Generic aversion and observational learning in the over-the-counter drug market[1]

Mariana Carrera[2]
Case Western Reserve University

Sofia Villas-Boas[3]
University of California- Berkeley

November 3, 2013

DRAFT.
Please do not cite without permission.

ABSTRACT.

In a set of four-week labeling interventions at six locations of one national retailer, we tested three hypotheses for consumer aversion to generic OTC drugs: lack of information on the comparability of generic and brand drugs, inattention to the price difference between generic and brand drugs, and biased priors about quality that can be shifted through information on peer purchase rates. We use a difference-in-differences approach to measure the average treatment effects of each type of label. We do not find evidence that consumers are unaware of the existence or the stated comparability of the generic products. Our results are consistent with the other two hypotheses. In particular, observational learning, through posted information on the purchases of other customers, increases generic purchase shares by twenty percent.

[1] We thank Stefano DellaVigna, Ulrike Malmendier, Mark Votruba, and seminar participants at U.C. Berkeley, Case Western, and Duke University for helpful comments. We thank the retailer, whom we cannot name, for sharing data and allowing us to conduct a field experiment.
[2] Department of Economics, 11119 Bellflower Road, Cleveland, Ohio. Mariana.carrera@case.edu
[3] Agricultural and Resource Economics, Berkeley, California. sberto@berkeley.edu

1. Introduction

There has never been greater enthusiasm for shared decision-making between patients and their health providers. The concept of patient engagement has attracted recent attention as a strategy for improving healthcare outcomes, with some going as far as calling it "the blockbuster drug of the century" (Roehr 2013). Little is known, however, about how increased consumer engagement might impact the health care cost curve (Hibbard and Greene 2013). When given more say over treatment choice, patients may reveal staunch preferences for more advertised products, or biases towards more costly products or procedures because they are presumed to be superior. It is thus important to understand how consumer preferences over health products are formed and updated.

The over-the-counter (OTC) drug market is a ripe setting for studying how consumers form preferences over pharmaceutical products. In contrast to the prescription drug market, in which physicians select drugs for their patients, consumers choose OTC drugs autonomously. To facilitate the comparison of different products, the FDA requires standardized "Drug Facts" labels to be posted on every package. Furthermore, visible price tags make price comparisons easy, in sharp contrast to the prescription drug market. Nevertheless, more than half of the sales of familiar household drugs are for branded versions which cost 40-60% more than their generic equivalents. [4]

In the realm of prescription drugs, generic substitution has long been promoted by policymakers as an uncontroversial way to reduce costs without compromising quality. After decades of state-level policy efforts to facilitate generic substitution in pharmacies, 89% of prescriptions for drugs with a generic substitute (with the same content of active ingredients)

[4] Authors' calculations at a national retailer described in Section 3.

now result in the purchase of a generic formulation.[5] The comparatively high brand market shares in the OTC market might reflect strong consumer preferences for the added consumption value of brand products (e.g. Advil's candy coating). However, they might also reflect low awareness of the existence of a generic substitute, a high degree of uncertainty regarding the safety or quality of the generic product, or the fact that some cognitive effort is required to compute the savings offered by the generic version.

We conducted a set of four-week labeling interventions at six locations of one national retailer, to test three hypotheses for consumer aversion to generic OTC drugs: (1) lack of awareness of the molecular equivalence of generic drugs to their brand counterparts, (2) inattention to the price difference between brand and generic drugs, and (3) significant uncertainty regarding other product attributes, such as safety or taste, which might be reduced by information on other customers' purchases. Using OTC sales data from previous years, we selected six treatment stores for the intervention. For each hypothesis, we designed a product-specific information label to display relevant facts to consumers. For two of the tests, we designed labels with two different ways of framing the relevant information. We chose a sample of OTC categories to be "treated" with a label attached to the shelf price tag, and each of the six selected stores was randomly assigned to one of five label types (three different types of information, two of which had two framing variations).

We use another six comparable stores as controls, allowing us to use a difference-in-differences approach to measure how consumers respond to the different labels. We estimate treatment effects on two outcomes of interest: total quantity purchased (of the brand and generic versions of a product, combined) and generic share of total quantity. Using household-level sales

[5] This measure refers to the generic fill rate of prescriptions written for a given off-patent drug molecule with at least one generically manufactured version available. Note that generic substitution is more narrowly defined than *therapeutic interchange*, or the use of one drug in place of another drug that is not molecularly identical, but has a similar therapeutic effect.

data, we are able to estimate different effects on consumers who have not yet been observed to purchase a generic drug, versus those who have already tried them in the past.

Our results are as follows.

First, information on the comparability of generics has no effect on their purchase rate relative to the national brand. Our labels included statements such as “The FDA certified [this product] as bioequivalent and therapeutically equivalent to [brand name],” as well as the simpler statement “This product contains the same active ingredient as [brand name].” This zero-effect is surprising, but serves to dispel the concern that our labels might increase generic purchase rates simply through a salience effect, i.e. by drawing customers’ attention to the existence of a labeled product. Our results on total quantity (of brand and generic products, combined) are consistent with a salience effect; the quantity of treated products sold—but not that of untreated products –increased in all treated stores. However, the fact that all labels were posted beneath generic products does not appear to have impacted the generic purchase rate *per se*.

Second, we find empirical evidence consistent with inattention to the price difference between brand and generic drugs. The average generic purchase share across treated products increased by 5.6 percentage points when the price difference between brand and generic drugs was posted in percentage terms. Third, we find evidence consistent with general uncertainty about product desirability, perhaps encompassing attributes beyond stated equivalence of active ingredients. Using pre-period sales data for each OTC product in each store, we calculated generic purchase rates and displayed them in three stores, which varied in their existing rates. We also introduced exogenous variation in the posted share (within store and product) by alternating weekly the length of pre-period time used to calculate the posted generic purchase rates. We find strong evidence of a positive consumer response to these labels: The generic share increases by 8.5 percentage points overall (20% of the pre-period level). In stores where the current generic purchase share was above 50%, disclosing these shares to customers raised the purchase share by 10 percentage points. This effect is equivalent to what we would expect from lowering the relative price of the generic from 70% to 40% of the brand price.

Aversion to generic drugs is not well understood. In recent work, Bronnenberg, Dube, Gentzgow and Shapiro (2013) find that medical professionals, especially pharmacists, are significantly more likely to purchase generic OTC drugs than other consumers, implying that if consumers had greater knowledge about drugs, they would purchase brands less often. They also find that generic purchasing is positively correlated with education, holding income constant. Our paper compares the effects of three parallel treatments testing different hypotheses for the typical consumer's aversion towards generic drugs. In doing so, we contribute to several literatures, which we outline in the next section.

The OTC market is a particularly interesting setting for testing how consumers respond to information on their peers' purchases. There are three main channels through which an individual's demand for a good might be affected by this information: (1) goods that are socially visible may confer status, (2) there may be network benefits or "social utility" to using the same product as a peer, and (3) signals of its quality.[6] The first two channels are absent in our context, since the consumption of pharmaceutical products is a private behavior.[7] The third channel—learning about quality—is likely to be the only one through which learning the purchases of other customers would affect one's purchasing decision. Furthermore, given that most consumers do not buy OTC drugs often, and that peers do not usually observe one another's drug purchases, even simple statistics of generic purchase rates may contribute new and useful information to the

[6] A fourth channel is the behavioral "nudge" that can come from identifying a particular behavior as an injunctive or descriptive norm. This channel is less relevant here for two reasons. First, purchasing generic OTC products is not really considered to be pro-social or an action for the collective good. Second, Cialdini (1984) and the related literature typically find that high existing shares are necessary for these messages to be effective.

[7] We conducted an in-store survey of customers, finding that 24% (of 294 respondents) do not know whether more of their friends use brand or generic OTC drugs. By contrast, only 5 percent respond that they do not know, when asked whether more of their friends drink diet or regular soda.

average consumer. However, it is not theoretically clear whether disclosing market shares to customers will lead to an increase in generic share, given that they are close to 50% in the stores that we study. We use consumer surveys to complement our field experiment by soliciting the priors of brand- and generic-buying consumers on the purchase shares of other consumers, and assessing their likelihood of being influenced by market shares in future purchasing decisions.[8]

The rest of the paper proceeds as follows. Section 2 describes the relevant literatures and institutional background. Section 3 explains a simple conceptual framework to frame the testable hypotheses. Section 4 describes the empirical setting and summarizes the data, while section 5 outlines the intervention, research design, and empirical strategy. Section 6 presents the results, Section 7 discusses robustness, and section 7 concludes.

2. Background

2.1 Relevant Literature

There exists a large and growing theoretical literature that examines the ways in which uncertainty regarding product quality affects consumer demand (see Akerlof, 1970; Nelson, 1970; Wiggins and Lane, 1983; and Wolinsky, 1995). A closely related empirical literature analyzes the extent to which changes in product quality information affects consumer behavior via a variety of information types and sources, including branding (Montgomery and Wernerfelt, 1992), mandatory product labeling (Jin and Leslie, 2003), and advertising (Ackerberg, 2001; Ackerberg, 2003).

There is a smaller but growing literature identifying consumer inattention to non-salient components of costs. With a labeling experiment of similar design as the present study, Chetty, Kroft and Saez (2009) find that the sales of taxable products at a grocery store are reduced when their tax-inclusive price is displayed in addition to the tax-exclusive price. Hossain and Morgan

[8] Results forthcoming.

(2006) find that eBay customers do not sufficiently take shipping costs into account when placing bids. Our hypothesis is that all consumers observe the price of the brand-name product, but may be inattentive to—and tend to underestimate—the savings offered by the generic product. Our second test also relates to a marketing literature on customer responses to the framing of discounts in absolute versus relative terms. This literature has generally found that when prices are high, dollar framing dominates (Chen 1998; Gendall 2006) but mixed results for the hypothesis that when prices are low, percent framing is more effective. A weakness of these studies is that they use stated purchase intentions rather than actual purchase data.

Our third test adds to a large literature on observational learning.[9] Other experimental studies have studied peer usage disclosure in settings including restaurant menu choice and music downloads. Cai et al. (2009) find that restaurant customers respond positively to the identification of the "Top 5" most popular menu items; demand for these items increases by 13-20 percent. Salganik et al. (2006) created an artificial ''music market'' in which participants chose unknown songs to download. When users are shown a song's download count, their probability of downloading responds positively to it. In both of these cases, however, there may be social motivations that play a role in addition to the updating of quality expectations. At the restaurant studied by Cai et al., groups of diners eat "family-style," so the orderer's utility may depend on the others' enjoyment of the items ordered. Similarly, when new music becomes popular, those aware of this music might enjoy greater social status; this could explain why

[9] On the theoretical side, Becker (1991) developed a formal model in which the demand for a good depends positively on its aggregate quantity demanded, i.e., on peer demand. McFadden and Train (1996) formalize consumer learning about a new good's quality through a rational, forward-looking decision process in which they learn from their own experience and from the experiences of their peers. Moretti (2011) derives predictions from a model of social learning for the time path of movie sales after their release. He analyzes how sales diverge over time for movies that perform well or poorly relative to prior expectations and finds a reinforcing pattern: When a movie exceeds expectations in its opening week, consumers update their expectations via input from their peers, leading to greater future sales.

people are more likely to download songs that have been downloaded by many others. We study observational learning in a setting which is more plausibly free of confounding social motives.[10]

Studies of learning in the context of health treatments are primarily focused on physician decision-making, though there are a few studies that examine how the choices of patients respond to information on the choices of others. For example, Zhang (2010) argues that patients in line to receive a kidney make negative inferences about the quality of available kidneys based on the fact that they have been rejected by each prior potential recipient in the queue.

2.2 OTC Drugs in the United States

The federal regulations concerning over-the-counter drugs in the U.S. are complex. For older drugs, including acetaminophen (Tylenol) and diphenhydramine (Benadryl), the FDA publishes a "Monograph" specifying strict regulations for production, packaging, and labeling, but does not actively examine and approve the formulations sold by each manufacturer.

For newer drugs, the process differs. After going through the FDA approval process, all new drugs receive patent protection for a certain length of time, around 12 years on average.

After patent expiration, other manufacturers must themselves obtain FDA approval prior to selling their versions of the product. In the prescription drug market, the FDA tests generics for bioequivalence to the brand, which means that beyond containing the same quantity of the active ingredient, the mode of delivery results in a similar time pattern of release and absorption into the blood stream. Of the drugs in our sample, many, but not all, were tested for bioequivalence prior to entering the OTC market.

In general, the FDA and clinical studies fail to find differences in safety or efficacy between versions of a drug produced by the original brand patent holder and generic entrants.

[10] Bursztyn et al. (2013) separately estimate the influences of "social utility" versus "social learning" (updating quality expectations) in the case of paired peers purchasing financial assets.

(Kesselheim 2008). An exception is drugs that have a “narrow therapeutic index (NTI)” meaning that patient responses can be affected by very small differences in the timing and speed of ingredient absorption. No over-the-counter drugs are considered NTI.

Despite this, the perceptions of consumers (and to a lesser degree, those of physicians) are that generic drugs are less desirable than the original brand product (Shrank et al. 2011, Shrank et al. 2009).

3. Conceptual framework

One can think of drugs as an experience good: individual *i* only learns her utility levels v_{ibx} (for the brand version of product *x*) and v_{igx} (for the generic version of product *x*) after having tried each of them. Assuming that since brands precede generics on the market, all buyers of product *x* have used the brand version in the past and know the quantity $v_{ibx.}$ Those who have not tried the generic version do not know v_{ig}, and will continue to purchase the brand version if:

$$v_{ibx} – E[v_{igx}] > - \alpha_i (p_{bx} – p_{gx})$$

where α_i represents sensitivity to price, $p_{bx} – p_{gx}$ is the price premium of the brand version of *x*, and $E[v_{igx}]$ is the individual’s expectation of how much she will value the generic version. We do not make the assumption that $E[v_{igx}]$ are unbiased expectations. While some consumers may have accurate information and unbiased expectations of v_{igx}, others might systematically underestimate their utility from the generic, based on a number of different factors including lack of information.[11]

Each of our interventions aims to shift either $E[v_{igx}]$ or $\alpha_{i.}$ By displaying information on drug similarity, we test whether $E[v_{igx}]$ is based on inaccurate knowledge. By displaying the typical price difference in percentage terms, we test whether increasing the salience of $p_{bx} – p_{gx}$

[11] For example, immigrants from countries where generic drugs are unregulated and often counterfeit may be unaware that U.S. generics are more trustworthy.

increases the consumer's response (α_i).[12] In our third test, we display the share of other customers who buy the generic (or brand) version, to test whether $E[v_{igx}]$ responds to the popularity of each version.

If we find that generic shares increase when popularity information is displayed, this would imply that (1) generic-buying customers put less weight on the behavior of their peers as a signal of a drug's quality and/or (2) brand-buying customers have priors farther from the posted generic shares.[13]

4. Empirical Setting and Data

In collaboration with a major supermarket chain, we were given permission to post labels at the point of purchase. We posted labels beneath the price tags of selected products (see Figure 1 below), and reposted them each week after price tags were updated by the retailer.[14]

Figure 1. Example of Label Placement

This retailer is a large national chain, covering a wide range of demographic areas in the US and in Canada. The shelf layout of OTC products is very uniform across stores, promotional

---

[12] Highlighting the "savings" associated with buying the generic might also increase generic purchases by increasing the perceived deal-value of the product (described as transaction utility by Thaler (1985).

[13] We are in the process of testing this directly, through surveys that ask customers to make inferences about the choices of other customers. Preliminary results indicate that the first explanation is correct.

[14] We thank Yann Pannasie, Raymond Gong, Karen Yao, Caitlin Crooks, Feyisola Shadiya, Brian Mitchell, Roni Hilel, Kyle Kennelly, Jonathan Arenas, Kathy Hua, Fanglin Sun for research assistance in gathering data and implementing the label experiment. We also thank Ishita Arora, Samantha Derrick, Kathy Hua, and Ye Zhong for helping us gather auxiliary data and perform in-store surveys.

efforts are common across stores, and prices are typically the same across the stores we selected for treatment. The retailer is mainly focused on food sales, although a large share of locations have in-store pharmacies that dispense prescriptions and offer other health services such as flu shots and blood pressure screenings. The pharmacist, in those stores, can be easily approached by consumers seeking advice on what drug to choose among competing OTC products. We conducted an in-store pre-treatment survey of pharmacists in selected store locations and found that in fact, less than 5% of customers seek the input of pharmacists.[15] Nevertheless, we made sure that all treated stores had in-store pharmacies and that the response of pharmacists to customer inquiries, according to our pre-treatment survey, was consistent in these stores.

4.1 Data

We obtain from the retailer two types of data. First, a panel data set of store-level weekly sales for OTC drugs at the product-version (UPC) level. Second, a transaction-level dataset identifying households via masked identifiers. We use the store-level sales to report treatment effects overall. The transaction-level dataset will allow us to investigate what types of consumers responded most strongly to each treatment: those who been observed to buy generic OTC products in the past, or those who have always been loyal to brands.

For simplicity, we use the term "product" to refer to a set of brand and generic formulations with the same active ingredient at the same dosage level. Thus, we refer to product "generic share" as the share of quantity purchased, within this set of equivalent products, for the products sold under the store's private label.

Store-level data

[15] Most pharmacists indicated that about 10 customers per week approach them to ask questions about OTC products. Our household-level sales data indicate that at these stores, the average number of loyalty card holders purchasing an OTC product each week is 190.

The data contain the gross revenue, net revenue (net of promotions) and the total quantity sold of a particular UPC in a given week in a given store, with which we calculate each item's gross price and net price.[16] Prices are similar within price divisions and all stores in the treatment and control below to one of two price divisions. Prices change at the same time each week for all stores in the same price region (and this was when we went to the stores to post the labels). For the rest of the paper, we use "price" to refer to the net or promotional price, given that this is the price faced by the majority of the customers.[17]

We collapse our data to the level of the active ingredient rather than the UPC. For example, if a drug is sold in quantities of 12, 30, and 100, and in gelcaps as well as tablets, we will combine all of these UPCs into one observation. Since promotions are sometimes offered on only one of the UPCs, we compute a weighted average of prices separately for brand and for generic UPCs of the same active ingredient ("product"), using weights that are held constant across weeks and stores based on the aggregate sales quantity of each UPC.

The relevant factor in the decision to buy brand or generic is the price difference between the two. We know, from conversations with product managers, that the retailer targets a specific percentage difference rather than a dollar difference when setting its prices relative to the national brand. However, this percentage difference varies significantly as either the brand or the generic product may be offered at temporary promotion prices. Thus, we compute "price difference, percent" as the percent of price savings offered by the generic (e.g. .30 if the generic is priced at 70% of the brand's price) and include this as a control when estimating treatment

[16] The revenue is obtained as two columns (net and gross) in the raw data that are equal to each other if the product was not on promotion during a certain week in a certain store. Those two revenue columns will differ if there are promotions: the net column will feature a smaller dollar value than the gross column. If we divide both revenue variables by the quantity sold, we obtain the gross shelf price and the promotional price.

[17] A loyalty card for this retailer is required to obtain the promotional price. Since our household level data is obtained through loyalty card purchases, we can compute the share of purchases that are made using a loyalty card, and thus, at the promotional price. Across the six treated stores, this share varies from 79% to 86% with a mean of 83%.

effects. We use the weighted average of prices among brand and generic formulations, as described above, to compute this variable.

OTC Product Classes

We focus our analysis on the twelve largest OTC drug categories that offer generic (store-label) products for the majority of the products. These twelve categories can be more broadly grouped into the categories of pain relief, allergy relief, digestive/stomach relief, and relief for acute ailments such as colds and cough. In selecting classes to treat, our intention was to keep closely substitutable categories either all treated or all untreated, but also, to leave at least one category from these four groups as untreated, for a more balanced representation.

Using the pre-treatment data from January-May, 2012, we compute summary statistics for the twelve drug categories. Because it is considered proprietary information by our retailer, we cannot share these statistics broken down by category. Table 1 displays class-level statistics for treated and untreated categories. The generic share of purchases is the ratio of generic products purchased over total number of purchases.

For two of our tests, we designed labels that share these statistics, at the store/product level, with consumers. In the test of price salience, we shared the mean percent savings, and in the test of observational learning, we shared the mean generic share. These tests will be described in greater detail in the following section.

Class Name(*)	Treated	Mean Generic Share	Mean Brand Price (in $)	Mean Percent Savings	Mean Savings per Dose	Average quantity by week
Pain relief						
Treated Category	yes	38.9%	7.4	15.2%	47.8%	11
Untreated Categories (avg.)	no	39.7%	8.9	14.5%	48.1%	26
Allergy Relief						
Treated Categories (avg.)	yes	49.7%	12.1	17.8%	32.9%	10
Untreated Category	no	51.8%	7.6	43.5%	53.4%	6
Digestive/Stomach relief						
Treated Categories (avg.)	yes	48.2%	14.3	19.2%	28.2%	5
Untreated Category	no	35.8%	8.2	21.6%	43.2%	3
Acute Ailments						
Treated Category	yes	49.5%	10.4	24.4%	37.8%	4
Untreated Category	no	36.1%	12.8	47.4%	57.0%	6

(*) When just one category name per class, its name was omitted due to proprietary nature of the data, and data was averaged within a category if more than one was present.

Table 1. Summary Statistics of OTC Product Classes

Choice of treatment and control stores

In choosing our six treatment stores, we first ruled out stores that were farther than 45 minutes' driving distance from our university, as we would be manually adding labels on the same morning each week at all six stores. The second criterion was high weekly sales quantities, to maximize the number of customers we would reach per label. Third, we only selected stores featuring an in-store pharmacy, to guarantee that the layout and the information available to customers (via pharmacist interaction) would be similar. We surveyed the pharmacists in the stores within driving distance and asked about their interactions with OTC consumers (the full survey can be found in the appendix). Of particular interest were pharmacists' estimates of how often people seek advice regarding OTC medications, and whether they recommend generic alternatives when asked. Responses were highly consistent across the pharmacists we surveyed (shown in appendix). During visits to each store we confirmed that OTC products were similarly organized and displayed on the shelves.

To verify the comparability of the treated stores, and to choose six similar control stores, we estimated an OLS regression of generic share at the week/store/product level, using pre-period data on purchases and prices, to test for store dummy effects on both the level and the price responsiveness of generic share. The results of the regression show that four of the six treatment stores have higher baseline generic shares than most other stores in the district. The other two treatment stores have significantly greater price sensitivity than most other stores in the district.[18] For a comparable control set of stores, we chose the 6 that matched most closely on these two characteristics, yielding the set of stores shown below:

	OTC sales statistics in the pre-treatment period		Coefficients from generic share regression,		Demographic characteristics of store locations		
	Generic share	Avg. qty per week	on a store dummy	on store-price interaction	Med. Income, in thousand $	Share Age 65+	Share Asian
Treated Store 1	0.44	12	0.06*	0.09	88	13%	36%
Treated Store 2	0.42	9.8	0.06*	-0.03	69	15%	11%
Treated Store 3	0.37	5.1	-0.01	0.13*	68	8%	29%
Treated Store 4	0.45	9.8	0.14*	-0.11	78	12%	9%
Treated Store 5	0.39	5.7	0	0.12*	88	13%	36%
Treated Store 6	0.44	5.7	0.12***	-0.06	57	12%	25%
Control Store 1	0.43	9.2	0.06*	-0.01	75	8%	6%
Control Store 2	0.38	12.7	0.06*	-0.01	58	17%	9%
Control Store 3	0.44	5.1	0.03	0.13*	75	8%	6%
Control Store 4	0.4	7.1	0.095***	-0.08	64	15%	11%
Control Store 5	0.41	4.2	-0.02	0.16*	57	16%	9%
Control Store 6	0.4	7.6	0.06***	-0.04	60	13%	6%

Table 2. Comparison of sales statistics and demographic characteristics, treated and control stores.

There are some differences between the treated stores and the control stores. First, the share of Asians is systematically higher in the neighborhoods surrounding treatment stores.[19] We

[18] Though the coefficients on these stores are positive, the explanatory variable is the percentage price savings offered by the generic. Thus, the positive coefficient indicates a stronger response to a greater price difference.
[19] Demographic characteristics are obtained from the U.S. Census 2000 at the zip-code level, and matched to stores by their zip codes.

found that across the 60 stores in these three divisions, *Share Asian* is weakly but positively correlated with price sensitivity to the brand-generic difference. The other difference is that the median household incomes are higher, on average, in the neighborhoods of treated stores. We found, however, that area household income is neither predictive of generic share nor of price-sensitivity across stores (results not shown).[20] Furthermore, as long as trends do not differ between the treatment and control stores, the difference-in-differences approach will eliminate any permanent differences across stores.

The table below shows the balance between treated stores and untreated stores on the statistics computed by treated and untreated products, and the subsequent table shows that there were no statistically significant differences in the trends of generic share for treated categories in treated stores.[21]

Table 3. Sales statistics across treated and control stores in both treated and untreated products.

	Treated Products		Untreated Products	
Variables	Control Stores	Treated Stores	Control Stores	Treated Stores
Generic Share	43.0%	43.2%	39.8%	41.1%
Average Brand Price (in $)	12.2	12.1	8.8	8.9
Percent Savings, average	22.4%	20.8%	29.9%	30.0%
Unit Savings, average	37.2%	36.8%	48.0%	46.7%
Weekly Quantity Sold	7.4	7.2	8.2	8.9

[20] This may be because the median income in the zip code where the store is located is a noisy proxy for the income of the store's shoppers. Unfortunately, we have no access to shopper demographics.
[21] There are, however, differing time trends for *Treated* and *Other* OTC product categories. This is due to the seasonality of some medicines included in the *Treated* group.

	All 54 other stores		Matched stores	
	Treated Products	Other OTC	Treated Products	Other OTC
Time Trend	-0.00256**	0.00319***	0.000145	0.00484***
	(0.000867)	(0.000725)	(0.00146)	(0.00129)
Treated store	0.0261	0.0505**	0.0346	0.0289
	(0.0187)	(0.0156)	(0.0233)	(0.0205)
Trend*Treated	0.000786	-0.000290	-0.00192	-0.00194
	(0.00157)	(0.00131)	(0.00190)	(0.00167)
N	374	374	192	192
R-sq	0.075	0.142	0.022	0.103

Table 4. Comparison of generic share trends, over a 19-week period prior to treatment, across treated and control stores for both treated and control products. (See Figures 2a, 2b for plot.)

Household-level data and purchase dynamics

We observe 20,561 households who make at least two purchases of the same drug (active ingredient) over the time period observed. OTC drugs are not purchased very frequently. The median number of weeks between purchases of the same active ingredient is 9, and the 25th percentile is 4. Of the 27,085 household-drug combinations that we observe with two or more purchases, 51% only ever buy the brand.[22] Of the remaining 49%, 80% buy the generic in the first purchase that we observe, and this group continues choosing the generic at a rate of 83% in future purchases.

[22] This share is not very responsive to increasing the minimum number of purchases. With a minimum of 5 purchases of the same product, 46% have only bought the brand, and with a minimum of 10 products, 42% have only bought the brand.

4. Intervention

The four-week intervention consisted of posting labels beneath the price tags of generic versions of products in treated categories in treated stores. We now explain the content of the labels for each of our five tests, which we number as tests 1, 2a, 2b, 3a, and 3b, corresponding to our three hypotheses, but with tests 2b and 3b offering variations on the information framing.

Test 1. To test the first hypothesis, lack of information on drug quality, we reported the similarities between brand and generic products as specifically as possible, with one of three labels displayed in Figure 3. The highest strength statement we used was: *"The FDA determined this product to be therapeutically equivalent and bioequivalent to* [corresponding brand product]," taken verbatim from the FDA approval letter, if accessible. The medium strength statement used was *"The FDA approved this product."* This was shown along with the reference number and date of FDA approval. This label appeared on products for which we found notices of FDA approval, but either no letter on the relevant webpage, or a letter that did not include any statement about bioequivalency. The lowest strength statement was *"This product contains the same active ingredient as* [corresponding brand product]*."* For older-generation drugs such as acetaminophen or aspirin, a special policy applies whereby companies need not seek approval from the FDA prior to marketing a generic, as long as they follow drug-specific formulation rules reported in an FDA monograph.

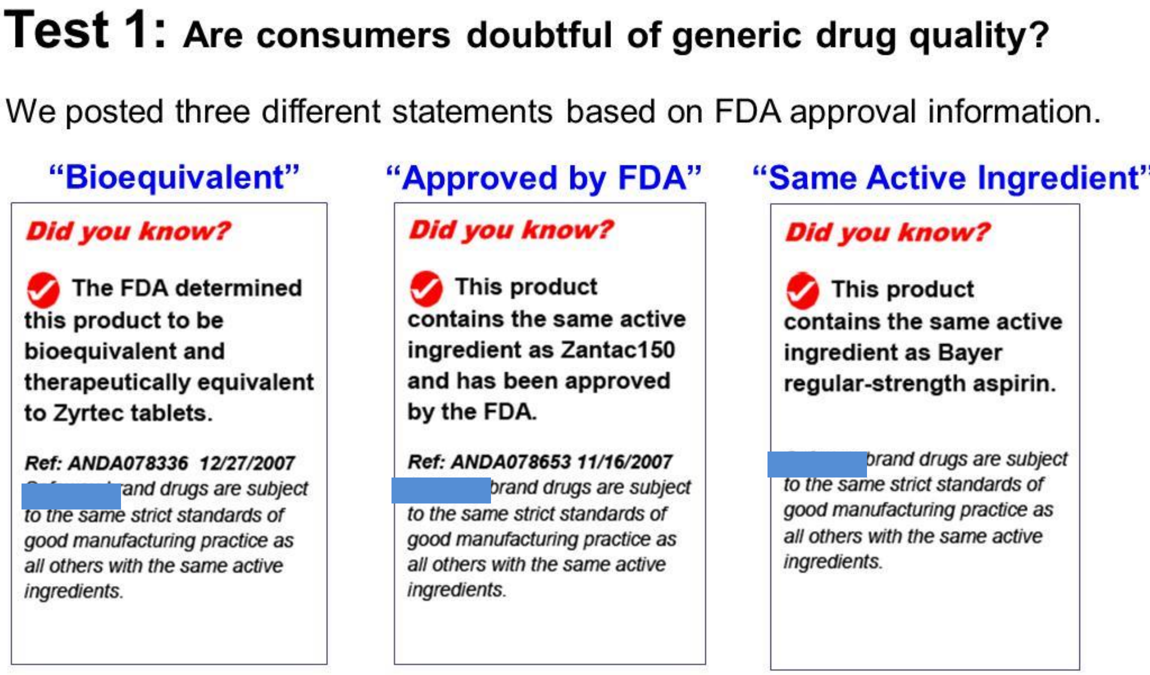

Figure 3. Labels highlighting quality

Test 2a. To test for inattention to price differences, we posted labels stating "*Customers who choose this product save X%* (relative to [corresponding brand product])." X ranged from 14% to 68% in the products labeled and an example of one such label is below.

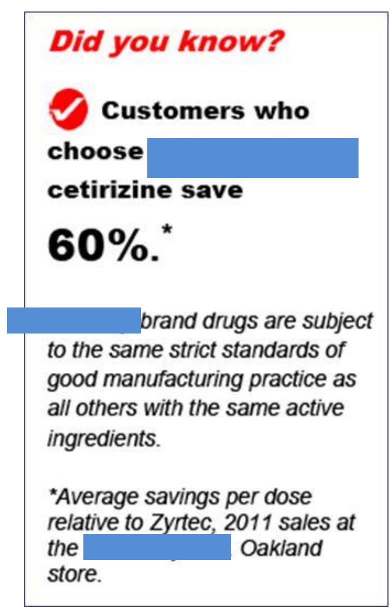

Figure 4. Label highlighting savings on generic OTC relative to brand

Test 2b. In another store, we highlighted the price differences in a different way, by stating "*Customers who choose* [corresponding brand product] *pay Y% more than the generic alternative*." Notice that for the same brand and generic prices, Y will be a larger number than X. In this type of label, the price difference is framed as a loss rather than a gain, and the percent value displayed is larger. For these reasons, we hypothesized that Test 2b would be more effective than Test 2a. However, the retailer only allowed us to place labels beneath their own (generic) items, so associating "pay more" with the item directly above the label might dampen the effect.

Test 3a. To test for observational learning, we posted labels stating *"X% of customers in this store choose this product instead of* [corresponding brand product]." The values of X were calculated for each product and each store, using either the previous year's sales data (Jan-Dec 2011) or the first three months of the current year (Jan-March 2012). To obtain quasi-exogenous variation in the value of X displayed, holding constant the product and the store, we alternated which method of calculation was used in each store's labels, each week. An example of one such label is below in Figure 4.

Figure 5. Label with Generic purchase rate of Peers

Test 3b. An alternate way to frame the information displayed in Test 3a is to report the share of customers who buy the brand product, e.g. *"Y% of customers in this store choose* [corresponding brand product]" *instead of this product."* If the mere act of bringing attention to the purchase of a specific product leads consumers to buy it, we would expect Test 3b to be less effective than Test 3a in raising the generic share of purchases.

5. Empirical Specifications and Results

5.1 Store-level analysis

The outcomes of interest are: total quantity sold $Q=(q_{gen} + q_{brand})$ and generic share $gs=(q_{gen} /(q_{gen} + q_{brand})$. With treated and untreated products, stores, and weeks, we used both differences-in-differences and a triple-difference approach to estimate the effect of each type of label on each of these two outcomes. Using a differences-in-differences (DD) specification, we examine the effect of the treatment on the treated products by comparing the change in the sales of treated OTC categories from the pre-treatment to treatment period in the treatment store to that in the control stores. We use the following difference-in-differences specification, first on only those products that received a label in the treatment store (treated products), and then on untreated products:

$$(1) \quad Y_{ist} = \beta_1 T_1 + \beta_2 T_2 + \beta_3 T_3 + \beta_4 t_t + \delta_i + \delta_s + \varepsilon_{ist}$$

where Y_{ist}, which is equal to Q_{ist} or g_{ist} in separate regressions, denotes the product i sold in store s in time t, δ_s denotes store fixed effects to control for store-specific constant factors, δ_i denotes product fixed effects, t_t is a "treatment time" dummy that is equal to one during the treatment month and equal to zero during the pre-treatment month, and T_1 T_2, and T_3 are interactions between t_{it} and the store indicators for stores treated with labels for Test 1, Test 2, and Test 3, respectively. The coefficients on T_s can be interpreted as average treatment store specific changes between the pre-treatment month and the treatment month, whereas t_{it} captures a time trend common to the control and treatment stores.

In a second specification, we include $p_{ist,}$ the price difference between the brand and generic versions of product *i*, expressed as a percentage of the brand price. Since the seasonality of different product categories can lead them to have varying trends from one month to the next, we also include separate interactions of t_t (treat time dummy) with dummies for each product category (allergy, pain relief, digestive, other acute ailments), and we allow each store to have not only a different baseline generic share (or total quantity) overall, but different baseline values for all product categories.[23] The estimated treatment effects remain essentially identical.

First, we estimate these models with Treatments 2a and 2b, and 3a and 3b, grouped together. Second, we estimate specifications in which we allow each treatment to have a different effect on each store, to compare the effects of varied ways of framing the relevant information (e.g. Test 3a vs. 3b, Test 2a vs. 2b).

Third, we interact treatment indicators T_1 T_2, and T_3 with product-store characteristics that were used to create the labels: the relative savings offered by the generic relative to the brand product, and the share of the store's customers that buy the generic version of each product. Since these are store-product specific variables, we include store X product fixed effects in these tests of heterogenous treatment effects.

In all regressions with generic share as the dependent variable, we weight each product's observations by its average quantity sold per store (across all stores) in the pre-treatment period.

[23] When comparing the point estimates with and without this slew of controls, we find no significant differences.

Results

Table 5 reports the results from our specifications (1)-(2), on our two outcomes of interest, for treated and untreated products. Note that in this table, we are constraining the effects of Test 2 and Test 3 to be equal across stores, despite the fact that we tested varied ways of framing the information contained in these tests, across stores.

Supporting the assumptions behind our experimental design, we find no significant differences between treatment and control stores in the pre-post changes in either outcome for the products which did not receive labels.[24] As the top-left portion of the table shows, both Test 2 and Test 3 had positive and significant effects on generic share. The bottom-left portion shows significant effects of all treatments on log quantity, suggesting that our treatments increased the quantity of sales in treated categories by 8-16%. The right-hand side of the table shows that the changes in untreated products' outcomes did not differ significantly between treated vs. untreated stores.

In Table 6, we disaggregate the results for Tests 2 and 3 across individual stores. The disaggregation of Test 2 shows that the two methods of framing the price difference ("save x%" relative to the brand or "pay additional y%" relative to the generic) had noisy but similar effects on generic purchase share. In contrast, varying the framing of Test 3 did not affect the impact of posted peer purchases on generic share. Instead, what drove variation in the magnitude of Test 3's effect might be the magnitude of the generic purchase rate that we posted on the labels, which was about 10 percentage points lower in Store 2 than in Stores 1 and 3.

Here, we discuss the specific results of each test in greater detail.

Test 1. On average, the Test 1 labels had no significant effect. However, a disaggregated analysis of this store's sales, by type of label (Figure 7) shows that products receiving the

[24] In theory, we could have found spillover effects from the labels, to other generic products in the same store.

"Bioequivalent" label (see Figure 3) had an average generic share increase of 7.5 percentage points, while the softer statements applied to other treated products had negative effects on generic shares. However, given the small sample size, the difference between these effects is not statistically significant.

Test 2. On average, posting the price difference led to a statistically significant 8.5 percentage point increase in generic share. Table 6 disaggregates these results by the two framing manipulations, demonstrating that neither store's results are statistically significant on their own. Table 7, however, shows that the treatment effect of Test 2 was significantly greater on products with a larger percentage price difference between the brand and generic versions. Test 2 labels appear to have increased the price sensitivity of consumers by about two-thirds, on average.

Test 3. These labels, which display the share of customers who purchase the generic version of a product, had a statistically significant positive effect of 8.5 percentage points on average, on the generic share of treated products. Disaggregating the results by store, we find that the point estimates were most significant and large (10 percentage points) in the two stores where generic shares were already greater than 50% at the outset, and therefore, in the labels we posted.

The interaction effects in Table 8 (boxed) do not show a general correlation between higher pre-existing store-level generic shares and the treatment effect in T_3 stores. However, Column 3 shows that when the posted generic share (from 2011) differed from the existing generic share, changes in generic purchasing were positively correlated with the magnitude of the posted share. [25]

Lastly, Table 7 shows that Test 3 labels had larger treatment effects on products with a larger percentage price difference. This suggests that consumers are more inclined to move their behavior in the direction of their peers when they stand to save money at the same time.

[25] As described in section 4.1, we created exogenous variation in the posted generic share by alternating, weekly, whether the posted share was calculated with 2011 data or with 2012q1 sales data.

Figure 8 graphically shows the correlation between the difference in posted share and the week-by-week differences in treatment effects for the same treatment stores and products. Each point in the scatterplot corresponds to one product (active ingredient) in one of the three stores that were treated with Test 3 labels.

4.2.2 Customer-level analysis

An important concern arises when the analysis is restricted to store-level totals: We cannot be sure whether existing buyers of OTC products within a store are shifting their purchases towards generics, or if the labels attract new customers, who may already be buyers of generics at other retailers. We use customer-level data to distinguish between these groups.

For this analysis, we focus on the purchase observations by individuals whose Club ID cards show a prior purchase of *any* OTC drug in our sample within the past six months (N=XX). We then group these individuals in two types: those who purchased at a store-label product (N=X) at least once, and those who purchased only brand products over the past six months (N=Y). We call these two types "Experienced with generic OTC" and "Inexperienced with generic OTC", respectively. Using a separate linear probability model for customers of each type, we estimate the effect of the experimental labels on the probability that a customer purchases a store-label product on their current visit, conditional on their past purchase history. We also estimate this model on two smaller, but more precisely defined subsamples: those who have been observed to purchase the generic version of the *same* product they are currently buying, and those who have only purchased the brand version of this product in the past six months.[26] We call these types "Experienced with this generic product" and "Inexperienced with this generic product."

$$(2) \quad Y_{ist} = XB + \beta_1 T_1 + \beta_2 T_2 + \beta_3 T_3 + \beta_4 t_t + \delta_{it} + \delta_s + \varepsilon_{ist}$$

[26] Customers who have purchased both generic and brand versions of this product in the past six months (N=XXX) are grouped with those who have only purchased the generic.

Controls in X include the number of OTC purchases in the past 6 months (total, total generic, same product, and same product generic), the price difference between brand and generic, in percentage terms, and dummies for whether brand, generic, or both versions of the product were on sale during a given week.

As in the store-level analysis, t_t identifies the four-week treatment period; T_1, T_2, and T_3 are the treatment period interactions with the stores used in tests 1-3; δ_{it} are product by month fixed effects, and δ_s are store fixed effects.

Results

The results from these customer type-specific linear probability models are shown in Table 7. First, a comparison of the mean outcome (the probability of purchasing a generic drug on a given purchase occasion) reveals the stark contrast between these types of customers. Only 19% of customers who have bought only brand OTC drugs in the past 6 months will buy a generic on their next visit, whereas 60% of customers who have some experience with OTC generics will purchase one. The difference is even larger when we compare customers who have been observed to buy either the brand or generic version of the same drug they are currently purchasing: only 9% of the brand-exclusive buyers will buy the generic next time, versus 85% of those who have bought the generic before.

We now turn our attention to the effects of our three labeling interventions on the generic purchase probability of each customer type. We discuss each of the three tests sequentially.

Consistent with the results of the store-level analysis, there is no evidence that customers (of any type) shopping at store 1 during the treatment period were more likely to purchase the generic version than they would be otherwise. The point estimates are consistently negative, but not significant.

For test 2 (price comparison), we find no evidence of an effect on customers who previously purchased brands, either of any OTC product (Column 1) or of the current product (Column 3). For customers who have prior experience with generic OTC drugs, however, Test 2 raised the probability of buying a generic by 6.5 percentage points (an 11% increase). For customers with prior experience buying this generic product, the point estimate is imprecise but suggests a similar magnitude increase (5.7 percentage points).

For test 3 (peer purchase information), we find a positive effect on the generic purchase probability of customers who were previously inexperienced with generic OTC drugs. The 4.4 percentage point increase in Column 1 indicates a 23% increase relative to the baseline probability of 19%, and the 4.8 percentage point increase in Column 3 indicates a slightly larger than 50% increase. The corresponding effects on customers with prior experience with the generic are smaller and statistically insignificant.

In sum the customer-level analysis rejects the possibility that our treatment effects are purely driven by new customers who already buy generic drugs elsewhere, or by existing buyers of generic drugs simply buying more of them. Among existing customers who were previously loyal to brands, we find that information on peer purchases of OTC drugs is the only treatment that increases the probability of buying a generic for the first time. This supports our intuition that these customers are wary of the overall quality of the generic products, but respond to the fact that many other customers find them acceptable.

5.3 Robustness

To verify that our results are not driven by seasonal shifts specific to the stores in our area, we run a placebo test. We use the sales data from the same months in 2010 and 2011 to estimate the placebo difference-in-differences effect on generic share in the stores we treated in 2012. Results are shown in Tables 9 and 10, revealing no discernible placebo effect.

6. Conclusion

Unlike prescription drugs, OTC drugs are purchased by consumers with direct access to price information. Nevertheless, a low responsiveness to price differences between near substitutes may result from biased beliefs about the differences between brand and generic drugs. Surveys have found that consumers prefer to receive a brand drug as a prescription, and we find that these preferences extend to the OTC market, despite high price premiums for the brand.

We implemented a labeling experiment at six locations of a national retailer to test three hypotheses for consumer aversion to generic OTC drugs: (1) lack of information regarding their similarity to the brand, (2) inattention to the price difference, and (3) biased beliefs regarding generic quality, that can be updated with information on their peers' purchases.

We find no evidence for the first hypothesis and some evidence for the second hypothesis. In the stores in which we tested hypothesis (2) by displaying price differences in percentage terms, we saw a 5.6 percentage point increase in generic purchase share.

We find the strongest evidence for hypothesis (3). Given that OTC drugs are consumed privately, individuals cannot easily observe the drugs that their peers purchase. Thus, it is perhaps not surprising that this information impacted behavior. In the two stores in which generic purchase shares of treated products averaged above 50%, labels that provided these shares, by product, led to a further increase of 10 percentage points. We conclude that some consumers do not buy generic products because they are unaware that these products are commonly purchased by shoppers like them, who apparently find them safe and effective.

Figures and Tables

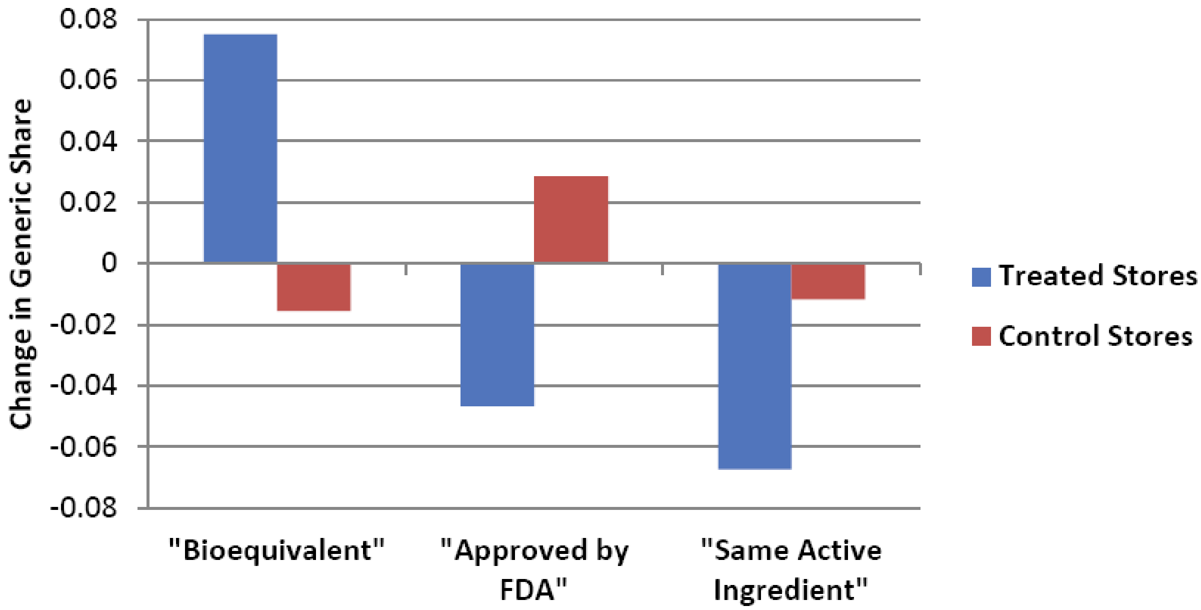

Figure 7. Raw change in generic share in Test 1 store, by label type.
3-4 products received each type of label for each of their packages.

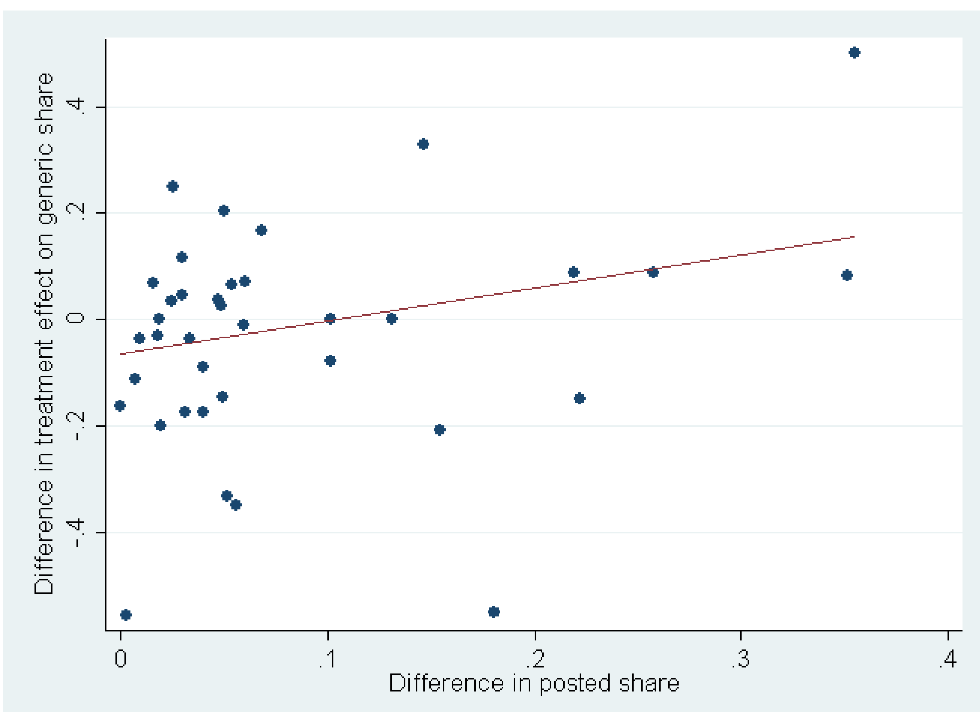

Figure 8. Effect of posted share of peers on current generic shares

Table 5: Average Treatment Effects

Panel A: **Y = Generic share by product, store, week**

	Treated products		*Untreated products*	
	(1)	(2)	(3)	(4)
T1: Comparability information	-0.00097	-0.00090	0.014	0.016
	(0.023)	(0.023)	(0.031)	(0.033)
T2: Price comparison	0.058*	0.059*	0.00085	0.0022
	(0.028)	(0.029)	(0.027)	(0.028)
T3: Peer purchase	0.088***	0.088***	-0.0015	0.00026
	(0.027)	(0.028)	(0.020)	(0.020)
Price difference, pct.		0.30***		0.51**
		(0.079)		(0.20)
N	1408	1408	1545	1545
R-sq	0.378	0.441	0.340	0.481

Panel B: **Y = Log quantity sold by product, store, week**

	Treated products		*Untreated products*	
	(5)	(6)	(7)	(8)
T1: Comparability information	0.17**	0.17**	-0.16	-0.16
	(0.077)	(0.080)	(0.10)	(0.10)
T2: Price comparison	0.088*	0.087*	0.094	0.094
	(0.047)	(0.048)	(0.093)	(0.094)
T3: Peer purchase	0.16**	0.16**	0.047	0.047
	(0.065)	(0.067)	(0.052)	(0.054)
Price difference, pct.		0.21		0.51***
		(0.18)		(0.13)
Store f.e. , product f.e.	X	X	X	X
Category x Treat Time		X		X
Category x Store f.e.		X		X
N	1408	1408	1545	1545
R-sq	0.764	0.790	0.746	0.780

Standard errors in parentheses; clustered by product.
Treatment T1 was only conducted in one store. T2 was conducted in two stores with different variations on the framing of the price difference. T3 was conducted in three stores with two framing variations.

Table 6: Average Treatment Effects, by disaggregated treatments across stores.

	Y = Generic share		Y = Quantity sold	
	(1)	(2)	(3)	(4)
	Treated Products	*Untreated Products*	*Treated Products*	*Untreated Products*
T1: Comparability information	-0.00090	0.016	0.17**	-0.16
	(0.023)	(0.033)	(0.080)	(0.10)
T2: Price comparison				
Framing: Savings	0.069	-0.0036	0.12	0.021
	(0.046)	(0.050)	(0.085)	(0.13)
Framing: Pay additional	0.050	0.0083	0.045	0.18*
	(0.046)	(0.022)	(0.056)	(0.098)
T3: Peer purchase				
Framing: "Generic" (Store 1)	0.098***	0.011	0.12	0.043
	(0.032)	(0.035)	(0.076)	(0.086)
Framing: "Generic" (Store 2)*	0.063	-0.0073	0.20*	-0.032
	(0.048)	(0.043)	(0.11)	(0.087)
Framing: "Brand"	0.10**	-0.0033	0.16*	0.14
	(0.037)	(0.024)	(0.089)	(0.13)
Price difference, pct.	0.30***	0.51**	0.21	0.52***
	(0.077)	(0.20)	(0.18)	(0.14)
N	1408	1545	1408	1545
R-sq	0.441	0.481	0.790	0.780

Standard errors in parentheses; clustered by product.

* In Store 2 of Test 3, Framing "Generic", posted generic shares averaged 43%, significantly lower than the average posted shares in the other Test 3 stores (53% and 54%).

Table 7: Customer-Level Regressions

	Y = 1 if generic version is purchased			
	Customers with previous OTC purchases		Customers with previous purchases of this product	
	All brand	Some generic	All brand	Some generic
	(1)	(2)	(3)	(4)
T1: Comparability information	-0.039	-0.047	-0.017	-0.055
	(0.040)	(0.032)	(0.022)	(0.046)
T2: Price comparison	-0.016	0.065**	-0.020	0.057
	(0.027)	(0.027)	(0.026)	(0.041)
T3: Peer purchase	0.044*	0.040	0.048*	0.010
	(0.024)	(0.035)	(0.023)	(0.030)
Price difference, pct.	0.16	0.20***	0.035	0.26***
	(0.12)	(0.059)	(0.050)	(0.079)
Generic on sale (0/1)	0.096***	0.073***	0.016	0.022**
	(0.031)	(0.0080)	(0.028)	(0.0094)
Brand on sale (0/1)	-0.055*	-0.081**	-0.034	-0.066**
	(0.026)	(0.031)	(0.020)	(0.025)
Both on sale (0/1)	-0.030	-0.0070	0.036	0.038
	(0.033)	(0.014)	(0.029)	(0.023)
Mean generic	0.19	0.60	0.09	0.85
N	10786	14011	6474	5868
R-sq	0.134	0.225	0.087	0.157

Standard errors in parentheses

Controls: # of OTC purchases in past 6 months (total, generic), store, product X month, sales promotion dummies

Table 7: Heterogeneity of treatment by price difference between brand and generic

Y = Generic share by product, store, week

	Treated products		*Untreated products*	
	(1)	(2)	(3)	(4)
Interactions of treatment dummies				
T1 (Comparability) x Price difference, pct.	-0.25	-0.26	-0.13	-0.13
	(0.26)	(0.26)	(0.19)	(0.19)
T2 (Price comparison) x Price difference, pct.	0.39*		0.00032	
	(0.19)		(0.13)	
T2a ("Savings" frame) x Price difference, pct.		0.65		-0.032
		(0.47)		(0.32)
T2b ("Pay more" frame) x Price difference, pct.		-0.045		0.035
		(0.30)		(0.18)
T3 (Peer purchases) x Price difference, pct.	0.37**	0.36**	0.061	0.066
	(0.14)	(0.15)	(0.10)	(0.11)
T1	0.040	0.041	0.053	0.053
	(0.059)	(0.059)	(0.054)	(0.054)
T2	-0.0071		0.0096	
	(0.044)		(0.037)	
T3	0.030		-0.0100	
	(0.037)		(0.032)	
Price difference, pct.	0.59***	0.60***	0.48**	0.48**
	(0.16)	(0.16)	(0.20)	(0.20)
Store x product fixed effects	X	X	X	X
Store-specific treatment dummies (T2a, T2b, T3a1, T3a2, T3b)		X		X
N	1408	1408	1545	1545
R-sq	0.558	0.559	0.555	0.555

Standard errors in parentheses, clustered by product

Table 8: Heterogeneity of treatment by existing and posted generic shares

Y = Generic share by product, store, week

	Treated products			*Untreated products*		
	(1)	(2)	(3)	(4)	(5)	(6)
T1 (Comparability) x GS 2012	-0.053	-0.054	-0.054	-0.13	-0.13	-0.13
	(0.24)	(0.24)	(0.24)	(0.23)	(0.23)	(0.23)
T2 (Price comparison) x GS 2012	-0.43*	-0.48**	-0.48**	0.051	0.071	0.071
	(0.21)	(0.22)	(0.22)	(0.25)	(0.25)	(0.25)
T3 (Peer purchases) x GS 2012	-0.029	-0.080	-0.37**	0.16***	0.17**	0.16
	(0.12)	(0.13)	(0.14)	(0.051)	(0.060)	(0.20)
T3 (Peer purchases) x Posted GS			0.31**			0.015
			(0.14)			(0.19)
T1	0.019	0.020	0.020	0.076	0.076	0.076
	(0.11)	(0.11)	(0.11)	(0.094)	(0.094)	(0.094)
T2	0.24**			-0.0091		
	(0.10)			(0.086)		
T3	0.10			-0.065*		
	(0.085)			(0.033)		
T3a ("Generic" frame) Store 1		0.14	0.13		-0.055	-0.056
		(0.098)	(0.093)		(0.048)	(0.046)
T3a ("Generic" frame) Store 2		0.096	0.082		-0.068*	-0.069
		(0.089)	(0.086)		(0.037)	(0.042)
T3b ("Brand" frame)		0.14	0.13		-0.082**	-0.083**
		(0.089)	(0.083)		(0.032)	(0.031)
Price difference, pct.	0.64***	0.65***	0.66***	0.47**	0.47**	0.47**
	(0.16)	(0.16)	(0.16)	(0.20)	(0.20)	(0.20)
Store x product fixed effects	X	X	X	X	X	X
Store-specific treatment dummies (T2a, T2b, T3a1, T3a2, T3b)		X	X		X	X
N	1374	1374	1374	1458	1458	1458
R-sq	0.561	0.562	0.563	0.563	0.563	0.563

Standard errors in parentheses; clustered by product

GS 2012 is the generic share of each product by package sales, calculated over the January-March 2012 period. Posted generic share is the information that was given in Test 3. In half of the weeks, this was determined by the store's 2011 generic share of the product. In the other half of the weeks, it was set to the store's 2011 generic share of the product.

Table 9: Placebo Treatment Effects, estimated with 2011 sales data, over the same months during which the treatments were implemented in 2012.

	Y = Generic share		Y = Quantity sold	
	(1)	(2)	(3)	(4)
	Treated Products	*Untreated Products*	*Treated Products*	*Untreated Products*
T1: Comparability information	-0.022	-0.0041	0.0089	-0.13*
	(0.038)	(0.043)	(0.054)	(0.072)
T2: Price salience				
Framing: Savings	-0.074*	-0.046	0.050	-0.042
	(0.038)	(0.027)	(0.086)	(0.12)
Framing: Pay additional	-0.037	-0.011	-0.17*	-0.058
	(0.048)	(0.048)	(0.081)	(0.12)
T3: Peer purchase				
Framing: "Generic" (Store 1)	0.053	0.0065	-0.0069	-0.18**
	(0.039)	(0.031)	(0.087)	(0.075)
Framing: "Generic" (Store 2)*	-0.00017	0.013	-0.060	-0.14
	(0.028)	(0.029)	(0.10)	(0.095)
Framing: "Brand"	-0.0074	0.029	-0.091	-0.13
	(0.071)	(0.038)	(0.094)	(0.099)
	-0.11	0.36***	-0.18	-0.14
Price difference, pct.	(0.18)	(0.12)	(0.29)	(0.33)
N	1356	1544	1356	1544
R-sq	0.474	0.436	0.745	0.763

Standard errors in parentheses; clustered by product.

Table 10: Placebo Treatment Effects, estimated with 2010 sales data, over the same months during which treatments were implemented in 2012.

	Y = Generic share		**Y = Quantity sold**	
	(1)	(2)	(3)	(4)
	Treated Products	*Untreated Products*	*Treated Products*	*Untreated Products*
T1: Comparability information	-0.025	0.068	-0.019	-0.097
	(0.038)	(0.045)	(0.14)	(0.091)
T2: Price salience				
Framing: Savings	0.0048	0.046	0.18*	0.067
	(0.034)	(0.061)	(0.095)	(0.098)
Framing: Pay additional	0.0077	0.025	-0.0085	0.069
	(0.035)	(0.051)	(0.13)	(0.086)
T3: Peer purchase				
Framing: "Generic" (Store 1)	0.013	-0.027	-0.0020	-0.13
	(0.030)	(0.044)	(0.086)	(0.12)
Framing: "Generic" (Store 2)*	0.0080	-0.015	0.039	-0.10
	(0.052)	(0.034)	(0.10)	(0.10)
Framing: "Brand"	-0.0050	-0.051	0.093	-0.022
	(0.061)	(0.055)	(0.12)	(0.097)
	0.098	0.42**	0.22	-0.10
Price difference, pct.	(0.092)	(0.17)	(0.29)	(0.32)
N	1230	1425	1230	1425
R-sq	0.363	0.454	0.777	0.780

Standard errors in parentheses; clustered by product.

References

Ackerberg, Daniel A., 2001. “Empirically Distinguishing Informative and Prestige Effects of Advertising,” The RAND Journal of Economics, vol. 32: 316- 333.

Ackerberg, Daniel A., 2003. “Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Empirical Examination,” International Economic Review, vol. 44: 1007-1040.

Anderson, Michael. and Jeremy. Magruder, 2012 "“Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database,” Economic Journal. 122(563): 957-989.

Akerlof, George A, 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," The Quarterly Journal of Economics, MIT Press, vol. 84(3), pages 488-500, August.

A.C. Nielsen. “The power of private label 2005: a review of growth trend around the world” 2005.

Banerjee, A. (1992). A simple model of herd behavior. Quarterly Journal of Economics, 107(3), 797-817.

Becker, G.S. (1991). A note on restaurant pricing and other examples of social influences on price. Journal of Political Economy, 99(5), 1109-1116.

Bikhchandani, S, Hirshleifer, D. and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy, 100(5), 992-1026.

Bronnenberg, Bart, Jean-Pierre Dubé, Matthew Gentzkow and Jesse Shapiro, “Do pharmacists buy Bayer? Sophisticated shoppers and the brand premium” University of Chicago mimeo.

Bursztyn, Leonardo, Ederer, Florian P., Ferman, Bruno and Yuchtman, Noam (2013) “Understanding Peer Effects in Financial Decisions: Evidence from a Field Experiment.” Available at SSRN: http://ssrn.com/abstract=2101391.

Cai, Hongbin, Yuyu Chen, and Hanming Fang. 2009. "Observational Learning: Evidence from a Randomized Natural Field Experiment." American Economic Review, 99(3): 864-82.

Chen, S S, Monroe, K B and Lou, Y (1998). The effects of framing price promotion messages on consumers' perceptions and purchase intentions. Journal of Retailing, 74(3), 353-372.

Chetty, Raj, Looney, Adam, and Kory Kroft, 2009. “Salience and Taxation: Theory and Evidence” American Economic Review 99(4): 1145-1177.

Eliashberg, Jehoshua, and Steven M. Shugan, 1997. “Film Critics: Influencers or Predictors?,” Journal of Marketing, vol. 61: 68-78.

Gendall, P., Hock, J., Pope, T., & Young, K. (2006).Message framing effects on price discounting.Journal of Product & Brand Management,15(7), 458-465.

Hibbard, Judith and Jessica Greene, 2013, “What The Evidence Shows About Patient Activation: Better Health Outcomes And Care Experiences; Fewer Data On Costs” Health Affairs, vol. 32:207-214

Hilger, James, Greg Rafert, Sofia B. Villas-Boas, 2011“Expert Opinion and the Demand for Experience Goods: An Experimental Approach,” Review of Economics and Statistics, 93 (4): 1289-1296

Hossain, T. and Morgan, J., 2006: “…Plus Shipping and Handling : Revenue (Non) Equivalence in Field Experiments on eBay,” B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy, 6, 2, 1-27.

Jin, Ginger Z., and Phillip Leslie, 2003 “The Effects of Information on Product Quality: Evidence from Restaurant Hygiene Cards,” Quarterly Journal of Economics, vol. 118: 409-451.

Kesselheim AS, Misono AS, Lee JL, et al. Clinical Equivalence of Generic and Brand-Name Drugs Used in Cardiovascular Disease: A Systematic Review and Meta-analysis. JAMA. 2008;300(21):2514-2526. doi:10.1001/jama.2008.758.

Khan, Romana, Kanishka Misra, and Vishal Singh. "Ideology and Brand Consumption." Psychological science (2013).

Lacetera, Nicola, Devin Pope, and Justin Sydnor, "Heuristic Thinking and Limited Attention in the Car Market" American Economic Review (2012) 102(5), 2206-2236.

Lee, Y.H. and Malmendier, U., 2007: "The Bidder's Curse," NBER Working Paper 13699.

McFadden, D. "The Measurement of Urban Travel Demand." Journal of Public Economics 3(1974):303-28.

McFadden, D.L. and Train, K. (1996). Learning from self and others. Journal of Political Economy, 104(4), 684-703.

Montgomery, Cynthia A & Wernerfelt, Birger, 1992. "Risk Reduction and Umbrella Branding," The Journal of Business, University of Chicago Press, vol. 65(1), pages 31-50, January.

Moretti, E. (2011). Social learning and peer effects in consumption: evidence from movie sales. The Review of Economic Studies, 78(1), 356-393.

Nease, Robert, Sharon Frazee, Larry Zarin, and Steven B. Miller (2013), "Choice Architecture is a Better Strategy than Engaging Patients to Spur Behavior Change," Health Affairs, vol. 32: 242-249.

Nelson, Phillip, 1970. "Information and Consumer Behavior," Journal of Political Economy, University of Chicago Press, vol. 78(2), pages 311-29, March-Apr.

Reinstein, David A., and Christopher M. Snyder, 2005. "The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics," Journal of Industrial Economics, vol. 53: 27-51.

Roehr B. Potential benefits of "well engaged patients" are akin to those from a blockbuster drug, say experts. BMJ 2013;346:f886

Salganik M.J., Dodds, P.S. and Watts, D.J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. Science, 311, 854-856.

Shrank, William H., et al. "Physician perceptions about generic drugs." The Annals of pharmacotherapy 45.1 (2011): 31-38.

Shrank, William H., et al. "Patients' perceptions of generic medications." Health Affairs 28.2 (2009): 546-556.

Train K.E. Discrete Choice Methods with Simulation. Cambridge: University Press, 2002.

Wiggins, Stephen. and W.J. Lane, 1983. "Quality Uncertainty, Search, and Advertising," American Economic Review, vol. 73: 881-894.

Wolinsky, Asher, 1995 "Competition in Markets for Credence Goods," Journal of Institutional and Theoretical Economics, vol. 151, 117-131.

Zhang, Juanjuan, 2010 "The Sound of Silence: Obsevational Learning in the U.S. Kidney Market," Marketing Science, vol 29, 315-335.