

von Mutius, Bernhard; Huchzermeier, Arnd

Article — Published Version

Customer-centric category selection for mobile and print promotions in loyalty reward programs

Naval Research Logistics (NRL)

Provided in Cooperation with:

John Wiley & Sons

Suggested Citation: von Mutius, Bernhard; Huchzermeier, Arnd (2021) : Customer-centric category selection for mobile and print promotions in loyalty reward programs, Naval Research Logistics (NRL), ISSN 1520-6750, John Wiley & Sons, Inc., Hoboken, USA, Vol. 68, Iss. 6, pp. 683-700,
<https://doi.org/10.1002/nav.21934>

This Version is available at:

<https://hdl.handle.net/10419/230248>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

Customer-centric category selection for mobile and print promotions in loyalty reward programs

Bernhard von Mutius¹ | Arnd Huchzermeier²

WHU—Otto Beisheim School of Management,
Vallendar, Germany

Correspondence

Bernhard von Mutius, WHU—Otto Beisheim
School of Management, 56179 Vallendar, Germany.
Email: bernhard.mutius@whu.edu

Abstract

In their loyalty reward programs, large retailers adopt category-specific targeted promotions as an effective means to increase shoppers' basket values. However, neither the literature nor observed marketing practice provides category selection methods that maximize the promotion's profitability. Hence, we provide a predictive customer-centric category selection approach based on the return on marketing investment measure that accommodates both fixed and variable promotion costs, captures cherry-picking effects, and encompasses the retailer's entire category assortment. We use a real-world promotional data set from a leading German retailer to show that our approach predicts customer responses to these promotions with high out-of-sample accuracy tested over time and also across promotion frequencies. We find that the most promising categories in *mobile* promotions maximize cross-category profits to curtail cherry-picking and boost the sales of nonpromoted items—that is, the profitability-driving part of the profit uplift. In contrast, the different cost structure of *print* promotions requires that categories achieve a high (but not too high) redemption rate, as cross-category profit declines, in order to recoup customer contacting costs. Our benchmark analysis reveals that current marketing practice fails to hit the profitability functions' "sweet spot" and can even work against the retailer by producing negative returns.

KEYWORDS

category selection, category-specific promotions, cross-category profit, targeted marketing

1 | INTRODUCTION

Targeted promotions are based on customer information and so enable retailers to maximize the impact of their marketing efforts (Rossi, McCulloch, & Allenby, 1996; Zhang & Krishnamurthi, 2004). Previous research has found that such marketing campaigns generate additional profits (Drèze & Hoch, 1998), prevent the alienation of loyal customers (Zhang & Breugelmans, 2012), and create lasting effects (Sahni, Zou, & Chintagunta, 2017). Category-specific promotions, or promotions that limit the customer benefit to just one product category, are widespread in retailers' loyalty reward programs

(LRPs). According to sample observations from June 2019, more than 30% of all received promotions in a nationwide LRP from three leading German grocery retailers are category-specific. This paper studies the category selection problem for these promotions from a multi-category retailer's perspective by addressing this practice-relevant question: which categories (eg, *wine*, *cheese*) should be emphasized in targeted promotions if the goal is to maximize the marketing campaign's profitability?

While the motivation for this research question originates in the context of grocery retailing, it is equally relevant for any online or offline multi-category retailer (eg, Amazon,

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Naval Research Logistics* published by Wiley Periodicals LLC.

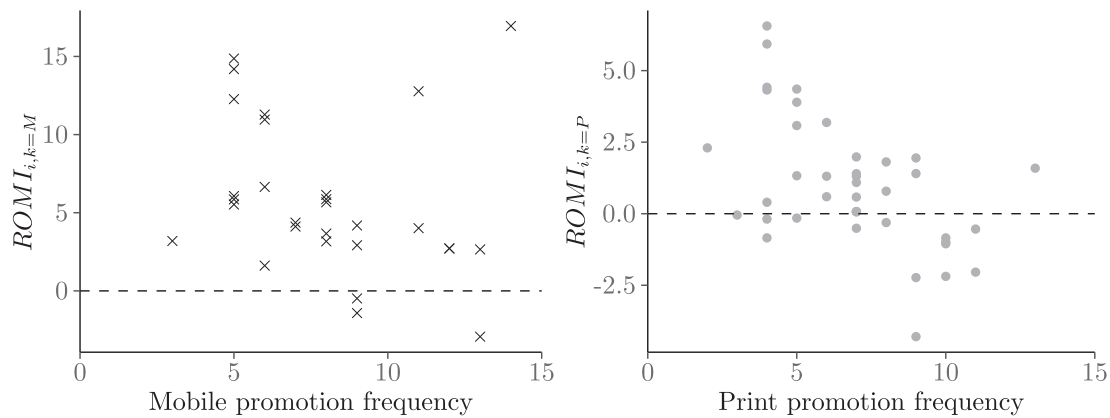


FIGURE 1 Profitability of category-specific promotions as a function of promotion frequency at this study's partnering retailer. Each data point represents a category; the underlying data set and costs are described in Section 3.3 and Section 4, respectively

eBay, JD.com, Walmart). In fact, already today all of these listed companies use category-specific (targeted) promotions. E-commerce players, thereby, have a significant head start. Due to the nature of their businesses, they have access to personalized customer information and, hence inherently have the opportunity to use targeted marketing and use the insights we provide in this paper. In contrast, traditional offline retailers must manage with tools such as a loyalty reward program to gather this “big data”.

The extant marketing literature does not adequately address our research question. Although researchers have examined customer responses to category-specific offers (Swaminathan & Bawa, 2005), investigated customer choice processes, and provided generalized recommendations for the related product selection problem (Ailawadi, Harlam, César, & Trounce, 2006; Dhar & Raju, 1998; Osuna, González, & Capizzani, 2016; Venkatesan & Farris, 2012), we are not aware of any published work on profit optimization models for the design of retailers' category-specific promotions. The data analytics research stream does provide the models proposed by Yang and Hao (2010) and Reutterer, Hornik, March, and Gruber (2017). Yet as we shall discuss, their models neglect cherry-picking effects and ignore such essential metrics as the redemption rate.

This problem has not been resolved in practice, either. We observe in the case of our partnering retailer, a major German hypermarket chain with both brick-and-mortar stores and a large online shop, that categories are selected either because they were frequently purchased by the targeted customers (at the segment level) compared to the entire customer base—that is, they were *distinctive* for them—or because the same promotion design performed well previously in generating a high increase in category spending (ie, “category spending lift” policy). A thorough search of the relevant literature yielded no other selection method that originated in practice. The approaches actually adopted fail to identify the most profitable categories; in fact, Figure 1 reveals that the frequently promoted categories at our focal retailer are those with low or even negative marketing profitabilities. This suboptimal category selection is due to managers who ignore not only the

relevant marketing key performance indicators (eg, the KPIs of redemption rate and marketing costs) but also the effects of cross-category spillovers.

In addition, the category spending lift policy is neither cost nor resource effective. To generate the required data, a retailer must empirically evaluate every customer segment's response to every category-specific offer. Overcoming this problem requires that marketing managers adopt an approach capable of predicting a category-specific promotion's profitability—even when that promotion has not yet been offered to a particular segment (or customer) or if the category was not promoted at all.

We address this research gap by developing a predictive customer-centric category selection approach based on the proven measure of *return on marketing investment*, or ROMI (Rust, Lemon, & Zeithaml, 2004). Our key objective is to understand the profitability drivers in the different promotion channels and to assess how current retail practice compares against it. We show that our approach can predict the ROMI's underlying variables with high out-of-sample accuracy when using only a limited data set of past segment-specific promotions as input. So in contrast to previous research and practice, our model can be scaled rapidly, deals with both fixed and variable promotion costs, incorporates cherry-picking effects, includes all of the retailer's category assortment, and can be deployed for the mass-market or (multiple) segments thereof.

The intuition behind our approach is the use of category-specific promotions to channel customer purchasing behavior toward high-profit baskets—in other words, so that customers will be incentivized to spend more in the promotion-redeeming transaction than they otherwise would. Following Chen, Hess, Wilcox, and Zhang (1999), we aim to select a category that also maximizes nonpromoted *cross-category profit*. The reason is that the profit uplift, or incremental profit, from this spending is “free” from the retailer's perspective (ie, it does not stress the marketing budget). Moreover, a category-specific promotion must achieve *redemption rates* high enough that the retailer can not only recover the fixed costs of contacting customers but

also reduce, as much as possible, its advertising costs in the promoted category (ie, *category spending*); the latter goal is a reflection of variable promotion costs being proportional to such category spending.

Analyzing the category selection problem requires that we make three assumptions. First, we consider the category selection decision independent of any other promotional design decisions (eg, targeted customers and promotion depth) and independent of previous promotions. That is, we assume those other design decisions to be pre-determined and equal for all categories and exclude those effects from past (category-specific) promotions.

Second, from a managerial viewpoint, it is crucial to maximize incremental spending, or the promotion's incremental profit (Bawa & Shoemaker, 1989). In an earlier study, Bawa & Shoemaker (1987b) show that most of the retailer's profit lift must stem from the transaction in which the promotion is redeemed. We therefore base our model on this profit lift only in promotion-redeeming transactions. The implication is that "exposure" effects (Bodapati, 2008; Venkatesan & Farris, 2012) and effects on future periods (Ailawadi et al., 2006; Sahni et al., 2017) are excluded. We assess incremental profit only for the single-category selection problem in a single-promotion setting. In practice, retailers can combine multiple categories for a special target group or events, and customers can redeem several promotions in one transaction. Yet all category selection approaches in these settings (multi-category, multi-promotion) and effects (exposure, future periods) are ultimately grounded in the single-category, single-promotion case of promotion-redeeming transactions—a connection that justifies the assumptions outlined here. In Appendix C, we outline how our approach can be extended to the multi-promotion case.

Third, we focus on determining the best category selection approach and leave any questions at the customer application level (mass-market, segment-specific, or personalized) open for future research (Zhang & Wedel, 2009). While the combinatorial challenge of selecting the right category for the right segment/customer can be significant, it represents a mere variant of the approach outlined in this paper. In short, all the analyses presented here address the mass-market level.

The rest of our paper proceeds as follows. Section 2 surveys the literature on customer-specific marketing and the category selection problem. In Section 3, we derive a customer-centric category selection approach and showcase our model's accuracy with respect to out-of-sample predictions. Section 4 explains, from the retailer's viewpoint, the profitability drivers of category-specific promotions in each of the two principal advertising channels: mobile and print. Then, in Section 5, we benchmark the relevant industry practice. We conclude in Section 6 with a summary of our results and some words on the outlook for future research.

2 | LITERATURE REVIEW

The classical marketing literature has studied various aspects of targeted marketing over the last decades (Drèze & Hoch, 1998; Khan, Lewis, & Singh, 2009; Sahni et al., 2017; Zhang & Krishnamurthi, 2004); yet research on the category selection problem has been scarce, and most studies cover the problem only in part. Venkatesan and Farris (2012) distinguish between *reward* and *cross-selling* coupons. A reward coupon promotes products that its recipients already buy frequently, whereas a cross-selling coupon is offered on products that the recipients rarely purchase. These authors analyze how the different coupon types affect shopping-trip redemption and exposure effects, but they are silent as regards how one should select the respective products or categories.

Dhar and Raju (1998), Ailawadi et al. (2006), and Osuna et al. (2016) are the first to provide such answers in the marketing literature for the product selection problem (eg, 10% discount off all *Coca-Cola* products). Ailawadi et al derive a net profit impact model for these offers and analyze the correlations among the factors of promotional type, brand, category, and store. Osuna et al. analyze a field experiment based on product-specific cash-out reward and cross-selling coupons; the authors offer general recommendations for product selection. Dhar and Raju similarly analyze the effects of cross-selling coupons on the sales of and profits derived from targeted brands. However, the category selection problem that we study has unique features (eg, no manufacturer subsidies, higher variable costs of promotions, lack of essential metrics on brand-specific items) that limit the application of models derived from the classical product selection problem.

The marketing scholars we have cited all employ explanatory models that use primarily brand, category, store, and promotion characteristics as input. However, these models are not applied to the retailer's entire category range—that is, so as to identify the overall most promising category (even if not yet promoted)—neither are the models' predictive accuracy verified for future promotions based on out-of-sample data. We shall address these issues and offer retailers a solution that uses the knowledge gained from past promotions to forecast the campaign profitability of all possible category-specific promotions.

A field of research related to category selection is that of *recommendation systems* for determining what item should be recommended to a customer. Such systems reflect either content-based or collaborative filtering approaches (Ricci, Rokach, Shapira, & Kantor, 2011). In the former, items are recommended as a function of the items' description and the customer's interest profile (Pazzani & Billsus, 2007); the latter approach recommends items enjoyed by similar customers (user-based method) or items similar to those the user is already known to like (item-based method); see Linden, Smith, and York (2003). Recommendation systems have been employed by researchers (eg, Bodapati, 2008) in the field of targeted marketing. Even though categories lack the

attributes (eg, color, size, brand) that item- and content-based approaches require, user-based collaborative filtering concepts can be adapted to the category selection problem. On the mass-market level, such adaptation involves promoting the most frequently purchased categories—which is similar to our partnering retailer’s approach of offering distinctive categories to different customers.

To the best of our knowledge, only two scholars have considered the stand-alone *category selection* problem: Yang and Hao (2010) and Reutterer et al. (2017). Each of these studies belongs to the data analytics research stream, and their contributions build on knowledge gained from studying the related *product assortment* problem (Brijs, Swinnen, Vanhoof, & Wets, 2004; Wong & Wang, 2005) and the data-mining methodology of “association rule mining.” The product assortment problem investigates what products to offer in the context of restricted retail space. Brijs et al. derive a profit-based product selection framework based on the “frequent itemsets” obtained from association rule mining. Their *PROFSET* model selects the best products while accounting for the promotional effects of complementary product relationships. Wong and Wang also address the product assortment problem; these authors propose a so-called *MPIS* model that, in essence, aims to minimize the profit contributions of products that are *not* assorted.

Building on the ideas of Brijs et al. (2004), Reutterer et al. (2017) develop a category selection framework for targeted promotions. They start by adopting a segmentwise approach in order to identify clusters of similar household composition. For each segment, they identify promising sets of categories that are purchased together. For each of these “category sets,” Reutterer et al. allocate profit by partially summing up the earnings of all transactions containing this itemset. A subsequent adjusted *PROFSET* optimization model identifies those θ categories that maximize overall profit under the constraint that, if a category set’s profit is unlocked, then all the categories contained in that set must be selected.

Yang and Hao (2010) employ a different model, one that is applicable to both the product and category selection problem. These authors ground their model on the product’s promotional effect, which has been studied extensively. Venkatesan and Farris (2012) find that customer-specific coupons for the categories *cereal* and *yogurt* have a positive effect on customer spending for a regional US supermarket chain, and Van den Poel, De Schampelaere, and Wets (2004) report similar results. Ailawadi et al. (2006) refer to this concept as the “halo” effect due to complementary products (Kumar & Leone, 1988); thus, for instance, a customer who buys bread will likely also purchase (say) butter or cheese. The spending based on these promotional effects is referred to as cross-category spending.

To quantify these effects for a single category, Yang and Hao (2010) calculate all promotional effects of the frequent categories (or items) and then determine their values as the

aggregate of those individual category or item effects. Thus the promotional effect of an item *A* on some other item *B* is defined as the *added value* (Tan, Kumar, & Srivastava, 2004) of this association—that is, the strength of the “ $A \rightarrow B$ ” rule *minus* the frequency with which *A* is purchased. In the end, those products with the highest added values are selected for promotion.

However, both the Yang and Hao (2010) and the Reutterer et al. (2017) category selection models have two downsides that inherently lead them to select suboptimal categories. First, they focus on the promotional effects of a category and ignore essential metrics (eg, the redemption rate, fixed costs of promotion). Second, these models assume similar customer behavior regardless of whether (or not) a shopping trip involves redeeming a promotion. Yet our results indicate that this assumption is invalidated by significant cherry-picking effects, which reduce the marketing campaign’s profitability.

We solve these issues by grounding our approach in the return on marketing investment (ROMI) financial performance indicator to address the first downside and, to address the second, by predicting the ROMI’s underlying components based on only a limited history of past promotions. It is surprising that the ROMI metric has not previously been used for the category selection problem—given that researchers have long recognized its utility (Blattberg & Deighton, 1996; Narayanan, Desiraju, & Chintagunta, 2004; Rust et al., 2004).

3 | CUSTOMER-CENTRIC CATEGORY SELECTION

To select the most profitable categories for promotion, we must first define how to *measure* profitability for category-specific promotions. The promotion schemes we study are category-specific points multiplier coupons (eg, 10× points on all *wines*) and category-specific discount coupons (eg, 5% off all *beers*). That is to say, we must not only identify which variables and parameters determine the return on marketing investment for each category but also understand how best to measure these variables accurately. For these purposes, we derive a ROMI equation based on past category-specific promotions. We can then use this knowledge when shifting our attention to the prediction of those variables for future promotions, which enables our development of a model for optimizing predictive category selection.

3.1 | Profitability of category-specific promotions

The notation we use to stipulate and analyze the ROMI equation is summarized in Table 1. The set of transactions is denoted $n \in N$; the set of categories, $i, j \in I$; and the set of promotion channels, $k \in K$. For each transaction, we know the amount spent by the customer in each of the categories and if a category-specific promotion was redeemed. If the latter is true, we further have the information on the promoted

TABLE 1 Notation for measuring the profitability of category-specific promotions

Sets	
$i, j \in I$	Categories
$k \in K$	Promotion channels
$n \in N$	Transactions
$n \in N_{ik}$	Transactions in which a category-specific promotion of i in channel k was redeemed
$n \in N'$	Nonpromotional transactions
Parameters	
C^f	Redemption-independent cost per coupon and customer
c^f	Fixed cost per coupon and customer in case of redemption
c^v	Variable cost margin per coupon
λ_i	Gross profit margin of category i
X_{ik}	Number of category-specific promotions of i distributed in channel k
s_{ni}	Spending in category i in transaction n
γ_{ik}	Redemption rate of a category-specific promotion of i in channel k
e_{ik}	Spending by customers in promoted category i in channel k
Φ_{ik}	Cross-category profit by customers on items from category i in channel k
E	Profit per nonpromotional transaction
$ROMI_{ik}$	Return on marketing investment for category i in channel k

category and the promotion channel. The spending per transaction and product category is written as s_{ni} , where $s_{ni} \geq 0$ for all $n \in N$ and $i \in I$. We split the set of all transactions N to obtain a set N' of *nonpromotional* transactions, or a set of transactions in which no promotions were redeemed. We also define sets of transactions N_{ik} in which *only* category-specific promotions of i in channel k were redeemed. Note that redeeming a promotion requires purchasing from a product category. It follows that $s_{ni} > 0$ for all $n \in N_{ik}$, all $i \in I$, and all $k \in K$.

To calculate the ROMI of a category-specific promotion, we need to measure the cost and the incremental profit that it generates. For a better interpretation, we calculate the $ROMI_{ik}$ as the average value per redeemed transaction; this interpretation is equivalent to the average value per customer if the promotion can be redeemed only once. Let us first consider the marketing costs, which have two per-customer components: a redemption-dependent part (eg, lost profit due to discounts, retailer spending to LRPs) and a redemption-independent part (eg, handling and mailing costs); the redemption-independent costs, denoted C^f , are constant for all customers.

The redemption-dependent cost can be divided into a fixed (ie, c^f) and a variable payment. This variable payment is a function of the customer's expenditure e_{ik} on the promoted category i in channel k , where such category spending $e_{ik} > 0$ for all $i \in I$ and all $k \in K$. This expenditure is defined as:

$$e_{ik} = \frac{1}{|N_{ik}|} \sum_{n \in N_{ik}} s_{ni} \quad \forall i \in I, \forall k \in K. \quad (1)$$

Therefore, e_{ik} (and hence Φ_{ik}) is conditional on the promotion stimulus and may differ from nonpromotional category spending e'_i . We address and quantify this “cherry-picking” effect in Section 3.2. Let c^v denote the variable cost margin and let γ_{ik} denote the redemption rate, where $\gamma_{ik} > 0$ for

all $i \in I$ and all $k \in K$. The redemption rate is defined in (2) as the ratio of the number of redemptions (ie, the cardinality of N_{ik}) over the number of promotions distributed in this channel (ie, X_{ik}). We shall derive $\gamma_{ik}(e_{ik}c^v + c^f)$ for the redemption-dependent costs.

$$\gamma_{ik} = \frac{|N_{ik}|}{X_{ik}} \quad \forall i \in I, \forall k \in K. \quad (2)$$

Category-specific promotions not only strain the retailers marketing budget, but they can also create incremental profits. To do so, the customers need to spend *more* in the promotion-redeeming transaction than they would otherwise. The incremental profit $\Delta Profit_{ik}$ consists of three parts, of which the first is due to the category spending e_{ik} multiplied by the category-dependent gross profit margin λ_i . The second part is the profit Φ_{ik} due to the spending in nonpromoted categories (ie, cross-category profit). This profit is associated with the promotional effect of category i ; that is, the nonpromoted categories are purchased *because* category i was promoted. Such cross-category profit is calculated as the average of all transaction profits excluding i but in which i was purchased:

$$\Phi_{ik} = \frac{1}{|N_{ik}|} \sum_{j \in I, j \neq i} \sum_{n \in N_{ik}} \lambda_j s_{nj} \quad \forall i \in I, \forall k \in K. \quad (3)$$

The third part of incremental profit is the average profit per transaction when no promotion was offered. This parameter, which we denote E , is defined as

$$E = \frac{1}{|N'|} \sum_{n \in N'} \sum_{i \in I} \lambda_i s_{ni}. \quad (4)$$

Thus incremental profit is the sum of $\lambda_i e_{ik}$ and Φ_{ik} net of E . The incremental margin is credited to *redeemed* promotions, which is why we multiply that margin by γ_{ik} . The $ROMI_{ik}$ of a category-specific promotion of category i in channel k is

defined as the incremental profit (a.k.a. profit uplift), net of the marketing costs, *divided by* these marketing costs. Hence we obtain the following equation for the $ROMI_{ik}$:

$$\begin{aligned} ROMI_{ik} &= \frac{\Delta Profit_{ik} - Cost_{ik}}{Cost_{ik}} \\ &= \frac{\gamma_{ik}(\lambda_i e_{ik} + \Phi_{ik} - E) - \gamma_{ik}(c^v e_{ik} + c^f) - C^f}{\gamma_{ik}(c^v e_{ik} + c^f) + C^f} \\ &= \frac{\gamma_{ik}(e_{ik}(\lambda_i - c^v) + \Phi_{ik} - E - c^f) - C^f}{\gamma_{ik}(c^v e_{ik} + c^f) + C^f} \\ &= \frac{\lambda_i e_{ik} + \Phi_{ik} - E}{c^v e_{ik} + c^f + \frac{C^f}{\gamma_{ik}}} - 1 \quad \forall i \in I, \forall k \in K. \end{aligned} \quad (5)$$

Statement (5) reveals three fundamental insights. First, as long as $\lambda_i > c^v$, the spending in promoted categories contributes to positive returns—though at the cost of reduced profitability. Second, the redemption rate γ_{ik} affects only C^f , the contacting cost per customer. If such costs are negligible (as in the case of mobile coupons), then $ROMI$ values are unaffected by the redemption rate. Thus the effectiveness of promotions in this setting depends only on e_{ik} and Φ_{ik} . If also the fixed cost c^f per coupon is negligible (as in the case of fully digital promotions), then the return on marketing investment depends only on the ratio $(\Phi_{ik} - E)/e_{ik}$. Third, in case $C^f > 0$, high $ROMI$ values require high redemption-independent costs C^f ; else C^f will dominate all other components in the ratio's denominator.

3.2 | Predicting marketing campaign profitability

So far we have presented a method for calculating the $ROMI$ of past category-specific promotions. As outlined in Section 1, our aim is to develop an approach that can be used to identify the most promising categories for future promotions in each promotion channel. This goal requires that one predict the $ROMI$ values underlying three components: $\hat{\gamma}_{ik}$, $\hat{\Phi}_{ik}$, and \hat{e}_{ik} . Mathematically, we have

$$\widehat{ROMI}_{ik} = \frac{\lambda_i \hat{e}_{ik} + \hat{\Phi}_{ik} - E}{c^v \hat{e}_{ik} + c^f + C^f / \hat{\gamma}_{ik}} - 1. \quad (6)$$

Prediction of the coupon redemption rate $\hat{\gamma}_{ik}$ has received considerable attention in the literature. Danaher, Smith, Ranasinghe, and Danaher (2015) find that the main determinants of the redemption likelihood are customers' coupon proneness (Bawa & Shoemaker, 1987a), their prior redemption history (Musalem, Bradlow, & Raju, 2008), and their category-specific usage rates (Swaminathan & Bawa, 2005). These usage rates are of particular interest to our problem. Swaminathan and Bawa (2005) model the odds of redeeming a category-specific coupon as a function of the customer's coupon proneness and the coupon's attractiveness; the authors use a customer questionnaire to quantify their model's parameters for the abstract notions of proneness and attractiveness.

Although customer surveys are a valuable source of knowledge about undifferentiated promotions, they yield only

limited insight for purposes of targeted marketing. Each such campaign requires a promotion design that is tailored to its intended targets. Generating new customer surveys for each receiver group is either costly and time-consuming or simply impossible owing to the lack of comparable panels.

Another straightforward approach to predicting redemption rates, category spendings, and cross-category profits is to field-test every possible category promotion and then to analyze empirically the $ROMI$ values that each promotion design generates in the redeemed transactions. However, time and cost constraints render this method infeasible for most retailers. Marketing managers (such as those at our partnering retailer) find themselves in a dilemma: they will not offer a new category because they do not know how it performed in the past; but then neither can they generate any historical data for that category because it has never been promoted.

Hence our goal is to develop a model capable of predicting, with a high level of accuracy, the underlying components of the return on marketing investment. We follow a two-step approach. First, we estimate regression equations using data from a few past category-specific promotions. That is, we use (sparse) promotional data to establish robust relationships between redemption rates and purchase frequencies and also between promotional and nonpromotional category spendings and cross-category profits. Second, we use the relationships determined in the first step—along with nonpromotional shopping trips, for which data are available to every retailer and for every category—to predict the $ROMI$ components of the retailer's category assortment. Using this method allows us to estimate the promotion profitability of categories, *even if* they have never been offered to customers, and to correct any bias that results from targeting only certain customer subgroups. As a consequence, information on just a few previous category-specific promotions is enough for a retailer to determine the $ROMI$ of its entire category range simply by reading the relationships in an “inverse” way. Our approach therefore provides marketing managers with a “fast to scale” means of prediction.

Thus we posit that the three variables of interest can be accurately determined if we use bivariate linear prediction models with the following independent variables:

- $\hat{\gamma}_{ik} = \beta_0^{\gamma k} + \beta_1^{\gamma k} pf'_i$, where pf'_i represents the customer's nonpromotional purchase frequency of items in category i ;
- $\hat{e}_{ik} = \beta_0^{ek} + \beta_1^{ek} e'_i$, where e'_i is the customer's nonpromotional category spending; and
- $\hat{\Phi}_{ik} = \beta_0^{\Phi k} + \beta_1^{\Phi k} \Phi'_i$, where Φ'_i denotes the customer's nonpromotional cross-category profit on category i .

The intuition behind the first prediction model (ie, $\hat{\gamma}_{ik}$) is that if a category is purchased often in nonpromotional shopping trips, we expect that customers redeem promotions of this category more often. The reasoning behind the second and third model is similar: if customers spend much in a category (or in cross-category products) in nonpromotional transactions, why should they not also in promotion-redeeming

transactions? Let $N'_i = \{n \in N' | s_{ni} > 0\}$. Then we can formalize these nonpromotional variables:

$$pf'_i = \frac{|N'_i|}{|N'|} \quad \forall i \in I, \tag{7}$$

$$e'_i = \frac{1}{|N'_i|} \sum_{n \in N'_i} s_{ni} \quad \forall i \in I, \tag{8}$$

$$\Phi'_i = \frac{1}{|N'_i|} \sum_{j \in I, i \neq j} \sum_{n \in N'_i | s_{nj} > 0} \lambda_j s_{nj} \quad \forall i \in I. \tag{9}$$

Ailawadi et al. (2006) and Osuna et al. (2016) report evidence of additional category characteristics (eg, storability, number of stock-keeping units, sales share of private labels) having explanatory power for their models of coupon redemption and incremental sales. Therefore, we tested if including more independent variables (in addition to those mentioned above: mean price and promotion frequency) by means of multiple linear and Lasso regression models further increased the already high predictive accuracy (as we shall establish) of the basic model. Appendix A reports the details of this analysis. While the in-sample performance metrics (R^2 and adjusted R^2) increased, the out-of-sample KPIs (WMAPE and MAPE) worsened. We attribute this behavior to overfitting. In consequence, we base our predictions solely on pf'_i , e'_i , and Φ'_i .

3.3 | Empirical study for estimating $\hat{\gamma}_{ik}$, $\hat{\Phi}_{ik}$, and \hat{e}_{ik}

To demonstrate the possibility of accurately predicting these determinative values based on already a small set of past promotions, we employ a data set provided by a leading German hypermarket chain. This partnering retailer’s targeted marketing vehicle is the issuance of coupons for its external loyalty reward program (cf. Zhang & Breugelmans, 2012), and about half of its customers are members of this program. Our data set consists of two parts: (a) all of the focal retailer’s targeted promotions via its LRP from January to July 2018, including a list of all the receivers and redeemers of those offers; and (b) the complete transaction history of a sample consisting of 100 000 customers—randomly selected from a pool of millions—who used the retailer’s LRP from January 2017 until July 2018. Overall, we account for more than 6 million transactions and 93 product categories that range from *wine* to *cosmetics* to *sports gear*. Among these categories, four are “historical” categories in which items are purchased at frequencies of less than 0.001% (ie, sales of remaining stock).

We start by describing the setup of our numerical study in Section 3.3.1 and by providing descriptive statistics in Section 3.3.2. In Sections 3.3.3 to 3.3.5 we train our prediction models and assess their predictive accuracy using out-of-sample test in Section 3.3.6. All subsequent analyses are performed in the programming language R (version 3.6.3).

3.3.1 | Setup of the numerical study

Analysis of sub-channel. We distinguish between our retailer’s two main promotion channels in its loyalty reward program.

The first is a *direct marketing mailing* in which the targeted customer receives only the retailer’s marketing effort; the second channel is the monthly *points-overview mailing* (mobile and print) of the LRP. Together with account balance information, the customer receives several coupons from all participants in the program (also direct competitors of the retailer). The presence of competitors entails that, in this case, the retailer’s customer contacting cost is but a fraction of the costs associated with its direct mailing channel. Our partnering retailer has the most experience with category-specific promotions in the second (points-overview mailing) promotion channel, with a category-specific promotion share exceeding 75%. Hence we analyze the customer response to category-specific coupons in this channel only. All these promotions are points multiplier coupons and have a 10× promotion “depth.”

Flat gross profit margin. Senior marketing executives highlighted to us that the calculation of category-dependent gross profit margins is subject to various assumptions and personal assessments. In practice, each category manager fine-tunes his or her λ_i so that these values are barely comparable. To avoid these complications, our partnering retailer deliberately sets a flat profit margin of about 30%. We follow this best practice in our empirical study and set $\lambda_i = 0.3 \quad \forall i \in I$. Appendix B shows that using nonflat gross profit margins does not worsen the prediction accuracy of our models.

Dealing with the multi-category and multi-promotion setting. In the past, our partnering retailer offered both single- and multi-category promotions. To consider all categories equivalently, we analyze its promotions at the most granular level: the single-category level. We determine a category’s redemption rate in a multi-category promotion by dividing that category’s number of promotional purchases by the number of coupons distributed. For the cross-category profit, we exclude for the dependent and independent variable those categories that were offered in the same promotion. To avoid conflating category and cross-category spending effects, which could result when multiple coupons are combined during a given shopping trip, we examine only the single-coupon redemption case when predicting Φ'_i and e'_i . Appendix C shows how extending our models to the multi-promotion setting impacts our empirical results.

Aggregation. We consider promotions at the aggregated level of unique promotional text, date, and channel. In order to account for different customer target groups, we determine the independent variables pf'_i , Φ'_i , and e'_i separately for each promotion by analyzing the nonpromotional purchasing behavior of only the targeted customers. Because customer habits change over time, we study only those transactions occurring within the 12-months *prior* to a promotion. We then aggregate all promotions of the same category to a single value by computing the mean of both the independent and dependent variables.

Robustness. For the sake of robustness, we restrict the analysis to promotions that more than 1,000 customers

received and we define a minimum threshold of 30 (single promotion-redeeming) transactions for each promoted category to calculate Φ_{ik} and e_{ik} . To avoid outliers distorting our prediction models, we omit the minimum and maximum values of γ_{ik} , Φ_{ik} , and e_{ik} when fitting the coefficients.

3.3.2 | Descriptive statistics

We obtain 307 individual category-specific promotions—including 32 unique promotion designs (eg, “10× points on all *frozen foods*”)—that were distributed in more than 4.9 million instances and were redeemed in some 140,000 transactions. These 32 unique promotion designs cover 75 product categories. Because our partnering retailer has applied its category selection scheme at the customer-segment level, it has already implemented a fairly large number (32) of different promotion designs. Each of the retailer’s 10 customer segments was targeted with an average of 8.1 promotion designs. It follows that each of the two category selection strategies outlined in Section 1 is responsible for $32/8.1 = 3.95 \approx 4$ promotion designs per segment.

For the 140,000+ category-specific promotion-redeeming transactions, mean customer spending is €57.9 (excluding bottle deposits and sales commissions); on average, customers purchase items from 9.5 categories. For the retailer’s 28,823 single promotion-redeeming transactions, the corresponding numbers are smaller: €43.3 and 7.3 categories. Mean spending is lower still in nonpromotional shopping trips, during which customers spend a mean of €33.5 and purchase from only 6.2 categories on average. The difference between promotion-redeeming and nonpromotional transactions—in terms of the size and value of market baskets—indicates that customer shopping is more “mindless” when redeeming coupons in the sense that purchase decisions are then guided less by informational considerations than by affect (Aydinli, Bertini, & Lambrecht, 2014). The average “hit rate” of coupons in our data set is 4.2%, although individual values range from 0.0019% to 24.5%.

We avoid overfitting by using the 222 promotions (72.3% of the total) initiated before June 10, 2018, as the training set and then validating the parameterization so obtained on the other 85 promotions (27.7%) until the end of July 2018. In Section 3.3.6 we present out-of-sample results for a training-testing data split based on promotion frequency; that is, we train the models using only categories that feature frequent promotions but test those models on categories that are promoted at a low frequency.

3.3.3 | Redemption rate $\hat{\gamma}_{ik}$

The uppermost plot in Figure 2 displays results based on the training data set. A strong linear correlation between the category purchase frequency and the redemption rate is observable for each of the promotion channels: $r_{pf'_i, \gamma_{i,k=M}} = 0.961$ and $r_{pf'_i, \gamma_{i,k=P}} = 0.879$ for the mobile and print channels, respectively (see Table 2). This graph also plots results from

equations of the linear models fitted by ordinary least-squares (OLS) linear regression. The regression’s high values of $R^2_{k=M} = 0.924$ and $R^2_{k=P} = 0.773$ indicate that the data confirm our hypothesis. These results highlight two characteristics of the promotion channels. First, they show that print promotions achieve approximately 74% higher redemption rates than do their mobile counterparts. Second, the y-intercepts of both linear regression equations are small in comparison with the slope parameters (0.001 vs 0.111, 0.002 vs 0.19). As a consequence, doubling the purchase frequency likewise nearly doubles the redemption rate. These slope parameters (and those for $\hat{\Phi}_{ik}$ and \hat{e}_{ik}) are statistically significant; see the *p*-values reported in Table 2.

3.3.4 | Cross-category profit $\hat{\Phi}_{ik}$

For both the print and mobile channels, we find a strong correlation between the cross-category profit per category in nonpromotional transactions and the cross-category profit in promotional transactions; see Table 2. The linear regression equations (fitted by OLS regression) reported in the middle graph of Figure 2 reveal that promotional cross-category profit is lower than nonpromotional cross-category profit: for €12, it is 6.1% lower in the mobile channel and 24.8% lower in the print channel. These outcomes are indicative of cherry-picking behavior: shoppers reducing their expenditures in nonpromoted categories. Such behavior is more prevalent in the print channel. So in comparison with the mobile channel, the print channel features customers redeeming more coupons but purchasing less from nonpromoted categories.

3.3.5 | Category spending \hat{e}_{ik}

Table 2 confirms also that customers’ category spending in nonredeeming transactions and in promoted categories are highly correlated: $r_{e'_i, e_{i,k=M}} = 0.901$ and $r_{e'_i, e_{i,k=P}} = 0.981$. The lowermost graph in Figure 2 plots the linear equations fitted by OLS regression. This graph shows that category spending is significantly higher when redeeming a category-specific promotion. Thus, for example, $\hat{e}_{i,k=M}$ is 36% (resp. 29.5%) higher than is e'_i in the mobile (resp. print) channel when $e'_i = €10$. This effect confirms our intuition that customers seek to maximize their benefit—namely, loyalty reward points—from the promotion (eg, via stockpiling) by buying more than usual from the offered category. However, this cherry-picking effect need not reduce the marketing campaign’s profitability; if steered intelligently, the effect can even be exploited to increase customer spending.

3.3.6 | Out-of-sample validation

In-sample accuracy measures are prone to overfitting. We are therefore motivated to assess the accuracy of our prediction models based on the testing data set by using the linear regression equations displayed in Figure 2. Given that

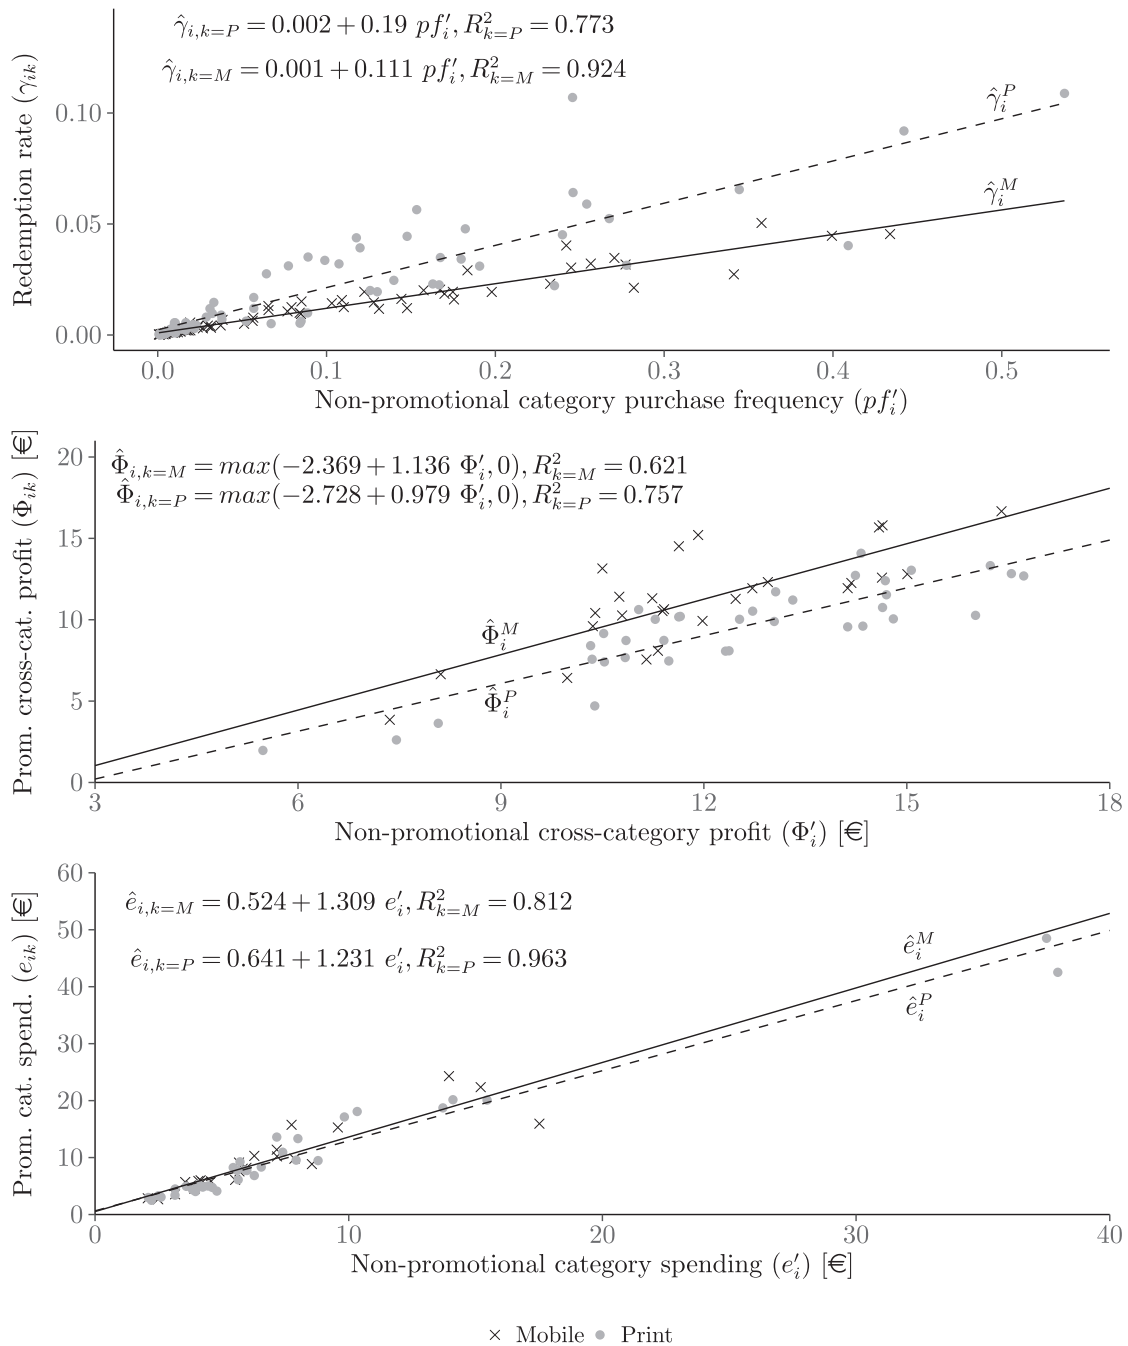


FIGURE 2 Prediction of redemption rate, promotional cross-category profits, and promotional category spending for category-specific promotions based on training data sets for the mobile and print promotion channels. The displayed values are based on the training data set that was split from the overall data set by date (ie, all category-specific promotions between January and June 10, 2018). In each plot and for each regression, the dependent variable's respective minimum and maximum values are omitted

redemption-rate values of zero will be involved, we use the *weighted mean absolute percentage error (WMAPE)* metric—also known as the MAD/mean ratio (Kolassa & Schütz, 2007). In addition, we report the out-of-sample *mean absolute percentage error (MAPE)* for our predictions of $\hat{\Phi}_{ik}$ and \hat{e}_{ik} .

Table 3 presents the out-of-sample results. Let us first focus on the testing data set that was split from the overall data set by date (viz., all promotions after June 10, 2018; see the sixth paragraph of Section 3.3). We achieve *low* WMAPE and MAPE values for both promotion channels, which indicates

that our predictions are accurate. The lowest mean absolute deviation percentages (17.9% and 16.5%) are reported for the cross-category profit prediction. The redemption rate prediction has a WMAPE of 25.6% in the mobile channel and 38.5% in the print channel, where the latter is the highest WMAPE value we find. For \hat{e}_{ik} , the WMAPE and MAPE scores are more accurate in the mobile than in the print channel; the respective scores are 23.8% and 29.2% (WMAPE) and 20% and 27.4% (MAPE).

In what follows, we test for whether this approach is robust to a different training-testing data split. Recall that our goal

TABLE 2 Correlations for γ_{ik} , Φ_{ik} , and e_{ik} ; P -values for OLS regressions

Correlation	Channel		Mobile channel			Print channel	
	Mobile	Print	P -value	Intercept	Slope Para.	Intercept	Slope Para.
$r_{p_i', \gamma_{ik}}$	0.961	0.879	$\hat{\gamma}_{ik}$	0.090	<0.001	0.216	<0.001
$r_{\Phi_{ik}', \Phi_{ik}}$	0.788	0.870	$\hat{\Phi}_{ik}$	0.294	<0.001	0.031	<0.001
$r_{e_{ik}', e_{ik}}$	0.901	0.981	\hat{e}_{ik}	0.594	<0.001	0.174	<0.001

TABLE 3 Out-of-sample validation of prediction models for training-testing data sets split by *date* and *promotion frequency*

	Data sets splits by date				Data sets splits by prom. Freq.			
	WMAPE		MAPE		WMAPE		MAPE	
	Mobile	Print	Mobile	Print	Mobile	Print	Mobile	Print
$\hat{\gamma}_{ik}$	25.6%	38.5%	—	—	19.1%	42.0%	—	—
$\hat{\Phi}_{ik}$	17.9%	16.5%	20.2%	22.1%	18.8%	15.7%	23.6%	16.8%
\hat{e}_{ik}	23.8%	29.2%	20.0%	27.4%	28.3%	27.9%	22.0%	20.5%

is to study the promotional customer response to only a few frequently promoted categories before predicting customer response to the remaining categories (ie, those that are not frequently promoted). We therefore use the 40 most frequently promoted categories, which account for 68% of all promotions in our data, as the training set and then evaluate the fit on the other 35 categories. We find that the reported out-of-sample performance (ie, Table 3's right side) is similar to that observed for the training–test data set split by date. Individual values change, of course, but the average of all WMAPE values is virtually identical in both cases (25.28% vs 25.3%). Although our models are simple and consist only of a linear function with one input parameter, it is clear that this approach yields a good fit to the data. Furthermore, it can predict customer response to new promotions and also to new categories—with high out-of-sample accuracy—even when relying on a relatively sparse data set of past promotions.

4 | PROFITABILITY DRIVERS OF CATEGORY-SPECIFIC PROMOTIONS

Having shown that the regression equations yield trustworthy predictions, we can now generalize our findings by using Figure 2's regressions to derive the target parameters of the entire category assortment—among those categories never offered to customers—while using abundant nonpromotional data. This approach helps us understand the profitability drivers of category-specific promotions in loyalty reward programs. We aim to describe the fundamental relationships between Equation (5)'s variables and the ROMI associated with the channels of mobile and print promotion. This knowledge would give multi-category retailers some guidance on how to direct their marketing budgets and optimize their promotions.

To achieve these goals, we turn to the data set provided by our partnering retailer and assume that an LRP-based marketing campaign (a) begins on June 10, 2018, and (b) targets all 100,000 customers in our sample using both the mobile and print distribution channels. For the reasons given in Section 3, for our model's input we use the nonpromotional transactions of all customers between June 11, 2017, and June 10, 2018—a total of 2,989,917 shopping trips. In these nonpromotional transactions, the average spending equals €32.81. Hence $E = 0.3 \times €32.81 = €9.84$. We ensure the robustness of this section's analysis by imposing a minimum purchase frequency of 0.001% in nonpromotional transactions for all categories.

In the loyalty reward program, customers receive a point for every €2 they spend and so 100 loyalty reward points are worth €1. Hence the 10× points multiplier that we consider is equivalent to a 4.5% discount, which means that $c^v = 0.045$. In the mobile distribution channel, the LRP charges the retailer a fixed amount of $c^f = €0.25$ if the coupon is redeemed; in the print channel, the retailer is charged a fixed customer contacting cost of $C^f = €0.0275$.

4.1 | Mobile promotion channel

Figure 3 plots predicted ROMI values for the mobile channel that are based on the regression equations from Figure 2. The middle graph reveals a strong positive correlation between cross-category profit and the return on marketing investment ($r_{\Phi_{ik}', \widehat{ROMI}_{i,k=M}} = 0.859$). Therefore, selecting categories with the highest Φ_{ik}' values is a simplified selection scheme with high ROMIs in the mobile channel. The rationale behind this profitability driver is the retailer's “free” profit lift from cross-category profit since it involves no variable costs. Note also that high cross-category profit curtails cherry-picking, which in turn increases the profitability of marketing promotions.

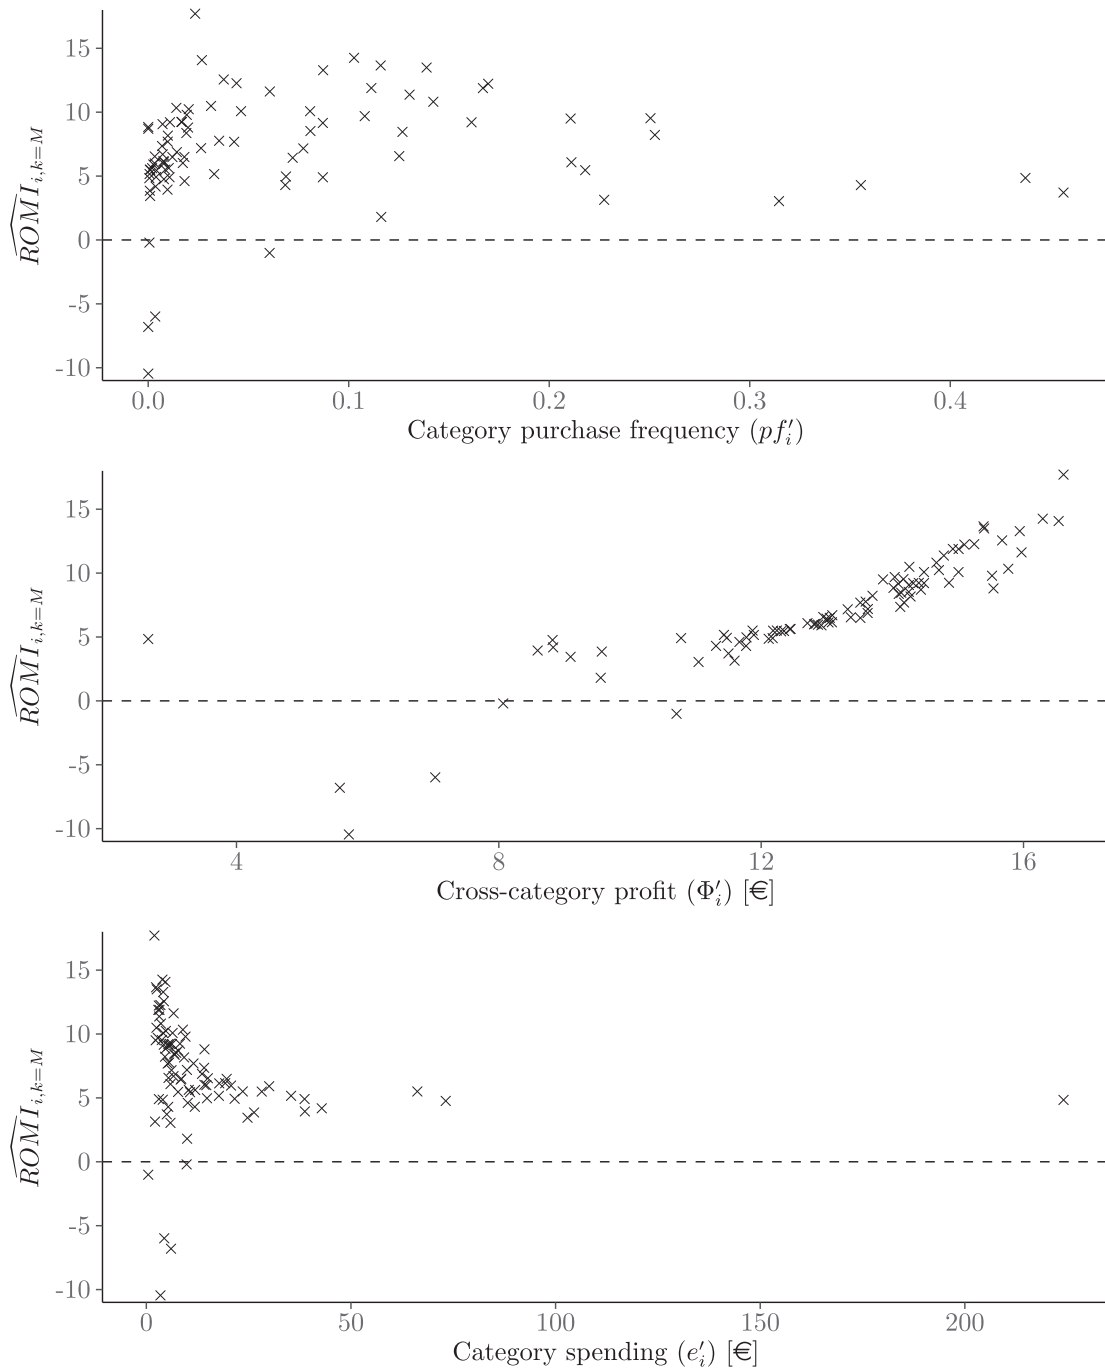


FIGURE 3 Mobile channel: value of $\widehat{ROMI}_{i,k=M}$ as a function of category purchase frequency pf'_i , nonpromotional cross-category profit Φ'_i , and nonpromotional category spending e'_i

However, cross-category profit is only one component of the ROMI equation. The lowermost scatter plot in Figure 3 indicates that low spending on the promoted category increases marketing profitability ($r_{e'_i, \widehat{ROMI}_{i,k=M}} = -0.162$). This negative correlation is intuitive because low category spending is directly related to low (variable) marketing costs. So in the mobile channel, the key to category-specific promotion profitability is selecting a category with a high Φ'_i value and a low e'_i value.

In this channel, the category purchase frequency (and its dependent variable, the redemption rate) are barely correlated with marketing profitability ($r_{pf'_i, \widehat{ROMI}_{i,k=M}} = 0.074$). In

contrast, the uppermost plot in Figure 3 shows that the mid- and low-redemption categories are the most promising ones for promotion. The low correlation between ROMI and both the purchase frequency and the redemption rate results from the zero fixed contacting cost per customer in the mobile channel. In this case, a promotion's hit rate plays no part in our expression for profitability; see Equation (5).

4.2 | Print promotion channel

If the redemption-independent contacting costs per customer are relevant, as they are in the print channel, then our




FIGURE 4 Print channel: value of $\widehat{ROMI}_{i,k=P}$ as a function of category purchase frequency pf'_i , nonpromotional cross-category profit Φ'_i , and nonpromotional category spending e'_i

understanding of the mobile channel's profitability drivers cannot be transferred in a straightforward way. Consider, for instance, the category purchase frequency (ie, the predicted redemption rate). The uppermost graph in Figure 4 shows that higher purchase frequencies tend to imply higher ROMI values, although only up to a certain purchase frequency boundary (the vertical dotted line). In our application example, that boundary appears at a purchase frequency of 28%. Until this point, ROMI and the purchase frequency exhibit a correlation of 0.596 (versus a correlation of 0.348 for all purchase frequencies). The reason is that the nonzero contacting cost necessitates a certain level of redemption so that the retailer

can recoup the redemption-independent cost. The higher the C^f , the stronger the correlation—until that cost dominates all other profitability drivers.

For purchase frequencies higher than 28%, the ROMI is either near zero or negative. Negative ROMI values for high-frequency categories are caused by the reduced cross-category profits associated with higher purchase frequencies (ie, the redemption rate). The more frequently that items from a category are bought, the lower its cross-category profits; see Figure 5. This figure's scatter plot establishes that, from a purchase frequency level of 2% onward (the vertical dotted line), nonpromotional cross-category profits reaches

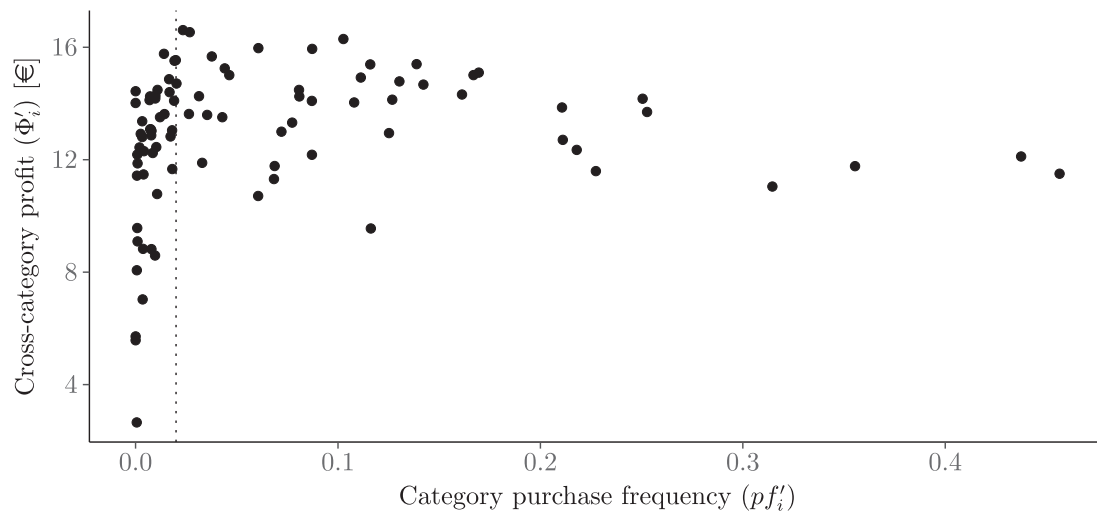


FIGURE 5 Cross-category profit Φ'_i as a function of category purchase frequency pf'_i for $\Phi'_i > \text{€}4.5$

its maximum at nearly €16.5. For higher category purchase frequencies, the Φ'_i scores decline to about €11.4. In combination with the stronger cherry-picking effect in the print channel, these low cross-category profit levels approximate (or fall short of) the average spending in nonpromotional shopping trips. The result is nearly zero or outright negative profit increases, which translate directly into negative ROMIs. In sum: a category's redemption rate must be high enough to recoup the redemption-independent cost yet still low enough to achieve a high cross-category profit. Because the trade-off between these two effects differs from one application to the next, promoting only the most frequently purchased items is not a viable option.

The middle graph in Figure 4 shows that, much as in the case of the mobile channel, high cross-category profit increases ROMI ($r_{\Phi'_i, \widehat{ROMI}_{i,k=P}} = 0.373$) in the print channel as well. However, the correlation is weaker here because the purchase frequency has a stronger effect (as explained previously). In contrast to the mobile channel, there is only a weak positive correlation in the print channel between category spending and ROMI ($r_{e'_i, \widehat{ROMI}_{i,k=P}} = 0.168$). This observation can be explained by the print channel's greater extent of cherry-picking behavior, for which the retailer compensates by higher category spending (provided $\lambda_i > c^v$; see Equation (5)). However, such increased spending entails a higher variable cost and lower overall promotion profitability.

We find that the ROMI varies significantly depending on the channel, categories selected, and type of coupons offered. The mobile promotion channel is more profitable than the print channel, and the former's high ROMI values reflect the low fixed costs ($c^f = \text{€}0.25$) of mobile points-overview mailings as well as higher category spendings and cross-category profits. Yet despite the print channel's relatively lower profitability, there could be strategic advantages to using it for promotions. For instance, it is unclear whether print coupons are necessary to trigger the redemption of mobile coupons—a situation akin to multi-channel shopping experiences. Assessing

and quantifying this effect is a promising avenue for further research.

In Appendices B and C, we test how extending our numerical study to include nonflat gross profit margins and the multi-promotion setting impacts our empirical results. Overall, we find that our key insights hold for both extensions. While the inclusion of category-dependent gross-profit margins results in almost identical correlations, the category spending in the print channel correlates strongly with the ROMI in the multi-promotion setting. We refer to Appendix C for a discussion and explanation of this phenomenon.

5 | OPTIMAL CATEGORY SELECTION AND RETAIL PRACTICE BENCHMARKING

In this section, we undertake a benchmark analysis to compare two common strategies in retail practice outlined in Section 1 (viz., the *distinctive categories* and the *category spending lift* policies) against the ROMI-maximizing category selection approach for each promotion channel. We also assess how the optimal ROMI is affected if the same categories in the mobile and print promotion channels must be selected—an extension motivated by retail executives who prefer to enforce consistency across channels. In the next paragraphs, we describe these four category selection approaches in more detail.

5.1 | Category selection approaches

ROMI-maximizing category selection (ROMI-max CS). This most profitable category selection regime aims to select the category with the highest \widehat{ROMI}_{ik} value for each promotion channel. However, if only one category per channel is promoted repeatedly, over time, customers will become accustomed to the “promotion stimulus.” Therefore, marketing managers want a list of several category suggestions for each channel from which they can alternate their choice of

high-profit categories to counteract the habituation effect. We model this requirement in the ROMI-maximizing category selection (ROMI-max CS) optimization model as follows:

$$\text{ROMI} - \max \text{CS} : \text{maximize } \frac{1}{\theta} \sum_{k \in K} \sum_{i \in I} y_{ik} \widehat{\text{ROMI}}_{ik} \quad (10)$$

$$\text{subject to } \sum_{i \in I} y_{ik} \leq \theta \quad \forall k \in K, \quad (11)$$

$$y_{ik} \in \{0, 1\} \quad \forall i \in I, \forall k \in K. \quad (12)$$

Thereby, constraint (11) ensures that the optimization model selects those θ categories per channel so that the average profitability of these categories in each channel is maximized in the objective function (10). The binary decision variable y_{ik} is set equal to 1 if category i is selected for promotion in channel k (and is set to 0 otherwise); see the constraint (12). The conditional inequality (11) ensures that only categories with positive returns are selected. The ROMI-max CS can be solved independently for each promotion channel, as the objective function is additive and a channel-linking constraint is absent. For the resulting knapsack problem for each channel, the following greedy procedure returns the optimal solution (as the item's weights are equal): first rank all categories according to $\widehat{\text{ROMI}}_{ik}$, then remove those categories with $\widehat{\text{ROMI}}_{ik} < 0$, and finally select the top θ categories.

Channel-synchronized ROMI-max CS. The ROMI-max CS chooses the best categories independently for each promotion channel. However, if we include constraint (13), then the ROMI-max CS can be forced to select the same categories in all promotion channels.

$$y_{ik_1} = y_{ik_2} \quad \forall i \in I, \forall k_1, k_2 \in K. \quad (13)$$

To solve the resulting “channel-synchronized” ROMI-max CS, $\widehat{\text{ROMI}}_i$ is computed and the same greedy solution procedure as in the previous paragraph is applied. Thereby,

$$\widehat{\text{ROMI}}_i = \frac{\sum_{k \in K} (\hat{\gamma}_{ik} (\hat{e}_{ik} (\lambda_i - c^v) + \hat{\Phi}_{ik} - E - c^f) - C^f)}{\sum_{k \in K} (\hat{\gamma}_{ik} (c^v \hat{e}_{ik} + c^f) + C^f)} \quad \forall i \in I.$$

Category spending lift policy. For the *category spending lift* policy, we determine the categories in the testing data set that achieved the highest relative increase in category spending. Thus we rank all categories according to their e_{ik}/e'_i ratio and select the top θ categories per promotion channel.

Distinctive categories policy. Recall that, at the mass-market level, the most distinctive categories are those from which items are most frequently purchased (an outcome identical to that of the adapted user-based collaborative filtering approach mentioned in Section 2). For the *distinctive categories* policy we select, from the testing data set, the θ categories with the highest purchase frequency pf'_i .

5.2 | Results

We treat the promotion channels independently and the benchmarking is based on the ROMI metric. Because we

do not have out-of-sample performance data for all product categories, the benchmarking is based on the *predicted* return on marketing investment values via the regression equations presented in Section 3.2. In our empirical study, we find that, on average, four categories are promoted per category selection regime. Hence we set $\theta = 4$. In other words, we calculate the average ROMI value ($\overline{\text{ROMI}}_k$) of the four most promising categories for each of the four benchmarking approaches.

Table 4 reports our results for both the mobile and print channels, which we compare by expressing all values as percentages of the highest value in the focal column. Including constraint (13) significantly reduces the mobile channel's profitability to 80.2% of its unconstrained value whereas the print channel's category choice is unaffected. This asymmetrical effect of the channel-synced ROMI-max CS stems from the print channel's higher absolute costs, which dominate the mobile channel's lower (yet more profitable) investments. That difference is driven, in turn, by the print channel's higher redemption rates (ie, higher variable cost) and higher redemption-independent costs. Therefore, retailers employing a channel-synchronized strategy actually consider only the print channel and ignore the mobile channel's potential. Abstracting this finding to the online and offline context, one could subsume that focusing too much on brick-and-mortar marketing holds retailers back from leveraging the upside of their digital channels.

In either channel, our focal hypermarket chain's basic distinctive categories policy fails to determine the most profitable categories: in the mobile channel it achieves only 26.6% of the maximal profitability; in the print channel it even results in a negative ROMI of -15% ; see Table 4's fourth row. This finding reflects the mobile channel's lack of correlation between ROMI and the purchase frequency as well as the print channel's negative ROMI values for high purchase frequencies.

The category spending lift policy gives the second-best results in both the mobile channel (40.4%) and the print channel (9.3%)—although these values are substantially lower, especially in the print channel, than the profitabilities enabled by our approach. The reason is that category spending is low in both channels, as described in Sections 4.1 and 4.2. Yet marketing managers who instead select categories that offer the greatest cross-category incremental profit (ie, categories with the highest values of Φ_{ik}/Φ'_i) would achieve better performance (62% in mobile and 29% in print). There is not much point in adopting such “lift” policies, however, because the simple strategy of selecting categories with the highest cross-category profit performs the best among all simplified policies (82.9% in mobile and 74.6% in print).

None of the commonly used category selection strategies hits the profitability function's sweet spot in both the mobile and print promotion channels. In particular: the mobile channel's category spending lift policy runs the risk of low

TABLE 4 Average ROMI values for the top four promotion categories in each framework: mobile channel versus print channel

	Mobile channel		Print channel	
	$\overline{ROMI}_{k=M}$	%	$\overline{ROMI}_{k=P}$	%
This paper				
• ROMI-max CS	14.92	100	2.21	100
• Channel-synced ROMI-max CS	11.97	80.2	2.21	100
Retail practice				
• Category spending lift	6.03	40.4	0.21	9.3
• Distinctive categories	3.98	26.6	-0.15	-6.8

Note: The “%” column reports each value as a percentage of the “ \overline{ROMI}_k ” column’s highest value (eg, 40.4% = 6.03/14.92).

profitability; and the print channel’s distinctive category approach yields an outright negative outcome (−6.8%). Thus extant implementations of category selection schemes can, in fact, result in promotions that undermine the retailer’s profitability.

6 | CONCLUSION AND FUTURE RESEARCH

This study details a new approach to the category selection problem and the first that targets promotions based on the metric of return on marketing investment. We document how the components of ROMI equations can be predicted, at high levels of out-of-sample accuracy, with reference to only a small set of past LRP promotions; we also show that prevailing category selection strategies can backfire on the retailer. For mobile promotions we find the most profitable categories to be those that exhibit the highest cross-category profits. For print promotions, the most profitable categories are those from which items are purchased at a high (but not too high) frequency. Since our data set’s distinctive features (eg, constant promotion depth, values of redemption-[in]dependent costs) only affect the strength of the correlations discussed in Section 4, not their direction, the empirical findings reported here can be generalized.

In addition to satisfying academic curiosity, we provide practitioners with a fast and easy-to-scale approach for the category selection problem even in cases where only a few categories have been promoted so far. Our research shows marketing managers the importance of *category selection* in any promotional campaign: it is an element that can channel customers’ purchase behavior toward higher basket values and greater promotion-based profits. Hence retailers should embrace the opportunities offered by category selection criteria and thereby maximize the impact of their marketing budgets.

Our approach could be made even more generalizable by relaxing the model’s simplifications (ie, those mentioned in Section 1’s penultimate paragraph). Three promising extensions of our work relating to these relaxations can be briefly

described as follows. First, extending our model to different promotion depths (eg, 3×, 5×, 20×) would allow offering different categories with different treatments (eg, 10× on sports shoes, but 20× on wines) to, for instance, increase redemption rates of categories with low \hat{e}_i in the print channel, as the higher returns (due to higher hit rates) could outweigh the increased variable marketing cost. However, this would require knowing these offers’ associated hit rates and (cross-)category spending levels. If one varies the depth of a few promotions, then the same approach as described here could be used to integrate promotion depth into the regression’s independent variables. As before, it requires only a small number of past promotions for retailers to predict customer response—here, to all possible multiplier coupons.

Second, extending the ROMI equation to accommodate temporal effects in a dynamic promotion setting, would require new approaches to distinguish each coupon’s effect. Formulating the ROMI-max CS dynamically allows optimizing the timing and sequence of the θ selected categories over time—a further step toward automated promotions.

Third, scaling our model to the fully personalized (or segment) level would lead scholars to additional topics of interest: customer-specific redemption rates and (cross-)category expenditures as well as management of increased marketing complexity. In that case, certain customers might not be offered any category-specific promotions, due to negative returns, while for others up to θ categories will be identified. If budget constraints are present, this results in an np-hard optimization problem that is difficult to solve for large customer bases. Providing fast exact algorithms or near-optimal heuristics for this problem would provide a significant contribution to the literature.

ACKNOWLEDGMENTS

We acknowledge the valuable feedback on the proposed methods from Marcel Uphues, Finance Director, and Stephan Länge, Senior Department Manager for Customer Information Management, of the German hypermarket chain Real.

ORCID

Bernhard von Mutius  <https://orcid.org/0000-0002-1274-8281>

Arnd Huchzermeier  <https://orcid.org/0000-0003-2178-3254>

REFERENCES

- Ailawadi, K. L., Harlam, B. A., César, 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.
- Aydinli, A., Bertini, M., & Lambrecht, A. (2014). Price promotion for emotional impact. *Journal of Marketing*, 78(4), 80–96.
- Bawa, K., & Shoemaker, R. W. (1987a). The coupon-prone consumer: Some findings based on purchase behavior across product classes. *Journal of Marketing*, 51(4), 99–110.
- Bawa, K., & Shoemaker, R. W. (1987b). The effects of a direct mail coupon on brand choice behavior. *Journal of Marketing Research*, 24(11), 370–376.
- Bawa, K., & Shoemaker, R. W. (1989). Analyzing incremental sales from a direct mail coupon promotion. *Journal of Marketing*, 53(7), 66–78.
- Blattberg, R. C., & Deighton, J. (1996). Manage marketing by the customer equity test. *Harvard Business Review*, 74(4), 136–144.
- Bodapati, A. V. (2008). Recommendation system with purchase data. *Journal of Marketing Research*, 45(1), 77–93.
- Brijis, T., Swinnen, G., Vanhoof, K., & Wets, G. (2004). Building an association rules framework to improve product assortment decisions. *Data Mining and Knowledge Discovery*, 8(1), 7–23.
- Chen, Y., Hess, J. D., Wilcox, R. T., & Zhang, Z. J. (1999). Accounting profits versus marketing profits: A relevant metric for category management. *Marketing Science*, 18(3), 208–229.
- Danaher, P. J., Smith, M. S., Ranasinghe, K., & Danaher, T. S. (2015). Where, when and how long: Factors that influence the redemption of mobile phone coupons. *Journal of Marketing Research*, 54(5), 710–725.
- Dhar, S. K., & Raju, J. S. (1998). The effects of cross-ruff coupons on sales and profits. *Management Science*, 44(11), 1501–1516.
- Drèze, X., & Hoch, S. J. (1998). Exploiting the installed base using cross-merchandising and category destination programs. *International Journal of Research in Marketing*, 15(5), 459–471.
- Khan, R., Lewis, M., & Singh, V. (2009). Dynamic customer management and the value of one-to-one marketing. *Marketing Science*, 28(6), 1063–1079.
- Kolassa, S., & Schütz, W. (2007). Advantages of the MAD/mean ratio over the MAPE. *Foresight*, 6(6), 40–43.
- Kumar, V., & Leone, R. P. (1988). Measuring the effect of retail store promotions on brand and store substitution. *Journal of Marketing Research*, 25(2), 178–185.
- Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76–80.
- Musalem, A., Bradlow, E. T., & Raju, J. S. (2008). Who's got the coupon? Estimating consumer preferences and coupon usage from aggregate information. *Journal of Marketing Research*, 45(6), 715–730.
- Narayanan, S., Desiraju, R., & Chintagunta, P. K. (2004). Return on investment implications for pharmaceutical promotional expenditures: The role of marketing-mix interactions. *Journal of Marketing*, 68(4), 90–105.
- 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.
- Pazzani, M. J., & Billsus, D. (2007). *Content-based recommendation systems*. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The adaptive web, lecture notes in computer science*. Berlin, Heidelberg: Springer.
- Reutterer, T., Hornik, K., March, N., & Gruber, K. (2017). A data mining framework for targeted category promotions. *Journal of Business Economics*, 87(3), 337–358.
- Ricci, F., Rokach, L., Shapira, B., & Kantor, P. B. (2011). *Recommender systems handbook*. New York: Springer.
- Rossi, P. E., McCulloch, R. E., & Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4), 132–340.
- Rust, R. T., Lemon, K. N., & Zeithaml, V. A. (2004). Return on marketing: Using customer equity to focus marketing strategy. *Journal of Marketing*, 68(1), 109–127.
- Sahni, N. S., Zou, D., & Chintagunta, P. K. (2017). Do targeted discounts offers serve as advertising? Evidence from 70 field experiments. *Management Science*, 63(8), 2688–2705.
- Swaminathan, S., & Bawa, K. (2005). Category-specific coupon proneness: The impact of individual characteristics and category-specific variables. *Journal of Retailing*, 81(3), 205–214.
- Tan, P. N., Kumar, V., & Srivastava, J. (2004). Selecting the right objective measure for association analysis. *Information Systems*, 29(4), 293–313.
- Van den Poel, D., De Schamphelaere, J., & Wets, G. (2004). Direct and indirect effects of retail promotions on sales and profits in the do-it-yourself market. *Expert Systems with Applications*, 27(1), 53–62.
- Venkatesan, R., & Farris, P. W. (2012). Measuring and managing returns from retailer-customized coupon campaigns. *Journal of Marketing*, 76(1), 76–94.
- Wong, R. C. W., & Wang, K. (2005). Data mining for inventory item selection with cross-selling considerations. *Data Mining and Knowledge Discovery*, 11(1), 81–112.
- Yang, Y., & Hao, C. (2010). Product selection for promotion planning. *Knowledge and Information Systems*, 29(1), 223–236.
- Zhang, J., & Breugelmans, E. (2012). The impact of an item-based loyalty program on consumer purchase behavior. *Journal of Marketing Research*, 49(1), 50–65.
- Zhang, J., & Krishnamurthi, L. (2004). Customizing promotions in online stores. *Marketing Science*, 23(4), 561–578.
- Zhang, J., & Wedel, M. (2009). The effectiveness of customized promotions in online and offline stores. *Journal of Marketing Research*, 46(2), 190–206.

How to cite this article: von Mutius B, Huchzermeier A. Customer-centric category selection for mobile and print promotions in loyalty reward programs. *Naval Research Logistics* 2021;68:683–700. <https://doi.org/10.1002/nav.21934>

APPENDIX A: MULTIPLE LINEAR AND LASSO REGRESSION MODELS

In this analysis, we test how multiple linear and Lasso regression models impact the prediction accuracy compared to the

TABLE A1 Mean out-of-sample WMAPE and MAPE values of the bivariate, multiple linear, and Lasso regression prediction models

Training-testing Data sets split by	Date		Promotion frequency	
	WMAPE	MAPE	WMAPE	MAPE
Bivariate linear regressions	25.3%	22.4%	25.3%	20.7%
Multiple linear regressions	26.4%	22.6%	34.1%	33.1%
Lasso regressions	25.8%	23.7%	32.4%	29.5%

basic model presented in Section 3.2. Besides using pf'_i , Φ'_i , and e'_i as inputs for all models, we also include five new independent variables for each prediction: the mean price per category ($price_i$), how fast the category perishes ($perish1_i$: not at all, $perish2_i$: slow, $perish3_i$: fast), the category’s promotion frequency ($promFreq_i$), the number of SKUs in the category ($nSKU_i$), and the category’s share of private labels ($nPrivate_i$). The linear prediction models can then be extended as follows: $\hat{y}_{ik} = \beta_0^{yk} + \beta_1^{yk} pf'_i + \beta_2^{yk} \Phi'_i + \beta_3^{yk} e'_i + \beta_4^{yk} price_i + \beta_5^{yk} perish2_i + \beta_6^{yk} perish3_i + \beta_7^{yk} promFreq_i + \beta_8^{yk} nSKU_i + \beta_9^{yk} nPrivate_i$. Thereby, y serves as a replacement for either γ , Φ , or e .

After checking for multicollinearity with the Farrar–Glauber test, we omit the independent variable $price_i$ as it can be explained to a large degree by e_i . We estimate the parameters of the multiple linear regressions using OLS and for the Lasso regression, we employ 10-fold cross-validation to obtain the best value for the regularization parameter. For the multiple linear regression models, the in-sample fitting measures improved significantly: the average R^2 of all models increased by 7 percent points to 0.878% and the adjusted R^2 increased by 6.3 percent points to 0.839% (data set splits by date). Analog to the analysis in Section 3.3, we calculate the out-of-sample WMAPE and MAPE metrics to compare the models’ predictive performance. The values we report in Table A1 are the mean WMAPEs and MAPEs across all predictions. We find that the bivariate linear regression models have the lowest out-of-sample prediction results. We attribute this finding due to overfitting of the other models on the testing data.

APPENDIX B: CATEGORY-DEPENDENT GROSS PROFIT MARGINS

In the following sections, we test if our empirical findings are robust to nonflat gross profit margins. We simulate this gross

profit margin as follows: $\lambda_i \sim 0.3 + \mathcal{U}(-0.1, 0.1)$. That is, we add an error term randomly selected between -0.1 and 0.1 to the average gross profit margin. Using these λ_i values, we recalculate Φ_{ik} and Φ'_i and update the cross-category profit prediction models. Out-of-sample performance remains high, that is, the WMAPE equals 18.2% in the mobile and 15.6% in the print channel (data sets split by date). We then determine the predicted ROMI values for each channel and uncover the correlation between the underlying variables and the ROMI. The two middle columns in Table B1 display the results. We find that no significant changes in the correlations occur, hence our findings in Section 4 are also valid for the category-dependent gross profit margin case (using the above described λ_i modification).

APPENDIX C: EXTENSION TO THE MULTI-PROMOTION SETTING

As outlined in the introduction, our focus in this paper is on the single-promotion setting. However, in reality, a significant part of the category-specific promotions might be redeemed together with other promotions. The overall profitability of promotions will then be lower, as the incremental profit effects are shared with other promotions, the costs, however, are not. In this chapter, we briefly discuss how our model can be extended to capture the multi-promotion setting.

In our data set, out of the 142,549 transactions in which at least one category-specific promotion was redeemed, 28,823 (20.2%) were single-promotion redeeming (case A; see section 3.1). 109,601 (76.9%) were redeemed together with at least one noncategory-specific promotion (eg, 10x points on all wines and 5x points on everything). This represents case B. Only 4,125 (2.9%) category-specific promotions were combined with one or more other category-specific or SKU-specific promotions. As the cases A and B cover over 97% of all transactions, we focus exclusively on the first-order

TABLE B1 Correlations for pf'_i , Φ'_i , and e'_i with \widehat{ROMI}_{ik}

Correlation	Base model		Cat.-dep. λ_i		Multi-prom.	
	Mobile	Print	Mobile	Print	Mobile	Print
$r_{pf'_i, \widehat{ROMI}_{ik}}$	0.074	0.348	0.045	0.287	-0.031	0.256
$r_{\Phi'_i, \widehat{ROMI}_{ik}}$	0.859	0.373	0.844	0.372	0.777	-0.023
$r_{e'_i, \widehat{ROMI}_{ik}}$	-0.162	0.168	-0.153	0.171	-0.012	0.524

Note: $r_{pf'_i, \widehat{ROMI}_{ik}}$ for $pf' < 0.28$: 0.596 (base model), 0.534 (cat.-dep. λ_i), and 0.318 (multi-prom.).

effects of these two cases in the subsequent analysis. Note that the coupons are not multiplicative, that is, in each category only the highest-valued coupon is used. Hence combined noncategory-specific promotions always have a lower promotion depth.

While the ROMI equation remains valid, the parameters $\Delta Profit_{ik}$ and $Cost_{ik}$ must incorporate the effects of both cases A and B. For the costs we can state $Cost_{ik}^B = \gamma_{ik}^B(c^y e_{ik}^B + c^f)$. Thereby, in both the definitions of γ_{ik}^B and e_{ik}^B the set N_{ik}^B only contains all multi-promotion redeeming transactions. For the incremental profits in case B, the cross-category profit is zero, as it is attributable to the noncategory-specific promotion (as we focus on first-order effects). The presence of other category-specific promotions is possible but irrelevant. In other words, the incremental profit is entirely due to the incremental spending in the promoted category multiplied by the gross profit margin of that category. We obtain $\Delta Profit_{ik}^B = \gamma_{ik}^B \lambda_i (e_{ik}^B - e'_i)$.

If we were to consider combining a category-specific promotion with one or several promotions on other categories, but no noncategory-specific one (case C; eg, 10× points on all *wines* and 10× points on *cheese*), the cost definition could trivially be extended. However, correctly capturing the cross-category profits (ie, avoid double-counting them) poses a significant challenge. In such a case, we propose to split the cross-category profit between all redeemed category-specific promotions. In the case of two, we can state: $\Delta Profit_{ik}^C = \gamma_{ik}^C \left(\lambda_i (e_{ik}^C - e'_i) + \frac{1}{2} (\Phi_{ik}^C - E^C) \right)$. Thereby, E^C denotes the average nonpromotional profit contribution of those categories in Φ_{ik}^C .

For case B, we build the prediction models analog to Section 3.2. The models together achieve similar (although slightly lower) prediction performance than those shown in the paper, that is, the average out-of-sample WMAPE equals 30.5% (data set splits by date). Then we use this information to calculate the predicted ROMI of the multi-promotion setting. As expected, the average ROI values are lower compared to the single-promotion setting: the mean ROMI of the top four categories is 1.19 in the mobile channel (14.92 in the single-promotion setting) and 0.46 in the print channel (2.21 in the single-promotion setting). The correlations between the underlying variables pf'_i , Φ'_i , and e'_i and the target variable \widehat{ROMI}_{ik} for both the single- and multi-promotions setting are shown in Table B2.

In general, we find that the key findings of the single-promotion setting are also valid in the multi-promotion setting, although the correlations are slightly weaker (see Table B2). That is, in the mobile channel, cross-category profits are essential while in the print channel, high, but not too high purchase frequencies are important. In the print channel, the category spending also becomes a strong profitability driver in the multi-promotion setting. The reasoning for this observation is similar to that for the redemption rate: if a large portion of promotions are redeemed together with other noncategory-specific promotions but for both promotions independent fixed contacting costs have to be paid, then only high category spendings can recoup these upfront investments. In the single-promotion setting, cross-category profits could also contribute to this goal, reducing the necessity for high e'_{ik} values.