



Marketing Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Effect of Retail Distribution on Sales of Alcoholic Beverages

Richard Friberg, Mark Sanctuary

To cite this article:

Richard Friberg, Mark Sanctuary (2017) The Effect of Retail Distribution on Sales of Alcoholic Beverages. Marketing Science 36(4):626-641. <https://doi.org/10.1287/mksc.2017.1038>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Marketing Science. Copyright © 2017, The Author(s). <https://doi.org/10.1287/mksc.2017.1038>, used under a Creative Commons Attribution License: <http://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2017, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Effect of Retail Distribution on Sales of Alcoholic Beverages

Richard Friberg,^a Mark Sanctuary^a

^a Stockholm School of Economics, SE-113 83 Stockholm, Sweden

Contact: richard.friberg@hhs.se,  <http://orcid.org/0000-0001-5951-1719> (RF); mark.sanctuary@hhs.se (MS)

Received: December 18, 2015

Revised: August 26, 2016; January 12, 2017

Accepted: February 7, 2017

Published Online in Articles in Advance:
June 13, 2017

<https://doi.org/10.1287/mksc.2017.1038>

Copyright: © 2017 The Author(s)

Abstract. We use monthly sales of all wines, beer, and spirits sold between 2006 and 2011 by Sweden's retail monopoly on alcohol to estimate the causal effect of retail distribution on market share by volume at the product level. Products are defined at the level of the stock-keeping unit. Two institutional features are key to identifying the causal effect: First, the monopolist uses four levels of retail distribution; a change in retail distribution is therefore associated with a discrete shift in the number of stores that carry a product in a given month. Second, the retailer is legally bound and monitored by the European Union to ensure that it acts in a non-discriminatory manner with all its suppliers. These features allow us to rule out many possible confounding factors in estimating the effect of distribution on sales volume. We find large and statistically significant effects from changes in retail distribution on market share by volume across all levels of retail distribution. The associated volume elasticity of retail distribution is convex; the wider the retail distribution the greater the marginal volume increase from further widening. In this market, wider distribution means reaching stores with successively smaller assortment. Our results indicate that the smaller assortment in smaller stores, coupled with a low resistance to compromise, is the main reason for the convex pattern. In other words, convexity appears to be generated by products achieving "a larger share of a smaller pie" as retail distribution expands.

History: Avi Goldfarb served as the senior editor and Harikesh Nair served as associate editor for this article.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "Marketing Science. Copyright ©2017, The Author(s). <https://doi.org/10.1287/mksc.2017.1038>, used under a Creative Commons Attribution License: <http://creativecommons.org/licenses/by/4.0/>."

Supplemental Material: Data are available at <https://doi.org/10.1287/mksc.2017.1038>.

Keywords: retail distribution • marginal returns to retail distribution • push-pull models

1. Introduction

Several studies have found a convex relation between retail distribution and sales in cross-sectional data (see, e.g., Reibstein and Farris 1995 or Wilbur and Farris 2014). The evidence further suggests that the links between retail distribution and sales are quantitatively important. Ataman et al. (2010) use rich French data to establish that the sales elasticity with respect to retail distribution (0.74) dwarfs the corresponding elasticity of advertising (0.13). However, despite the apparent importance of retail distribution for sales there is scant evidence of the magnitude of the causal effect of retail distribution on sales. Establishing such evidence may, for instance, be valuable for determining the willingness to pay slotting fees (see Bloom et al. 2000) or to determine damages in cases of exclusionary conduct in retail channels.

Identifying the causal effect is complicated by the fact that causality clearly runs both ways: A product is stocked because retailers expect it to sell well; it sells well because it is available to many consumers.

There are likely to be confounding factors as well. For instance, various sales promoting activities may make it easier to get onto retailers' shelves at the same time as sales expand. Yet some of these promotional activities may be hard for the researcher to observe.

Our aim in this article is to provide a clean identification of the causal effect of changes in retail distribution on volumes sold. We have access to rich monthly data for all products in the Swedish market for alcoholic beverages covering the years 2006 to 2011 inclusive. Observations are at the level of stock-keeping unit (SKU), which we henceforth refer to as a *product*. The particulars of retail distribution for alcohol in Sweden create a quasi-experimental setting that allows us to trace the causal effect of distribution on sales in an unusually transparent way. At the retail level, the state-owned Systembolaget has a monopoly. The suppliers to Systembolaget are independent, private firms, as is typical of retail markets. Because a monopoly could be used to favor some products, a number of mechanisms are in place to ensure that Systembolaget provides a

level playing field for different producers. This entails non-differential treatment of Systembolaget's suppliers and also non-differential in-store treatment of products. Thus, we can rule out many possible confounding factors for estimation of the effect of distribution on sales volume.¹

For its regular assortment, Systembolaget has four discrete levels of retail distribution; each product is assigned to a distinct level of distribution. The larger the store, the larger the assortment of products. All products that are available at smaller stores are also available in the larger stores. Changes in the retail distribution of products can be made twice per year, in April and October. Systembolaget bases these decisions on three backward looking observable criteria, i.e., availability, markup, and sales volume. Products are compared against others in the same segment (whiskey is an example of a segment). We match each product that experiences a shift in retail distribution to the most similar product (in terms of the three criteria) in the same segment that does not shift, and use the bias-adjusted matching estimator provided in Abadie and Imbens (2011) to estimate the causal effect of changes in retail distribution on volume. "Treatment" thus comes in the form of discrete shifts in the number of stores that carry a particular product as the product is moved from one level of distribution to another. We use logit regressions to verify that the criteria identified by Systembolaget also have predictive power for the decision to change. We further use the estimated probabilities to examine the "overlap" assumption of matching estimators, i.e., that the probability of treatment is bounded away from 0 and 1. Systembolaget makes its decision on retail distribution two months before the changes take effect. This provides us with a "pseudo" outcome, i.e., the impact on sales in the two months after selection into treatment has been made, but before it has taken effect. The pseudo-outcome results support the "unconfoundedness" assumption of the matching estimators: Conditional on covariates the assignment to treatment is independent of the outcome.

Our goal is not only to determine whether a causal relationship between sales volume and distribution exists but also to determine its *shape*. Nuttall (1965) argued that there were decreasing marginal returns to retail distribution, which imply a concave relation between retail distribution and sales. On the other hand, the modeling by Farris et al. (1989) and Reibstein and Farris (1995) suggests a convex relation in many situations; the wider the distribution the greater the marginal gains from further distribution. Note that a convex relation in the cross-section does not preclude a concave causal relationship of distribution on sales. Intuitively, Absolut Vodka may be available in almost all stores and sell well whereas more obscure vodkas

are available in very few stores and sell very little, which would appear as a convex pattern in the cross-section. If one of these more obscure vodkas gains wider distribution it is likely to sell more, but would clearly not reach the same intrinsic valuation as Absolut just by having the same retail distribution. Brand and product characteristics remain important for sales.

We find that widening distribution has a statistically and economically significant effect on volumes sold. For wine, we estimate that market share (by volume) of a product increases for each discrete step across the four levels of distribution, from narrow to broad, by 50%, 18%, and 17%, respectively. Expressed in terms of a distribution elasticity (percentage change in volume over percentage change in turnover weighted retail distribution) the results suggest that sales are a convex function of distribution. A widening in the retail distribution for wine across the four levels of distribution for wine yields estimated distribution elasticities of 0.12, 0.21, and 0.62, respectively. This convexity holds for widening and narrowing retail distribution. The quantitative effects and convexity of sales in distribution are similar for beer and spirits.

We examine "push" and "pull" factors (Farris et al. 1989, Reibstein and Farris 1995) that can explain the convexity of sales in distribution. The institutional setting largely rules out "push" factors apart from distribution per se. On the "pull" side, we find that expanding retail distribution reaches local markets with fewer consumers per store and where consumers on average have lower disposable income. This would suggest that sales are a concave function of distribution, i.e., lower volume expansions as distribution reaches additional stores. On the other hand, the assortment is narrower in these additional stores, which raises the volume effect of gaining distribution in these stores. Our results suggest that a narrower assortment in smaller stores is the key mechanism explaining the convexity of sales volume in value weighted retail distribution.

A few articles have sought to identify the causal effect of retail distribution on sales. To our knowledge, our paper is the first to use discrete shifts in retail distribution to identify the effects of distribution on volume. Bronnenberg et al. (2000) use data from the U.S. market for ready-to-drink tea to study the positive feedback between sales and distribution. They establish that such feedback shocks are important at early stages of a new product's life: Short-term changes may generate larger long-term changes. Bucklin et al. (2008) use cross-sectional variation in closeness to car retailers for California consumers to establish that dealer accessibility has a strong effect on market share for the respective brand.

The next section presents the data and institutional setting. Section 3 specifies the matching model and details how Systembolaget's practices inform the

Table 1. Summary Statistics

Variable	Mean	Median	S.D.	N
Wines				
<i>Liters</i>	13,895.20	4,878.75	28,674.82	65,163
<i>Price</i>	126.23	98.67	103.01	65,163
<i>Advertising</i>	69.52	.00	331.30	65,163
Beers				
<i>Liters</i>	80,469.88	21,818.11	166,645.02	15,106
<i>Price</i>	37.61	35.45	13.15	15,106
<i>Advertising</i>	24.08	.00	241.53	15,106
Spirits				
<i>Liters</i>	5,053.78	1,849.05	12,999.40	22,871
<i>Price</i>	401.43	355.71	152.29	22,871
<i>Advertising</i>	63.14	.00	462.00	22,871

Notes. The table reports descriptive statistics for the regular assortment of Systembolaget. Monthly data at the product-level are from January 2006 through November 2011. Advertising is given in thousands of SEK. In December 2008, 1 Euro is the equivalent of 10.7 SEK and 1 U.S. dollar is the equivalent of 8.0 SEK.

dimensions in which we search for the products that are to serve as controls for the treated products. Section 4 presents the results on the volume effects of changes in retail distribution; in Section 5 we discuss the patterns of distribution elasticities in light of a push-pull framework. Section 6 presents concluding remarks.

2. The Data and the Institutional Setting

We use monthly data at the product-level on retail sales of wines, beer, and spirits in Sweden from January 2006 through November 2011. Our data source is Systembolaget, the Swedish state-owned monopoly retailer for alcoholic beverages. We include all products in Systembolaget's regular assortment in the data set that we use for analysis.² For each product and month we observe the number of units sold and the price in Swedish kronor (SEK), as well as product characteristics such as container volume, alcohol content, region or country of origin, and product category such as light lager or sparkling wine. Table 1 provides descriptive statistics on the key variables for the sample. By far, most observations are for wine.

Data on advertising expenditure at the product level was obtained from SIFO/Research International.³ On average across years, the advertising intensity (i.e., advertising as a share of retail sales) equals 4.4%.⁴ Advertising for alcoholic beverages is done by the wholesalers and brewers that are suppliers to Systembolaget. Systembolaget does not undertake advertising of alcoholic beverages itself.

Systembolaget's purpose is to sell alcoholic beverages in such a way that alcohol-related damages are minimized.⁵ Moreover, Systembolaget is bound to rules that seek to ensure that it does not abuse its monopoly to favor some producers. "Systembolaget

shall, in its capacity as a purchaser, act in a non-discriminatory and brand neutral manner [...] The principle of equal treatment means that Systembolaget must apply the same terms and conditions to all beverage suppliers at any given time, and is not at liberty to negotiate the wording of them with one or several individual suppliers."⁶ In its relations with suppliers, Systembolaget is charged with providing a competitive, transparent, and uncomplicated setting. For example, Systembolaget's markup is determined by the Swedish parliament and it uses the same percentage markup for all products, which was lowered from 23% to 19% in August 2006 and remained at this level for the remainder of the sample period. Retail prices are therefore a deterministic function of wholesale prices; strategic interaction between Systembolaget and its suppliers on issues such as quantity rebates and trade promotions are mute.⁷

Systembolaget is instructed by the Swedish Government to operate efficiently, but is forbidden from taking some actions in which a profit maximizing retailer might engage. In-store, Systembolaget is, for instance, not permitted to provide differential treatment of products: There are no end-of-aisle displays; all beverages are sold at room temperature; on the shelves, products are organized according to type of product (wines, beer, spirits), country of origin, and price; there are no in-store tastings or postings of reviews.⁸ Prices are the same across the country; there are no temporary sales or any form of rebates for particular consumers (e.g., customer loyalty cards). Sales expanding efforts of this kind may be hard to control for in other settings. This aspect of the institutional setting is useful for interpreting the patterns that we later establish.

Systembolaget operates around 410 stores across Sweden during the sample period. The stores are spread across Sweden, a country roughly the size and shape of California and with around 9.4 million inhabitants at the time covered by the data. Each store is classified into one of four categories; extra-large, large, medium or small. *Extra-large stores* are in large towns and large shopping centers and offer the widest selection of products. *Large stores* are in large or medium-sized towns. *Medium stores* are in small towns or as "complements" to stores in medium or large towns. *Small stores* are in sparsely populated municipalities. By comparison, with many other retail markets there are relatively few Systembolaget outlets in Sweden and people are likely to shop at their local retail store; this is particularly true for the smallest stores where distances to the next store are large.⁹ Table 2 describes the share of stores belonging to each category as well as the share of turnover that accrues to each. Because Systembolaget sells alcoholic beverages only, the latter is a measure of what is often called weighted Product Category Volume (PCV).

Table 2. Number of Products Offered by Store Type and Product Category

Store size	Share of stores percent	Share of sales percent	Product range	No. of products offered		
				Wine	Beer	Spirits
Extra-large	23	33	Base, T1, T2, T3	918	212	322
Large	24	44	Base, T1, T2	714	167	267
Medium	31	17	Base, T1	495	117	201
Small	23	6	Base	253	65	119

Notes. This is the regular assortment of Systembolaget; the products offered are average number across periods in the sample as observed in the data; shares of sales and stores for 2010.

Systembolaget follows a store-class dependent stocking rule. All products in the regular assortment are categorized as belonging to one of four levels of retail distribution, which Systembolaget terms modules, and which are denoted: Base, T1, T2, and T3.¹⁰ The distribution level with the lowest coverage, T3, is distributed only in extra-large stores. When a product moves to the next tier, T2, it is distributed in the large and extra-large stores. A move to T1 means that a product is distributed in medium, large, and extra-large stores. A product in the Base module is distributed in all stores. The number of products offered by each store type is summarized in Table 2.

Any store will provide any product in the full assortment for delivery within two days; ordering can be done by phone, fax or in the store. A consumer in a sparsely populated area willing to wait two days thus has access to the full assortment. A relatively low share of sales is made up of such pre-orders: In the first half of 2012 they made up 1.6% of total sales (Konkurrenserket 2012).¹¹

Figure 1 illustrates, in cross-section, the relationship between a product's sales volume and the breadth of retail distribution. The most broadly distributed products sell more on average; the pattern appears especially marked for the products with the broadest retail distribution. Such a convex relation between sales and distribution in the cross-section has been documented by researchers studying a number of different markets, see, for instance, Wilbur and Farris (2014). Our key interest is how changes in retail distribution can be used to identify the causal effect of distribution on sales volume. After presenting the empirical model, we provide a closer examination of Systembolaget's practices as to changes in retail distribution.

3. Applying the Matching Method

To estimate the causal effect of a change in retail distribution on sales volume we use matching methods, comparing each “treated” product to a “control” product that is similar to the treated product in relevant dimensions. Our choice of dimensions in which to match products is guided by the explicit rules that Systembolaget is obliged to follow.

Following standard notation in applications of matching methods (see, e.g., Imbens and Rubin 2015) we are interested in the outcome variable of product i , Y_i . Define the market share of product i in month t in category $c \in (\text{wines}, \text{spirits}, \text{beer})$ as $Y_{it} = \text{Liters}_{it} / (\sum_{j \in c} \text{Liters}_{jt}) \times 100$. We use the average of Y_{it} in each four month treatment period as our outcome variable.¹² A product is “treated” if it changes distribution level in month t . We denote treatment status with W_i , which is equal to 0 for a product that does not change distribution level and equal to 1 for a product that changes distribution level. There are 12 occasions in our sample when a product can change distribution level. We use r to denote each of these occasions. Thus,

$$Y_{ir} \equiv \begin{cases} Y_{ir}(0), & \text{if } W_{ir} = 0, \\ Y_{ir}(1), & \text{if } W_{ir} = 1. \end{cases}$$


For a treated product we observe $Y_{ir}(1)$ but we do not observe the potential outcome $Y_{ir}(0)$. For each treated product we determine the “most similar” untreated product $Y_{jr}(0)$ and use the outcome for this product as our measure of $Y_{ir}(0)$. We use Mahalanobis weighting of three continuous variables as described below to find the closest product. We use the matrix X to denote these three variables.

Abadie and Imbens (2006) show that differences in covariates between the treated products and the controls can lead to a bias of matching estimators, even in large samples. Abadie and Imbens (2011) propose a bias-corrected matching estimator where the $Y_{jr}(0)$ is adjusted for differences in the values of the covariates between the treated products and the controls. Define $\mu_1 = E(Y(1) | X = x)$ and use ordinary least squares (OLS) on the matched sample to estimate $\hat{\mu}_1$. Following Abadie and Imbens (2011) we thus examine

$$\tilde{Y}_{ir}(W_{ir}) = \begin{cases} Y_{jr} + \hat{\mu}_1(X_{ir}) - \hat{\mu}_1(X_{jr}), & \text{if } W_{ir} = 0, \\ Y_{ir}, & \text{if } W_{ir} = 1, \end{cases}$$

and estimate the average treatment effect on the treated product, which is defined as¹³

$$ATET = E[\tilde{Y}_{ir}(1) - \tilde{Y}_{ir}(0) | W_{ir} = 1]. \quad (1)$$

Figure 1. (Color online) Average Monthly Sales Volume for Products by Level of Distribution, Systembolaget 2006–2011

Two key assumptions underlie the consistent estimation of treatment effects using matching estimators. One is that the probability of treatment, the propensity, is bounded away from 0 and 1 in support of the covariates such that each product has a positive probability of being treated. This is sometimes referred to as the *overlap* assumption: $0 < \Pr(W_{ir} = 1 | X_{ir} = x) < 1$.

The second key assumption is *unconfoundedness* (or conditional independence, selection on observables). Conditional on the covariates X , the assignment to treatment is independent of the outcome: $(Y_{ir}(0), Y_{ir}(1)) \perp W_{ir} | X_{ir}$.

Below we describe how the decision to change distribution of a product is made in Systembolaget and whether the assumptions on overlap and unconfoundedness are likely to hold in the current setting.

3.1. The Decision to Change Distribution Tier in Systembolaget

According to its rules, Systembolaget evaluates products based on three criteria only: volume, availability, and markup. First, we will describe the timing of events; then we describe the variables.¹⁴

Twice a year, on April 1 and October 1, changes in the distribution take effect. Whether a product changes its distribution level depends on its performance in the *pre-change period*, which ends two months before the possible change in distribution. For instance, Systembolaget bases the April 1 change in distribution on variables measured up to January 31 of that year as illustrated in Figure 2. While the decision to expand


the distribution is forward-looking, the specified criteria are essentially backward-looking; in particular the measure of volume is historical rather than a projection. The reason for the backward-looking criteria is that projections and expectations are hard to verify and more open to argument.

We measure the market share by volume per month for the four months following the change and call this the *treatment period*. We index treatment periods by r ; $r - 1$ refers to the pre-change, or selection, period for r . The outcome variable is thus $\tilde{Y}_{ir}$. The treatment period for one change will also serve as the selection period for the next change. The fact that Systembolaget decides on retail distribution two months before it takes effect implies that we observe sales for two months after treatment has been determined but before it has been implemented. The gap in time between selection and treatment, what we call the *pseudo outcome* period, will be used to assess unconfoundedness. We denote the outcome in this period by Y'_{ir} .

Next, we discuss the three criteria on which Systembolaget matches products (volume, availability, and markup) and how we operationalize those criteria. In our matching estimator, we include average liters per month in the pre-change period as the measure of *volume*. We expect that a higher volume is associated with a higher probability to move up in the hierarchy and with a lower probability to move down.

Availability aims to penalize products that run out of stock. According to interviews with Systembolaget,

Figure 2. The Timing of Changes; the Example of an April Change



this is relatively rare but it was nevertheless a criterion for retail distribution during the time of this study. To capture this, we create a measure of *volatility*, the standard deviation of sales volume in the pre-change period divided by the mean sales volume in the pre-change period. We expect products with a higher volatility to be less likely to move up and more likely to move down.

All products have the same percentage markups; more expensive products therefore have a higher absolute markup measured in SEK. A product with a higher markup is more likely to have a wider distribution since Systembolaget wants to avoid having low price/high volume products become too dominant. In our matching, we include average *markup* in the pre-change period. This is calculated as a percentage markup on the wholesale price plus a small fixed markup that depends on container size.¹⁵

We estimate the effect of six different changes in retail distribution, i.e., three steps for widening distribution by one tier (from T3 to T2, T2 to T1, and T1 to Base) and three for shrinking distribution by one tier (from Base to T1, T1 to T2, and T2 to T3). For each respective shift we set the sample to the products that have the same distribution level in the pre-change period. Thus, for examining the treatment effect of widening distribution from T3 to T2 we set the sample to all products that are T3 in the pre-change period $r - 1$. We assign $W = 1$ for a product that changes from T3 to T2 at r and $W = 0$ for the products that remain at T3 in r . In our estimation of treatment effects, we pool all 12 treatment dates but apply nearest neighbor matching within each pre-change period $r - 1$.

Products are compared to each other within segments and we consider 20 such segments. Eight segments are for red and white wines respectively, i.e., wines in bottles with a price below 70 SEK; wines in bottles with a price between 70 and 100 SEK; wines in bottles with a price above 100 SEK; and wines sold as

bag-in-box. One segment is for sparkling wines. Seven segments are for beer: ales; dark lagers; light lagers in bottles; imported light lagers in cans; domestic light lagers in cans; specialty beer (such as Belgian Kriek); and Weissbeer. Four segments are for spirits: Whiskey; Vodka; Schnapps; and other spirits. For wine and beer our definitions coincide well with Systembolaget's segment but for spirits we use somewhat wider segment categories; narrowing segments makes it harder to find products that can serve as controls.¹⁶

There are instances where a product jumps across several distribution modules (for instance, from T2 to Base). We have also estimated the treatment effect of these larger changes. Results conform to expectations: Widening distribution is associated with higher volumes, and conversely for shrinking distribution; effects are typically larger than when changing by just one tier. However, 91% of changes are to the nearest tier. Because of the possibility that the products selected as controls differ from the products that receive such exceptional treatment, we refrain from further analysis of changes across more than one distribution tier.

As in most retail markets, the product assortment is affected by considerable entry and exit.¹⁷ Because distribution is based on sales in the pre-treatment period we do not include current entrants in the estimations; an entrant in, for instance, $r - 2$ will be included in the sample in $r - 1$.¹⁸ Products can exit through various mechanisms. Products that sell too little are dropped (typically from T3 but may also exit from higher distribution levels). Exit may also occur for other reasons such as shifting strategy by the wholesaler or cost shocks. We do not try to examine the exit decision in detail. Products that exit in period r are not included in the sample in $r - 1$, but before this they are part of the sample. Because product entry and exit may lead to selection bias, we stress that here we only examine the effects of “marginal” changes in retail distribution.

Table 3. Logit Regressions for Wine

Variables	(1) T3 to T2	(2) T2 to T1	(3) T1 to Base	(4) Base to T1	(5) T1 to T2	(6) T2 to T3
<i>Volume</i>	0.0338** (0.0149)	0.0519*** (0.0101)	0.0242*** (0.00545)	−0.0123*** (0.00371)	−0.00744 (0.00538)	−0.0249** (0.0118)
<i>Markup</i>	0.00777* (0.00461)	0.00935* (0.00540)	0.0295*** (0.00897)	−0.0385*** (0.00949)	−0.0312*** (0.0102)	−0.00694 (0.00589)
<i>Volatility</i>	−0.00503 (0.00786)	−0.0165* (0.00891)	−0.00522 (0.00829)	−0.0417*** (0.00943)	−0.0341*** (0.00856)	−0.0343*** (0.00793)
<i>Rank</i>	−0.0472*** (0.00674)	−0.0309*** (0.00652)	−0.0485*** (0.00987)	0.0884*** (0.00836)	0.0883*** (0.00764)	0.0322*** (0.00349)
<i>Growth</i>	−0.667** (0.315)	0.00103 (0.0526)	−0.000476 (0.00701)	0.000744 (0.0287)	−0.000287 (0.00250)	0.00860 (0.0272)
<i>Advertising</i>	0.109 (0.108)	−0.0165 (0.0888)	0.00434 (0.0620)	−0.110* (0.0644)	−0.266*** (0.0908)	−0.276** (0.132)
Constant	−1.148* (0.671)	−4.479*** (0.713)	−3.566*** (0.597)	−1.128** (0.441)	−1.008*** (0.383)	−1.440*** (0.416)
Observations	1,591	2,181	1,882	2,488	2,014	2,305
Pseudo R2	0.0837	0.0763	0.0658	0.129	0.120	0.0871

Notes. Fixed effects for each segment are included in the specification. Standard errors are in parentheses. Volume in units of 1,000 liters. Advertising in units of 1,000 SEK. Volatility measured as (standard deviation of liters)/(mean liters) × 100. Dependent variable equal to 1 if a product moves as indicated by column headings. All explanatory variables measured for the pre-change period.

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Conceptually we may think of a curve relating distribution to volume; we use matching methods to identify the slope of that curve at three points.

A causal interpretation of results from a matching estimator relies on the assumptions of “overlap” and “unconfoundedness.” Next, we assess these assumptions.

3.2. The Probability of Being Treated and Assessing Overlap


The “overlap” assumption requires that treated products and controls should have overlapping probabilities of being treated. One common way to evaluate this is to examine the balance of covariates between treated products and controls. We report this in Table A.3 in the appendix, although it does not raise concerns. We can also estimate the probability of treatment using logit regressions and compare results for the treated products with the population. The logit regressions are also useful for evaluating the criteria that we examine and their link to treatment, i.e., a check that Systembolaget’s practices are in line with their stated policy. The estimated probability of treatment (propensity score) can also be used as a basis for matching, but the Mahalanobis matching on three criteria corresponds closer to Systembolaget practices and performed better in examinations of unconfoundedness.

In Table 3 we report the results from logit estimations on wine. The dependent variable is in each case 1 if a product is treated in period r and 0 otherwise. We also include the mean sales rank in the pre-change period.

The product with the highest sales in a segment and distribution level has rank 1, the second highest selling product has rank 2, and so forth. When estimating the probability of treatment the rank is a useful measure of volume to include.¹⁹ A priori, it cannot be ruled out that Systembolaget also considers other variables when shifting retail distribution. If Systembolaget did a wider evaluation than that inferred by the three criteria above, two other criteria suggest themselves. One is advertising expenditure. To the extent that advertising expenditure raises expected sales, this should be associated with a higher probability for wider distribution. We also calculate product-specific trends in market shares for the pre-change period to examine the possibility that trends are an important determinant. Apart from the reported coefficients we also include fixed effects for the different segments.

In line with Systembolaget’s stated policy, markup and volume are positively related to widening distribution and negatively to narrowing distribution. Our proxy for possible stock-outs, volatility, shows little relation to widening distribution but a negative relation to the probability of narrowing distribution. The latter is counter to our predictions but it must be noted that volatility is likely to be a rather weak proxy for the likelihood of stock-outs. In line with expectations rank has strong predictive power: A higher rank (for instance having rank 10 rather than rank 1 among red wines sold in Bag-in-Box in $r - 1$) lowers the probability of achieving wider distribution and raises the

Figure 3. (Color online) The Probability Density Functions for Treated (Solid Line) and Nontreated Wines (Dashed Line)



probability of narrowing distribution. With one exception pre-change growth is not significantly related to the probability of moving up or down in the distribution hierarchy. Overall the findings suggest that Sys-tembolaget follows its policies closely: While products that sell well are more likely to move up, it is the *level* and rank of volume rather than their product specific trends that appear to be driving movements. Although advertising seems to have little relation to movements up in the hierarchy, estimates for downward movements are negative and significant. This suggests that advertising may be playing a defensive role for wine, helping to avoid a downgrade. Our estimates of the causal effect that follow were essentially unaffected by the inclusion of advertising, however. Qualitative results are similar for beer and spirits but significance levels are lower, possibly reflecting fewer observations for those categories.

One way to illustrate the logit regressions and evaluate the overlap assumption is to plot the estimated probability distributions for the treated and nontreated wine (see Figure 3). While the estimated probability density function for the treated is to the right of the nontreated there is substantial overlap, suggesting that it is plausible that for each treated product we can find another control product that might have been treated but was not.²⁰

3.3. Assessing Unconfoundedness

The second key assumption in the matching framework is that assignment to treatment is independent of the potential outcomes, conditional on the pre-change variables. One way to assess unconfoundedness is to examine whether an alternative, “pseudo,” outcome, which should not be affected by treatment, indeed appears independent of treatment. A variable well suited for such an analysis is average market share in the two months after the distribution level has been determined but before the change has taken effect as discussed above. In Table 4 we present the estimates for this pseudo outcome for wine. Point estimates of the effect of treatment on the pseudo outcome are small. To interpret the results, we use the fact that on average 12.8 million liters of wine are sold each month. In one of six cases, the point estimate for the difference between treated product and control for the “pseudo” outcome is statistically significant: For wine moving from T3 to T2 it is around 256 liters per month ($(0.002/100) \times 12.8$ million). This difference is small, however, when set in relation to the average volume per wine of 13,900 liters as seen in Table 1. The other five changes in distribution have no statistically significant effect on the “pseudo” outcome and the estimated effects are very small, for instance the estimated effect of moving from T2 to T1 corresponds to 77 liters per

Table 4. Assessment of Unconfoundedness for Wine: The Effects of Treatment on Pseudo Outcome (the Market Share of a Product After a Decision to Treat But Before Treatment Has Been Implemented)

Change in distribution module	Market share (percent)			Pseudo-treatment effect	<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $\tilde{Y}'_{ir}(0)$	Treated $\tilde{Y}'_{ir}(1)$	Controls $\tilde{Y}'_{ir}(0)$			
T3 to T2	0.0234 (0.0079)	0.0408 (0.0914)	0.0388 (0.0885)	0.0020 (0.0096)	2.84 0.005	191
T2 to T1	0.0526 (0.0061)	0.1084 (0.1873)	0.1078 (0.1871)	0.0006 (0.0291)	0.27 0.783	199
T1 to Base	0.1116 (0.0175)	0.1704 (0.2795)	0.1669 (0.2830)	0.0035 (0.0428)	1.16 0.245	198
Base to T1	0.2567 (0.0113)	0.1562 (0.2247)	0.1565 (0.2230)	−0.0003 (0.0366)	−0.13 0.900	269
T1 to T2	0.1079 (0.0215)	0.1024 (0.1482)	0.1020 (0.1463)	0.0003 (0.0249)	0.23 0.818	308
T2 to T3	0.0535 (0.0064)	0.0370 (0.0615)	0.0366 (0.0588)	0.0004 (0.0113)	0.63 0.530	313

Notes. The treatment effect is *ATET* defined in Equation (1). Outcome variable $\tilde{Y}'$ is the average market share (in percent) in the two months after the decision to treat has been taken but before it has been implemented. *N* denotes the number of treatments in the estimation (same products as in Table 5). Standard deviations are in parentheses.

month. Also, for beer and spirits, the differences are typically not statistically significant and economically small, as seen in Tables A.1 and A.2 in the appendix.

We conclude that the setting appears well suited for applying matching methods. We have the largest number of observations for wine, which is the category where our examinations of overlap and unconfoundedness look particularly good. Although we do not see any grave concerns for beer and spirits, we believe that including these categories provides a richer picture than a unique focus on wine.

4. Results

In Table 5 we present the outcome of the analysis for wines. The first column shows the average market share by volume in the treatment period for products that are part of the population but which were not matched. Thus, for example, wines that had T3 distribution in the pre-change period and were not matched had an average market share of 0.027%. The second column shows the market share for the treated ($\tilde{Y}_{ir}(1)$); the third column gives the bias-adjusted market share for the products that act as controls ($\tilde{Y}_{ir}(0)$). We see that the (bias-adjusted) market share for the controls is substantially higher than for the non-matched as we consider movements up; products that are similar to those that are moved up in the distribution chain do better than more “ordinary” products. For instance, in the move from T2 to T1, the controls have a market share of 0.104%, which is substantially higher than the market share for the non-matched, at 0.058%.

Table 5. Wine: The Effects of Changes in Distribution on Market Share

Change in distribution module	Market share (percent)				<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $\tilde{Y}_{ir}(0)$	Treated $\tilde{Y}_{ir}(1)$	Controls $\tilde{Y}_{ir}(0)$	Treatment effect		
T3 to T2	0.0267 (0.0085)	0.0518 (0.1059)	0.0345 (0.0759)	0.0174 (0.0331)	7.24 0.000	191
T2 to T1	0.0578 (0.0077)	0.1228 (0.2086)	0.1038 (0.1787)	0.0190 (0.0723)	3.70 0.000	199
T1 to Base	0.1179 (0.0179)	0.1826 (0.2815)	0.1561 (0.2555)	0.0266 (0.0681)	5.49 0.000	198
Base to T1	0.2323 (0.0125)	0.1395 (0.2047)	0.1510 (0.2165)	−0.0115 (0.0709)	−2.67 0.008	269
T1 to T2	0.0972 (0.0182)	0.0824 (0.1209)	0.0977 (0.1439)	−0.0153 (0.0446)	−6.02 0.000	308
T2 to T3	0.0473 (0.0067)	0.0225 (0.0402)	0.0347 (0.0574)	−0.0122 (0.0220)	−9.82 0.000	313

Notes. The treatment effect is *ATET* defined in Equation (1). *N* denotes the number of treatments in the estimation. Standard deviations are in parentheses.

To test whether the treatment effect is statistically significant we use a *t*-test; *P*(*t*) gives the significance level. As seen, the treatment effects are statistically significant at the 1% level across all treatments. The treatment effect is also economically important in all cases. Moving from T3 to T2 raises the volume by 50% (0.0174/0.0345). Similarly, moving from T2 to T1 raises the volume by some 18%, and from T1 to Base by some 17%. As expected, widening distribution raises volumes and shrinking distribution lowers volumes. We discuss the pattern in Section 5.

In Table 6 we present the corresponding outcomes for beer. The results are as for wine: An increase in retail distribution raises volumes and shrinking retail distribution lowers volumes. The effects are large and statistically significant, with the exception of moves between Base and T1, which are imprecisely estimated. This is one instance where the bias adjustment due to Abadie and Imbens (2011) is binding. Not adjusting for the bias of a matching estimator would, for this case, result in an estimated treatment effect of 0.35 for movement from T1 to Base and −0.11 for demotion from Base to T1, both significant at the 5% level. Although, overall the bias adjustment does not affect the qualitative pattern, it is intuitive that it is binding for precisely this kind of change, when there are relatively few products available to use as controls (i.e., beer that is sold in all or almost all stores only) and the treated products differ from the controls in terms of covariates.

In Table 7 we present the corresponding outcomes for spirits. Again, the patterns are as for wine and beer: Wider distribution raises volumes sold; shrinking distribution lowers volumes sold. The effects are economically large and statistically significant.

Table 6. Beer: The Effects of Changes in Distribution on Market Share

Change in distribution module	Market share (percent)			Treatment effect	<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $Y_{jr}(0)$	Treated $\tilde{Y}_{jr}(1)$	Controls $\tilde{Y}_{jr}(0)$			
T3 to T2	0.0793 (0.0348)	0.1371 (0.2028)	0.0930 (0.1425)	0.0441 (0.0778)	3.16 0.004	31
T2 to T1	0.1722 (0.0349)	0.3227 (0.3522)	0.2154 (0.2448)	0.1073 (0.1628)	3.55 0.001	29
T1 to Base	0.4210 (0.0557)	1.1616 (1.4599)	1.1551 (1.5295)	0.0065 (0.3082)	0.11 0.915	26
Base to T1	1.0058 (0.0685)	0.5478 (0.5989)	0.5797 (0.6324)	−0.0318 (0.2602)	−0.86 0.396	49
T1 to T2	0.3899 (0.0643)	0.2193 (0.2213)	0.2664 (0.2598)	−0.0471 (0.0766)	−4.34 0.000	50
T2 to T3	0.1438 (0.0305)	0.0800 (0.0913)	0.1157 (0.1315)	−0.0357 (0.0570)	−4.84 0.000	60

Notes. The treatment effect is *ATET* defined in Equation (1). *N* denotes the number of treatments in the estimation. Standard deviations are in parentheses.

Table 7. Spirits: The Effects of Changes in Distribution on Market Share

Change in distribution module	Market share (percent)			Treatment effect	<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $Y_{jr}(0)$	Treated $\tilde{Y}_{jr}(1)$	Controls $\tilde{Y}_{jr}(0)$			
T3 to T2	0.0387 (0.0092)	0.0793 (0.0691)	0.0492 (0.0427)	0.0301 (0.0331)	6.04 0.000	44
T2 to T1	0.0871 (0.0091)	0.1651 (0.1817)	0.1195 (0.1273)	0.0456 (0.0615)	4.57 0.000	38
T1 to Base	0.2028 (0.0612)	0.3011 (0.2937)	0.2664 (0.2666)	0.0347 (0.0652)	3.54 0.001	44
Base to T1	0.6143 (0.0553)	0.2233 (0.2292)	0.2377 (0.2513)	−0.0144 (0.0777)	−1.67 0.098	81
T1 to T2	0.1874 (0.0321)	0.1014 (0.0848)	0.1160 (0.0963)	−0.0146 (0.0278)	−4.45 0.000	72
T2 to T3	0.0800 (0.0136)	0.0408 (0.0361)	0.0529 (0.0441)	−0.0121 (0.0240)	−4.14 0.000	68

Notes. The treatment effect is *ATET* defined in Equation (1). *N* denotes the number of treatments in the estimation. Standard deviations are in parentheses.

5. The Sales Volume Elasticity of Retail Distribution

Panel I of Table 8 shows the percentage change in market share by volume set in relation to the percentage change in turnover weighted distribution levels. To aid comparisons to other market settings and comparisons between widening and shrinking distribution we calculate all percentage changes at the midpoint.²¹ The elasticities summarized in Panel I are computed using the matching estimates presented in Tables 5–7. For comparison, Panel II shows purely cross-sectional “elasticities,” relating the percentage

change in mean volume (as reported in Figure 1) to the percentage change in turnover weighted distribution levels. Panel III relates the causal estimate of the percentage change in market volume to the percentage change in the *number* of stores that carry a product for a change distribution.

First, we observe that the cross-sectional elasticities in Panel II are much higher than the corresponding causal estimates in Panel I. Our estimates of the causal effects are, roughly speaking, an order of magnitude lower than the cross-sectional estimates. This highlights the fact that caution should be exercised if one is to give cross-sectional elasticities a causal interpretation in other markets.

The elasticities shown in Panel I of Table 8 suggest that sales are convex in turnover weighted retail distribution; the wider the distribution, the greater the percentage change in sales volume from further expansion. Whether there are increasing (“convex”) or decreasing (“concave”) returns to increases in retail distribution has been the subject of substantial theoretical interest. We examine this in the next section.

5.1. Push and Pull Factors as a Possible Source of Convexity

The academic literature shows that a convex relation can result from the interplay of “push” and “pull” factors (Farris et al. 1989). On the push side, Farris et al. (1989) highlight two factors, i.e., *distribution* and *in-store attractiveness*. Distribution is our key variable of interest: It is the impact of changes in this factor that we measure. On in-store attractiveness, there is little or no scope for a systematic relationship between a product’s in-store attractiveness and firm strategies in the market that we study: Because the rules governing Systembolaget limit differential product treatment, these mechanisms are largely shut down in our setting. More generally, other push side factors are likely to be weak in this setting. In many markets in-store attractiveness can be affected by competitive interactions between retailers selling the same product. In our setting, the retail monopoly has relatively few stores charging the same prices: Inter-store competition is limited. “Push” factors, apart from distribution per se, are weak. The explanation for the convex pattern should therefore come from “pull” factors.

On “pull” factors, Farris et al. (1989) consider *unmodified preferences* and *resistance to compromise*. Here, we consider unmodified preferences first. If consumers have a strong preference for a product, then that product is more likely to sell well; in turn, preferences may be affected by retail distribution. One possible mechanism is that preferences are partly shaped by habit. We may, for instance, grow accustomed to particular brands or develop an appreciation for wines of a certain origin or type.²² For the time period that we study,

Table 8. The Volume Elasticity of Retail Distribution

I. Marginal effect (turnover weighted)						
	Wine		Beer		Spirits	
	Up	Down	Up	Down	Up	Down
Between T3 and T2	0.12	−0.13	0.12	−0.11	0.15	−0.08
Between T2 and T1	0.21	−0.21	0.51	−0.24	0.39	−0.16
Between T1 and Base	0.62	−0.32	N.S.	N.S.	0.47	−0.25
II. Cross-sectional “elasticity” (turnover weighted)						
	Wine		Beer		Spirits	
Between T3 and T2	1.06		0.94		1.04	
Between T2 and T1	3.14		4.08		3.79	
Between T1 and Base	11.54		13.41		13.39	
III. Marginal effect (number of stores)						
	Wine		Beer		Spirits	
	Up	Down	Up	Down	Up	Down
Between T3 and T2	0.14	−0.15	0.14	−0.13	0.17	−0.09
Between T2 and T1	0.08	−0.09	0.20	−0.18	0.16	−0.06
Between T1 and Base	0.15	−0.08	N.S.	N.S.	0.12	−0.26

Notes. The elasticities are computed by taking the ratio between the change in sales from a change in distribution (from Tables 5–7) and the appropriate change in distribution (from Table 2). Panels I and II are computed using turnover weighted retail distribution; Panel III is computed using the number of stores. N.S. denotes the case where the change in volume was not statistically significant from 0 at the 10% level.

the four months following a change in retail distribution, it seems implausible that the effects of retail distribution on unmodified preferences via habit formation are strong enough to be the main drivers of the convexity of sales that we see in this market.

In simple economic models of choice, consumers are aware of all products and know the intrinsic value that they attach to any of them. This simplification, while convenient, may be especially inappropriate for products such as alcoholic beverages, a large number of which can be characterized as experience products. As noted by Nelson (1970): Consumers only know their valuation of such a product after having tried it. However, consumers may still have beliefs about the intrinsic valuation of a product. Such beliefs may be affected by changes in their information set which, in turn, might be triggered by changes in retail distribution. For instance, expert reviews or reports in social media could be triggered by widening distribution. Friberg and Grönqvist (2012) examine the impact of expert newspaper reviews on wine sales using Systembolaget data for 2002 to 2006. They find that positive reviews have a significant positive effect on sales, while the effect of neutral or negative reviews is small. If widening distribution triggers a positive review, this would then be a conduit by which widening distribution increases sales volume. Evaluated at the means, their estimates imply that volume in the treatment period is 3.4% higher because of a favorable review.²³ Although

favorable reviews could contribute to the positive effect of widening distribution, we do not expect this to be the main mechanism driving the convexity of sales. In particular, there is evidence that sales are convex in shrinking distribution as well; it is harder to envision reviews as a key driver of that result. Advertising responses by wholesalers/producers could in principle also shape preferences and be an important explanation for the patterns. However, in specifications that include advertising expenditure, qualitative patterns are the same. Moreover, as noted by Friberg and Grönqvist (2012), reviews are a much stronger force on short run demand in this market than advertising.²⁴

Resistance to compromise, i.e., the willingness of consumers to substitute their ideal product for another when their ideal product is not available, is a mechanism that could be driving the convexity of sales in distribution. In the words of Reibstein and Farris (1995, p. 193) “added distribution provides access not only to the customers who prefer the brand, which would be the linear impact, but also to the other unavailable brands’ consumers.” In the market for alcoholic beverages in Sweden, there are a large number of products and it is plausible that consumers have a relatively low resistance to compromise.²⁵ The stocking rule that stores follow is key to the interaction between the curvature of the distribution elasticity and resistance to compromise. Next, we provide a closer examination of this interplay.

5.1.1. The Stocking Rule and Convexity. Systembolaget applies a store class dependent stocking rule. When a product achieves wider distribution, it becomes available at successively smaller stores with smaller assortments. The richness of our data allows us to directly observe Systembolaget's stocking rule. For example, the number of U.S. red wines stocked by store type in July 2010 is increasing with store size. There are 4, 8, 23, and 30 U.S. red wine products sold in small, medium, large, and extra-large stores, respectively.

We find that the stocking rule is sufficient to generate a sales volume function that is convex in turnover weighted retail distribution. Application of the stocking rule means that as a product's distribution widens, it competes with successively fewer products. As we document below, widening distribution means the product accesses markets where there are fewer consumers with lower income. A product that moves to a wider distribution takes a larger share of a smaller pie. When the effect of fewer competing products dominates the effect of fewer, low-income consumers, we find that sales can be a convex function of distribution, other push-pull factors notwithstanding.

To illustrate these countervailing effects for wine, consider a standard model of demand in a monopolistically competitive market, as shown by Dixit and Stiglitz (1977). Assume that at each store there are a number of k consumers who each spend E_f on wine, where f denotes the size of the store.

Dixit–Stiglitz demand for a single consumer for a particular product i is

$$D_{if} = \frac{p_i^{-\sigma}}{\sum_{j \in \Omega_f} p_j^{1-\sigma}} E_f, \quad j \in \Omega_f,$$

where j denotes a product in the set of available products Ω_f in a store of size f ; σ denotes the elasticity of substitution; and p denotes prices.

Suppose all varieties are equally priced across all stores so that $p_j = p$ and there are n_f products in store f . The quantity demanded for each product in store f is then given by

$$kD_f = \frac{p^{-\sigma}}{n_f p^{1-\sigma}} E_f k = \frac{E_f k}{n_f p}.$$

This latter demand equation describes how volume sold varies with store size and consumer expenditures. Product demand is increasing in expenditure and the number of consumers, but decreasing in the number of products offered by the store.

For Systembolaget, the smaller the municipality, the smaller the store, and the smallest stores are in sparsely populated areas where per capita incomes are lower. Table 9 provides summary statistics of category level

Table 9. Summary Statistics on Annual Store Level Data

Store size	Variable	Wine mean	S.D.	Beer mean	Spirits mean
XL	Liters sold per store	772,379	366,681	875,953	104,619
	By municipality				
	Liters per capita	23.85	32.52	25.33	3.28
	Income per capita	225.91	31.17		
L	People per store	31,237	11,666		
	Liters sold per store	412,828	133,750	546,514	66,090
	By municipality				
	Liters per capita	18.86	7.77	23.58	2.85
M	Income per capita	220.43	24.71		
	People per store	26,851	9,774		
	Liters sold per store	220,199	170,781	343,375	44,786
	By municipality				
S	Liters per capita	17.59	27.61	25.09	3.34
	Income per capita	204.68	18.57		
	People per store	17,844	8,788		
	Liters sold per store	86,326	33,461	213,609	24,237
	By municipality				
	Liters per capita	11.32	4.77	27.71	3.14
	Income per capita	195.53	11.86		
	People per store	9,586	5,299		

Notes. The summary statistics are for annual store-level data for July 2003 to July 2004 using data from Asplund et al. (2007). The statistics per municipality are computed contingent on a given store type being in the municipality. Income is average annual income in '000 SEK for persons older than 20 for 2004.

Source. Systembolaget and Statistics Sweden.

sales at the store level by municipality.²⁶ The first row for each store class reports outcomes on volume sold during a year. Larger stores sell more wines, beer, and spirits on average. For instance, an extra-large store on average sold 772,379 liters of wine in a year, which can be compared to 412,828 liters for large stores, 220,199 liters for medium stores, and 86,326 liters for small stores. This conforms to expectations.

We aggregate sales for each store class in a municipality and report the number of liters sold per inhabitant in the municipality. Smaller stores tend to be in municipalities where less wine is sold per capita. By contrast, liters of beer and spirits sold per capita in the municipality are roughly independent of store size. To the extent that this is due to preferences, the diminishing sales of wine per capita in smaller stores would suggest that wine sales would be concave in distribution, not convex.

We also report the income per capita and number of consumers per store in the respective municipalities. Average income is lower for smaller stores in the municipality. Previous studies indicate that alcoholic beverages are normal goods, in Sweden and elsewhere (see, e.g., Asplund et al. 2007). Lower average incomes in municipalities with smaller stores would lead to lower demand per consumer in smaller stores. Finally,

we observe that there are fewer consumers per store in the municipalities with smaller stores. Both of these features would tend to work against a convex pattern, i.e., progressively lower demand in the progressively smaller stores, or a smaller pie to divide in the smaller stores. Expressed in terms of the Dixit–Stiglitz model: Lower expenditure per capita E_f ; and fewer consumers per store k . All else equal, these features suggest that sales for a product D_f should be a concave function of widening distribution.

However, product level sales in a store also depend on the number of products stocked (n_f). Systembolaget's stocking rule is such that smaller stores stock fewer products. The effect on sales of widening distribution to smaller stores is convex if the stocking rule (n_f) dominates the effects of expenditure (E_f) and consumers per store (k). With Dixit–Stiglitz demand, fewer consumers and lower demand per consumer in small stores can be offset by fewer competing products in these smaller stores. For instance, in the extra-large stores there are on average 918 wines whereas in the small stores there are 253 wines as seen in Table 2. Other things equal, achieving 1/918 of category sales is less than receiving 1/253 of category sales. Each product is thus likely to receive a larger share of the pie in the smaller stores.

Panel III in Table 8 shows distribution elasticities in terms of percentage changes in the *number of stores* that carry a product. These elasticities suggest a flat pattern. Around a quarter of the stores are in each size class and the flat elasticity with respect to the number of stores suggest that the number of products offered shrinks so as to broadly cancel out the effects of lower store-level demand. However, the smallest 23% of stores only account for some 6% of sales. With a more common measure of the distribution elasticity (which takes into account how much different stores sell) the relation becomes convex. In Section 6, we offer some observations about the extent to which these findings are likely to generalize to other retail markets.

On a final note, we acknowledge that consumers reached in the small stores may have different preferences than those reached in the larger stores. Although we abstracted from this in the discussion above, and in the sketch of Dixit–Stiglitz demand, preference heterogeneity may influence the shape of the relation between retail distribution and sales. We hope to have shown that a convex relation between retail distribution and sales volume can be explained by Systembolaget's stocking rule, rather than by the number of consumers reached and their respective incomes as retail distribution expands to successively smaller stores. Systematic differences in preferences may affect the strength of this observed pattern. This would be an interesting question to investigate in future research if researchers can gain access to product sales at the store level.

6. Concluding Remarks

We use detailed data provided by the Swedish retail monopolist for alcohol, Systembolaget, and the particular institutional features of this market, to identify and characterize the causal effect of retail distribution on the volume of alcohol sales. Wider retail distribution has a large, statistically significant effect on volumes sold. While the relation is intuitive, discrete shifts in retail distribution are not a common feature of retail markets. Identifying the causal effect in other markets may then need to rely on stronger assumptions: Evidence presented here may serve as a useful complement to evidence from such markets.

We highlight three features of our study that are likely to be important in taking our findings to other retail markets. First, retail elasticities are likely to differ according to product type: Per se, there is little to suggest that effects should be the same for relatively cheap convenience products such as wine or beer versus high-priced items such as a new car where consumers are likely to engage in a substantial search before purchasing.²⁷ Second, to an uncommon degree, our estimates exclude any measures related to in-store attractiveness that correlate with retail distribution. An endogenous response by retailers and brand owners to expanded distribution is likely to further increase the volume effect of expanded distribution, suggesting larger predicted effects in markets where such responses are important. Third, we measure the change in sales volume in the four months following a change in retail distribution. We thus capture the relatively short run. Long-run elasticities are likely to be higher as they also capture habit formation as consumers grow accustomed to particular brands.

The results indicate that a stocking rule where distribution starts in the stores with the largest assortments and expands to stores with progressively smaller assortments is sufficient to generate a sales volume function that is convex in turnover-weighted retail distribution. Is this finding likely to translate to other retail markets? Although discussing the generalizability of our results is somewhat speculative, we note that some of the features of Swedish retail distribution of alcohol are commonly observed in other retail markets. Key to the convexity in the present setting is that a few brands are available in all stores and large stores carry a wider assortment, which includes products found in smaller stores. Such patterns are common: Reibstein and Farris (1995, p. 190), for instance, state that “in the typical convenience goods distribution system there are a few large outlets that stock many brands and numerous smaller outlets that stock the leading brands only.” Hwang et al. (2010) establish that for U.S. supermarkets there is a set of top selling products that account for

half of all sales and that are available at almost all stores and that larger stores carry larger assortments within a given category. The factors that generate the convexity in the setting we study are therefore also likely to play a role in many other retail markets.

One aspect that could affect the portability of results is that Sweden has a relatively low income dispersion and has long been seen as relatively homogeneous. This would point to sales that are relatively similar across stores, which would dampen the extent to which smaller stores address very different consumer groups than larger stores. An indication that such differences play a limited role in some other settings comes from Briesch et al. (2009) who note that in many markets for convenience goods the links between total category sales in a store and the number of products available in that category are weak.

Acknowledgments

Financial support from the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged. The authors thank Therese Elmgren and Ulf Sjödin at Systembolaget for help with understanding Systembolaget's assortment policies, Lars Näsström for providing the advertising data, Kerem Cosar and Thomas Seiler for valuable comments, and two anonymous referees and the associate editor for valuable comments.

Appendix

A.1. Assessment of Unconfoundedness for Beer and Spirits

Table A.1. Assessment of Unconfoundedness for Beer: The Effects of Treatment on Pseudo Outcome (the Market Share of a Product After a Decision to Treat But Before Treatment Has Been Implemented)

Change in distribution module	Market share (percent)			Pseudo-treatment effect	<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $\bar{Y}'_{jr}(0)$	Treated $\bar{Y}'_{ir}(1)$	Controls $\bar{Y}'_{ir}(0)$			
T3 to T2	0.1146 (0.0355)	0.1138 (0.1686)	0.1011 (0.1508)	0.0127 (0.0261)	2.71 0.011	31
T2 to T1	0.1953 (0.0396)	0.2636 (0.2871)	0.2438 (0.2632)	0.0198 (0.0479)	2.22 0.035	29
T1 to Base	0.4627 (0.0618)	1.1637 (1.4921)	1.1458 (1.4637)	0.0179 (0.2266)	0.40 0.690	26
Base to T1	1.1436 (0.0841)	0.5988 (0.6607)	0.5784 (0.6684)	0.0204 (0.1829)	0.78 0.439	49
T1 to T2	0.4469 (0.0654)	0.2519 (0.2457)	0.2587 (0.2550)	−0.0068 (0.0581)	−0.83 0.412	50
T2 to T3	0.1857 (0.0399)	0.1214 (0.1354)	0.1243 (0.1413)	−0.0029 (0.0232)	−0.97 0.335	60

Notes. The treatment effect is *ATET* defined in Equation (1). Outcome variable $\bar{Y}'$ is the average market share (in percent) in the two months after the decision to treat has been taken but before it has been implemented. *N* denotes the number of treatments in the estimation (same products as in Table 6). Standard deviations are in parentheses.

Table A.2. Assessment of Unconfoundedness for Spirits: The Effects of Treatment on Pseudo Outcome (the Market Share of a Product After a Decision to Treat But Before Treatment Has Been Implemented)

Change in distribution module	Market share (percent)			Pseudo-treatment effect	<i>t</i> -statistic <i>P</i> (<i>t</i>)	<i>N</i>
	Not matched $\bar{Y}'_{jr}(0)$	Treated $\bar{Y}'_{ir}(1)$	Controls $\bar{Y}'_{ir}(0)$			
T3 to T2	0.0444 (0.0108)	0.0630 (0.0570)	0.0586 (0.0576)	0.0045 (0.0125)	2.36 0.023	44
T2 to T1	0.0962 (0.0078)	0.1510 (0.1756)	0.1444 (0.1530)	0.0066 (0.0295)	1.37 0.178	38
T1 to Base	0.2078 (0.0230)	0.3071 (0.3315)	0.2982 (0.3184)	0.0089 (0.0371)	1.59 0.119	44
Base to T1	0.6710 (0.0331)	0.2506 (0.2813)	0.2431 (0.2792)	0.0075 (0.0424)	1.59 0.116	81
T1 to T2	0.2154 (0.0258)	0.1233 (0.1122)	0.1186 (0.1147)	0.0047 (0.0156)	2.57 0.012	72
T2 to T3	0.0968 (0.0099)	0.0608 (0.0567)	0.0590 (0.0541)	0.0018 (0.0144)	1.04 0.303	68

Notes. The treatment effect is *ATET* defined in Equation (1). Outcome variable $\bar{Y}'$ is the average market share (in percent) in the two months after the decision to treat has been taken but before it has been implemented. *N* denotes the number of treatments in the estimation (same products as in Table 7). Standard deviations are in parentheses.

A.2. Covariate Balance

Matching methods require that treated and nontreated products have positive probabilities of being treated. One way to evaluate this is to examine if the covariates are relatively balanced between the treated products and the products

Table A.3. Normalized Difference in Covariates (Δ_X) Between Treated Products and the Full Sample and Between Treated and Control Products

	Liters		Markup		Volatility	
	Treated-pop.	Treated-control	Treated-pop.	Treated-control	Treated-pop.	Treated-control
I. Wines						
T3 to T2	0.22	0.08	−0.08	0.03	−0.07	0.08
T2 to T1	0.26	0.11	0.01	0.04	−0.16	0.03
T1 to Base	0.20	0.10	0.09	0.04	−0.08	0.02
Base to T1	−0.24	−0.05	0.03	−0.01	−0.08	0.02
T1 to T2	−0.02	−0.02	−0.05	−0.01	−0.12	0.02
T2 to T3	−0.19	−0.06	0.10	0.05	−0.12	−0.05
II. Beers						
T3 to T2	0.15	0.17	−0.01	0.07	−0.28	0.11
T2 to T1	0.25	0.07	−0.04	0.02	−0.40	−0.28
T1 to Base	0.47	0.25	−0.10	−0.09	−0.29	−0.08
Base to T1	−0.32	−0.09	−0.01	−0.01	0.01	0.09
T1 to T2	−0.36	−0.24	0.02	0.03	−0.12	0.22
T2 to T3	−0.24	−0.16	0.08	0.06	−0.21	0.11
III. Spirits						
T3 to T2	0.36	0.32	−0.10	−0.02	−0.08	0.06
T2 to T1	0.28	0.25	−0.07	−0.02	−0.15	0.05
T1 to Base	0.22	0.06	0.03	0.02	−0.04	0.00
Base to T1	−0.49	−0.19	0.07	−0.01	0.09	0.01
T1 to T2	−0.37	−0.18	0.17	0.02	−0.07	−0.02
T2 to T3	−0.38	−0.10	0.09	0.08	−0.17	−0.02

that serve as controls. Following the suggested roadmap for implementing a matching estimator in Imbens (2015) we may therefore assess covariate balance. If the covariates are relatively balanced between treated products and controls, this is an indication that they are similar enough for the controls to live up to their name. We thus calculate the normalized differences in covariates X where T denotes treated, C controls, s standard deviations, and an overbar denotes the mean

$$\Delta_X = \frac{\bar{X}_T - \bar{X}_C}{\sqrt{(s_T^2 + s_C^2)/2}}.$$

Imbens and Rubin (2015) suggest that if the normalized difference is below 0.25 then the covariate balance is not a major reason for concern. In Table A.3 we report the normalized differences between treated on one hand and nontreated at large as well as the control products chosen by our Mahalanobis matching. While differences are sometimes substantial relative to all the nontreated, they are generally small relative to the products that act as controls; only in 2 of 54 cases are they strictly above the 0.25 threshold (for movements from T2 to T1 for beer and from T3 to T2 for spirits).

Endnotes

¹ A number of these mechanisms were adopted in conjunction with Swedish accession to the European Union (EU) in 1995 to ensure that the retail monopoly could not be used to distort competition between producers and would be compatible with the internal market. The Swedish Competition Authority monitors Systembolaget on behalf of the EU Commission and produces two reports every year on Systembolaget's compliance with the rules on non-differential treatment.

² The regular assortment makes up between 95.7% (in 2006) to 97.4% (in 2011) of total volume sold in Systembolaget. The bulk of the remaining sales stem from the temporary assortment, mostly seasonal products or small shipments of high end wines. Systembolaget also maintains a catalog of products that are available on order only, which accounts for less than 1% of sales volume across years.

³ In cases where several products are affected (an advertisement for Absolut Vodka will, for instance, affect all of the products sold under this brand) we split the advertising expenditures according to sales volume.

⁴ As one comparison, for 2008 *Advertising Age* reports a U.S. advertising to sales ratio of 6.7% for the beverage industry (SIC 2080) and 6.2% for wine, brandy, and spirits (SIC 2084).

⁵ For more information on the purpose of Systembolaget, see <http://www.systembolaget.se/english/>. The description of Systembolaget's practices is based on discussions with representatives for Systembolaget but also on information from their yearly reports on "launch plans" (see Systembolaget 2012) and the bi-annual reports from the Swedish Competition Authority (see Konkursverket 2012).

⁶ <http://www.systembolaget.se/english/purchasing-process/> (accessed November 13, 2015).

⁷ See, for instance, Nijs et al. (2010) on the effects of trade promotions.

⁸ For evidence on the effects of such in-store measures to increase sales, see, for instance, Nakamura et al. (2014) on end-of-aisle displays or Hilger et al. (2011) on posting reviews in stores.

⁹ On average there are some 0.44 stores per 10,000 inhabitants. We may compare this to the corresponding store density for two U.S. states with somewhat similar alcoholic monopolies: Utah (0.2 stores per 10,000 inhabitants) and Pennsylvania (1.4 stores per 10,000 inhabitants) as reported in 2010 by <http://www.healthindicators.gov/>.

Sweden has about twice the land area of Utah and four times that of Pennsylvania.

¹⁰ Until October 2009, there was one additional level of distribution in the regular assortment: T4, which was distributed in 45 stores. The share of this category was at its largest in 2006, when it accounted for 1.6% of volume. A large number of these products were switched to T3 in October 2009, which makes it harder to identify the effect of distribution on sales from the switch between T4 and T3. We therefore do not use this source of variation and include only products in the Base to T3 distribution.

¹¹ Since February 2011, consumers can also pre-order using the Internet.

¹² Normalizing volume by category sales eases interpretation but, qualitatively, results are the same if we instead define the outcome variable Y_{it} in terms of liters of product i sold.

¹³ We use the *teffects nmmatch* command in STATA and apply the *biasadj* option.

¹⁴ The practices that we describe here changed in 2013. The basic principles were maintained but changes are now possible four times per year and the number of store formats have increased.

¹⁵ The relation between consumer price per liter (p) and wholesale price (p_{wh}) is as follows: $p = p_{wh} \times (1.25 \times (1 + \text{percm})) + (\text{percomm} + \text{taxperlitter}) \times 1.25$ where 0.25 is the value added tax; *percm* is Systembolaget's percentage markup which was 23% until August 2006 after which it was lowered to 19%; *percomm* is a fixed per-container charge of minor quantitative importance: For a regular bottle of wine, for instance, it is 3.5 SEK, amounting to less than 50 cents (U.S.) per bottle. The alcohol tax per liter is calculated by volume; tax rates differ depending on the alcohol content and whether the product is wine, beer or spirits.

¹⁶ We also do not include cases in which there are fewer than four products that can serve as controls for a particular change.

¹⁷ In total there are 2,589 products in the data set. Seven hundred eighty-seven products are available throughout the sample; 1,356 products enter, and 943 exit.

¹⁸ Most entrants enter in T3 but for new entrants Systembolaget does make a forward-looking estimate of sales volume; thus, new products can also enter at higher distribution levels.

¹⁹ In Mahalanobis matching we do not include rank, however, as it is essentially an alternative way to measure sales volume. We want to find a close relative to a product that is moved up in the distribution chain. Minimizing distance with respect to volume is a more direct measure of closeness than rank.

²⁰ Based on Figure 3 it may seem that a substantial share of estimated probabilities of change are 0, which would cast doubt on whether the overlap assumption holds. While in some cases the lowest probabilities are close to 0, such estimates are few. For instance, at the fifth percentile, the estimated probabilities for treated products are, respectively, 0.05, 0.03, and 0.04 in the top row of Figure 3, and 0.06, 0.07, and 0.07 in the bottom row. The corresponding probabilities for the products later chosen as controls by our nearest neighbor matching estimator are: 0.04, 0.03, 0.04, 0.03, 0.04, and 0.06.

²¹ Thus, for instance, the change for wine moving from T3 to T2 is calculated as $0.0174 / ((0.0533 + 0.0345)/2) / ((77 - 33) / ((77 + 33)/2))$ where the denominator uses values from Table 5 and the numerator uses values from Table 2.

²² Stigler and Becker (1977) provide an influential formalization of preferences with such "consumption capital." Bronnenberg et al. (2012) use the migration history of households and consumer scan data on a large number of packaged goods categories to establish that some 40% of geographic variation in market shares across the United States can be attributed to persistent preferences.

²³ Friberg and Grönqvist (2012) use weekly data; the effect peaks in the week after the review and gradually diminishes to an accumulated effect of a favorable review that is equivalent to 4,418 liters in

higher sales. With average sales of 8,218 liters per week (Friberg and Grönqvist 2012 only examine the T1 and Base modules) the total effect of a favorable review can thus be calculated as $4,418/(8,218 \times 16 \text{ weeks})$.

²⁴For example, we estimated a regression adjustment specification in our matching framework where we included advertising expenditure: The point estimates of the average treatment on the treated products were only marginally different than in our baseline specification.

²⁵Consumers can pre-order any product to their local store. As noted, however, volumes of pre-orders are relatively low, which is consistent with a relatively low resistance to compromise for most consumers.

²⁶While Systembolaget were forthcoming and willing to help with discussions of practices, their current policy is not to share store level data. Previously they were willing to share such data, however; Table 9 uses data from Asplund et al. (2007) for the last 12 months in that data set, covering July 2003 to July 2004. While clearly some variation may occur between years, we have no reason to believe that the overall pattern is not qualitatively the same for the period covered by our main data set.

²⁷The level at which we observe products may also matter. Previous evidence is typically at the brand rather than SKU-level. Indeed, Wilbur and Farris (2014, p. 155) note that “all prior evidence [on the convex relation between market share and distribution] used brand-level data.”

References

- Abadie A, Imbens GW (2006) Large sample properties of matching estimators for average treatment effects. *Econometrica* 74(1): 235–267.
- Abadie A, Imbens GW (2011) Bias-corrected matching estimators for average treatment effects. *J. Bus. Econom. Statist.* 29(1):1–11.
- Asplund M, Friberg R, Wilander F (2007) Demand and distance: Evidence on cross-border shopping. *J. Public Econom.* 91(1): 141–157.
- Ataman MB, Van Heerde HJ, Mela CF (2010) The long-term effect of marketing strategy on brand sales. *J. Marketing Res.* 47(5): 866–882.
- Bloom PN, Gundlach GT, Cannon JP (2000) Slotting allowances and fees: Schools of thought and the views of practicing managers. *J. Marketing* 64(2):92–108.
- Briesch RA, Chintagunta PK, Fox EJ (2009) How does assortment affect grocery store choice? *J. Marketing Res.* 46(2):176–189.
- Bronnenberg BJ, Dubé J-PH, Gentzkow M (2012) The evolution of brand preferences: Evidence from consumer migration. *Amer. Econom. Rev.* 102(6):2472–2508.
- Bronnenberg BJ, Mahajan V, Vanhonacker WR (2000) The emergence of market structure in new repeat-purchase categories: The interplay of market share and retailer distribution. *J. Marketing Res.* 37(1):16–31.
- Bucklin RE, Siddarth S, Silva-Risso JM (2008) Distribution intensity and new car choice. *J. Marketing Res.* 45(4):473–486.
- Dixit AK, Stiglitz JE (1977) Monopolistic competition and optimum product diversity. *Amer. Econom. Rev.* 67(3):297–308.
- Farris P, Olver J, De Kluyver C (1989) The relationship between distribution and market share. *Marketing Sci.* 8(2):107–128.
- Friberg R, Grönqvist E (2012) Do expert reviews affect the demand for wine? *Amer. Econom. J.: Appl. Econom.* 4(1):193–211.
- Hilger J, Rafert G, Villas-Boas S (2011) Expert opinion and the demand for experience goods: An experimental approach in the retail wine market. *Rev. Econom. Statist.* 93(4):1289–1296.
- Hwang M, Bronnenberg BJ, Thomadsen R (2010) An empirical analysis of assortment similarities across U.S. supermarkets. *Marketing Sci.* 29(5):858–879.
- Imbens GW (2015) Matching methods in practice: Three examples. *J. Human Resources* 50(2):373–419.
- Imbens GW, Rubin DB (2015) *Causal Inference in Statistics, Social, and Biomedical Sciences* (Cambridge University Press, New York).
- Konkurrensverket (2012) Övervakning av det svenska detaljhandelsmonopolet för alkoholdrycker—Rapport till europeiska kommissionen. Technical report, Swedish Competition Authority, Stockholm. <http://docplayer.se/18864521-Overvakning-av-det-svenska-detaljhandelsmonopolet.html>.
- Nakamura R, Pechey R, Suhrcke M, Jebb SA, Marteau TM (2014) Sales impact of displaying alcoholic and non-alcoholic beverages in end-of-aisle locations: An observational study. *Soc. Sci. Medicine* 108:68–73.
- Nelson P (1970) Information and consumer behavior. *J. Political Econom.* 78(2):311–329.
- Nijs V, Misra K, Anderson ET, Hansen K, Krishnamurthi L (2010) Channel pass-through of trade promotions. *Marketing Sci.* 29(2):250–267.
- Nuttall C (1965) The relationship between sales and distribution of certain confectionery lines. *Commentary* 7(4):272–285.
- Reibstein DJ, Farris PW (1995) Market share and distribution: A generalization, a speculation, and some implications. *Marketing Sci.* 14(3suppl):G190–G202.
- Stigler GJ, Becker GS (1977) De gustibus non est disputandum. *Amer. Econom. Rev.* 67(2):76–90.
- Systembolaget (2012) Lanseringsplan. Technical report, Stockholm. <https://www.systembolaget.se/imagelibrary/publishedmedia/4w2o5i7szhsm31mpobo1/2012-lanseringsplan.pdf>.
- Wilbur KC, Farris PW (2014) Distribution and market share. *J. Retailing* 90(2):154–167.