In-Store Coupons: A Large-Scale Field Experiment

Sebastian Gabel[*], Duncan Simester[†], and Artem Timoshenko[‡]

[*]Rotterdam School of Management, Erasmus University
[†]MIT Sloan School of Management
[‡]Kellogg School of Management, Northwestern University

Abstract

New technologies are disrupting the distribution of CPG coupons. Many customers now obtain digital coupons while in-store, using retail kiosks or mobile apps. These technologies are designed to prompt unplanned (impulse) decisions and enable physical retailers to distribute promotions at the very bottom of the purchasing funnel.

We investigate the effectiveness of in-store coupons using a large-scale field experiment involving 101 CPG brands. On average, in-store coupons triple sales of the promoted brands. This is not merely cannibalization from future sales or stealing share from other brands in the category. Instead, in-store coupons drive substantial category expansion.

Category expansion benefits both the focal brand and other brands in the same category. We show that in-store coupons *increase* sales of competing brands, particularly among customers who had previously purchased the competing brands. This is the first paper to document positive competitive spillovers for price promotions. Notably, when customers respond to an in-store coupon by purchasing a competing brand, they choose to forgo the coupon discount on the promoted products.

Keywords: in-store marketing, coupons, field experiments, causal effects, incrementality

1. Introduction

New technologies are transforming the way retailers promote consumer packaged goods (CPG). Retailers such as CVS, Winn Dixie, Food Lion, and Stop & Shop now distribute coupons via in-store kiosks (Hamstra 2023). Kroger's, Target, Publix, and Best Buy have introduced mobile apps that consumers can use to browse coupons in-store (Berthiaume 2022). Other emerging technologies include digital in-store and intelligent shopping carts (Grewal et al. 2023). These technologies are disrupting the $430 billion CPG coupon industry by delivering promotions to customers at the very bottom of the purchasing funnel, when they are already in-store and about to make a purchase. Amazon's sponsored search and on-site advertising share the same characteristic. The effectiveness of Amazon's lower funnel interventions has helped it become the world's third-largest media company (Statista 2022). We might expect that promotions delivered at the bottom of the purchasing funnel in physical stores would also be very effective, but little is known about how customers respond to this new generation of promotions.

We investigate the effectiveness of in-store coupons using a large-scale field experiment at a supermarket chain in Germany. The retailer delivers coupons via kiosks located near store entrances. The experiment was conducted in a business-as-usual environment, by distributing random coupons using the retailer's regular couponing system. The random coupon assignments exogenously varied the promoted brands across different customers. We use the retailer's loyalty program to link coupon exposures to consumers' individual purchases and report the incremental (causal) effects across 101 CPG brands spanning 44 categories.

In-store coupons dramatically increase purchases of the promoted brand. On average, receiving an in-store coupon causes purchase incidence to increase from 0.71% to 2.11% on the current visit. Even though in-store coupons increase unit sales, we may be concerned that they do not increase revenue, because discounts offset the sales lift. Our findings assuage this concern. The sales increases in our study were more than enough to offset the discounts: the average revenue increased from €15.64 to €47.41 per mille.[1] We did not observe a statistically significant reduction in revenue for any of the 101 brands.

In-store coupons are a versatile tool that brands can use for different marketing goals. We show that they increase sales of the promoted brand by past customers of the promoted brand, by customers who had previously purchased other brands in the category, and by customers who have not recently purchased in the category. Some brands are focused on increasing topline sales. Others want to attract customers away from competitors or induce trial by customers who do not regularly purchase in the category. In-store coupons can accomplish each of these objectives.

Three mechanisms could explain the sales increase for the promoted brands: (a) forward buying, (b) brand switching, and (c) category expansion. It is possible that in-store coupons merely pull demand from the future, with customers taking advantage of the discounts by stockpiling for future demand. We do not find evidence of either a post-promotion dip or stockpiling. Extending the measurement window to include future visits reveals no evidence of a post-promotion dip in either units purchased or revenue. Moreover, the sales increase does not vary significantly according to whether or not the category is suitable for stockpiling.

An alternative explanation is that incremental sales of the promoted brand are driven by stealing market share from competitors. Brand switching is frequently cited in the managerial literature and academic research as a primary objective of CPG coupons and price promotions (Inmar 2023; Neslin 1990). Contrary to this expectation, we find that coupons for the promoted brand cause an *increase* in competitors' sales. This result is primarily driven by customers who had previously purchased the competing brands. For example, purchase incidence on the current trip increases by 0.28% for the competing brands, and

[1] Per mille is a standard industry measure describing the outcome per 1,000 impressions. In our analysis one impression is a single customer receiving a coupon for the promoted brand on a given day.

this treatment effect doubles when we focus on the competitors' past customers.

The positive spillovers across CPG brands from competitors' price promotions is a surprising finding. Three studies have documented positive spillovers in other markets: Sahni (2016) studies restaurant advertising, Sahni et al. (2017) study email promotions for event tickets, and Anderson and Simester (2013) study private label apparel advertising. Compared to the positive spillovers in advertising, positive competitive spillovers from price promotions are even more striking: By purchasing a competing brand instead of the promoted brand, customers choose to *forgo the coupon discount*.

We find strong and consistent evidence supporting a third mechanism for why sales of the promoted brand increased: category expansion. This result extends across each of our outcome measures (incremental purchase incidence, quantity and revenue), and across 41 of the 44 categories, including non-food categories such as laundry care, oral care and pet food. The category expansion effect is significantly stronger for customers who had previously purchased in the category, but is also significant for customers who had not recently purchased in the category.

Although we do not measure store traffic, category expansion and positive competitive spillovers can be explained by more customers visiting the aisle. Presumably, these customers intend to purchase the promoted brand when they get to the aisle. However, when they arrive, some of this incremental traffic decides to purchase a competing brand instead. This effect is sufficient to offset any brand switching from the competing brands to the promoted brand due to the coupon discount. By increasing traffic to the aisle, in-store coupons increase purchases of all brands.

Past studies on price promotions have not reported large category expansion effects (Gupta 1988; Bell et al. 1999). One explanation is that traditional price promotions are more effective at driving customers to stores than they are at prompting impulse purchases when customers are in a store. In contrast, in-store promotions are primarily designed to drive impulse (unplanned) purchases (Iyer et al. 2020). If in-store coupons lead to more impulse purchasing, we would expect larger treatment effects in categories that are more susceptible to impulse purchasing. To investigate this prediction, we measured the tendency to purchase a category on impulse using survey responses to Narasimhan et al. (1996) two-item impulse scale. The findings confirm that in-store coupons cause larger sales lifts in high-impulse categories (compared to low-impulse categories). This finding extends to the promoted brand, competing brands and the overall category.

Our study is conducted in a business-as-usual environment, and so the results are indicative of what brands can expect when adopting in-store coupons. In-store coupons can effectively drive demand for the promoted brand by expanding the entire category. Amazon, Home Depot, Walmart, and almost every other large retailer now provide brands with options for delivering promotions to customers when the customers are in the retailers' digital or physical stores (Gabel et al. 2024). Our findings help to justify these investments. While Amazon's success confirms that targeting customers with marketing actions at the bottom of the purchasing funnel is effective in a digital channel, we demonstrate that it is also strikingly effective in physical retail.

Outline of the Paper

The paper proceeds with a review of the related research in Section 2. We describe the data, experimental setting, and identification strategy in Section 3. We document outcomes for the promoted brand in Section 4. In Section 5, we investigate forward buying, brand switching, and category expansion, and in Section 6 compare outcomes in high and low impulse categories. The paper concludes in Section 7.

2. Related Literature

There is a rich literature studying how consumers respond to coupons and price promotions. We first review the literature on in-store coupons, and then discuss research investigating three mechanisms that can explain the increase in sales of the promoted brand.

2.1 In-Store Coupons and Impulse Purchasing

Previous investigations of in-store promotions include coupons distributed at the shelf, such as coupons dispensed from J-hooks located directly in front of the promoted product (Dhar and Hoch 1996) or peel-off coupons placed on individual packages (Raju et al. 1994). Dhar and Hoch (1996) conducted two field experiments at 86 Chicago supermarkets. They find that on-shelf coupons yield sales increases similar to price discounts, but because only 55% of purchasers redeemed coupons, the revenue and profit increases were much larger for the coupons. Raju et al. (1994) conducted a series of quasi-experiments at a UCLA student store in which they placed coupons on 14oz and 20oz soda-fountain drinks. They conclude that on-pack coupons lead to larger market share and category sales lifts than peel-off and in-pack coupons.

Two studies consider coupons delivered as customers enter a store. Heilman et al. (2002) distributed printed coupons to 105 customers intercepted at the entrance of two St. Louis grocery stores. In a similar study, Hui et al. (2013) distributed printed coupons to 90 customers entering a grocery store in Pittsburgh. Both studies document the impact of the coupons on "unplanned" purchases (in categories not on the shopping list). The interactions between researchers and customers introduce a potential limitation to both studies, by introducing a risk that the interactions themselves could impact customers' behavior in the store (Pollay 1968).[2]

A key motivation for retailers' in-store marketing activities, including in-store coupons, is to increase impulse purchasing (Abratt and Goodey 1990; Zhou and Wong 2004). Impulse purchases are unplanned (see the discussion in Shapiro 2001), and often include a "powerful and persistent urge to buy something immediately" (Rook 1987, p. 191). Bucklin and Lattin

[2] In addition, Chiou-Wei and Inman (2008) study a setting in which "consumers browse the Web at home for any suitable coupons available on the store coupon pages and print them out. In the case of forgetting and/or misplacing, they can also get access to the coupon pages by running their shopping cards through terminals at the store and printing out the coupons" (p. 297). The browsing and printing at home is very similar to clip-at-home coupons. An important distinction is that we study the impact of in-store coupons using data from a field experiment.

(1991) propose a two-state model of purchase incidence and brand choice in which consumers shop in two ways. Shoppers who have planned their shopping trips are unresponsive to in-store promotions. In contrast, shoppers who have not planned their trips make impulsive decisions in the store and respond more strongly to in-store promotions.

Our paper contributes to this literature by investigating the effectiveness of in-store coupons using a large-scale unobtrusive field experiment in a business-as-usual environment. The loyalty program data helps us link the coupon distribution with purchase outcomes and evaluate incremental sales for both the promoted brand and competing brands across customer segments. The findings demonstrate that in-store coupons are effective for driving sales of the promoted brands and can lead to category expansion sufficiently large to *increase* sales of competing (non-promoted) brands.

2.2 Forward Buying, Brand Switching and Category Expansion

We investigate three mechanisms for why sales increase for the promoted brand: forward buying, brand switching and category expansion. All three explanations have received attention in the literature.

In the case of forward buying, Neslin et al. (1985) were among the first to study how coupons affect purchase timing. They find that coffee and bathroom tissue coupons impact purchase quantities but have relatively little impact on purchase timing. They attribute this to the long expiration dates on the coupons. Subsequent studies of household-level purchasing have often reported a dip in purchasing following a price promotion (Gupta 1988; Chiang 1991; Grover and Srinivasan 1992; Bell et al. 1999). However, this does not always translate into measurable dips in weekly store data (Litvack et al. 1985; Moriarty 1985; Blattberg et al. 1995; Neslin and Schneider Stone 1996). Van Heerde et al. (2000) and Hendel and Nevo (2003) both investigate this apparent paradox and conclude that post-promotion dips can be detected in store-level data and are sometimes over-stated in household-level data.

A separate but closely related literature studies consumer stockpiling in response to price promotions. Hendel and Nevo (2004, 2006) study the implications of household stockpiling on demand models. They confirm that failing to account for household inventory can yield static demand estimates that overestimate price sensitivity. Hendel and Nevo (2013) argue that household stockpiling can facilitate price discrimination.

Several studies have investigated whether variation in how suitable items are for stockpiling explains variation in the response to price promotions. Narasimhan et al. (1996) propose a two-item "suitability for stockpiling" scale, which has been widely used for this purpose. The findings have been mixed, with several studies reporting a positive relationship between promotional response and suitability for stockpiling (Narasimhan et al. 1996; Bell et al. 1999; Ailawadi et al. 2006; Fok et al. 2006), while others report a negative or non-significant relationship (Nijs et al. 2001; Lim et al. 2005; Osuna et al. 2016). Reconciling these conflicting results is challenging because the studies use a wide variety of outcome measures and different identification strategies. We provide a more detailed discussion of these papers in Web Appendix A.

The second mechanism is brand switching, which has been identified as a primary cause of the sales response to coupons and price promotions (Neslin 1990; Chiang 1995; Bell et al. 1999). In-store coupons have been shown to be an effective mechanism for driving brand switching. In particular, the past studies of peel-off and J-hook coupons (Raju et al. 1994; Dhar and Hoch 1996) confirm that coupons delivered *at the shelf* can lead to large market share gains.

There is also strong evidence that *at-home* coupons (such as FSI and direct mail coupons) can increase market shares, although the magnitude of the effect is not always large. Neslin (1990) concludes that coupons for instant coffee have pronounced effects on market share, but for many brands, the incremental sales effect is not large enough to yield additional short-run profits. Bawa and Shoemaker (1987b) report that coupon redemption rates are higher among customers who have previously purchased the promoted brand, and that customers quickly return to their favored brand after their redemption purchase. Srinivasan et al. (1995) also study FSI coupons, and find evidence of brand-switching in just two of the six categories they study.

The third mechanism we study is category expansion. One of the first studies to investigate category expansion was Chiang (1995). He chose to study laundry detergent because overall household consumption is not expected to vary with price. His findings confirm that laundry detergent coupons do not contribute to category expansion. Ailawadi and Neslin (1998) study how price promotions change yogurt and ketchup consumption and conclude that yogurt's usage rate increases as household inventory expands, but ketchup usage is much less sensitive to the available inventory. Nijs et al. (2001) study the impact of consumer price promotions on CPG category sales and conclude that price promotions can cause category expansion in the first 10 weeks (see also Dekimpe et al. 1998; Mela et al. 1998).

Past studies of coupons distributed in-store do not reveal the extent to which they contribute to category expansion. The Heilman et al. (2002) and Hui et al. (2013) papers both focus on customer basket sizes rather than total category sales. In the Dhar and Hoch (1996) study, the effects on category expansion are confounded by concurrent promotional activity on competing brands during the study period. The Raju et al. (1994) also explicitly acknowledges that it does not investigate category expansion.

Our study systematically investigates how in-store coupons impact sales of the promoted brand, sales of competing brands, and overall category sales. The scale and scope of the study allow us to disentangle the roles of forward buying, brand switching, and category expansion. In the next section, we introduce our empirical setting, discuss the design and implementation of the field experiment, and describe our identification approach for estimating treatment effects.

3. Data Overview

The data for our study was provided by a large German grocery retailer. The stores use a traditional supermarket format, and during the experiment, the retailer operated over 500 stores. The product range included over 30,000 SKUs spanning food, beverages, and non-food products, such as laundry detergents and paper towels. The stores use a HiLo price format, and the average price point is similar to its moderately priced competitors. The retailer maintains a loyalty program to collect customer data and distribute in-store coupons. Our dataset contains the purchases of all customers who shopped in the retailer's 147 stores in one of Germany's federal states, including over 250,000 customers in the retailer's loyalty program.[3]

We extracted the brand definitions from the retailers' product descriptions. Examples include Ben's Original rice, Lindt chocolate, and Johnnie Walker whisky. We use the category definitions provided by the retailer. The definitions are relatively narrow and include, for example, frozen pizza, cat food, deodorant, and juice. On average, the categories include 103 SKUs, and the promoted brands in our experiment represent 3.8% of the revenue in their corresponding categories. In Appendix A, we provide a complete list of the 44 categories included in the study.

The retailer communicates with loyalty program members through self-service kiosks (see Figure 1). The kiosks allow customers to print in-store coupons and check their loyalty point balances. During the study, the kiosk system was the only touchpoint between the retailer and its loyalty program customers. On average, customers who participate in the loyalty program obtain coupons from an in-store kiosk on 40% of their shopping trips.

Customers typically interact with the kiosks and obtain in-store coupons when entering the store (see Figure 1). After customers present their loyalty cards, the kiosk prints seven coupons, each of which offers a discount on a single brand. When customers show their

[3]Loyalty program participants account for 21.1% of the retailer's total revenue. The penetration tends to be larger in suburban residential areas with lower household incomes (up to 43.7%).

Figure 1: In-Store Kiosk and Coupon Printout

The figure depicts the self-service kiosk (left) and the coupon printout (right). The printout contains seven coupons valid for the same day, regular and discounted prices, and the discount rate in percent.

loyalty card at the checkout, they automatically receive these discounts if they purchase any of the brands represented by the coupons. If customers use the kiosk system multiple times in one day, they receive the same coupons; the coupon assignments are determined the night before and are valid for a day.

The retailer allocates coupons to customers based upon their past spending. For example, the targeting system prioritizes coupons for the brands that a customer purchased before. We will label coupons that the system chooses to distribute to individual customers as *targeted coupons*. The system ensures customers' daily coupons do not include multiple targeted coupons for the same or competing brands (in the same category).

3.1 Experimental Design

We obtained data from an experiment implemented by the retailer as part of its standard practice. The experiment distributed random coupons to both provide training data and evaluate the performance of the couponing system. We use this experimental variation to

Figure 2: Experimental Design with Random Coupons

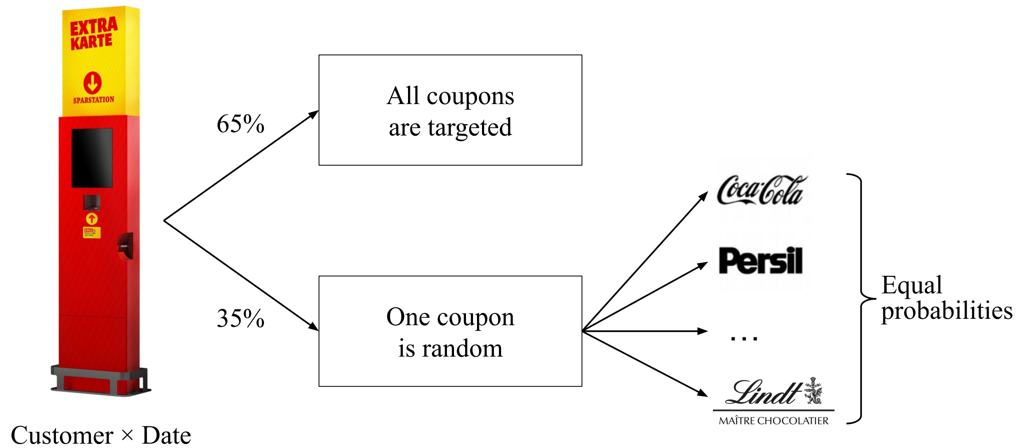

measure brand-level treatment effects. Next, we discuss the experimental data and then describe the identification approach that we use to calculate the treatment effects.

We illustrate the experimental design in Figure 2. Our unit of observation is a customer $\times$ date combination with at least one kiosk visit. By default, all seven coupons are targeted. In the experiment, in 35% of kiosk visits, the couponing system replaces one of the targeted coupons with a random coupon. The random coupons feature a randomly selected brand with a random discount depth ranging from 5% to 50%. The coupon assignments and randomization are implemented daily so that, on a given day, each customer has a 65% chance of receiving only targeted coupons, and a 35% chance of receiving a random coupon.

If a customer receives a random coupon, the random coupon replaces one of the seven targeted coupons. The targeted coupon that is replaced is randomly selected. The only exception is that if a printout includes a targeted coupon for the same brand as the random coupon, the random coupon replaces this targeted coupon. This ensures that the printout never includes two coupons for the same brand.

We emphasize that the experiment only includes customers who are part of this retailer's loyalty program, and we do not seek to generalize the results to customers who are not part of the loyalty program. The retailer price discriminates by only offering coupons to customers in its loyalty program, and customers who have not chosen to join the loyalty program are intentionally precluded from receiving coupons.

The retailer's couponing system records which coupons are distributed to customers, including the featured brands and discounts. The system also logs which coupons are targeted and which coupons are random (if any). However, the retailer does not record which targeted coupons are replaced by the random coupons, meaning there is no counterfactual policy logging in the system (Johnson et al. 2017).

3.2 Descriptive Statistics

Table 1 reports descriptive statistics. The dataset includes 1,503,000 kiosk visits across 158 different dates (November 2015 to May 2016). We focus on 88,840 random coupons distributed for 101 CPG brands, which had regular sales prior to the experiment, distributed at least 400 random coupons, and had no eligibility criteria.[4] We provide additional details about sample construction and randomization checks in Appendix B. The final sample spans 44 product categories, including food, drinks, and non-food CPG categories. The average discount per random coupon was 23.95%, and the redemption rate for random coupons was 1.68%.

Table 1: Experimental Data Descriptive Statistics

Number of brands	101
Number of categories	44
Number of random coupons	88,840
Number of shopping trips with kiosk usage	1,503,000
Average discount per random coupon	23.95%
Redemption rate for random coupons	1.68%
Average face value per redeemed random coupon	€0.86

Our analysis focuses on incremental changes in purchase incidence, quantity, and rev-

[4]Some brand managers restricted the distribution of coupons to a subset of customers (e.g. customers who had recently purchased in the category). To avoid selection effects and facilitate our comparison of treatment effects, we excluded brands that used such eligibility criteria.

enue (see Table 2). Brand managers who execute and evaluate campaigns typically focus on quantity and revenue, and the retailer reports these outcomes in its campaign dashboard. The revenue measure offers the advantage of accounting for price discounts. We consider incidence because we can measure treatment effects for this binary variable with high precision. Comparing multiple metrics also allows us to evaluate the robustness and consistency of our findings.

Table 2: Performance Metrics

Metrics	**Description**	**Type**
Purchase Incidence	Did a customer purchase or not?	$\{0, 1\}$
Quantity	How many units did a customer purchase?	$\mathbb{N}^0$
Revenue	How much did the customer spend?	$\mathbb{R}^0$

Our primary analysis focuses on the day that the customer used the kiosk to obtain in-store coupons (see Section 4), and we extend the measurement window to include purchases in subsequent weeks in Section 5. We next describe an identification approach that estimates the average treatment effect for each brand.

3.3 Identification Approach

Our goal is to measure the average treatment effect (ATE) of in-store coupons for the loyalty program customers on days that they use the kiosk (customer×date observations). The ATE captures the average impact of providing a coupon for the focal brand, compared to not providing a coupon for the focal brand:

$$\text{ATE}_b \equiv \mathbb{E}[Y(C = 1) - Y(C = 0)], \tag{1}$$

where $Y(C = 1)$ and $Y(C = 0)$ indicate potential outcomes with and without a coupon, and the expectation integrates over customer×date observations. Y refers to the performance metrics described in Table 2. The ATE does not confound in-store coupon effectiveness with the retailer's targeting algorithm.[5]

To estimate Equation (1), we can interpret the random sample of customers who received a random coupon for the focal brand as a "treatment" sample. Outcomes for these customers directly provide an estimate of $\mathbb{E}[Y(C=1)]$. The remaining customers are equivalent to the treatment sample and did not receive a *random* coupon for the focal brand. However, we cannot use their average outcomes to directly estimate $\mathbb{E}[Y(C=0)]$, because some of these customers received a *targeted* coupon for the focal brand.

To account for the focal brand's targeted coupons, our identification approach leverages random coupons for <u>non-focal</u> brands to infer the ATE for the <u>focal</u> brand. We summarize our identification approach here and provide additional details in Appendix C. We emphasize that empirically, the adjustment for targeted coupons has little impact on the results. If we forgo this adjustment and use the simple difference-in-means estimator, we obtain similar ATE estimates and our substantive conclusions are unchanged.

For each brand, we define three groups of customer×date observations:

1. $Focal_b$: Observations in which a customer received a random coupon for brand b. This is the treatment sample.
2. $NonFocal_b$: Observations in which a customer received a random coupon for any brand other than b (non-focal brands).
3. $NoRandom_b$: Observations in which a customer did not receive any random coupons (all coupons are targeted).

The experimental variation ensures that, for each brand, the three groups are equivalent and representative of the overall customer×date population. The $Focal_b$ observations always received a coupon for the focal brand, due to the random assignment. The $NonFocal_b$ and $NoRandom_b$ observations do not receive a random coupon for the focal brand in the experiment, but they could receive a targeted coupon for the focal brand.

[5] As an alternative, we could study the treatment effects of targeted in-store coupons (ATT), but these findings would be specific to the firm's current targeting algorithm.

We estimate the average treatment effect for brand b using the following equation:

$$\widehat{\text{ATE}}_b = \underbrace{\bar{Y}(\text{Focal}_b)}_{\mathbb{E}[Y(C=1)]} - \underbrace{\bar{Y}(\text{NoRandom}_b) - 1/p \cdot (\bar{Y}(\text{NoRandom}_b) - \bar{Y}(\text{NonFocal}_b))}_{\mathbb{E}[Y(C=0)]} \quad (2)$$

where $\bar{Y}(\cdot)$ indicates the average outcome in the corresponding group, and p is the probability that a random *non-focal* coupon replaces a targeted *focal* coupon. This probability is defined over observations in which the customer is assigned to receive both a random non-focal coupon and a targeted focal coupon. In our setting, there are seven targeted coupons that could be replaced by the non-focal coupon, so there is a one in seven chance that the targeted coupon for the focal brand is the one that is replaced ($p = 1/7$).

The treatment effect estimator in Equation (2) contains two components. The first component captures the average outcome in the treatment sample, $Focal_b$, to estimate $\mathbb{E}[Y(C = 1)]$. The second component estimates $\mathbb{E}[Y(C = 0)]$ using the outcomes in the $NoRandom_b$ and $NonFocal_b$ conditions. Intuitively, the experimental design ensures the equivalence of observations in these two conditions. However, in $NonFocal_b$, random non-focal coupons replace p proportion of the targeted coupons, so there are fewer targeted coupons in $NonFocal_b$ than in $NoRandom_b$. We know this is the only cause of different outcomes between the two conditions. We can thus recover the effect of the targeted coupons in $NoRandom_b$ by scaling the difference in outcomes between $NoRandom_b$ and $NonFocal_b$ by $1/p$. We obtain $\mathbb{E}[Y(C = 0)]$ by subtracting this effect from $NoRandom_b$. In Appendix C, we provide the identification assumptions and characterize the variance of our ATE estimator in Equation (2).

The estimator is unbiased under the assumption that treatment effects do not extend across brands. This is a common assumption in the causal inference literature (Angrist et al. 1996; Johnson et al. 2017). However, we will present evidence in Section 5 that treatment effects spill over to competing brands. In Appendix C, we recognize that competing brands' targeted coupons can co-occur with random focal coupons, which can introduce a bias. We empirically investigate this concern and demonstrate that it does not affect our conclusions. First, only a small fraction of customers received competing coupons, and the magnitudes of the spillover effects are small compared to the main effects on the focal brand. Second, the pattern of results and our conclusions are both unchanged when we focus on a subset of brands with zero active coupon campaigns by competing brands.

Our identification approach can also be applied to calculate Conditional Average Treatment Effects (CATEs) that measure treatment effects for subgroups of customers. In Sections 4 and 5, we use these CATEs to compare treatment effects across different customer segments.

4. Brand Treatment Effects

We report the distribution of ATEs for in-store coupons across 101 brands. We then investigate brand-level factors that explain variation in the ATEs and analyze treatment effect heterogeneity across customer segments.

4.1 Distribution of Brand ATEs

In Figure 3, we report the treatment effects for the 101 brands. Each data point represents a treatment effect estimate for a single brand on the day that the customer used the kiosk (on the current trip), and the error bars indicate 95% confidence intervals. The purchase incidence treatment effects (Panel A) measure the percentage increase in purchase incidence with versus without a coupon for a brand. The quantity and revenue treatment effects (Panels B and C) are measured as the average incremental units purchased and the euro revenue lift per 1,000 coupons (per mille). For example, an effect size of €10 indicates that the focal brand's average revenue per 1,000 coupon impressions was €10 higher than the average revenue per 1,000 kiosk visits without a coupon for the focal brand.

In-store coupons drive incremental purchase incidence, incremental purchase quantity, and incremental revenue for almost all brands. The purchase incidence treatment effects are positive for 99 of the 101 brands, and 56.4% of the effects are statistically significant

Figure 3: Distribution of Average Treatment Effects (ATEs)

(A) Purchase Incidence

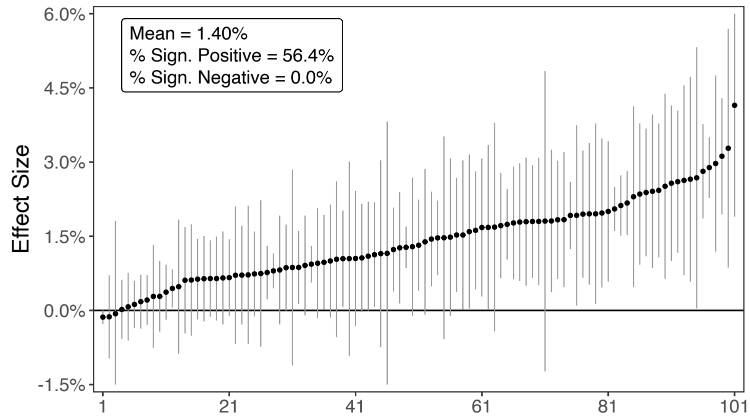

(B) Quantity

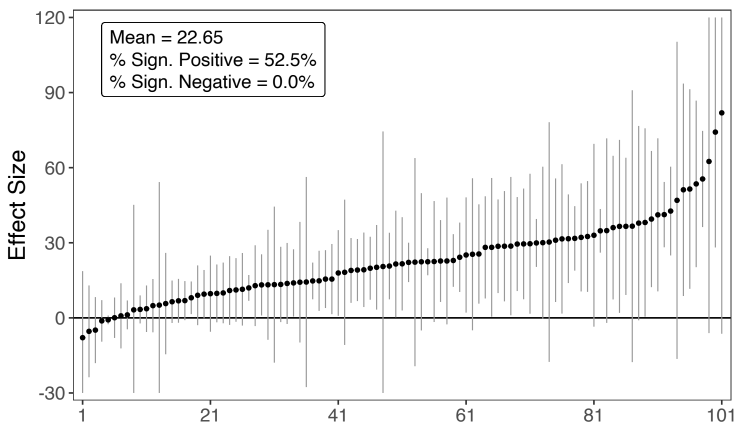

(C) Revenue

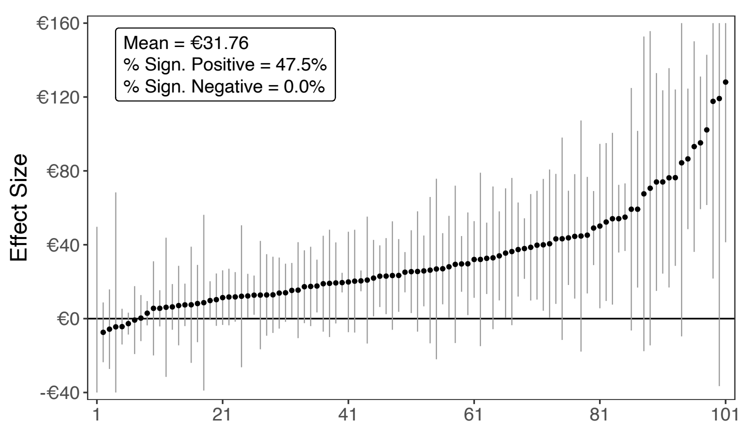

The figure summarizes the distribution of average treatment effects for the 101 brands. Panel A reports the incremental purchase incidence for each brand (in percentage points). Panel B reports the incremental quantity sold (in units per 1,000 coupons). Panel C reports the incremental revenue (in € per 1,000 coupons). Error bars indicate 95% confidence intervals.

($p < 0.05$). We observe a very similar pattern for the quantity and revenue treatment effects. Overall, the average ATE across the 101 brands is 1.40% incremental purchase incidence, 22.65 incremental units, and €31.76 additional revenue per 1,000 coupons. We do not find any negative treatment effects that are statistically significant for any of the 101 brands across the three outcomes. Notice that we measure treatment effects at the brand level, so all outcome measures account for substitution between packaging sizes.

We highlight that the revenue measure incorporates the coupon discount, and in-store coupons in our sample provide an average discount of 23.95%. A revenue breakeven analysis suggests the coupons need to increase volume by at least 1.31x to offset discounting. This is a substantial threshold, which might make some manufacturers hesitant to allocate budget to this channel. However, despite this breakeven threshold, in-store coupons yield revenue gains for essentially every brand in our sample. In-store coupons drive a 3.0x average increase in purchase incidence, a 3.1x increase in quantity sold, and a 3.0x increase in revenues.

We do not observe costs and so cannot calculate profit margins. However, we can calculate profit ATEs under different assumptions about the profit margin. In Web Appendix B we report plots of the profit ATE when using profit margins of 40%, 50% or 60%. The findings reveal that the in-store coupons yield an average profit increase of €4.80, €7.90, and €11.10 per mille when using these three profit margins (respectively).

To provide concrete examples that help interpret the scale and variation in the treatment effects, we report baseline outcomes and treatment effects for six globally recognized brands in Table 3. The brands are Ben's Original rice, Coke soft drinks, Johnnie Walker whisky, Lindt chocolate, Maggi spice mixes, and Schauma shampoo. In-store coupons are effective in driving sales for all six brands. For example, in-store coupons increase revenues for Coke by over 75%, and yield even more substantial gains for other brands, with sales lifts for Lindt chocolate exceeding 22x for quantity and 11x for revenue. Notably, revenue lifts are generally smaller than quantity lifts, reflecting the lower prices paid in the coupon

Table 3: Treatment effects for six global brands

	Incidence		**Quantity**		**Revenue**	
	No Coupon	ATE	No Coupon	ATE	No Coupon	ATE
Ben's Original rice	0.48%	1.23%	6.53	14.80	€16.25	€19.33
Coke soft drinks	5.93%	1.81%	84.00	46.95	€89.97	€67.57
Johnnie Walker whisky	0.16%	0.97%	1.82	9.48	€23.62	€84.36
Lindt chocolates	0.24%	1.92%	1.41	31.75	€7.62	€86.49
Maggi spice mix	0.87%	1.81%	12.52	51.47	€12.88	€43.18
Schauma shampoo	0.22%	0.91%	2.39	11.16	€5.57	€17.43

conditions due to the discount.

The Coke example is particularly interesting. Customers purchase Coke on 5.93% of shopping trips without Coke coupons, which is among the highest purchase rates for brands in our sample. This creates a risk that in-store coupons provide discounts to customers who would purchase anyway. We find that despite the high baseline purchase rate, the incremental demand for Coke is large enough to offset the discounting.

The brands in Table 3 represent different product categories, including chocolates, alcoholic and non-alcoholic beverages, and hair products. For completeness, we report treatment effects separately for each of the 44 categories in Appendix A. To systematically explore the variation in the treatment effects across the full sample of 101 brands, we next use the retailer's loyalty card data to construct four brand features and investigate whether they explain variation in the ATEs.

4.2 Correlation Between ATEs and Brand Features

In Table 4, we report the pairwise correlation between the three treatment effects (incidence, quantity, and revenue) and four brand features that have appeared prominently in the price promotions literature (Narasimhan et al. 1996; Bell et al. 1999; Ailawadi and Neslin 1998):

Brand Penetration: Percentage of customers that purchased at least one item from brand b in the data period.

Brand Loyalty: Spending on brand b as a % of total spending in the category. Calculated first at the customer level (among customers of the focal brand), and then averaged across customers.

Price Position: Average price paid for brand b, divided by average price paid in the category.

Interpurchase Time: Average (across customers) in the number of days between each customer's sequential purchases of brand b.

These features are all constructed using the retailer's transaction data in the three months prior to the in-store coupon experiments. Summary statistics and pairwise correlations between the four measures are reported in Web Appendix C.

Table 4: Pairwise Correlation Between Treatment Effects and Brand Features

	Incidence ATE	**Quantity ATE**	**Revenue ATE**
Brand Penetration	0.50 **	0.54 **	0.27 **
Brand Loyalty	-0.18 †	-0.24 *	0.08
Price Position	-0.08	-0.15	0.17 †
Interpurchase Time	0.06	-0.10	0.17 †

The table reports the Spearman pairwise correlations between each brand feature and purchase incidence, purchase quantity, and revenue treatment effects (on the current trip). The unit of analysis is a brand, and the sample size is 101. Standard errors are in parentheses. ** indicates significantly different from zero, $p < 0.01$. * indicates significantly different from zero, $p < 0.05$. † indicates significantly different from zero, $p < 0.10$.

Brand Penetration explains more variation in in-store coupon treatment effects than brand loyalty, price position, or interpurchase time. Brands that are purchased by more customers tend to have larger ATEs. One explanation for the effect is that some categories are not relevant for certain customers, and coupons do not change this relevance. If few customers have pets, then we would expect both a low penetration rate for pet food brands, and a relatively small response to pet food coupons.

Purchase incidence and quantity purchased treatment effects are larger for brands with less loyal customers. A possible interpretation is that when *Brand Loyalty* is low, customers in the focal category are more price-sensitive or have lower switching costs, so they frequently switch brands during a promotion. For these customers, the coupon discount is sufficient to increase preferences for the brand and prompt a purchase.

We observe a positive relationship between a brand's relative *Price Position* and revenue treatment effects. This presumably reflects the value of each incremental sale (revenue scales with price). Revenue treatment effects are also larger for brands with longer *Interpurchase Times*.

To illustrate which brands can expect the largest sales lifts from in-store coupons, we calculated treatment effects for the brands in each feature's top and bottom decile (see Table 5). For the ten brands with the highest penetration, the average quantity purchased treatment effect is 41.82 incremental units, compared to just 5.72 for the ten brands with the lowest penetration. For the ten brands with the highest and lowest loyalty measures, the ATE varied from 9.98 (high loyalty) to 28.38 (low loyalty) incremental units.

Table 5: Which Brands Can Expect the Largest Treatment Effects?

		Incidence ATE	**Quantity ATE**	**Revenue ATE**
Brand Penetration	Highest	2.15%	41.82	€44.81
	Lowest	0.48%	5.72	€28.43
Brand Loyalty	Highest	0.88%	9.98	€36.67
	Lowest	1.66%	28.38	€21.57
Price Position	Highest	1.09%	14.90	€25.51
	Lowest	1.07%	24.99	€17.75
Interpurchase Time	Highest	1.27%	14.60	€29.04
	Lowest	0.80%	22.56	€23.13

The table reports the average ATE (on the current trip) for the ten brands with the highest values and the ten brands with the lowest values on each brand feature.

Previous studies have documented a negative relationship between loyalty and the response to promotion (Bawa and Shoemaker 1987a; Bell et al. 1999; Lim et al. 2005). Other studies investigated whether brand penetration, price position and interpurchase time explain variation in the response to different types of price promotions (Narasimhan et al. 1996; Bell et al. 1999; Ailawadi and Neslin 1998). In Web Appendix A, we compare our findings to this past literature. However, we caution that past studies use different identification methods (including correlations in historical data) and different outcome measures, making direct comparisons difficult.

4.3 Heterogeneity in Treatment Effects Across Customers

Many retailers allow brands to choose whether in-store coupons are distributed to past brand customers, customers of competing brands, or customers who have not recently purchased in the category. Treatment effects may vary depending on customers' past purchase behavior. To investigate this, we use loyalty data in the three months before the coupon experiment to group customers into three customer segments: *Not Category Customers* (60%), *Category Customers with Promoted Brand* (6%), and *Category Customers without Promoted Brand* (34%). We then calculate conditional average treatment effects (CATEs) separately for each segment, which we then average across the 101 brands. In Figure 4, we report the CATEs for the promoted brand (on the current trip) by each segment.

Figure 4: Heterogeneity in Treatment Effects for the Promoted Brand

(A) Purchase Incidence (B) Quantity (C) Revenue

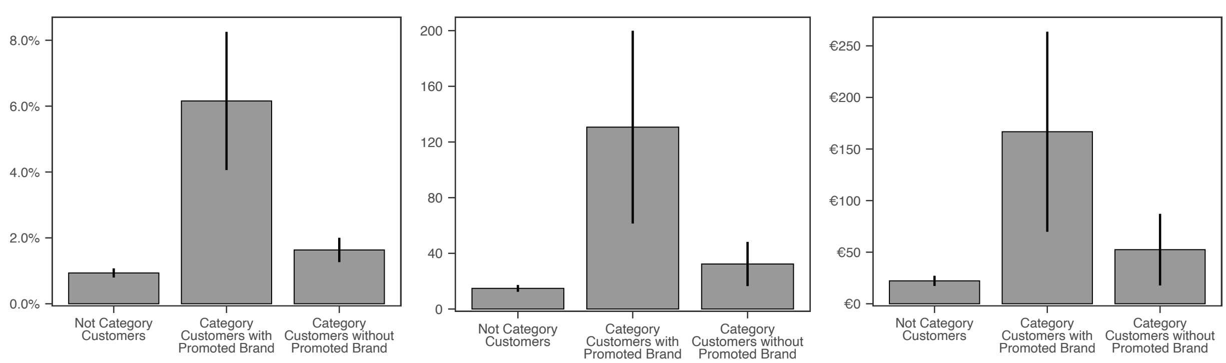

The figures report the average ATEs (on the current trip) averaged across the 101 focal brands. Each column represents a different customer segment. Error bars indicate 95% confidence intervals.

The largest treatment effects occur among the *Category Customers with Promoted*

Brand segment. Among these customers, the promoted brand's revenue increased by over €167 per mille. However, we also see significant increases in revenue among the larger segments of customers who had only purchased other brands in the category in the previous three months (€52 per mille) or had not purchased in the category in that prior period (€22 per mille). Recall that the *Category Customers with Promoted Brand* segment is over 5 times larger than the *Category Customers without Promoted Brand* segment, and the *Not Category Customers* segment is ten times larger. As a result, the aggregate revenue lift is higher among the customers who had not purchased the promoted brand.

4.4 Discussion

Our findings reveal a consistently large increase in purchases of the promoted brand on the visit on which customers received the coupon. The increase in sales volumes is large enough to offset the coupon discount, and so the coupons resulted in revenue increases for almost all of the 101 brands. The consistency of the revenue result is notable. In-store coupons can be exercised immediately, which results in high redemption rates. In our setting, 100% of the customers that received a focal brand coupon and purchased the focal brand received the coupon discount. The revenue increase occurs despite this 100% redemption rate.

While in-store coupons increase purchases and revenue for essentially every brand in our study, the magnitude of the effects varies. The largest treatment effects were observed on brands that many customers had purchased the brand in the past (high brand penetration). Purchase incidence and sales quantity treatment effects were also larger on brands with lower brand loyalty. Revenue treatment effects did not vary significantly with brand loyalty, but they were significantly larger on brands with higher prices and longer interpurchase times.

Our findings also highlight the versatility of in-store coupons. Different brands have different objectives for their promotions. Some are focused solely on expanding top-line sales, while others want to induce brand switching or attract new customers into the category. In-store coupons are an effective tool for achieving these different goals. The sales increases for the promoted brand are not limited to the promoted brand's past customers. Sales of the promoted brand also increase for past customers of other brands, and for customers who had not recently purchased in the category.

5. Why Do Sales Increase for the Promoted Brand?

In this section, we investigate three mechanisms for the increase in purchases of the promoted brand on the current trip: forward buying, brand switching, and category expansion.

5.1 Did the In-Store Coupons Pull Demand from the Future?

Previous work in the price promotions literature has documented a "post promotion dip" for CPG products in the days or weeks following a price promotion (see for example Neslin et al. 1985; Neslin and Schneider Stone 1996; Van Heerde et al. 2000; Hendel and Nevo 2003; Macé and Neslin 2004). To investigate whether in-store coupons also result in a post-promotion dip, we repeated our analysis using longer measurement windows that combine outcomes on the current visit with outcomes on future visits.

In Figure 5, we report the treatment effects averaged across brands for six different measurement windows: the day the coupon is distributed ("Immediate"), the distribution date and the next six days ("One Week"), the distribution date and the next thirteen days ("Two Weeks"), etc. For each measurement window, we calculate purchase incidence (Panel A), quantity purchased ATEs (Panel B), and revenue ATEs (Panel C) for the promoted brands. The columns present the average treatment effects across the 101 brands; the error bars indicate 95% confidence intervals.

For all three outcome measures, we find that the treatment effects remain positive and statistically significant, even when evaluating the cumulative outcome in the five weeks after the coupon is distributed. The purchase incidence treatment seems to decrease slightly as the length of the measurement window is extended farther into the future. This suggests that some of the customers in the treatment condition who purchased on the current visit would have purchased in future weeks. However, the decrease in the cumulative treatment effect is small compared to the overall size of the treatment effects.

We do not observe any decrease for purchase quantity or revenue. If anything, the quantity and revenue treatment effects trend upwards slightly, although the increase is not statistically significant. We conclude that in-store coupons do not simply pull demand forward in time.

To further explore forward buying, we analyzed whether the effect is larger for items more suitable for stockpiling. Previous research of post-promotion dips has generally attributed this phenomenon to stockpiling, whereby customers take advantage of the discounted price to meet future consumption needs. In our setting, the in-store coupons are only valid for a day, which may tend to encourage stockpiling. We construct a category-level stockpiling measure following Narasimhan et al. (1996). The measure is survey-based and includes a two-item stockpiling scale, which has become a standard for measuring the suitability of a category for stockpiling (see for example Ailawadi and Neslin 1998; Bell et al. 1999). We obtained responses to both items from 63 German Prolific participants and averaged the responses across the two items and participants to construct a category level *Stockpiling* measure.[6]

We do not find evidence of a significant relationship between suitability for stockpiling and the response to in-store coupons. The pairwise correlations between *Stockpiling* and the purchase incidence, purchase quantity, and revenue treatment effects are 0.01, -0.08, and 0.12 (respectively). None of these correlations are significantly different from zero.

We conclude that the increase in the promoted brand sales cannot be attributed solely to in-store coupons pulling demand forward. We next investigate whether the in-store coupons pulled demand from competing brands.

[6]The two items include "It is easy to store extra quantities of this product in my home" and "I like to stock up on this product when I can." Both items are measured using a 5-point Likert agreement scale. In Web Appendix D, we describe additional details of this survey, including copies of the stimuli.

Figure 5: Cumulative Treatment Effects Over Time

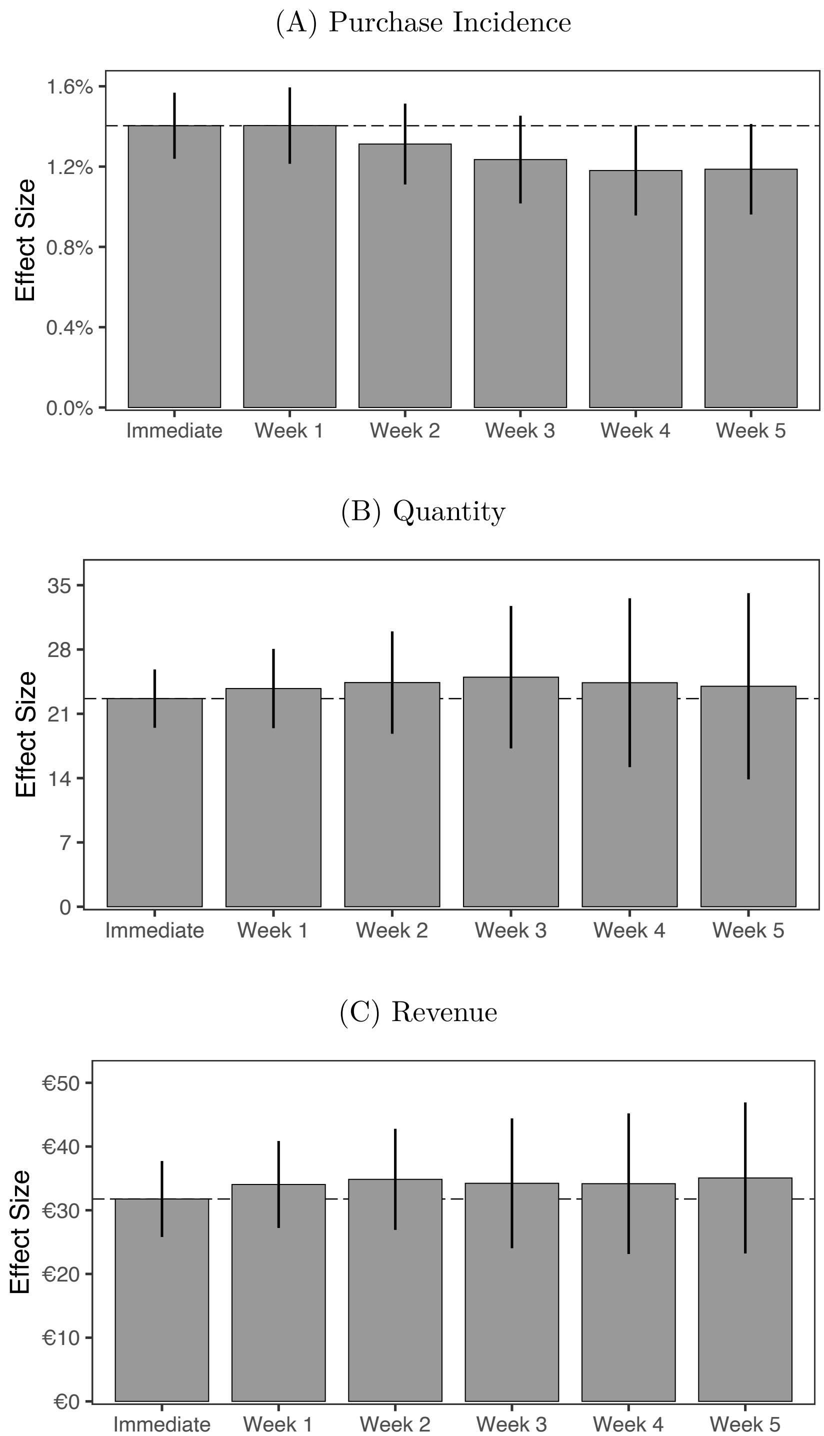

The figure reports ATEs averaged across the 101 focal brands. Each column represents a different measurement window for the outcome measure. "Week 1" includes the coupon distribution date and the next 6 days. "Week 2" includes the distribution date and the next 13 days. The other columns are similarly defined. The dashed horizontal line represents the immediate treatment effect. Error bars indicate 95% confidence intervals.

5.2 Did the In-Store Coupon Pull Demand from Competing Brands?

In Figure 4 we reported that in-store coupons cause sales increases across different types of customers, including those who had previously purchased competing brands. The evidence that the promoted brand's sales increase for the competing brand's customers suggests that the coupons contributed to brand switching. This interpretation is consistent with previous research on price promotions, which has found that brand switching accounts for more than 75% of the demand increases for the promoted brands (Gupta 1988; Bell et al. 1999). Similarly, the couponing literature has emphasized that coupons attract brand switchers (Neslin 1990). In this section, we investigate how in-store coupons affect sales of competing brands in the same category as the promoted brand.

In Table 6, we report the average treatment effects on the current visit for *competing* brands. The treatment effects are first calculated at the promoted brand level and then averaged across the 101 promoted brands. The results measure how a coupon for a promoted brand changes sales of the competing brands, aggregated across all competitors in the category.

The in-store coupons for the promoted brand have *positive* spillovers. The competing brands benefit from the promoted brands' in-store coupons, without taking any marketing action themselves, and without lowering margins through discounts. The positive spillovers to the competing brands are larger among the customers who had previously purchased a competing brand (*Competing Brand Customers*) than among those who had not (*Not Competing Brand Customers*).

The findings in Table 6 measure the outcomes on the day that customers received the in-store coupons. It is possible that the positive spillovers to the competing brands on the current visit are offset by negative spillovers to the competing brands on future visits. We investigated this possibility by measuring the change in sales of competing brands on future visits. We did not find evidence of a drop in competing brands sales on future visits.

Table 6: Treatment Effects for Sales of the Competing Brands

	Incidence ATE	**Quantity ATE**	**Revenue ATE**
All Customers	0.28% * (0.13%)	6.99 * (3.39)	€1.64 (€4.58)
Customer Segment			
Competing Brand Customers	0.64% * (0.31%)	23.92 ** (8.81)	€9.26 (€13.17)
Not Competing Brand Customers	0.24% † (0.14%)	2.81 (3.73)	-€0.20 (€4.53)
Difference	0.39% (0.33%)	21.12 * (9.26)	€9.46 (€13.06)

The table reports treatment effects for the competing brands on the current visit. The results are reported separately, according to whether customers had purchased the competing brands in the three months prior to receiving the random in-store coupon. Standard errors are in parentheses. ** indicates significantly different from zero, $p < 0.01$. * indicates significantly different from zero, $p < 0.05$. † indicates significantly different from zero, $p < 0.10$.

This represents some of the first documented evidence of positive competitive spillovers for price promotions in the CPG sector. The extensive literature investigating the outcome of coupons distributed at the checkout or to customers outside the store, including FSI and direct mail coupons, does not report evidence of positive competitive spillovers. Outside the CPG setting, Anderson and Simester (2013) document positive competitive spillovers for private-label apparel. However, those spillovers are for competitors' advertising instead of competitors' price promotions. What makes positive competitive spillovers striking for price promotions is that customers who make incremental purchases of the competing brands *forgo the discounts* available for the promoted brand. They respond to the promoted brand coupons by purchasing the competing brand and forgoing the coupon.

We might have expected that brand switching could have resulted in reduced sales of the competing brands. We observe the opposite outcome; for people who were customers of the competing brand in the past, the in-store coupons cause an increase in competing brand sales on the current trip. We caution that we should not conclude there is no brand switching. Indeed, the findings in Figure 4 suggest that at least some of the competing brands' past customers respond to the in-store coupon by increasing their purchases of the promoted brand. Instead, sales of the competing brands increase despite some customers switching from the competing brand to the promoted brand. We next examine a third mechanism that could explain why sales increase for both the promoted brand and the competing brands: category expansion.

5.3 Category Expansion

In Table 7, we report findings measured at the category level. This analysis uses the same identification approach that we described in Section 3.3, but the outcome measure focuses on category-level outcomes instead of outcomes for either the promoted brand or competing brands. There is strong evidence of category expansion. The effects are largest for customers who had previously purchased in the category (*Category Customers*), but also extend to customers who had not purchased in the category (*Not Category Customers*).

Table 7: Category Expansion Treatment Effects

	Incidence ATE	**Quantity ATE**	**Revenue ATE**
All Customers	1.45% ** (0.14%)	29.64 ** (3.96)	€33.40 ** (€5.58)
Customer Segment			
Category Customers	2.44% ** (0.31%)	62.55 ** (14.46)	€71.26** (€19.02)
Not Category Customers	1.10% ** (0.15%)	17.92 ** (4.22)	€22.61 ** (€5.57)
Difference	1.34% ** (0.35%)	44.63 ** (14.90)	€48.65 * (€18.92)

The table reports the category-level treatment effects (on the current visit). The results are reported separately, according to whether customers had purchased any product in the promoted category in the three months prior to receiving the random in-store coupon. Standard errors are in parentheses. ** indicates significantly different from zero, $p < 0.01$. * indicates significantly different from zero, $p < 0.05$.

We report findings separately for each of the 44 categories in Appendix A. The category-level treatment effects reveal that the in-store coupons caused a lift in category units purchased even for laundry care. We might wonder how category demand for laundry care can expand, given we do not expect usage to vary with price. Our category-level sales measures focus on purchases from a single chain of stores and does not measure purchases from competing stores. The results may indicate that the in-store coupons pulled demand from competing stores (Venkatesan and Farris 2012).

For retailers, this category expansion makes in-store coupons an attractive promotional tool. The lift in demand for the promoted brand is not at the expense of either future sales or sales of competing brands. Instead, the retailer enjoys unit sales increases for the category, and some of this demand increase is on competing brands for which there is no price reduction.

5.4 Discussion

We have investigated the response to a new generation of price promotions: in-store coupons distributed to customers at the bottom of the purchasing funnel (as they enter the store). Category expansion is the primary mechanism for why in-store coupons increase sales of the promoted brand. The increase in category sales extends beyond the promoted brand to other brands in the category, resulting in a significant lift in sales of the competing brands. These positive spillovers to the competing brands are the opposite of what we would expect if in-store coupons merely drive brand-switching. For in-store coupons, purchase acceleration and stockpiling also appear to play only a minor role.

The findings for this new generation of price promotions contrast with previous research on traditional price promotions. For example, Gupta (1988) investigated the short-term effects of price promotions on coffee and concluded that 84% of the effect is due to brand switching, 2% is due to forward buying (stockpiling), and 14% is due to an increase in category purchase quantities. Bell, Chiang, and Padmanabhan (1999) study thirteen CPG categories and revise this breakdown to 75% brand switching, and 25% to short-term primary demand expansion (category expansion and forward buying). One explanation for why these results contrast with our findings is that traditional price promotions are effective at driving customers to stores, while in-store coupons drive customers to the shelf (when they are in the store).

In Section 3.3, we recognized that positive spillovers on competing brands breach the identifying assumption in our ATE estimator that treatment effects do not extend across brands. However, our results are robust to this concern. First, competing coupons are relatively rare (for the median brand, customers receive competitors' targeted coupons in 5.4% of observations), and the spillover effects are small compared to the main effect of in-store coupons on the promoted brand (see Figure 3 and Table 6). Second, we repeat all of the analyses in this section when restricting attention to brands for which there are no in-store coupons for competing brands. The pattern of findings is impressively robust even when using this smaller set of brands (Web Appendix E). Average ATEs remain large and significant for incidence, quality and revenue. There is little evidence of a post-promotion dip, or that in-store coupons pulled demand from the future. We continue to see strong evidence of positive competitive spillovers among past purchasers of the competing brands. We also consistently observe large category expansion effects. The presence or absence of competing in-store coupons does not change any of our conclusions.

One explanation for why we observe category expansion is that the coupons may increase traffic in the store aisle in which the promoted brand is located. Although this incremental aisle traffic may initially intend to purchase the promoted brand, when these customers arrive at the aisle, they see not just the promoted brand, but also all the competing brands. The availability of the competing brands, together with the packaging and other on-shelf stimuli, may prompt some customers to purchase a competing brand instead, which could contribute to positive spillovers to competing brands on the current trip. We next investigate this explanation by evaluating the role of impulse purchasing in the response to in-store coupons.

6. Impulse Purchasing and the Response to In-store Coupons

In the previous section, we documented that in-store coupons contribute to category expansion, by increasing the number of customers that purchase both the promoted brand and competing brands in the same category. The aisle traffic explanation that we proposed for the observed treatment effects is akin to President John F. Kennedy's idiom that "A rising tide lifts all boats." In-store coupons generate more traffic in the aisle and benefit all the brands in the category. This interpretation suggests that customers were willing to change their shopping plans. They decide to visit the aisle when they would not otherwise have done so, thereby changing their in-store shopping route. By switching from the promoted brand to a competing brand when they arrived at the shelf, some customers again chose to change their plan.

The decision to visit the aisle when they would not otherwise have done so, and the switch to purchase a competing brand when arriving at the aisle would both be interpreted as unplanned, or impulsive decisions. Past studies have established that purchases in some categories tend to be planned, while purchases in other categories are often unplanned (see, for example, Narasimhan et al. 1996; Bell et al. 1999; Ailawadi et al. 2006). If the in-store coupons in our study prompted customers to impulsively visit the store aisle and some of these customers impulsively switch to a competing brand when they arrived, we would expect more category expansion and larger competitive spillovers in categories for which purchases tend to be impulsive (versus planned).

To investigate these predictions, we asked the same 63 German participants who responded to the *Stockpiling* questions (see earlier discussion) to rate the categories using a two-item *Impulse* scale from Narasimhan et al. (1996).[7] We averaged the response across

[7] The two items are "I often buy this product on a whim when I pass by it in the store." and "I typically like to buy this product when the urge strikes me." Both items are measured on 5-point Likert scales. See Web Appendix D for a screenshot of the survey stimuli and a discussion of the survey design and implementation.

items and participants within each category to calculate a category-level measure of customers' tendency to buy items on impulse (*Impulse*). In Table 8, we report the Spearman pairwise correlation between the three treatment effects (purchase incidence, purchase quantity, and revenue) and this *Impulse* measure. The correlations are reported separately when measuring outcomes for the promoted brand, competing brands, and the category.

Table 8: Correlation between the ATEs and the Impulse Measure

Outcome	**Incidence ATE**	**Quantity ATE**	**Revenue ATE**
Promoted Brand	0.274 **	0.252 *	0.309 **
Competing Brands	0.184 †	0.135	0.176 †
Category (all brands)	0.158	0.185 †	0.170 †

The table reports the Spearman pairwise correlation between the ATEs and the Impulse measure. The unit of analysis is a category and the sample size is 101. ** indicates significantly different from zero, $p < 0.01$. * indicates significantly different from zero, $p < 0.05$. † indicates significantly different from zero, $p < 0.10$.

The positive correlations indicate larger treatment effects in high impulse categories. This pattern is consistent when measuring outcomes for the promoted brand, for the competing brands and for the category. In-store coupons cause larger demand increases in categories that tend to be purchased on impulse compared to categories where purchases tend to be planned. In Figure 6, we illustrate the relationships in Table 8 by comparing the purchase incidence ATEs in categories with high and low *Impulse* scores. We select the thirty brands with the highest *Impulse* scores and the thirty brands with the lowest *Impulse* scores and report the average ATEs within these groups of brands.

We see that among the thirty high-impulse brands, the incremental purchase incidence is 1.50% for the promoted brand, 0.80% for the competing brand, and 1.92% for the category. In contrast, among the thirty low-impulse brands, these ATEs are just 0.91%, 0.11%, and 0.93%. The difference between these averages is statistically significant ($p < 0.05$) for all three outcome measures.

Figure 6: Purchase Incidence ATEs in High and Low Impulse Categories

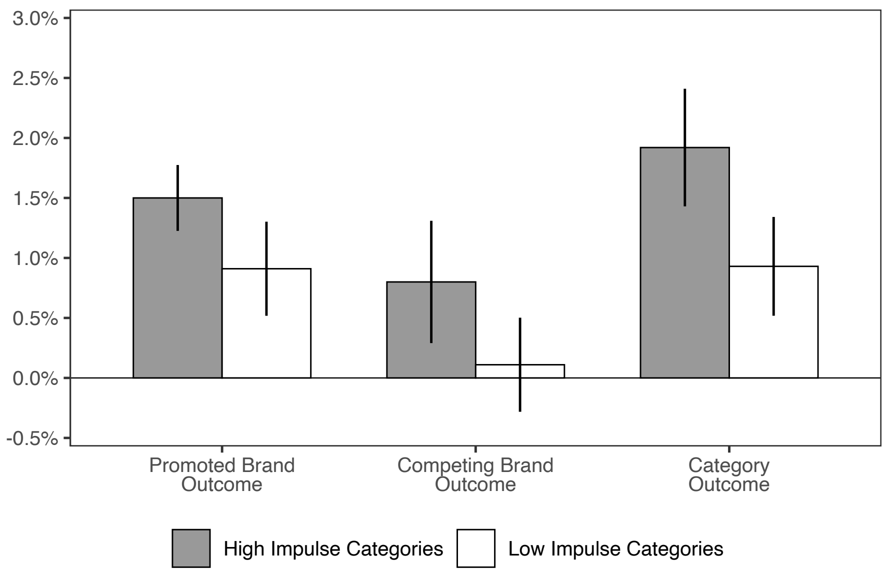

The figure reports the average purchase incidence ATEs for the 30 brands with the highest Impulse measure and the 30 brands with the lowest Impulse measure. The unit of analysis is a brand, and the sample size for each bar is 30.

Among the 30 low-impulse brands, the (average) competitive spillovers are essentially zero. Yet we see significant increases in both promoted brand and category purchase incidence (light grey bars). This suggests that in low-impulse categories, in-store coupons also drive additional traffic to the aisle, but when these customers arrive, they purchase the promoted brand with a discount, rather than switching to a competing brand and forgoing the discount.

In contrast, positive spillovers to the competing brands are prominent among the 30 high-impulse categories. We also see positive increases in both promoted brand and category purchase incidence (dark grey bars). The implication is that in high impulse categories, the in-store coupons are again effective at driving additional traffic to the aisle, but when these customers arrive, some of them purchase the promoted brand with a discount, and some of them switch to a competing brand and forgo the discount.

We report analogous figures for the quantity and revenue treatment effects in Web Appendix F and observe a similar pattern of results. However, the standard errors are larger for these two outcome measures, and the differences between the high and low-impulse ATEs are not statistically significant.

We could not find previous studies that have compared the impact of CPG coupons in high and low-impulse categories. Perhaps the closest example is a study by Heilman et al. (2002), who study how coupons impact planned and unplanned purchases. They show that coupons for planned purchases do not change planned purchases but can prompt customers to make additional unplanned purchases outside the promoted category. If we interpret planned categories as low-impulse categories, we can compare the null result in their promoted category with the positive results we observe for the promoted brand in both low and high-impulse categories (see the first two bars in Figure 6).[8]

Our finding also contrasts with previous research on advertised price promotions that has reported negative relationships between incremental sales and impulse measures.[9] This is the reverse of the positive relationships we report in Table 8 and Figure 6. A notable difference is that advertised promotions are received before a store visit, and so customers may use these promotions to help plan future shopping trips. This could contribute to a different relationship between treatment effects and the susceptibility of brands to impulse purchasing.

Finally, the finding that the treatment effects of in-store coupons are larger in high-impulse categories is consistent with the model proposed by Bucklin and Lattin (1991). They argue that opportunistic shoppers, who have not planned their shopping trips in advance, make more impulsive decisions in the store and are therefore more responsive to in-store promotions. In contrast, customers making planned purchases do not respond to promotions. In our empirical setting, customers receive coupons after visiting the kiosk system, arguably putting them into the opportunistic state, in which coupons affect category and brand choice.

[8] These positive results for the promoted brand in low-impulse categories also extend to the purchase quantity and revenue ATEs (see Web Appendix F).

[9] Bell et al. (1999) and Ailawadi et al. (2006) report a significant negative relationship between promotion treatment effects and *Impulse*. Narasimhan et al. (1996) do not find a significant relationship. All three studies use the same *Impulse* measure that we use.

7. Conclusions

A long-standing challenge for retailers is to find a mechanism for delivering promotions to customers when customers are about to purchase. This problem is easier to solve in digital channels than in physical retail channels. Digital retailers have introduced sponsored search and on-site advertising. However, in physical channels, firms have resorted to sending promotions to customers' homes by direct mail, distributing FSIs in newspapers and magazines, advertising through weekly circulars, or handing customers coupons at checkout. None of these distribution channels reach customers at the beginning of a store visit. Retailer apps and in-store kiosks represent a breakthrough that can overcome this long-standing challenge. Customers can now receive coupons during their visits to physical stores.

We investigate the effectiveness of in-store coupons using data from a large-scale field experiment. The coupons were distributed through kiosks, which customers in the retailer's loyalty program typically visited at the start of their shopping trips. The data provides a causal measure of the impact of in-store coupons on customer purchasing. We summarize the findings in Table 9.

In-store coupons are very effective at increasing sales of the promoted brand. The sales increases are large enough to offset the discounts, so that revenue for the promoted brand also increases. The sales increases come from different types of customers, including past purchasers of the promoted brand, past purchasers of competing brands, and customers without past category purchases. This makes in-store coupons a versatile marketing tool, that can be used to increase topline sales, induce trial by customers of competing brands and attract customers to the category.

The increase in sales of the promoted brand is not at the expense of either future sales or sales of competing brands. We see little evidence of forward buying or stockpiling. Moreover, rather than decreasing sales of competing brands, in-store coupons cause sales of competing brands to *increase*. What makes these positive competitive spillovers particularly surprising

Table 9: Summary of Key Findings

Sales Increase for the Promoted Brand: In-store coupons increase purchase incidence and unit purchases of the promoted brand. The unit sales increases are large enough to offset the coupon discount, so that revenue also increases for the promoted brand.
Accomplishes Multiple Goals: In-store coupons are a versatile marketing tool that can be used to increase top-line sales, attract purchases by past customers of competing brands, and induce purchasing by customers who do not regularly purchase in the category.
Category Expansion: The increase in sales of the promoted brand reflects category expansion rather than merely forward buying or brand switching.
Positive Competitive Spillovers: In-store coupons *increase* purchase incidence and purchase quantities for competing brands, particularly among past purchasers of the competing brands. For these customers, category expansion is large enough to offset brand switching.
Stimulate Impulse Purchasing: The sales lift for the promoted brand and competing brands are larger in categories that customers tend to purchase on impulse.

is that customers make incremental purchases of the competing brands without receiving the coupon discounts. The in-store coupons prompt customers to buy the competing brands, but in doing so, they forgo the opportunity to purchase the promoted brand at a discount. Because purchases of the competing brands are not discounted, the value of these transactions is magnified for both the retailer and the competing brands.

Our findings suggest that when customers receive an in-store coupon, some customers are prompted to make two unplanned decisions. First, they choose to visit the aisle in which the promoted product is located. Second, at least some of them decide to purchase a competing brand instead of the promoted brand when they get there. We might expect both outcomes to be more likely in categories in which customers tend to purchase on impulse (rather than making planned purchases). Our findings are consistent with this. We observe larger sales lifts for both the promoted brand and competing brands in high impulse categories. Consequently, there is also more category expansion in these categories.

We conclude that in-store coupons provide a very effective marketing tool for increasing sales of the promoted brand. The sales increase does not merely reflect forward buying or brand switching. Instead, in-store coupons drive category expansion, which benefits the promoted brand, competing brands, and the retailer itself.

Limitations and Future Research

A limitation in our study is that we do not observe the wholesale prices that the retailer pays to manufacturers for their brands. This means we are unable to report how in-store coupons affect the incremental profits earned by the retailer or the manufacturers. Retailers and manufacturers could combine the treatment effects reported in this paper with their own wholesale price data to estimate changes in profits. A related question is how in-store coupons increase overall channel profits, and how the incremental profits are shared between retailers and manufacturers.

Our data describes the response to CPG coupons delivered via in-store kiosks. Many retailers have developed alternative in-store media channels to deliver promotions to customers, including retail apps, in-store video and displays. Examples can be found in consumer electronics (Best Buy), department store (Saks), fast food (Starbucks), home improvement (Home Depot), beauty (Ulta), and other retailers. Future research could study the performance of retail media in these settings.

Future research could also investigate the value of targeting in-store coupons. Many firms are interested in using experimental data to design targeting policies that vary marketing actions across different customer segments. For example, retailers allow manufacturers to design programmatic marketing campaigns, in which marketing actions (such as coupons) are restricted to a subset of customers. Retailers also often target different marketing actions to different individual customers. We would need additional data in order to study the targeting of in-store coupons.

References

Abratt R, Goodey SD (1990) Unplanned buying and in-store stimuli in supermarkets. *Managerial and Decision Economics* 11(2):111–121.

Ailawadi KL, Harlam BA, Cesar J, Trounce D (2006) Promotion profitability for a retailer: the role of promotion, brand, category, and store characteristics. *Journal of Marketing Research* 43(4):518–535.

Ailawadi KL, Neslin SA (1998) The effect of promotion on consumption: Buying more and consuming it faster. *Journal of Marketing Research* 35(3):390–398.

Anderson ET, Simester D (2013) Advertising in a competitive market: The role of product standards, customer learning, and switching costs. *Journal of Marketing research* 50(4):489–504.

Angrist JD, Imbens GW, Rubin DB (1996) Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* 91(434):444–455.

Bawa K, Shoemaker RW (1987a) The coupon-prone consumer: some findings based on purchase behavior across product classes. *Journal of Marketing* 51(4):99–110.

Bawa K, Shoemaker RW (1987b) The effects of a direct mail coupon on brand choice behavior. *Journal of Marketing Research* 24(4):370–376.

Bell DR, Chiang J, Padmanabhan V (1999) The decomposition of promotional response: An empirical generalization. *Marketing Science* 18(4):504–526.

Berthiaume D (2022) Survey: Online shoppers search for coupons. https://chainstoreage.com/survey-online-shoppers-search-coupons, accessed: 2025-01-16.

Blattberg RC, Briesch R, Fox EJ (1995) How promotions work. *Marketing science* 14(3_supplement):G122–G132.

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

Chiang J (1991) A simultaneous approach to the whether, what and how much to buy questions. *Marketing Science* 10(4):297–315.

Chiang J (1995) Competing coupon promotions and category sales. *Marketing Science* 14(1):105–122.

Chiou-Wei SZ, Inman JJ (2008) Do shoppers like electronic coupons?: A panel data analysis. *Journal of Retailing* 84(3):297–307.

Dekimpe MG, Hanssens DM, Silva-Risso JM (1998) Long-run effects of price promotions in scanner markets. *Journal of econometrics* 89(1-2):269–291.

Dhar SK, Hoch SJ (1996) Price discrimination using in-store merchandising. *Journal of Marketing* 60(1):17–30.

Fok D, Horváth C, Paap R, Franses PH (2006) A hierarchical bayes error correction model to explain dynamic effects of price changes. *Journal of Marketing Research* 43(3):443–461.

Gabel S, Simester D, Timoshenko A (2024) How retailers became ad platforms. Harvard Business Review, URL https://hbr.org/2024/06/how-retailers-became-ad-platforms, accessed: 2025-01-16.

Grewal D, Ahlbom CP, Noble SM, Shankar V, Narang U, Nordfält J (2023) The impact of in-store inspirational (vs. deal-oriented) communication on spending: The importance of activating consumption goal completion. *Journal of Marketing Research* 60(6):1071–1094.

Grover R, Srinivasan V (1992) Evaluating the multiple effects of retail promotions on brand loyal and brand switching segments. *Journal of marketing research* 29(1):76–89.

Gupta S (1988) Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing research* 25(4):342–355.

Hamstra M (2023) Stop & shop pilots kiosks for coupons and more. https://www.supermarketnews.com/finance/stop-shop-pilots-kiosks-for-coupons-and-more, accessed: 2025-01-16.

Heilman CM, Nakamoto K, Rao AG (2002) Pleasant surprises: Consumer response to unexpected in-store coupons. *Journal of Marketing Research* 39(2):242–252.

Hendel I, Nevo A (2003) The post-promotion dip puzzle: what do the data have to say? *Quantitative Marketing and Economics* 1:409–424.

Hendel I, Nevo A (2004) Intertemporal substitution and storable products. *Journal of the European Economic Association* 2(2–3):536–547.

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

Hendel I, Nevo A (2013) Intertemporal price discrimination in storable goods markets. *American Economic Review* 103(7):2722–2751.

Hui SK, Inman JJ, Huang Y, Suher J (2013) The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing* 77:1–16.

Inmar (2023) Report shows grocery shoppers will switch brands for coupons or discounts. Produce Blue Book, URL https://www.producebluebook.com/2023/03/01/report-shows-grocery-shoppers-will-switch-brands-for-coupons-or-discounts/, accessed: 2025-01-16.

Iyer GR, Blut M, Xiao SH, Grewal D (2020) Impulse buying: a meta-analytic review. *Journal of the Academy of Marketing Science* 48:384–404.

Johnson GA, Lewis RA, Nubbemeyer EI (2017) Ghost ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research* 54(6):867–884.

Lim J, Currim IS, Andrews RL (2005) Consumer heterogeneity in the longer-term effects of price promotions. *International Journal of Research in Marketing* 22(4):441–457.

Litvack DS, Calantone RJ, Warshaw PR (1985) An examination of short-term retail grocery price effects. *Journal of Retailing* 61(3):9–25.

Macé S, Neslin SA (2004) The determinants of pre-and postpromotion dips in sales of frequently purchased goods. *Journal of Marketing Research* 41(3):339–350.

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

Moriarty MM (1985) Retail promotional effects on intrabrand and interbrand sales performance. *Journal of Retailing* 61(3):27–47.

Narasimhan C, Neslin SA, Sen SK (1996) Promotional elasticities and category characteristics. *Journal of Marketing* 60(2):17–30.

Neslin SA (1990) A market response model for coupon promotions. *Marketing Science* 9(2):125–145.

Neslin SA, Henderson C, Quelch J (1985) Consumer promotions and the acceleration of product purchases. *Marketing Science* 4(2):147–165.

Neslin SA, Schneider Stone LG (1996) Consumer inventory sensitivity and the postpromotion dip. *Marketing Letters* 7:77–94.

Nijs VR, Dekimpe MG, Steenkamp JBE, Hanssens DM (2001) Tracing the impact of price promotions across categories. *Tilburg University the Netherlands Working paper* .

Osuna I, González J, Capizzani M (2016) Which categories and brands to promote with targeted coupons to reward and to develop customers in supermarkets. *Journal of Retailing* 92(2):236–251.

Pollay RW (1968) Customer impulse purchasing behavior: A reexamination. *Journal of Marketing Research* 5(3):323–325.

Raju JS, Dhar SK, Morrison DG (1994) The effect of package coupons on brand choice. *Marketing Science* 13(2):145–164.

Rook DW (1987) The buying impulse. *Journal of Consumer Research* 14(2):189–199.

Rubin DB (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5):688.

Sahni NS (2016) Advertising spillovers: Evidence from online field experiments and implications for returns on advertising. *Journal of Marketing Research* 53(4):459–478.

Sahni NS, Zou D, Chintagunta PK (2017) Do targeted discount offers serve as advertising? evidence from 70 field experiments. *Management Science* 63(8):2688–2705.

Shapiro MF (2001) Impulse purchases. *NPN, National Petroleum News* 93(7):40–40.

Srinivasan SS, Leone RP, Mulhern FJ (1995) The advertising exposure effect of free standing inserts. *Journal of Advertising* 24(1):29–40.

Statista (2022) Leading media companies worldwide by revenue 2022. URL https://www.statista.com/statistics/261827/leading-media-companies-worldwide, accessed: 2025-01-16.

Van Heerde HJ, Leeflang PS, Wittink DR (2000) The estimation of pre-and postpromotion dips with store-level scanner data. *Journal of Marketing Research* 37(3):383–395.

Venkatesan R, Farris PW (2012) Measuring and managing returns from retailer-customized coupon campaigns. *Journal of Marketing* 76(1):76–94.

Zhou L, Wong A (2004) Consumer impulse buying and in-store stimuli in chinese supermarkets. *Journal of International Consumer Marketing* 16(2):37–53.

A. Category-Level Average Treatment Effects

	Category	Incidence ATE	Quantity ATE	Revenue ATE
Drinks	Alcoholic mixers (e.g., alcopops, long drinks)	0.52%	5.62	€6.81
	Carbonated soft drinks (e.g., Coke, Fanta)	0.92%	9.27	-€54.34
	Juice	0.84%	22.22	€23.09
	Spirits (e.g., vodka, gin, rum)	0.56%	7.37	€54.37
	Water	0.38%	-7.44	-€94.60
Food	Baking mixes	-0.61%	5.93	€3.00
	Cake (packaged)	0.01%	2.00	€8.97
	Cheese	2.66%	40.97	€16.93
	Chocolate bars	2.37%	59.77	€65.18
	Chocolates	0.84%	19.08	€56.36
	Coffee	1.62%	24.80	€67.83
	Cold cuts	2.34%	56.62	€70.79
	Cookies	1.70%	16.45	€20.79
	Dairy desserts (e.g., puddings, jello)	1.00%	49.94	€31.98
	Fish and other frozen seafood	1.44%	27.71	€110.06
	Fries and other frozen potato products	2.16%	23.30	€20.62
	Frozen vegetables	1.81%	21.29	€35.41
	Gummy candy	2.00%	81.67	€67.83
	Jams and spreads (e.g., peanut butter)	1.14%	17.01	€46.87
	Milk	1.69%	42.51	€22.93
	Oats	1.01%	23.42	€22.59
	Olive oil and vegetable oil	1.42%	22.80	€52.13
	Pasta	1.11%	29.36	€65.81
	Pizza (frozen)	2.49%	55.31	€118.92
	Ready meals	1.65%	44.03	€35.01
	Rice	1.43%	25.64	€36.67
	Salty snacks (e.g., potato chips, nuts)	2.42%	46.77	€52.81
	Sauces	-0.17%	-13.17	€3.28
	Sausages boiled (packaged)	0.56%	17.24	€3.00
	Sausages cooked (packaged)	1.25%	12.47	€7.30
	Soups	1.83%	53.16	€66.99
	Spice mixes	0.94%	40.90	€41.58
	Toast	2.65%	34.61	€35.86
	Yogurt and cream cheese	1.72%	33.71	€42.84

	Category	Incidence ATE	Quantity ATE	Revenue ATE
Nonfood	Deodorant	1.02%	15.83	€16.94
	Dishwashing products	0.29%	-1.48	€8.27
	Hair care (e.g., shampoos, conditioners)	1.01%	14.14	€21.05
	Laundry care	1.82%	23.42	€29.19
	Oral care products (e.g., toothpaste)	1.20%	10.26	€7.51
	Other hair care (e.g., hair gel, hair spray)	0.79%	7.53	€16.13
	Soap	1.48%	17.81	€25.89
	Toilet paper, tissues, and paper towels	-0.02%	-12.73	-€12.75
Pet care	Cat food	1.27%	60.57	€57.06
	Dog food	0.57%	23.79	€27.27

B. Sample Selection Criteria

Our analysis uses two sample selection criteria. The first criterion focuses on estimation precision. We restrict our analyses to the brands that had regular sales prior to the experiment and distributed at least 400 random coupons. This criterion eliminates new products and short campaigns (i.e., single weekend).

The second sample selection criterion recognizes that brands could define eligibility criteria for their coupons. For example, a pet food brand only distributed coupons to customers who had purchased pet food before. Random coupons were not distributed to non-eligible customers, so we cannot measure *average treatment effects* for these brands. Not all brands used eligibility criteria, and the retailer did not record the eligibility criteria defined by brand managers.

We detect which brands had eligibility criteria by comparing the pretreatment variables across experimental conditions. In particular, we regress coupon assignments on three pretreatment variables, and identify brands for which pretreatment variables are not jointly significant in an F-test with 95% confidence level. The pretreatment variables include the number of store visits, revenue, and the number of unique products purchased during 90 days prior to the experiment (from August to October 2015). For each brand, we estimate two regression models that compare assignment into *Focal vs. NonFocal* and *Focal vs. NoRandom* (see Section 3.3), and we require that neither F-test is statistically significant.

We recognize that the F-test approach for detecting eligibility criteria is typically used as a randomization check. In our setting, the randomization is implemented programmatically in the system and was separately verified with the engineering team. Reassuringly, when we run randomization tests using a *different* set of pretreatment variables, the campaigns that remained in our sample pass these tests. The pretreatment variables for randomization checks include the total spending, recency and average basket size in 90 days prior to obtaining a coupon (specific to the customer×date observation). We report the distribution of p-values across the brands for *Focal vs. NonFocal*, *Focal vs. NoRandom*, and *NonFocal vs. NoRandom* in Figure 7. The distributions are not statistically different from the uniform distribution in the Kolmogorov-Smirnov test.

Figure 7: Distribution of p-values in the Randomization Tests

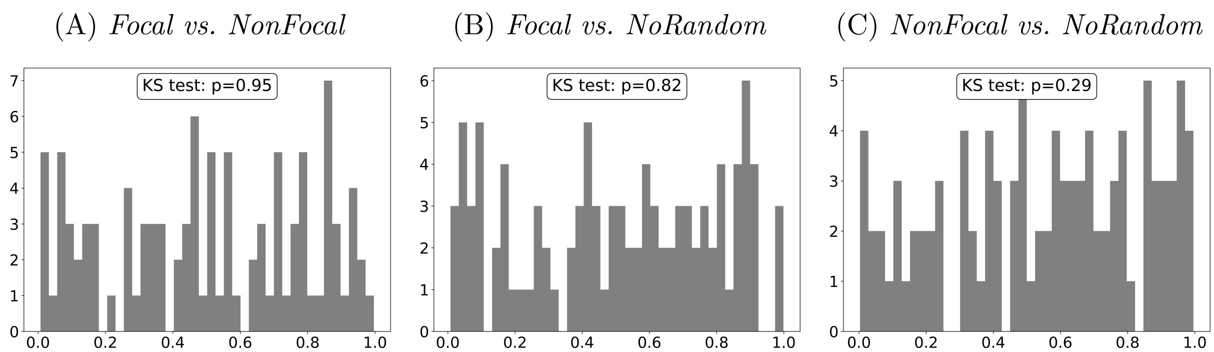

The figure reports histograms of the distribution of p-values in randomization tests. The sample size in each histogram is 101 (brands). The y-axis is a count of the number of brands and the x-axis are p-values in the F-test with three pretreatment variables.

C. Identification of Treatment Effects

Recall that in the main text, we define an *observation* as a customer×day combination, and we restrict our analysis to customer×day combinations in which the customer used the kiosk system (at least once). We denote a customer×day combination by $i \in 1, .., N$ and a brand by $b \in 1, .., B$. We also define $Y_{ib}(1)$ and $Y_{ib}(0)$ as potential outcomes for a customer×day observation i for brand b, with and without a coupon for brand b (Rubin 1974). Our goal is to estimate the average treatment effect (ATE) for brand b: $\text{ATE}_b = \mathbb{E}[Y_{ib}(1) - Y_{ib}(0)]$.

The proposed identification approach relies on two assumptions:

Assumption 1 (Stable Unit Treatment Value Assumption; *SUTVA*): The potential outcomes for an observation do not vary with the treatment assignments for other observations.

Assumption 2: Purchases of a focal brand are not affected by coupons from other brands.

Our first assumption recognizes that the same customer could appear in multiple customer×date observations in our data, and we treat these observations as independent. The second assumption implies that only focal coupons are relevant for estimating ATEs (instead of all seven coupons on the printout). Replacing a focal coupon with a random non-focal coupon moves from a potential outcome $Y_{ib}(1)$ to a potential outcome $Y_{ib}(0)$. These assumptions are standard in the causal inference literature.

Proposition: The following estimator provides an unbiased estimate of the average treatment effect for brand b:

$$\text{ATE}_b = \Big(\bar{Y}(\text{Focal}_b) - \bar{Y}(\text{NoRandom}_b)\Big) + 1/p \cdot \Big(\bar{Y}(\text{NoRandom}_b) - \bar{Y}(\text{NonFocal}_b)\Big) \tag{3}$$

where $\bar{Y}(\cdot)$ indicates the average outcome in the corresponding group, and p captures the probability that a random *non-focal* coupon replaces a targeted *focal* coupon in the NonFocal_b group.

Proof. We can estimate the first term in the ATE_b as an average monetary outcome in the Focal_b condition. Focal_b observations are representative of the customer population, and in our experiment, they always received a coupon for the focal brand:

$$\mathbb{E}[Y_{ib}(1)] = \mathbb{E}[Y_{ib}|\text{Focal}_b] = \bar{Y}(\text{Focal}_b) \tag{4}$$

To estimate the expected potential outcome without a focal coupon, $\mathbb{E}[Y_{ib}(0)]$, it is helpful to consider a group NoFocal_b, in which all targeted coupons for a focal brand are replaced by random coupons. Our data does not contain a group NoFocal_b. However, if we had such data, we could estimate $\mathbb{E}[Y_{ib}(0)]$ as follows:

$$\mathbb{E}[Y_{ib}(0)] = \mathbb{E}[Y_{ib}|\text{NoFocal}_b] = \bar{Y}(\text{NoFocal}_b)$$

The distribution of outcomes in NonFocal_b is a mixture of outcomes from NoFocal_b and outcomes from NoRandom_b, where a share of NoFocal_b observations is p. Indeed, in the NonFocal_b, random coupons replace share p of targeted coupons in the printouts, and the remaining random coupons replace targeted coupons for non-focal brands, which does not change the outcomes for the focal brand. We can thus estimate $\mathbb{E}[Y_{ib}(0)]$ as follows:

$$\mathbb{E}[Y_{ib}|\text{NonFocal}_b] = p \cdot \mathbb{E}[Y_{ib}|\text{NoFocal}_b] + (1-p) \cdot \mathbb{E}[Y_{ib}|\text{NoRandom}_b] \tag{5}$$

$$\mathbb{E}[Y_{ib}|\text{NoFocal}_b] = \mathbb{E}[Y_{ib}|\text{NoRandom}_b] - 1/p \cdot (\mathbb{E}[Y_{ib}|\text{NoRandom}_b] - \mathbb{E}[Y_{ib}|\text{NonFocal}_b])$$

$$\mathbb{E}[Y_{ib}(0)] = \bar{Y}(\text{NoRandom}_b) - 1/p \cdot (\bar{Y}(\text{NoRandom}_b]) - \bar{Y}(\text{NonFocal}_b)) \tag{6}$$

We obtain Equation (3) by combining Equations (4) and (6). $\square$

Variance

Under the SUTVA assumption, the observations in groups Focal_b, NonFocal_b, and NoRandom_b are independent, so we can write the variance of the proposed estimator:

$$\mathbb{V}(\text{ATE}_b) = \mathbb{V}(\bar{Y}(\text{Focal}_b)) + \left(\frac{p-1}{p}\right)^2 \mathbb{V}(\bar{Y}(\text{NoRandom}_b)) + \left(\frac{1}{p}\right)^2 \mathbb{V}(\bar{Y}(\text{NonFocal}_b)) \quad (7)$$

Competitors' Targeted Coupons

Our identification approach assumes that coupons for one brand do not affect sales of other brands. In Section 5, we present evidence that there are spillovers between brands within the same category. The spillovers between competing brands have implications for $\mathbb{E}[Y_{ib}(1)]$ and $\mathbb{E}[Y_{ib}(0)]$ in our identification approach.

First, in the presence of competitive spillovers, we can construct the NonFocal_b group using observations in which a customer received a random coupon for any brand *outside* the category of the focal brand (non-focal category). By doing so, we partial out the effects of targeted coupons from the estimation of $\mathbb{E}[Y_{ib}(0)]$ for both focal and competitors' brands. Equation (6) still holds because the couponing system in our experiment ensures that (1) any printout does not contain multiple targeted coupons from the same category, and (2) random coupons from outside the focal category replace one of the targeted coupons at the random position in the printout.

Second, observations in the Focal_b group receive the random coupon for the focal brand, but they may also receive a targeted competitors' coupon on the same printout. This can introduce a difference between the estimation of $\mathbb{E}[Y_{ib}(1)]$ and $\mathbb{E}[Y_{ib}(0)]$ because targeted competitors' coupons affect $\mathbb{E}[Y_{ib}(1)]$, but this effect is partialled out from $\mathbb{E}[Y_{ib}(0)]$ (see the previous paragraph). This difference can introduce a bias to our estimate of the ATE. Notice that the source of the bias is that observations in the treatment group can receive a targeted competitors' coupon *in addition* to the focal coupon. We expect the bias to be orders of magnitude smaller than the main effect. We validate that our findings are robust to this bias in Web Appendix E, where we replicate our analysis with brands that faced no competitors' coupons.