bar{s}] \cdot v_j$ is the *realized value* of bidder j. Corresponding low-type profits are

$$\pi_0(E) \quad \equiv \quad \alpha(\underline{v})\bar{s}^{N-1}\underline{v} - p_0(E),$$

where $p_0(E)$ are mechanism-determined expected payments of a low-type bidder given entry structure E.

Finally, note that under Assumption 5 we have

$$p_0(E) \equiv \bar{s}^{N-1}\alpha(\underline{v})\underline{v} - \int_0^{\bar{v}} \alpha(y)dy - E[\rho|\bar{s}, N],$$

where $E[\rho|\bar{s}, N]$ is defined as in Equation I.4. The characterization in Proposition I.1 follows immediately.

□

Proof of Lemma I.1. Let M be any RS mechanism satisfying Assumptions 1-5, and suppose that M involves award rule $\alpha(\cdot)$ and induces low-type payment $p_0(\mathbf{z}_{-i})$. Define $\hat{\rho}(\mathbf{z}_{-i})$ as follows:

$$\hat{\rho}(\mathbf{z}_{-i}) \equiv p_0(\mathbf{z}_{-i}) - \mathbf{1}[n=1]\left\{\alpha(\underline{v})\underline{v} - \int_0^{\underline{v}} \alpha(y)dy\right\}$$

Now define $\hat{p}(z_i; \mathbf{z}_{-i})$ for $z_i \leq \underline{v}$ as in Lemma I.1:

$$\hat{p}(z; \mathbf{z}_{-i}) = \mathbf{1}[n=1]\left\{\alpha(z_i)z_i - \int_0^{z_i} \alpha(y)dy\right\} + \hat{\rho}(\mathbf{z}_{-i}).$$

By construction of $\hat{\rho}(\mathbf{z}_{-i})$, $\hat{p}(\cdot)$ induces the same realized low-type payments as $p_0(\cdot)$ under truthful revelation:

$$\begin{aligned} \hat{p}(\underline{v}; \mathbf{z}_{-i}) &= \mathbf{1}[n=1]\left\{\alpha(\underline{v})\underline{v} - \int_0^{\underline{v}} \alpha(y)dy\right\} + \hat{\rho}(\mathbf{z}_{-i}) \\ &\equiv p_0(\mathbf{z}_{-i}). \end{aligned}$$

It remains only to show that $\hat{p}(\cdot)$ induces truthful revelation from an entrant with type $\underline{v}$. The initial mechanism induced an equilibrium, so $z_i > \underline{v}$ cannot be optimal. For $z_i \leq \underline{v}$, the new mechanism induces profit

$$\begin{aligned} \pi(z_i; \mathbf{z}_{-i}) &= \mathbf{1}[n=1]\alpha(z_i)\underline{v} - \hat{p}(z; \mathbf{z}_{-i}) \\ &= \mathbf{1}[n=1]\left[\alpha(z_i)[\underline{v} - z_i] + \int_0^{z_i} \alpha(y)dy\right] - \hat{\rho}(\mathbf{z}_{-i}). \end{aligned}$$

The derivative of this function with respect to the report z is

$$\pi'(z;\mathbf{z}_i) = \mathbf{1}[n=1]\alpha'(z)[\underline{v}-z] \geq 0 \forall z \leq \underline{v}.$$

Thus a low-type bidder can do no better than to report $z_i \equiv \underline{v}$. Equivalence of $\hat{p}(z;\mathbf{z}_{-i})$ with $p_0(\mathbf{z}_{-i})$ then follows immediately from above.

□

Proof. Consider first the social optimum when $r=0$. By definition, social welfare at entry threshold $\bar{s}$ is the expected value of the object to the highest bidder less total expected entry costs:

$$\begin{aligned} S(\bar{s}) &\equiv E[Y_{1:N}|\bar{s}] - N(1-\bar{s})c \\ &= \int_0^{\bar{v}} y g^*_{1:N}(y;\bar{s})dy - N(1-\bar{s})c \\ &= \bar{v} - \int_0^{\bar{v}} G^*_{1:N}(y;\bar{s})dy - N(1-\bar{s})c \\ &= \bar{v} - \int_0^{\bar{v}} [\bar{s} + \int_{\bar{s}}^1 F(y|t)dt]^N dy - N(1-\bar{s})c \qquad \text{(A.1)} \end{aligned}$$

where the third line follows from integration by parts and the last line follows by definition of $G^*_{1:N}(\cdot)$.

Differentiating this expression with respect to $\bar{s}$ yields a FOC characterizing socially optimal entry s^*:

$$S'(\hat{s}) = -N\int_0^{\bar{v}} F^*_w(y;\hat{s})^{N-1}[1-F(y|\hat{s})] + Nc \equiv 0.$$

Let $\bar{s}_e$ be the entry threshold corresponding to entry fee e. By Proposition I.2, $\bar{s}_e$ satisfies

$$\int_0^{\bar{v}} F^*_w(y;\bar{s}_e)^{N-1}[1-F(y|\bar{s}_e)] \equiv c+e.$$

When $e=0$, we thus get $\int_0^{\bar{v}} F^*_w(y;\bar{s}_0)^{N-1}[1-F(y|\bar{s}_0)] \equiv c$, which in turn implies $S'(\bar{s}_0)=0$. Further, by the proof of Proposition I.2, we know $\frac{\partial}{\partial s}\{F^*_w(y;s)^{N-1}[1-F(y|s)]\} \geq 0$ for any y, so $S''(s) \leq 0$ and $\hat{s}=\bar{s}_0$ is the unique global optimum. Hence when $r=0$, social welfare is maximized by setting $e=0$.

Finally, note that $r \neq 0$ can never improve social welfare. Reserve prices affect surplus in two ways: by shifting the equilibrium threshold $\bar{s}$, and by affecting the allocation of the object being sold. The first effect can be offset via an appropriate entry subsidy (a transfer, hence welfare-neutral). The second effect must decrease allocative efficiency (and hence social welfare). For any $(\bar{s}, r)$, we thus conclude:

$$S(\bar{s}; r) \leq S(\bar{s}; 0) \leq S(\hat{s}; 0).$$

Hence the social optimum obtains when the seller sets no reservation price or entry fee: $\hat{e} = 0$ and $\hat{r} = 0$.

□

Proof of Proposition I.5. Assume $r = 0$, and focus on second-price auctions WLOG. Since there is a one-to-one correspondence between the entry fee e and the entry threshold s, we can derive optimal policy in terms of either. When $r = 0$, expected revenue under equilibrium threshold s is

$$R^*(s) = \int_0^{\bar{v}} y \, dG^*_{2:N}(y; s) + N(1-s) \cdot e(s),$$

where $e(s) \equiv \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy - c$ is the entry fee required to produce equilibrium threshold s. Integrating by parts and substituting for $e(s)$ then gives

$$\begin{aligned} R^*(s) &= \bar{v} - \int_0^{\bar{v}} G^*_{2:N}(y; s) dy + N(1-s) \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy - N(1-s)c \\ &= \bar{v} - N \int_0^{\bar{v}} [1 - F^*_w(y; s)] F^*_w(y; s)^{N-1} - \int_0^{\bar{v}} F^*_w(y; s)^N dy \\ &\quad + N(1-s) \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy - N(1-s)c \\ &= \bar{v} - N(1-s) \int_0^{\bar{v}} [1 - F^*(y; s)] F^*_w(y; s)^{N-1} - \int_0^{\bar{v}} F^*_w(y; s)^N dy \\ &\quad + N(1-s) \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy - N(1-s)c \\ &= \left\{ \bar{v} - \int_0^{\bar{v}} F^*_w(y; s)^N dy - N(1-s)c \right\} - \left\{ N(1-s) \int_0^{\bar{v}} [F(y|s) - F^*(y; s)] F^*_w(y; s)^{N-1} dy \right\} \end{aligned}$$

where the first term represents expected total surplus (see Equation (A.1)) and the second represents net profit accruing to inframarginal bidders.

Differentiating this expression yields the necessary FOC for the seller's optimal threshold s^*:

$$\begin{aligned} R^*_s(s) &= \left\{ -N \int_0^{\bar{v}} [1 - F(y|s)] F^*_w(y; s)^{N-1} dy + Nc \right\} \\ &\quad - \frac{\partial}{\partial s} \left\{ N(1-s) \int_0^{\bar{v}} [F(y|s) - F^*(y; s)] F^*_w(y; s)^{N-1} dy \right\}. \end{aligned} \tag{A.2}$$

In the "regular" case, net bidder profit should be decreasing in total entry costs. In this case, total bidder profits will also be decreasing in the entry threshold s (since e and s are monotonically related), so the derivative of the second term will be negative. We know (from I.4) that the first term vanishes at the social optimum $\hat{s}$. Hence in the regular case we expect

$$R_s^*(\hat{s}) = -\frac{\partial}{\partial s}\left\{N(1-\hat{s})\int_0^{\bar{v}}[F(y|\hat{s}) - F^*(y;\hat{s})]F_w^*(y;\hat{s})^{N-1}dy\right\} > 0,$$

and a revenue-maximizing seller will set $s^* > \hat{s}$ (or equivalently $e^* > 0$).

Analytically, however, it is not clear that net bidder profit must be decreasing in total entry cost for all possible fundamentals. Consequently, we derive an explicit form for $R_s^*(s)$ as well:

$$\begin{aligned}
R_s^*(s) &= \left\{-N\int_0^{\bar{v}}[1-F(y|s)]F_w^*(y;s)^{N-1}dy + Nc\right\} \\
&\quad +N\int_0^{\bar{v}}[F(y|s) - F^*(y;s)]F_w^*(y;s)^{N-1}dy \\
&\quad -N(1-s)\int_0^{\bar{v}}F_s(y|s)F_w^*(y;s)^{N-1}dy \\
&\quad +N(1-s)\int_0^{\bar{v}}[\frac{\partial}{\partial s}F^*(y;s)]F_w^*(y;s)^{N-1}dy \\
&\quad -N(1-s)\int_0^{\bar{v}}[F(y|s) - F^*(y;s)]\cdot(N-1)F_w^*(y;s)^{N-2}[1-F(y|s)]dy \\
&= \left\{-N\int_0^{\bar{v}}[1-F(y|s)]F_w^*(y;s)^{N-1}dy + Nc\right\} \\
&\quad - \{\,N(1-s)\int_0^{\bar{v}}F_s(y|s)F_w^*(y;s)^{N-1}dy \\
&\quad +N(1-s)\int_0^{\bar{v}}[F(y|s) - F^*(y;s)]\cdot(N-1)F_w^*(y;s)^{N-2}[1-F(y|s)]dy\,\}
\end{aligned}$$

where the second and fourth lines of the first equality cancel since

$$\begin{aligned}
\frac{\partial}{\partial s}F^*(y;s) &\equiv \frac{\partial}{\partial s}\left\{\frac{1}{1-s}\int_s^1 F(y|t)dt\right\} \\
&= \left\{\frac{1}{1-s}\right\}^2\int_s^1 F(y|t)dt - \frac{1}{1-s}F(y|s) \\
&= -\frac{1}{1-s}[F(y|s) - F^*(y;s)].
\end{aligned}$$

We thus obtain $R_s^*(s)$ given in Lemma I.5. By construction, the seller's optimal threshold will satisfy $R_s^*(s^*) \equiv 0$, but without being able to sign the derivative of bidder profit the relationship of s^* to $\hat{s}$ may be ambiguous in general.

□

Proof of Corollary I.2. Though Corollary I.2 is a special case of Lemma I.5, the special features of the S model make it easier to establish directly. Entry in the S model will involve a value threshold $\bar{y}$ such that bidder i will enter if and only if $v_i \geq \bar{y}$.[1] A bidder with value $v_i = \bar{y}$ wins only when no other bidder enters, so given entry fee e this threshold satisfies the following breakeven condition:

$$\bar{y}F(\bar{y})^{N-1} \equiv c + e.$$

As above, seller revenue corresponding to threshold $\bar{y}$ is the expected value of the second-highest entrant plus total expected entry fees:

$$\begin{aligned}
R^*(\bar{y}) &= \bar{v} - \int_0^{\bar{v}} G^*_{2:N}(y;\bar{y})dy + N[1-F(\bar{y})][\bar{y}F(\bar{y})^{N-1} - c] \\
&= \bar{v} - \int_{\bar{y}}^{\bar{v}} \left\{N[1-F(y)]F(y)^{N-1} + F(y)^N\right\} dy \\
&\quad - \int_0^{\bar{y}} \left\{N[1-F(\bar{y})]F(\bar{y})^{N-1} + F(\bar{y})^N\right\} dy \\
&\quad + N[1-F(\bar{y})][\bar{y}F(\bar{y})^{N-1} - c] \\
&= \bar{v} - \int_{\bar{y}}^{\bar{v}} \left\{N[1-F(y)]F(y)^{N-1} + F(y)^N\right\} dy \\
&\quad - \bar{y}F(\bar{y})^N - N[1-F(\bar{y})]c,
\end{aligned}$$

where the second equality follows since in the S case

$$\begin{aligned}
F^*_w(y;\bar{y}) &\equiv F(\bar{y}) + [1-F(\bar{y})]\mathbf{1}[y \geq \bar{y}] \cdot \frac{F(y) - F(\bar{y})}{1 - F(\bar{y})} \\
&= F(\bar{y}) + \mathbf{1}[y \geq \bar{y}] \cdot [F(y) - F(\bar{y})].
\end{aligned}$$

Differentiating $R^*(\bar{y})$ and simplifying yields:

$$R^*_y(\bar{y}) = N[1-F(\bar{y})]F(\bar{y})^{N-1} - N[\bar{y}F(\bar{y})^{N-1} - c]f(\bar{y}).$$

The seller's optimum y^* will satisfy $R^*_y(y^*) \equiv 0$. Meanwhile, from Proposition I.4 we know the social optimum $\hat{y}$ satisfies $\hat{y}F(\hat{y})^{N-1} \equiv c$, so we must have $R^*_y(\hat{y}) \geq 0$. We thus conclude $y^* \neq \hat{y}$ in general, and under standard regularity conditions we will have $y^* > \hat{y}$.

□

[1] In particular, under the normalization assumption above, we would have $\bar{y} \equiv F_y^{-1}(\bar{s})$.

Proof of Proposition I.6. Seller revenue corresponding to reserve price $r \geq 0$ at entry threshold s is given by

$$\begin{aligned} R(r;s) &= E^*[Y_{2:N}|Y_{2:N} \geq r; s] \Pr(Y_{2:N} \geq r|s) + r \cdot \Pr(Y_{1:N} \geq r \cap Y_{2:N} \leq r|s) \\ &= yG^*_{2:N}(y;s)|_r^{\bar{v}} - \int_r^{\bar{v}} G^*_{2:N}(y;s)dy + r[1 - F^*_w(r;s)]F^*_w(r;s)^{N-1} \\ &= \bar{v} - r\left\{[1 - F^*_w(r;s)]F^*_w(r;s)^{N-1} + F^*_w(r;s)^N\right\} \\ &\quad - \int_r^{\bar{v}} G^*_{2:N}(y;s)dy + r[1 - F^*_w(r;s)]F^*_w(r;s)^{N-1} \\ &= \bar{v} - \int_r^{\bar{v}} G^*_{2:N}(y;s)dy - rF^*_w(r;s)^{N-1}. \end{aligned}$$

As in the proof of Proposition I.5, we frame the problem as a choice of the optimal entry threshold: $R^*(s) \equiv R(r(s), s)$, where $r(s)$ is the reserve price inducing equilibrium entry s. Differentiating with respect to s then yields the seller's necessary FOC:

$$R^*_s(s^*) \equiv R_r[(r(s^*), s^*]r'(s^*) + R_s[r(s^*), s^*] = 0.$$

We obtain $R_r(r, s)$ and $R_s(r, s)$ directly from above:

$$\begin{aligned} R_r(r,s) &= N(1-s)F^*_w(r,s)^{N-1}[1 - F^*(r;s)] - r \cdot N(1-s)F^*_w(r;s)^{N-1}f^*(r;s), \\ R_s(r,s) &= -N(N-1)(1-s)\int_r^{\bar{v}} [1 - F^*(y;s)]F^*_w(y;s)^{N-2}[1 - F(y|s)]dy. \end{aligned}$$

We obtain $r'(s)$ by differentiating the Stage 1 entry equilibrium condition (I.7) completely with respect to s:

$$\int_{r(s)}^{\bar{v}} [1 - F(y|s)]F^*_w(y;s)dy \equiv c,$$

which in turn implies

$$\begin{aligned} r'(s) \cdot [1 - F(r|s)]F^*_w(r;s)^{N-1} &= (N-1)\int_r^{\bar{v}} [1 - F(y|s)]F^*_w(y;s)^{N-1}[1 - F(y|s)]dy \\ &\quad - \int_r^{\bar{v}} F_s(y|s) \cdot F^*_w(y;s)^{N-1}dy. \end{aligned}$$

We could combine all these elements to obtain an explicit form for the seller's FOC, but in general this condition is likely to be unwieldy and unintuitive. We therefore focus instead on the derivative at $\hat{s}$. By Proposition I.4, we know $r(\hat{s}) \equiv 0$, so the relevant terms are

$$\begin{aligned} R_r(0,\hat{s}) &= N(1-\hat{s})\hat{s}^{N-1} \\ R_s(0,\hat{s}) &= -N(N-1)(1-\hat{s})\int_0^{\bar{v}}[1-F^*(y;\hat{s})]F_w^*(y;\hat{s})^{N-2}[1-F(y|\hat{s})]dy \\ r'(\hat{s})\hat{s}^{N-1} &= (N-1)\int_0^{\bar{v}}[1-F(y|\hat{s})]F_w^*(y;\hat{s})^{N-1}[1-F(y|\hat{s})]dy \\ &\quad -\int_0^{\bar{v}} F_s(y|\hat{s})\cdot F_w^*(y;\hat{s})^{N-1}dy. \end{aligned}$$

Finally, combining these terms and simplifying gives

$$\begin{aligned} R_s^*(\hat{s}) \;\propto\; & \int_0^{\bar{v}}(-1)F_s(y|\hat{s})\cdot F_w^*(y;\hat{s})^{N-1}dy \\ & -(N-1)\int_0^{\bar{v}}[F(y|\hat{s})-F^*(y;\hat{s})]F_w^*(y;\hat{s})^{N-2}[1-F(y|\hat{s})]dy. \end{aligned}$$

Inspection reveals that this expression is identical to the last two terms of Equation (I.12) of Proposition I.5. Hence the seller can gain by setting a positive reserve price under the same conditions that the seller can gain by setting a positive entry fee. By the argument in the proof of Proposition I.5, a sufficient condition for such gain to be feasible is for potential bidders to strictly prefer lower total entry costs.

□

Proof of Corollary I.4. By arguments similar to those in the proofs of Proposition I.6 and Corollary I.2, we eventually obtain:

$$\begin{aligned} R(r,\bar{y}) &= \bar{v}-\bar{y}\left\{N[1-F(\bar{y})]F(\bar{y})^{N-1}+F(\bar{y})^N\right\}+rN[1-F(\bar{y})]F(\bar{y})^{N-1} \\ &\quad -\int_{\bar{y}}^{\bar{v}}\left\{N[1-F(y)]F(y)^{N-1}+F(y)^N\right\}dy, \end{aligned}$$

again (repeatedly) taking advantage of the fact that

$$F_w^*(y;\bar{y}) = \begin{cases} F(\bar{y}) & \text{if } y \le \bar{y} \\ F(y) & \text{if } y \ge \bar{y}. \end{cases}$$

Again proceeding as in Corollary I.2, we find that $\bar{y}$ is related to r by the breakeven condition:

$$(\bar{y}-r)F(\bar{y})^{N-1} \equiv c.$$

Taking appropriate partial and total derivatives, we thus conclude

$$\begin{aligned} R_r(r,\bar{y}) &= N[1-F(\bar{y})]F(\bar{y})^{N-1} \\ R_y(r,\bar{y}) &= -(\bar{y}-r)\cdot N(N-1)[1-F(\bar{y})]F(\bar{y})^{N-2}f(\bar{y}) - r\cdot NF(\bar{y})^{N-1}f(\bar{y}) \\ r'(\bar{y}) &= 1+(N-1)(\bar{y}-r)F(\bar{y})^{-1}f(\bar{y}). \end{aligned}$$

Finally, putting everything together yields the FOC given:

$$R_y^*(\bar{y}) = N[1-F(\bar{y})]F(\bar{y})^{N-1} - rNF(\bar{y})^{N-1}f(\bar{y}).$$

□

A.2 Proofs from Chapter 2

Proof of Lemma II.1. We establish claims for $\check{F}^+(v|\bar{s})$; the argument for $\check{F}^-(v|\bar{s})$ is analogous.

By construction, if $s^-(\bar{s}) \notin \mathcal{S}$ then $s^-(\bar{s}) \equiv 0$, and if $s^-(\bar{s}) = 0$ and $0 \notin \mathcal{S}$ then $\check{F}^+(v|\bar{s}) \equiv 1 \geq F(v|\bar{s})$. Hence we focus on the case $s^-(\bar{s}) \in \mathcal{S}$.

When $s^-(\bar{s}) \in \mathcal{S}$, there are two possible subcases:

- $\bar{s} = s^-(\bar{s})$: By construction of $s^-(\bar{s})$, this occurs when $\bar{s} \in \text{int}(\mathcal{S})$, which implies that there exists an open neighborhood of identified thresholds $t \in \mathcal{S}$ around $\bar{s}$. Consequently, we can identify the function $(1-t)F^*(v;t)$ at points arbitrarily close to $\bar{s}$, and the limit defining $\check{F}^+(v|\bar{s})$ converges to the corresponding derivative:

$$\lim_{t \uparrow s^-(\bar{s})} \left\{ \frac{(1-t)F^*(v;t) - (1-\bar{s})F^*(v;\bar{s})}{\bar{s}-t} \right\} = -\frac{\partial}{\partial \bar{s}}(1-\bar{s})F^*(v;\bar{s}) \equiv F(v|s).$$

  Hence $\check{F}^+(v|\bar{s}) = F(v|s)$, so $\check{F}^+(v|\bar{s})$ is a distribution and $F(v|s)$ is exactly identified.

- $\bar{s} > s^-(\bar{s})$: By construction, $s^-(\bar{s})$ is then the nearest lower neighbor of $\bar{s}$ in $\mathcal{S}$

(but separated by an open interval). In this case,

$$\begin{aligned}
\lim_{t\uparrow s^-(\bar{s})}\left\{\frac{(1-t)F^*(v;t)-(1-\bar{s})F^*(v;\bar{s})}{\bar{s}-t}\right\} &= \frac{[1-s^-(\bar{s})]F^*(v;s^-(\bar{s}))-(1-\bar{s})F^*(v;\bar{s})}{\bar{s}-s^-(\bar{s})} \\
&= \frac{1}{\bar{s}-s^-(\bar{s})}\left\{F^*(v;s^-(\bar{s}))-F^*(v;\bar{s})\right\} \\
&= \frac{1}{\bar{s}-s^-(\bar{s})}\left\{\int_{s^-(\bar{s})}^1 F(v|t)dt-\int_{\bar{s}}^1 F(v|t)dt\right\} \\
&= \frac{1}{\bar{s}-s^-(\bar{s})}\int_{s^-(\bar{s})}^{\bar{s}} F(v|t)dt \\
&= F(v|S_i\in[s^-(\bar{s}),\bar{s}]).
\end{aligned}$$

Line 1 implies $\check{F}^+(v|\bar{s})$ is identified (since it depends only on identified components), Line 5 implies that $\check{F}^+(v|\bar{s})$ is a distribution, and Line 4 implies that $\check{F}^+(v|\bar{s})$ bounds $F(v|s)$:

$$\begin{aligned}
\frac{1}{\bar{s}-s^-(\bar{s})}\int_{s^-(\bar{s})}^{\bar{s}} F(v|t)dt &\geq \frac{1}{\bar{s}-s^-(\bar{s})}\int_{s^-(\bar{s})}^{\bar{s}} F(v|\bar{s})dt \\
&= \frac{\bar{s}-s^-(\bar{s})}{\bar{s}-s^-(\bar{s})}F(v|\bar{s})=F(v|\bar{s}),
\end{aligned}$$

where the first inequality follows since affiliation implies $F(v|s')\leq F(v|s)$ for $s'\geq s$ (see Milgrom and Weber [1982]).

Taken together, the cases above establish all claims in Lemma II.1.

□

Proof of Proposition II.1. We establish claims for $F^+(v|s)$; the argument for $F^-(v|s)$ is analogous.

Identification of $F^+(v|t)$ follows from (i) $s^-(t)\in\{\mathcal{S},1\}$ by construction, (ii) identification of $\check{F}^+(v|s)$ for $s\in\mathcal{S}$, and (iii) $\check{F}^+(v|s)\equiv 1$ for $s=0$ if $s\neq\mathcal{S}$. Hence $F^+(v|t)$ depends only on objects recoverable from process $\mathcal{L}$.

The distribution and exact identification properties of $F^+(v|t)$ are inherited directly from the corresponding properties of $\check{F}^+(v|t)$.

Finally, to establish bounds, we consider cases:

- If $t\in\mathcal{S}$, then $F^+(v|t)\equiv\check{F}^+(v|t)\geq F(v|t)$.
- Otherwise, $F^+(v|t)\equiv\check{F}^+(v|s^-(t))\geq F(v|s^-(t))\geq F(v|t)$, where the last inequality follows by the stochastic-dominance property of affiliation.

Taken together, the cases above establish all claims in Proposition II.1.

□

Proof of II.2. Identification of $c^+(z)$ and $c^-(z)$ and the inequalities $c^+(z) \geq c(z) \geq c^-(z)$ follow immediately from identification of $F^-(y|\bar{s}_N(z))$ and $F^+(y|\bar{s}_N(z))$ and $F^+(y|\bar{s}_N(z)) \geq F(y|\bar{s}_N(z) \geq F^-(y|\bar{s}_N(z))$, with exact equality obtaining when $F^{\pm}(y|\bar{s}_N(z)) = F(y|\bar{s}_N(z))$.

□

Proof of Proposition II.3. To establish Statement 1, suppose $z \in \text{int}(Z)$ and $\bar{s} \in (0, 1)$. Then there exists an open ϵ-ball $B_\epsilon(z) \subset Z$ around z. By Assumption 7, $c(\cdot)$ is continuous and monotonic in continuous components of z, and hence maps open sets to open sets. Thus $\mathcal{L}$ involves an open ϵ-ball $B_\epsilon(c(z)) \subset R^+$ of costs around $c(z)$. Finally, by Proposition I.2, the equilibrium threshold $\bar{s}(\cdot, N)$ is continuous and monotonic in $c(\cdot)$ for $\bar{s} \in (0, 1)$. Hence $\bar{s}(\cdot, N)$ also maps open sets to open sets, so $\bar{s}(c(z), N) \in \text{int}(\mathcal{S}(\mathcal{L}))$. Exact local identification then follows from Propositions II.1 and II.2.

To establish Statement 2, suppose the range of $c(\cdot)$ is as given. Then for an appropriate choice of z we can produce any $c(z) \in [0, \bar{v}]$. The former will ensure universal entry, and the latter will exclude all potential entry. The rest of the interval will produce every intermediate case. Consequently $\mathcal{S}(\mathcal{L}) = [0, 1]$ and full identification follows.

□

Proof of revenue characterization in Lemma II.3. For any $(s; N)$ pair, expected seller revenue at allocation rule α is given by

$$R_\alpha(\bar{s}; N) = AV_\alpha(s; N) - N\Pi^*_\alpha(s; N),$$

where $AV_\alpha(\cdot)$ is *ex ante* expected allocation value of the object being auctioned and $\Pi^*(s; N)$ is expected *ex ante* equilibrium profit for an arbitrary bidder.

To obtain $AV_\alpha(\cdot)$, let $Y_{1:N}$ be max realized value among N potential bidders. Then net value created is$Y_{1:N}$ if sale, v_0 if no sale. Conditional on $Y_{1:N}$, expected allocation value is thus

$$\alpha(Y_{1:N})Y_{1:N} + [1 - \alpha(Y_{1:N})]v_0.$$

Integrating with respect to $Y_{1:N}$, we obtain *ex ante* expected allocation value:

$$\begin{aligned} AV_\alpha(s;N) &= s^N v_0 + \int_{\underline{v}}^{\bar{v}} \{\alpha(y)y + [1-\alpha(y)]v_0\}\, g^*_{1:N}(y;s)dy \\ &= \int_{v_0}^{\bar{v}} \{\alpha(y)y + [1-\alpha(y)]v_0\}\, dG^*_{1:N}(y;s), \end{aligned}$$

where $g^*_{1:N}(y;s) \equiv NF^*_w(y;s)^{N-1} f^*_w(y;s)$ is the density of $Y_{1:N}$ on $[\underline{v}, \bar{v}]$ given entry threshold s, $G^*_{1:N}(y;s) = F^*_w(y;s)^N$ is the corresponding distribution on $[v_0, \bar{v}]$, and $\alpha(v_0) \equiv 0$ by Assumption I.1.

To obtain $\Pi^*(s;N)$, we start from the result in Proposition I.1:

$$\begin{aligned} \pi_\alpha(v;s,N) &= \pi_0^\alpha(s,N) + \int_{\underline{v}}^{v} \alpha(t)\cdot F^*_w(t;s)^{N-1}dt \\ &= \int_{v_0}^{\underline{v}} \alpha(t)s^{N-1}dt + \int_{\underline{v}}^{v} \alpha(t)\cdot F^*_w(t;s)^{N-1}dt - \rho \\ &= \int_{v_0}^{\underline{v}} \alpha(t)\cdot F^*_w(t;s)^{N-1}dt - \rho \\ &= \lambda_\alpha(v;s,N)\cdot \alpha(v)F_w(v;s)^{N-1} - \rho, \end{aligned}$$

where the second equation follows from Assumption I.1 and

$$\lambda_\alpha(v;s,N) \equiv \begin{cases} 0 & \text{if } \alpha(v)=0; \\ \int_{v_0}^{v} \frac{\alpha(t)}{\alpha(v)} \cdot \frac{F^*_w(t;s)^{N-1}}{F^*_w(v;s)^{N-1}}dt & \text{otherwise.} \end{cases}$$

gives the average incremental profit (above $-\rho$) a bidder of type v receives per win.

Integrating over the distribution $F^*_w(y;s)$ then gives $\Pi^*(s;N)$

$$\Pi^*(s;N) = \int_{\underline{v}}^{\bar{v}} \lambda_\alpha(y;s,N)\cdot\alpha(y)F^*_w(y;s)^{N-1}f^*_w(y;s)dy - (1-s)\rho$$

and multiplying by N yields

$$\begin{aligned} N\Pi^*(s;N) &= \int_{\underline{v}}^{\bar{v}} \lambda_\alpha(y;s,N)\alpha(y)\cdot NF^*_w(y;s)^{N-1}f^*_w(y;s)dy - N(1-s)\rho \\ &= \int_{\underline{v}}^{\bar{v}} \lambda_\alpha(y;s,N)\alpha(y)\cdot g^*_{1:N}(y;s)dy - N(1-s)\rho \\ &= \int_{\underline{v}}^{\bar{v}} \lambda_\alpha(y;s,N)\alpha(y)\, dG^*_{1:N}(y;s)dy - N(1-s)\rho. \end{aligned}$$

Combining the results above gives a final expression for seller revenue:

$$\begin{aligned} R_\alpha(s;N) &= \int_{v_0}^{\bar{v}} \{\alpha(y)y + [1-\alpha(y)]v_0\}\, g^*_{1:N}(y;s)dy \\ &\quad - \int_{\underline{v}}^{\bar{v}} \lambda_\alpha(y;s,N)\alpha(y)\, dG^*_{1:N}(y;s) + N(1-s)\rho \\ &= \int_{v_0}^{\bar{v}} \{\alpha(y)[y-\lambda_\alpha(y;s,N)] + [1-\alpha(y)]v_0\}\, dG^*_{1:N}(y;s)dy + N(1-s)\rho. \end{aligned}$$

where the second equality follows because $\int_{v_0}^{\underline{v}} \lambda_\alpha(y;s,N)\, dG^*_{1:N}(y;s) = 0$: $\lambda_\alpha(v_0;s,N) \equiv 0$ and $g^*_{1:N}(y;s) \equiv 0$ for $y \in (v_0, \underline{v})$.

□

Proof of Lemma II.3. Identification of $R_\alpha(s;N)$ for $s \in \mathcal{S}$ follows directly from Equation II.4: $R_\alpha(\cdot)$ depends only on mechanism components (α, ρ, v_0) (known by hypothesis) and distributions $F^*_w(\cdot;s)$ and $G^*_{1:N}(\cdot;s)$ (identified for $s \in \mathcal{S}$). Thus it only remains to show $R_\alpha(s;N)$ is decreasing in s. Equation (II.4) implies that s affects seller revenue through (at most) three channels: the per-win profit function $\lambda_\alpha(v;s,N)$, the distribution $G^*_{1:N}(\cdot;s)$, and the residual term $N(1-s)\rho$. We show that each of these partial effects is negative.

First, consider effects through the per-win profit function $\lambda_\alpha(v;s,N)$. Note that

$$\frac{\partial}{\partial s}\lambda_\alpha(v;s,N) = \int_{v_0}^{v} \frac{\alpha(t)}{\alpha(v)} \cdot \frac{\partial}{\partial s}\left\{\frac{F^*_w(t;s)^{N-1}}{F^*_w(v;s)^{N-1}}\right\} dt.$$

By algebra,

$$\begin{aligned} \frac{\partial}{\partial s}\left\{\frac{F^*_w(t;s)^{N-1}}{F^*_w(v;s)^{N-1}}\right\} &= \frac{(N-1)F^*_w(t;s)^{N-2}\frac{\partial}{\partial s}F^*_w(t;s)}{F^*_w(v;s)^{N-1}} - \frac{(N-1)F^*_w(t;s)^{N-1}\frac{\partial}{\partial s}F^*_w(v;s)}{F^*_w(v;s)^{N}} \\ &= (N-1)\frac{F^*_w(t;s)^{N-2}}{F^*_w(v;s)^{N-1}}\left\{[1-F(t|s)] - \frac{F^*_w(t;s)}{F^*_w(v;s)}[1-F(v|s)]\right\} \\ &\geq 0\, \forall t \leq v, \end{aligned}$$

since $t \leq v$ means $F^*_w(t;s) \leq F^*_w(v;s)$ and $F(t|s) \leq F(v|s)\ \forall s$. Thus $\lambda_\alpha(v;s,N)$ is increasing in s for all v, so the effect of s on R through $\lambda_\alpha(v;s,N)$ is negative.

Next, consider effects through the distribution $G^*_{1:N}(\cdot;s)$. It is easy to show that $G^*_{1:N}(v;s)$ is increasing in s for any v, hence $s' \geq s$ means $G^*_{1:N}(\cdot;s)$ first-order stochastically dominates $G^*_{1:N}(\cdot;s')$. Thus if the integrand

$$\{\alpha(y)[y-\lambda_\alpha(y;s,N)] + [1-\alpha(y)]v_0\} \tag{A.3}$$

is increasing in y, an increase in s will involve taking the expectation of an increasing function with respect to a stochastically dominated distribution, which must imply a decrease in revenue. It is therefore is sufficient to show that the integrand (A.3) is increasing in y.

- First, note that $[y - \lambda_\alpha(y; s, N)]$ is increasing in y:

$$
\begin{aligned}
\frac{\partial}{\partial y}[y - \lambda_\alpha(y; s, N)] &\equiv 1 - \frac{\partial}{\partial y}\int_{v_0}^{y} \frac{\alpha(t)}{\alpha(y)} \cdot \frac{F_w^*(t; s)^{N-1}}{F_w^*(y; s)^{N-1}} dt \\
&= 1 - \frac{\partial}{\partial y}\frac{1}{\alpha(y) F_w^*(y; s)^{N-1}} + 1 \\
&= -\frac{\partial}{\partial y}\frac{1}{\alpha(y) F_w^*(y; s)^{N-1}} \geq 0
\end{aligned}
$$

  since $\alpha(y) F_w^*(y; s)^{N-1}$ is increasing in y by construction.

- Second, note that $[y - \lambda_\alpha(y; s, N)] \geq v_0$ for $y \geq v_0$:

$$
\begin{aligned}
[y - \lambda_\alpha(y; s, N)] &\equiv [y - \lambda_\alpha(y; s, N)]|_{v_0} + \int_{v_0}^{y} \frac{\partial}{\partial t}[t - \lambda_\alpha(t; s, N)] dt \\
&= v_0 + \int_{v_0}^{y} \frac{\partial}{\partial t}[t - \lambda_\alpha(t; s, N)] dt \\
&\geq v_0
\end{aligned}
$$

  since we know $\frac{\partial}{\partial y}[y - \lambda_\alpha(y; s, N)] \geq 0$.

- Finally, note that (by construction) $\alpha(y)$ is increasing in y.

Hence increasing y has two effects on the function (A.3): it increases $[y - \lambda_\alpha(y; s, N)]$ and shifts weight from v_0 to $[y - \lambda_\alpha(y; s, N)]$ (through $\alpha(y)$). Since $[y - \lambda_\alpha(y; s, N)] \geq v_0$, both these effects are positive, so (A.3) is increasing in y. It follows that increasing s leads to taking an expectation of an increasing function with respect to a stochastically dominated distribution. Hence the effect of s on R through the distribution $G_{1:N}^*(y; s)$ is negative.

Finally, note that $\rho_\alpha \geq 0$ by construction. Hence an increase in s implies a decrease in $(1 - s)N\rho$.

Combining these observations, we conclude that seller revenue $R_\alpha(s; N)$ is decreasing in s for any N.

□

Proof of Lemma II.4. We establish claims for $s_\alpha^+(z, N)$; the argument for $s_\alpha^-(z, N)$ is analogous. Suppose $\bar{s}$ is an equilibrium at (z, N, α). Then Proposition I.2 implies

$\Pi_\alpha(\bar{s}, N; F) \equiv c(z)$. Since $\Pi_\alpha(s, N; \tilde{F})$ is increasing in s and decreasing in F, it follows that

$$c^+(z) \geq c(z) \equiv \Pi_\alpha(\bar{s}, N; F) \geq \Pi_\alpha(\bar{s}, N; F^+)$$

Hence $\Pi_\alpha(s', N; F^+) > c^+(z)$ implies $s' \neq \bar{s}_\alpha(z, N)$ and (in particular, by monotonicity of $\Pi_\alpha(\cdot)$ in s) $s' > \bar{s}_\alpha(z, N)$. Taking the smallest such s' identified by $\mathcal{L}$ (or the uninformative bound 1 if no such s' exists) yields $s^+_\alpha(z, N)$ defined above.

□

Proof of Corollary II.4. See Haile and Tamer [2003] for proof that $r^*_+(z, N)$ and $r^*_-(z, N)$ bound $r^*(z, N)$. The final equality statement follows from exact identification of counterfactual revenue when $s_{r^*}(z, N) \in \text{int}(\mathcal{S}(\mathcal{L}))$.

□

A.3 Constructing 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.*[2]
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.*

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

Figures III.2, A.1, and A.2 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.[3]

Figure A.1: Price vs Regprice for Purex 100oz, IRI store 683960 (2002-2008)

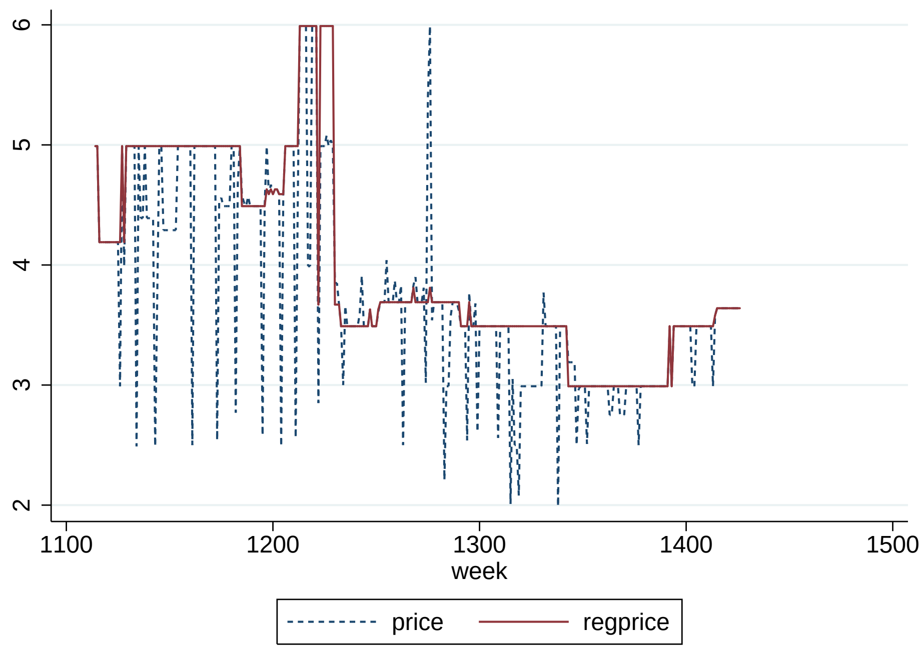

[3]In preliminary work, I also explored a regular price filter based on rolling modes. Experiments suggest the rolling median filter is more robust to potential noise in the price data. However, qualitative results were vere similar in both cases.

Figure A.2: Price vs Regprice for All 100oz, IRI store 683960 (2002-2008)

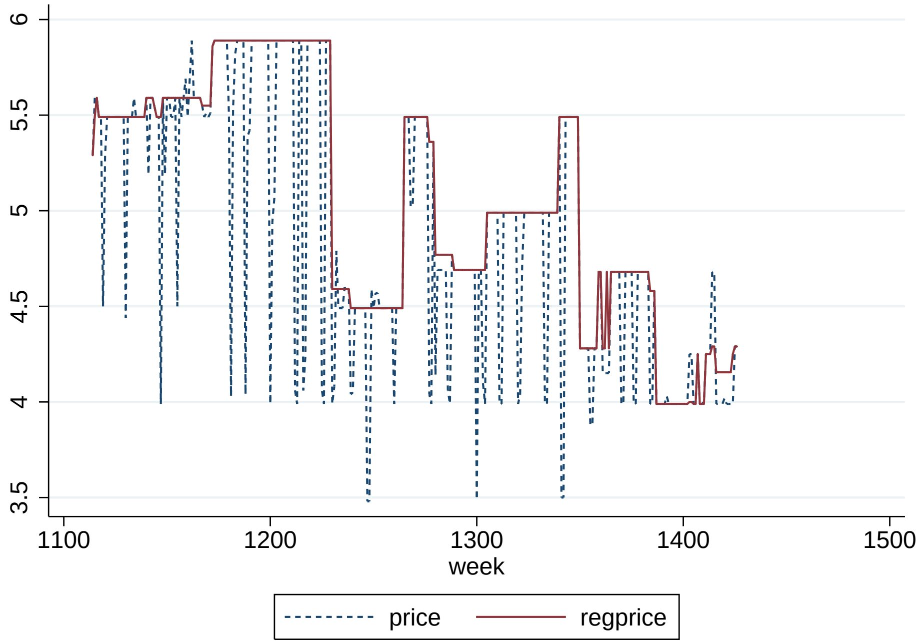

A.4 Simulating market shares

Consider first simulation of the search probability $\pi_s(\mathbf{u};\mathbf{p}_d)$ for a known utility vector $\mathbf{u}$. The key component in this probability is expected gain from search:

$$g_s(\mathbf{u};\mathbf{p}_d) = E_{p_s}[\max\{\mathbf{u} - \alpha \cdot (\mathbf{p}_d, \mathbf{p}_s)\}|\mathbf{p}_d] - \bar{v}_d(\mathbf{u}, \mathbf{p}_d).$$

Under Assumption 16, resampled discounts from the empirical sale distribution $\hat{F}_{\Delta p}$ can be used to produce a consistent simulator for the the conditional search probability above:

$$\hat{\pi}_s(\mathbf{u};\mathbf{p}_d) \equiv F_c\left[\hat{g}_s(\mathbf{u};\mathbf{p}_d);\theta\right],$$

where

$$\hat{g}_s(\mathbf{u};\mathbf{p}_d) \equiv \frac{1}{R}\sum_{r=1}^{R} \max\left\{\mathbf{u} - \alpha \cdot (\mathbf{p}_d, \mathbf{r}_s + \Delta\mathbf{p}_s)\right\} - \bar{v}_d(\mathbf{u}, \mathbf{p}_d).$$

For simplicity, the resampling algorithm I employ here further assumes that sale realizations Δp_j are independent across products. As noted in Section III.4, however, this additional restriction is not essential, and will be relaxed in future work.

In turn, combining the resampling simulator $\hat{\pi}_s(\mathbf{u};\mathbf{p}_d)$ with a standard GHK probit simulator gives a consistent, always positive simulator of the overall market share function $\sigma_j(\cdot)$. Note first that

$$\Pr(\text{choose } j) \equiv \Pr(\text{search} \cap \text{choose } j) + \Pr(\text{no search} \cap \text{choose } j).$$

Now consider simulation of $\Pr(\text{search} \cap \text{choose } j)$ as follows. First, following GHK, draw a size-R sample of utilities $\mathbf{u}_r$ such that good j will be chosen for each $\mathbf{u}_r$, and calculate standard GHK probability weights $\Phi_r(\theta)$ for these utility draws. Second, simulate search probabilities $\hat{\pi}_{sr}(\theta)$ for each $\mathbf{u}_r$ using the resampling simulator $\hat{\pi}_s(\mathbf{u};\mathbf{p}_d)$ as above, and reweight standard GHK probabilities Φ_{jr} by these search probabilities. Finally, averaging these reweighted probabilities across simulation draws yields a simulator for $\Pr(\text{search} \cap \text{choose } j)$ at parameters θ:

$$\hat{\Pr}_{sj}(\theta) = \frac{1}{R}\sum_{r=1}^{R} \Phi_{jr}(\theta) \cdot \pi_{sr}(\theta)$$

This simulator is computationally faster than its accept / reject analogue and inherits smoothness and positivity from the underlying GHK simulator. To show consistency, note that

$$\begin{aligned} \Pr(\text{search} \cap \text{choose } j) &\equiv \int_{\Omega} \pi_s(\mathbf{u}|\theta) \cdot \iota_j(\mathcal{J};\mathbf{u}|\theta)\, dF_{\mathbf{u}}(\mathbf{u}|\theta) \\ &= \int_{A_j(\theta)} \pi_s(\mathbf{u}|\theta)\, f_{\mathbf{u}}(\mathbf{u}|\theta)\, d\mathbf{u}, \end{aligned}$$

where $A_j(\theta)$ is the set of $\mathbf{u}$'s such that j is chosen (i.e. the set such that $\iota_j(\mathcal{J};\mathbf{u}|\theta) \equiv 1$). Further, by Assumption 15, $f_{\mathbf{u}}(\mathbf{u}|\theta)$ has a known joint normal form. GHK simulation involves approximation of the second integral via importance sampling from a corresponding truncated normal distribution $f_{A_j}(\mathbf{u}|\theta)$ defined on the acceptance region $A_j(\theta)$:

$$\begin{aligned} \Pr(\text{search} \cap \text{choose } j) &= \int \pi_s(\mathbf{u}|\theta)\, \frac{f_{\mathbf{u}}(\mathbf{u}|\theta)}{f_{A_j}(\mathbf{u}|\theta)} \cdot f_{A_j}(\mathbf{u}|\theta)\, d\mathbf{u} \\ &= \int \pi_s(\mathbf{u}|\theta)\, \frac{f_{\mathbf{u}}(\mathbf{u}|\theta)}{f_{A_j}(\mathbf{u}|\theta)} \cdot f_{A_j}(\mathbf{u}|\theta)\, d\mathbf{u} \\ &= \int \pi_s(\mathbf{u}|\theta) \cdot \Phi_j(\theta)\, f_{A_j}(\mathbf{u}|\theta)\, d\mathbf{u}, \end{aligned} \tag{A.4}$$

where $\Phi_j(\theta) \equiv \Pr(\mathbf{u} \in A_j(\theta)|\theta)$ and the last equality follows since $f_{A_j}(\mathbf{u}|\theta) \equiv f_{\mathbf{u}}(\mathbf{u}|\theta)/\Phi_j(\theta)$. The simulator above is consistent for the RHS of (A.4), and hence consistent for $\Pr(\text{search} \cap \text{choose } j)$.

A simulator for $\Pr(\text{no search} \cap \text{choose } j)$ can be constructed similarly, and summing over disjoint events yields a simulator for the overall vector $\sigma_j(\theta)$. I thus obtain a computationally tractable, always positive, and smooth-in-parameters simulator of the market demand vector $\sigma(\theta)$. A step-by-step description of the resulting algorithm is given below.

Algorithm (Simulating market shares)**.**

1. *Simulate search-plus-purchase probabilities.*

   (a) *Draw a size-R_1 sample of utilities $\mathbf{u}_r$ such that $j \equiv \arg\max_k\{u_{ik} - \alpha p_k\}$ following the standard GHK algorithm, then save the corresponding sampling probabilities $\Phi_r \equiv \Phi(\mathbf{u}_r|\theta)$ for future reference.*

   (b) *For each $\mathbf{u}$ drawn in Step 1(a), simulate implied search value V_{sr} using price resampling, and obtain corresponding search probability π_{sr}:*

      i. *Calculate maximum display-set utility $\bar{v}_r$ based on current-period market characteristics and known utility draw $\mathbf{u}$.*

      ii. *Draw R_2 resampled vectors of prices-to-search q_r from the empirical price distribution $F_p(\cdot)$, where $q \equiv \{p_j|j \notin \mathcal{D}\}$. This can be done using any model of $F_p(\cdot)$, and could in principle incorporate correlation between displayed and nondisplayed prices.*

      iii. *Simulate V_s via resampled average: $\hat{V}_s \equiv \frac{1}{R_2}\sum_{r=1}^{R} \max\{\mathbf{u} - \alpha p_r, \bar{v}_r\} - \bar{v}_r$, and define$\hat{\pi}_s \equiv F_c(\hat{V}_s|\theta)$.*

   (c) *Reweight Φ_s by π_s and average to obtain simulated probability of event "search and purchase j":*

$$\hat{\Pr}(search \cap j) = \frac{1}{R_1}\sum_{r=1}^{R_1} \Phi_r \cdot \pi_r.$$

2. *If $j \in \mathcal{D}$: Simulate no-search-plus-purchase probabilities.*

   (a) *As above, draw a size-R_1 sample of utilities $\mathbf{u}_r$ such that $j \equiv \arg\max_k\{u_{rk} - \alpha p_k|k \in \mathcal{D}\}$ following the standard GHK algorithm, and recover the corresponding sampling probabilities $\Phi_r \equiv \Phi(\mathbf{u}_r|\theta)$.*

   (b) *As above, simulate search values $\hat{V}_{sr}$ via resampling, and obtain corresponding search probabilities $\hat{\pi}_{sr}$.*

(c) *Reweight* Φ_r *by* $(1 - \hat{\pi}_{sr})$ *and average to obtain simulated probability of event "don't search and purchase* j*":*

$$\hat{\Pr}(no\ search \cap j) = \frac{1}{R_1} \sum_{r=1}^{R_1} \Phi_r \cdot (1 - \hat{\pi}_{sr}).$$

3. *Combine two conditional estimators above to obtain final simulated market share:*

$$\hat{\sigma}_j \equiv \hat{\Pr}(\boldsymbol{search} \cap j) + \hat{\Pr}(\boldsymbol{no\ search} \cap j).$$

This algorithm yields a consistent simulator of the true market share σ_j, which in addition is always positive and continuous in model parameters θ.

APPENDIX B

SUPPLEMENTAL TABLES AND FIGURES

Table B.1: 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 III.1.

Figure B.1: Price history for Purex 100oz, IRI store 683960 (2002-2008)

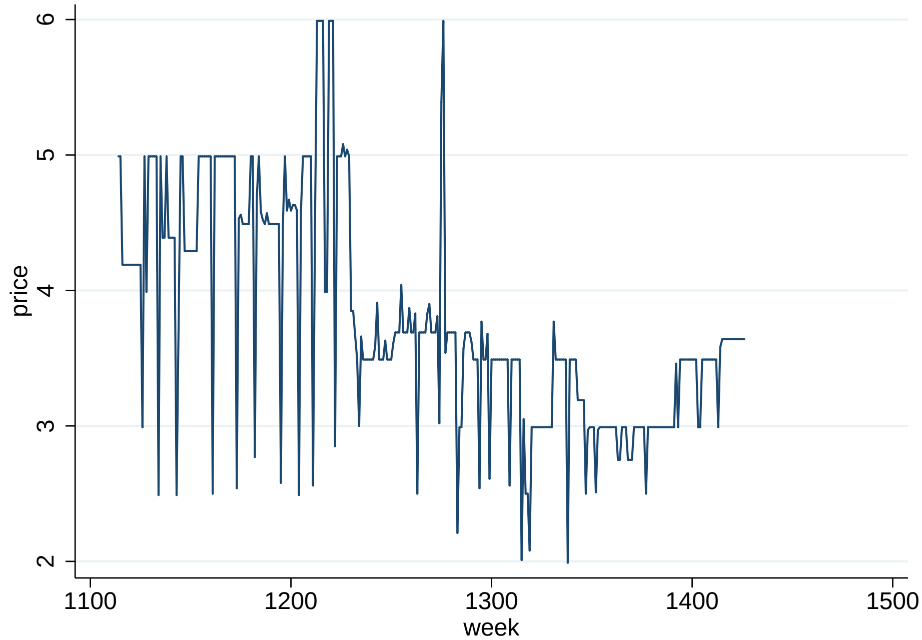

Figure B.2: Price history for All 100oz, IRI store 683960 (2002-2008)

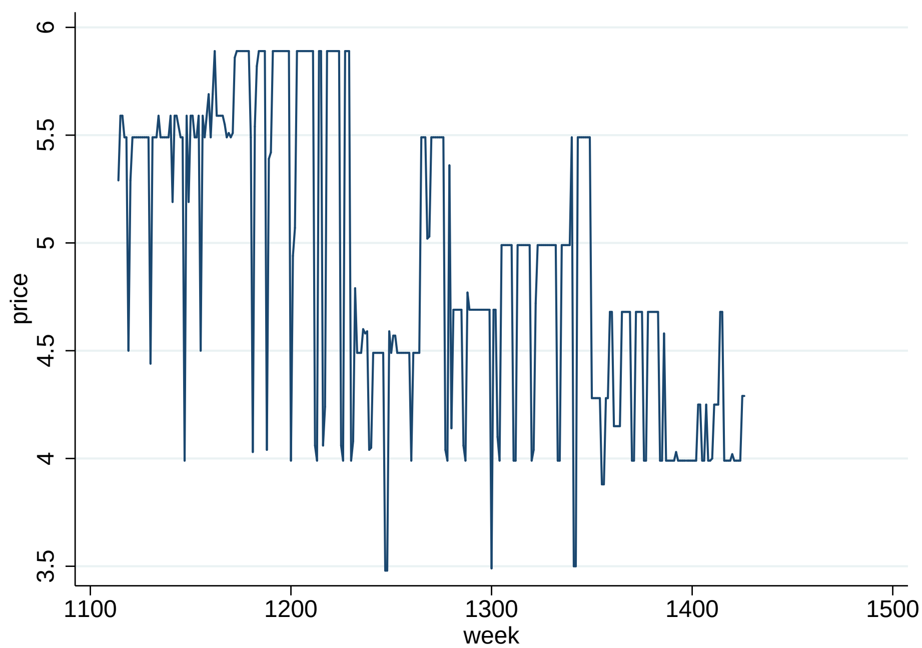

Figure B.3: Alternative price aggregates for Purex 100oz, IRI store 683960

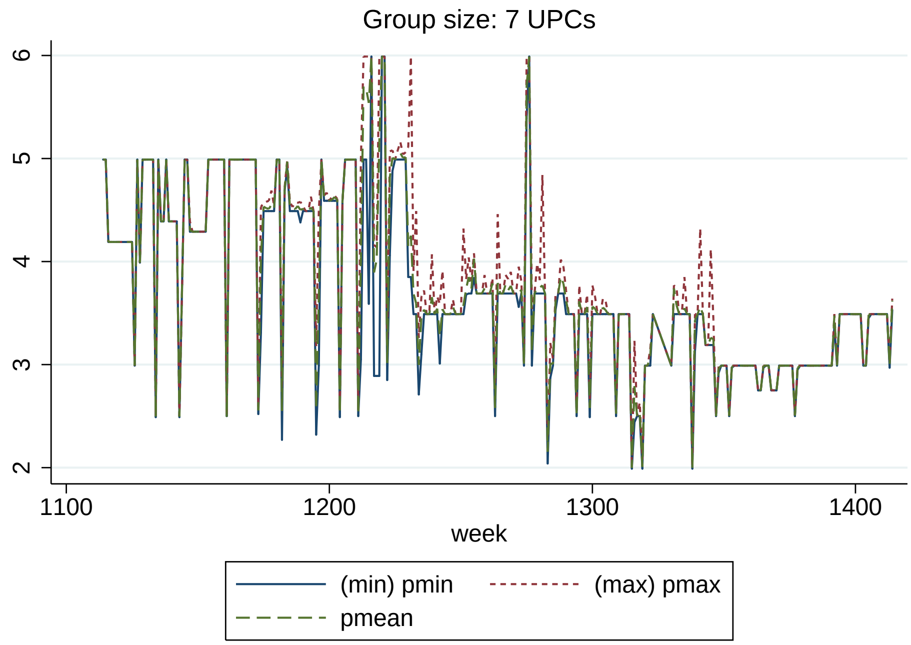

Figure B.4: Alternative price aggregates for All 100oz, IRI store 683960

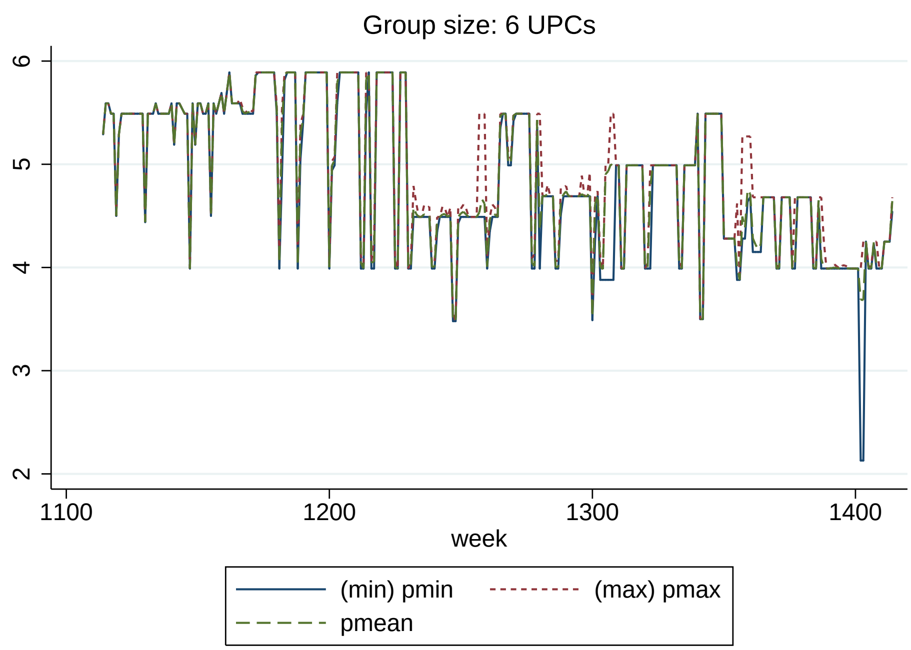

Figure B.5: Predicted N_t for IRI store 683960

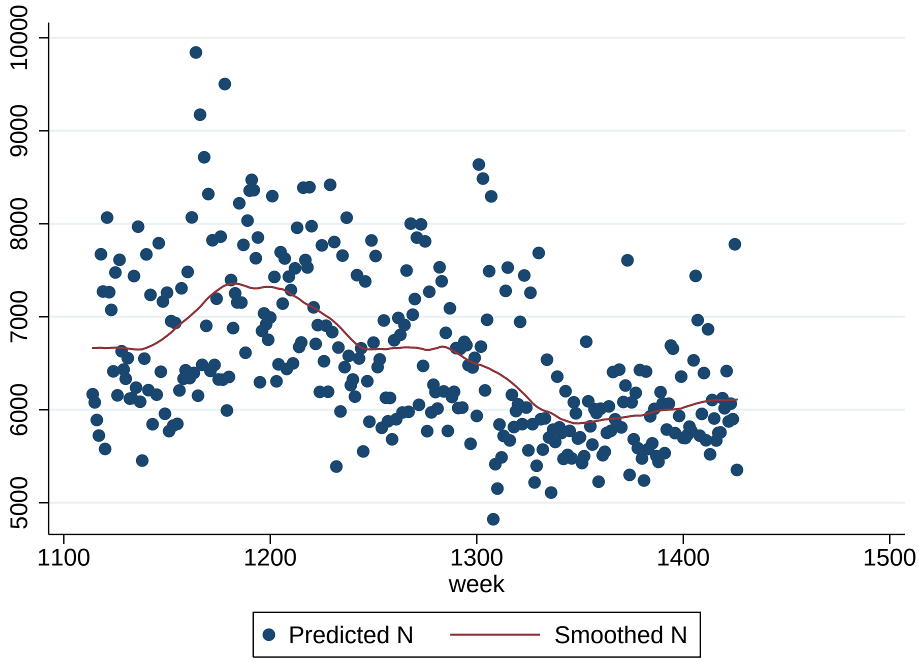

Table B.2: Promotional effects in Atlanta detergent markets, alternative specifications

	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 B.3: 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 B.4: 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 B.5: 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 B.6: Structural parameters: All stores, broad N_t

Parameter	243785	250094	263568	266100	653369	683960
α	0.325	0.262	0.375	0.343	0.299	0.307
λ	0.726	0.658	0.361	0.433	0.508	0.776
γ	4.44	0.82	1.187	0.633	0.245	>35
lag_1	-0.009	0.001	-0.016	-0.019	-0.005	-0.008
lag_2	-0.003	-0.008	-0.016	-0.011	-0.015	0
lag_3	-0.014	-0.006	-0.002	-0.016	-0.008	-0.009
lag_4	-0.007	0.004	0.017	0.002	-0.004	-0.005
$ftaste$	0.162	0.267	0.142	0.118	0.187	0.197
Objective	1861.8	2091.6	2128.6	2107.2	1946.3	2304.6

Table B.7: Selected structural own-price elasticities, all stores (narrow N_t)

Prod/Store	653369	250094	263568	266100	243785	683960
SURF100	-2.96	-2.40	-3.17	-3.32	-3.47	-3.57
ALL100	-2.29	-1.51	-1.72	-1.96	-2.42	-2.53
A&H125	-1.28	-1.46	-1.40	-1.33	-1.33	-1.99
CHEER80	-1.42	-1.45	-2.35	-1.38	-1.23	-2.83
ERA100				-1.25		
FAB100	-2.38		-2.75	-2.91		
GAIN100	-2.55	-1.84	-1.81	-2.21	-3.24	-2.50
PUREX100	-1.74	-1.82	-1.74	-1.82	-2.43	-2.35
PUREX200	-3.30		-3.32	-3.35		-4.33
TIDE100	-1.49	-1.22	-1.57	-1.43	-1.16	-1.87
TIDE200	-6.49	-5.36	-6.32	-5.43	-5.21	-6.27
WISK100	-3.02	-2.15	-3.20	-2.83	-2.70	-3.06

Note: Elasticities relative to baseline market with no discounts or promotions.

Table B.8: Cross-price elasticities, store 266100 (narrow N_t)

	ALL100	A&H125	CHR80	FAB100	GAIN100	PRX100	PRX200	TIDE100	TIDE200
ALL100	**1.96**	-0.12	-0.09	-0.03	-0.14	-0.18	-0.05	-0.21	-0.08
A&H125	-0.1	**1.33**	-0.09	-0.03	-0.14	-0.17	-0.42	-0.13	-0.43
CHR80	-0.09	-0.11	**1.38**	-0.03	-0.12	-0.16	-0.04	-0.19	-0.1
FAB100	-0.14	-0.16	-0.11	**2.91**	-0.2	-0.28	-0.08	-0.26	-0.07
GAIN100	-0.11	-0.12	-0.09	-0.03	**2.21**	-0.2	-0.06	-0.22	-0.08
PRX100	-0.15	-0.15	-0.12	-0.05	-0.2	**1.82**	-0.09	-0.26	-0.12
PRX200	-0.1	-0.95	-0.09	-0.03	-0.14	-0.19	**3.35**	-0.14	-0.42
TIDE100	-0.06	-0.06	-0.06	-0.02	-0.08	-0.1	-0.02	**1.43**	-0.71
TIDE200	-0.04	-0.26	-0.04	-0.01	-0.05	-0.06	-0.13	-1.45	**5.43**

Note: Elasticities simulated relative to baseline market with no displays, features or discounts.

Figure B.6: Model validation: actual vs predicted market shares, Store 683960

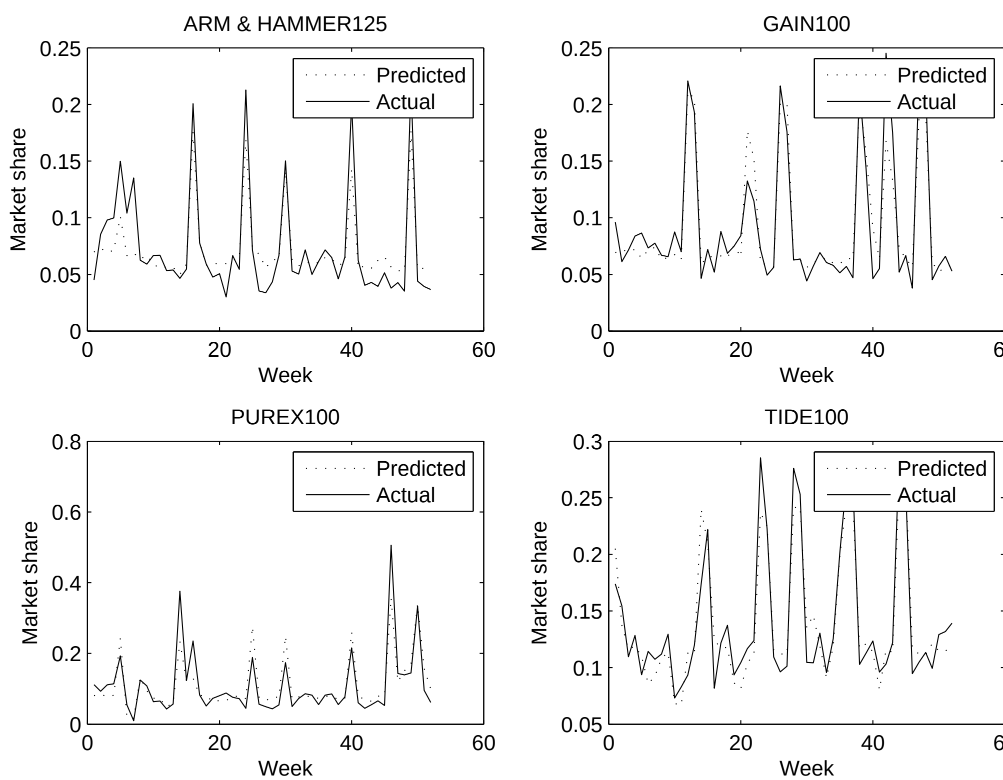

Table B.9: Cross-specification comparison, store 266100 (narrow N_t)

Parameter	Naive	Promo FX	Search
α	0.577	0.345	0.444
λ	–	–	0.354
γ	–	–	0.652
lag_1	-0.029	-0.06	-0.028
lag_2	-0.004	-0.025	-0.01
lag_3	-0.015	-0.006	0.008
lag_4	-0.012	0.028	0.003
$ftaste$		0.114	0.226
$dtaste$		0.398	–
Objective	-2310.8	-2149.2	-2123.7