

Automated adaptive selling

Maurits Kaptein

*Jheronimus Academy of Data Science, Den Bosch, The Netherlands and
Tilburg University, Tilburg, The Netherlands*

Richard McFarland

Department of Marketing, ESSEC Business School, Cergy-Pontoise, France, and

Petri Parvinen

*FAF University of Helsinki, Helsinki, Finland and DIEM Aalto University,
Helsinki, Finland*

Automated
adaptive
selling

1037

Received 30 August 2016
Revised 8 February 2017
6 September 2017
23 November 2017
Accepted 19 December 2017

Abstract

Purpose – This paper aims to develop and test a method of automating, for online retailers, the practice of adaptive selling, which is typically used by salespeople in face-to-face interactions. This method customizes persuasive messages for individual customers as they navigate a retailer's website.

Design/methodology/approach – This paper demonstrates a method for the online implementation of automated adaptive selling using sales influence tactics. Automated adaptive selling is compared to nonadaptive selling in three e-commerce field studies.

Findings – The results reveal that adaptive selling is more effective than nonadaptive selling. The click-through rates increased significantly when adaptive selling was used.

Research limitations/implications – This paper highlights the effectiveness of existing theories concerning adaptive human-to-human selling and their utility to online selling. The authors demonstrate the added value of adaptive selling in e-commerce, thereby opening up a novel area of research into adaptive selling online. While the paper focuses on the adjustment of sales influence tactics, other factors could be investigated for adjustment in future research (e.g. prices).

Practical implications – The methods, described in detail, are readily available for implementation by online retailers. The implementations are timely and increasingly valuable as e-commerce expands into interpersonal channels (e.g. instant messengers and social media).

Originality/value – To the authors' knowledge, this paper is the first to formally implement automated adaptive selling as described in the ISTE model in an e-commerce setting.

Keywords Online retailing, Adaptive selling, e-selling, ISTE model, Sales influence tactics

Paper type Research paper

1. Introduction

The practice of Web customization is ubiquitous among online retailers. Marketing scholars have made notable contributions to this field, including in the areas of product recommender systems (Ansari *et al.*, 2000; Bodapati, 2008; Ying *et al.*, 2006), customized design and content of e-mail communications (Ansari and Mela, 2003) and morphing Web content and banner ads based on customers' cognitive preferences (Hauser *et al.*, 2009; Urban *et al.*, 2013). Despite this foundational research, one form of customization that has received comparatively little attention is the customization of persuasive messages (Kaptein and Eckles, 2012). This is a major theoretical gap, given the fundamental finding in the

Screenshots of the participating websites and the implementations of the influence tactic are available from the corresponding author upon request.



marketing literature that tailoring persuasive messages to individuals pays off (Griskevicius *et al.*, 2009; McFarland *et al.*, 2006; Spiro and Weitz, 1990). We contribute to this Web customization literature by developing an automated method that selects from among multiple persuasive messages regarding any product by experimentally determining the specific message that has the highest probability of success for each, individual customer in real time. To do so, we implement the adaptive sales process from the professional selling literature in an online context.

Adaptive selling, defined as the unique ability to adapt sales messages and persuasive arguments by salespeople, is one of the most prominent areas of study within the marketing domain (Franke and Park, 2006; Román and Iacobucci, 2010). As “inherently a dynamic influence process” (Weitz *et al.*, 1986, p. 187), personal selling provides marketers with “the opportunity to tailor-make presentations for each individual customer” (Sujaan *et al.*, 1988, p. 9) and update these presentations in real time in response to customer reactions. Other forms of marketing communication, including the current practices of Web customization, still tend to be geared toward delivering “messages targeted toward a “typical” customer within a market segment” (Spiro and Weitz, 1990, p. 61), which is modified at discrete intervals of time rather than continuously. By comparison, the method we introduce adapts continuously and on a customer-by-customer basis.

The ability to customize on a customer-by-customer basis represents a meaningful advance over existing approaches currently used by online retailers. Under current practices, for example, when a customer initially visits an online retailer, such as a bookseller, he/she is presented with multiple books on the homepage. As she navigates through the website, she receives product recommendations based on product ratings or purchases by similar groups of customers (such as Amazon.com’s popular “customers who bought this also bought [...]” recommendations). In this type of approach, preferences are determined off-line on a static basis, and the live component simply serves to assign customers to homogenous segments of similar customers (Ansari *et al.*, 2000; Bodapati, 2008).

In this paper, we propose a different Web customization approach, which we refer to as *automated adaptive selling*. Continuing our bookseller example, the homepage of the online store still presents multiple books, but each product now also includes one of several possible sales influence tactics: for example, a message listing a book as either a “bestseller,” or as being “highly rated by experts.” A customer choosing to click on a book listed as a bestseller versus one listed as being highly rated by experts conveys information about which of those messages are most effective (Kaptein and Eckles, 2012; Kaptein *et al.*, 2015). As the customer browses through the website, we determine with increasing accuracy which influence tactics will be most effective for this specific customer by experimentally presenting him/her with different influence tactics using an efficient algorithm and observing her clicking behavior. Our method thus differs in its focus (argument type versus product) and granularity (offline versus continuously augmented) from typical Web customization methods.

Given the importance of adaptive selling and the increasing proportion of retail sales moving from brick-and-mortar to online retailers, our method of automating the adaptive selling process for online retailers has the potential to improve customers’ shopping experiences and increase online retailers’ sales revenues. Our method is also less expensive and easier to implement than the existing methods (Hauser *et al.*, 2009; Urban *et al.*, 2013). Creating an automated method of adaptive selling also provides researchers with opportunities for advancing and testing the adaptive selling theory in situations that would be difficult in person-to-person settings. For example, systematically manipulating variations in sales tactics across customers is cumbersome in face-to-face settings, but

within reach in online selling (Ansari and Mela, 2003). In addition, we contribute to the adaptive selling literature by automating each step of the adaptive selling process, giving greater insight into each of these steps.

In the remainder of this paper, we start by reviewing the literature and introducing the five-stage model of adaptive selling developed by Weitz (1978). Next, we review the use of influence tactics for adaptive selling and then develop our hypothesis. Following that section, we introduce our method of automating adaptive selling. Using both a simulation study (presented in the supplementary materials) and multiple field experiments, we subsequently test the effectiveness of our method against current practices used by online retailers. According to our findings, automated adaptive selling is more effective than existing, competitive versions of online selling in real-world implementations. Finally, we discuss the theoretical and managerial implications of our work and highlight future research directions.

2. Literature review

In this section, we first review the adaptive selling literature with a focus on the five-stage process model developed by Weitz (1978), and we relate this model directly to more recent developments in *online* selling. Second, we review work on the role of influence tactics in the adaptive selling literature. On the basis of this review of the literature, we develop our hypotheses in the subsequent section.

Adaptive selling is defined as the alteration of sales tactics based on the characteristics of the selling situation and of the individual buyer (Weitz, 1981; Weitz *et al.*, 1986). The customization method we present and evaluate in this paper focuses on adapting the content of persuasive messages for each product presented to individual customers browsing an online store. Following the literature on adaptive selling, we refer to these persuasive messages as *sales influence tactics*, or more simply as *influence tactics* (McFarland *et al.*, 2006).

2.1 The ISTEА adaptive selling process model

Adaptive selling can be considered one of the most influential theories in marketing, as well as one of the few theories native to the field. The practice of adaptive selling has been demonstrated to improve sales outcomes (e.g. see meta-analyses by Franke and Park, 2006 and Verbeke *et al.*, 2011). Evans (1963) has been credited for the earliest research mentioning “adaptive selling.” Weitz (1978) was the first to develop a process model of adaptive selling – that is, a multistage model of the individual steps *within* the adaptive selling process:

- impression formation (I);
- strategy formulation (S);
- message transmission (T);
- evaluation of the effect of the sales attempt (E); and
- subsequent adjustment (A).

The resulting *ISTEA* model represents a dynamic process through which salespeople cycle, making adjustments at multiple times during sales encounters; see Figure 1 for an overview.

Next, we describe each of the five stages of the *ISTEA* model in greater detail and relate these stages to recent developments in online retailing.

2.1.1 Impression formation. In this first stage, the seller builds an impression of the customer by “examining past experiences with the target customer [or] other customers [and] by observing the target customer during an interaction” (Weitz, 1978, p. 502). The

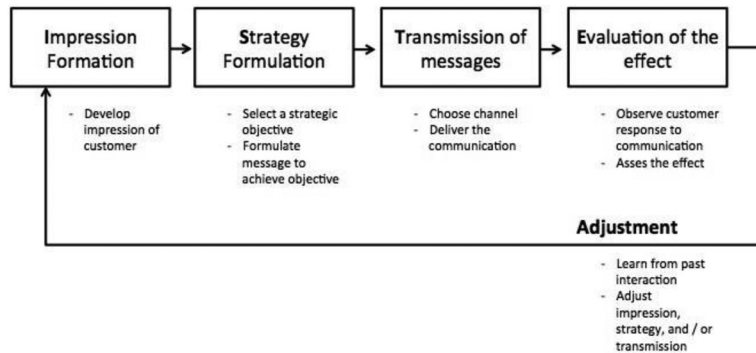


Figure 1.
Simplified
representation of the
ISTEA model

Notes: Simplified representation of the ISTE model (adopted from Weitz, 1978). Displayed are the “impression formation”, “strategy formulation”, message “transmission”, “evaluation” and “adjustment” stages

seller further explicitly adjusts impressions upon acquiring further information through interactions with the customer. Several dimensions on which the impression is formulated in personal selling are known to be important in adaptive selling, including buyer communication styles, customer needs and preferences, customer emotions and the quality of the relationship with the seller (Hall *et al.*, 2015; McFarland *et al.*, 2006; Mullins *et al.*, 2014).

Much of the personalization effort in online marketing and selling relies on similar dimensions. Most notably, the extensive work on recommender systems is geared primarily toward identifying user needs and preferences through the examination of online browsing behavior (Montgomery and Smith, 2009; Gretzel and Fesenmaier, 2006). The aim in *website morphing* is also to adapt online product presentations to consumers' preferred *cognitive style* (as determined by browsing behavior), resulting in visually expressive versus verbally detailed presentations (Hauser *et al.*, 2009).

2.1.2 Strategy formation. During the second stage (strategy formation), “the salesperson analyzes his impression of the customer and develops a communication strategy” (Weitz, 1978, p.502). Weitz (1978) argued that the ability to choose effective influence tactics based on accurate customer impressions has a substantial impact on sales performance. Providing empirical support for this argument, McFarland *et al.* (2006) found that more effective salespeople match categories of influence tactics to the communication style of buyers. Spiro and Weitz (1990, p. 62) argued that adaptive selling ability is linked to “an ability to collect information to facilitate the process of matching sales situations to categories in memory.”

A critical aspect of the resulting communication strategy is deliberate and conscientious experimentation with different sales messages (influence tactics) to develop additional knowledge about which messages are the most suitable (Sujan *et al.*, 1988). Thus, the idea of experimentation and learning both within and across sales interactions based on the responses of customers is an explicitly critical component of successful adaptive selling. Similarly, in the context of online selling, it is theoretically possible to learn which influence tactics are most effective for each customer by presenting them with specific influence tactics and registering their responses as they navigate the website; on pages featuring multiple products, each product could be accompanied by its own influence tactic (Kaptein *et al.*, 2015).

2.1.3 Transmission. In the third stage, transmission, “the salesperson delivers the message,” according to Weitz (1978, p. 502). Although the adaptive selling literature tends to assume that this stage is limited to face-to-face communication, in the online selling process, more communication channel choices are possible. Thus, one could actively experiment with the channels used to deliver specific messages at different points in the sales process (Ansari and Mela, 2003). Specifically, during the transmission stage, an influence tactic can be presented in text on the customer’s computer or mobile phone screen or by using an image when the customer visits the online store. Visual demonstrations of the transmission stage in the online selling context can be provided by the corresponding author upon request; showing that distinct influence tactics can be presented with each product.

2.1.4 Evaluation. According to Weitz (1978), in the fourth (evaluation) stage of the process, the salesperson observes the customer’s response to the message transmitted and may also solicit the customer’s opinions. This implies an active and iterative process of information exchange in which, through multiple criteria, the effects of the sales pitch, including influence tactics, are assessed and continuously adapted. Despite the importance of this process, the methods used tend to be implicit in the sales literature. Although the ability to “read” the customer is acknowledged as one of the core ingredients of successful adaptive selling (Hall *et al.*, 2015; Kidwell *et al.*, 2007; Mullins *et al.*, 2014), few methods are explicitly provided for doing so. In the process of automating adaptive selling, however, the effectiveness of the messages is relatively easily monitored by observing the clicking behaviors of customers (Hauser *et al.*, 2009).

2.1.5 Adjustment. In the final stage of the model, salespeople make adjustments on the basis of the evaluations from the previous step (Weitz, 1978). For example, evaluation of a sales attempt, including the use of sales influence tactics, will lead to an adjustment of the impression formed about a specific customer (Figure 1). Initial impressions can be updated in line with new evidence. This process is already implemented in online customization; for example, a click on a recommended product updates the customer’s preference profile (Montgomery and Smith, 2009).

2.2 Adaptive selling using influence tactics

Online retailers commonly use influence tactics on a static (i.e. nonadaptive) basis. For example, persuasive messages such as “almost out of stock” illustrate the scarcity influence tactic, and “bestselling product” uses the social proof influence tactic (Cialdini, 1993). Messages like these are frequently encountered in e-commerce stores although not all online retailers use influence tactics, and those who do, do so non-adaptively (Kaptein and Eckles, 2012).

In the personal selling context, customers are more likely to make purchases when salespeople adapt their use of influence tactics based on customer differences than when they refrain from doing so (McFarland *et al.*, 2006). Thus, this literature suggests that customers individually differ in terms of which influence tactic is most effective. This heterogeneity in responses to influence tactics across customers has proven stable within customers; with them tending to respond best to similar (or the same) influence tactics over time, making such influence tactics a prime candidate for customization (Kaptein and Eckles, 2012).

This study adopts Cialdini’s (1993) taxonomy of six influence tactics, of which we focus on the following three (Kaptein *et al.*, 2015):

- (1) *Social proof* tactics appeal to people’s desire to act in ways that are consistent with social norms (Ajzen and Fishbein, 1980; Goldstein *et al.*, 2008; Zhang and Wedel, 2009): lists of bestsellers are examples of such messages.

- (2) The *scarcity* tactic is based on the assumption that scarce products must be desirable, as shown in phrases, such as “limited time offer” and “while supplies last” (Lynn, 1991; Parker and Lehmann, 2011).
- (3) The *authority* tactic exploits the desire to make the “right” choice based on the opinions of experts; examples include expert reviews or endorsements from influential sources (Kelman and Hamilton, 1989; Martin and Hewstone, 2003).

3. Hypotheses

Based on the information presented above, in this section we develop specific hypotheses. Influence tactics, especially the social proof, scarcity and authority tactics, are easily implemented in online selling (Kaptein *et al.*, 2015). There is a *positive main effect* of the use of influence tactics (Ajzen and Fishbein, 1980; Goldstein *et al.*, 2008; Zhang and Wedel, 2009). Thus, on average, products pitched with an influence tactic are more likely to be purchased than those pitched without an influence tactic. Finally, the literature on offline selling finds heterogeneity in the effectiveness of influence tactics across customers in interpersonal settings (McFarland *et al.*, 2006). However, to the best of our knowledge, no prior work has exploited such heterogeneity in online selling.

On the basis of our literature review, we formulate three increasingly more stringent hypotheses. The purpose of our first, “baseline” hypothesis is to verify that the static (nonadaptive) use of influence tactics associated with the presentation of products on retailer webpages is more effective than the presentation of products without the use of influence tactics at all[1].

- H1.* The use of (static) influence tactics associated with products displayed on webpages of online retailers will outperform (in terms of the click-through rates) products displayed without the use of influence tactics.

This first hypothesis does not evaluate adaptive selling, but rather intends to replicate earlier work investigating the average positive effects of the static use of influence tactics (always presenting the same influence tactic with the same product) as compared to the average effects of presenting products without influence tactics.

Next, we hypothesize that the adaptive presentation of influence tactics will also outperform presenting products without influence tactics:

- H2.* Adapting the presentation of influence tactics associated with individual products displayed on webpages of online retailers to individual customers will outperform (in terms of the click-through rates) products displayed without the use of influence tactics.

Finally, we specify a more stringent hypothesis that predicts that the adaptive presentation of influence tactics with products will outperform the static presentation of influence tactics with products.

- H3.* Adapting the presentation of influence tactics associated with individual products displayed on webpages of online retailers to individual customers will outperform (in terms of the click-through rates) products displayed with influence tactics, but on a static basis.

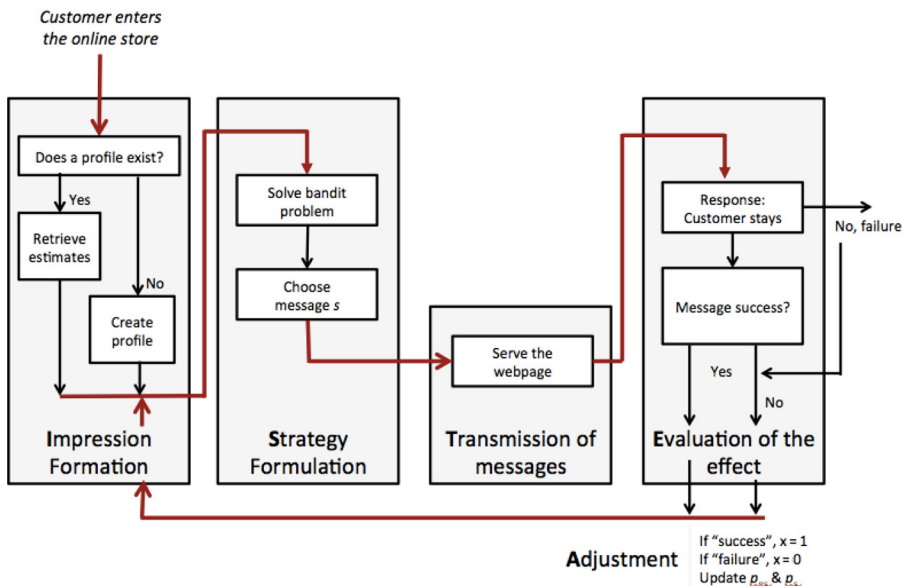
The second and third hypotheses constitute incrementally more challenging tests of the utility of adapting influence tactics in e-commerce. Below, each of these three hypotheses is

tested in different e-commerce settings in three field experiments. However, before discussing our field experiments, we first introduce the method by which online sellers can (automatically) adapt influence tactics to individual customers.

4. Automated adaptive selling algorithm

In this section, we describe our novel method of automated adaptive selling (our new algorithm) with reference to the ISTE model (Figure 2). Here, we provide a high-level overview of implementation at each stage of the process; the mathematical details of our method are provided in the supplementary materials. We automate the five stages of the ISTE model as follows:

- (1) *Impression formation*. Our initial impression consists of an estimate of the probability of success for each of the distinct influence tactics used on a customer entering the retail platform (Figure 2, top-left). On the customer's first visit to an online store, these estimates are derived from the responses of other customers. We call these estimates of the effectiveness of distinct influence tactics for a specific customer a *profile*; the profile describes which influence tactics are most likely to influence the customer. On the customer's subsequent visits to the website, this personal profile is used. More formally, for each influence tactic, $s = 1, \dots, s = 4$ (where, in our case, 1 = "social proof," 2 = "scarcity," 3 = "authority," and 4 =



Notes: Overview of automated adaptation of influence tactics using the ISTE model. For a new customer (top-left), the initial estimates of the susceptibilities to different tactics are derived from the average-level estimates (see Supplementary Materials S1). Next, based on the estimates, a trade-off between learning and earning is made (see S2 for mathematical details), a message is selected and the effect of the message is logged. Finally, the estimates for the current customer are dated and used in subsequent interactions

Figure 2.
Implementation of
automated adaptive
selling

“neutral,” e.g. not explicitly displaying an influence tactic[2]), and for each customer $i = 1, \dots, i = n$ the estimate is composed of a point estimate p_{is} (the estimated probability that the influence tactic will be successful) and a measure of certainty (denoted n_{is} in Supplement S1).

- (2) *Strategy formation.* After forming the impression (operationalized by the set of estimates, p_{i1}, \dots, p_{i4} , for customer i), we formulate a strategy. The strategy at this stage involves either choosing the influence tactic that we believe to be the most effective (e.g. the one with the highest p_{is}) or experimenting with and displaying a different influence tactic, the effectiveness of which we are more uncertain about. A naïve strategy is to always present the type of persuasive message that has had the highest success rate in the past. More formally, this is known as an *exploitation* strategy. There is, however, a benefit to experimenting with alternative messages that appear to have a lower success rate, but whose probability estimates are uncertain (and ultimately may prove more effective). This experimentation is known as *exploration*. Balancing exploitation and exploration for long-term gains is a delicate task (Gittins, 1979). To address this, we implement a heuristic approximation to *Thompson sampling* (Thompson, 1933; Scott, 2010; Agrawal and Goyal, 2012); mathematical details of this approach are given in S2. Because there are other methods for optimizing these payoffs, in the methodology section below, we test our method for balancing exploitation versus exploration against common, but more computationally intensive, alternative methods.
- (3) *Transmission.* The adapted message is transmitted via a sales influence tactic displayed on top of, or next to, the products displayed in the online retail platform. Details of the webpages are provided in the descriptions of the field experiments (discussed below).
- (4) *Evaluation.* In each of the field experiments, the success of the displayed sales influence tactic is evaluated in the evaluation stage by monitoring the click-through behavior of the visiting customer. The tactic is deemed successful if the visitor clicks on the product (to view more information or to add the product to his/her shopping cart).
- (5) *Adjustment.* We update the estimates obtained in the impression-formation step in line with the observed clicking behavior. This adjustment is made at the level of the individual customer.

Our implementation of the ISTE model results in the continuous adjustment of sales influence tactics in online selling at the individual customer level. The initial profile determines the probability with which specific tactics are selected (if, for example, social proof works well on a specific platform, it is more likely to be displayed), after which the behavioral responses of individual customers continuously influence the probabilities of selection for that customer. For example, if an individual customer clicks on (again, for example) authority messages and ignores others, for this specific customer, the probability of seeing more authority messages will dynamically increase. Thus, the behavioral responses of customers are used directly to infer their susceptibility to distinct influence tactics (and hence customer traits). This method of inferring traits based on recurrent behavior is known as an *implicit* measure of a customer trait as opposed to *explicit* (survey-based) measures that are more common in the marketing literature (for a discussion of explicit versus implicit measures of customer traits, see Kaptein *et al.*'s study, 2015).

5. Methodology

In this section, we first discuss the two simulation studies we conducted to verify the effectiveness of our method of optimizing the exploitation-exploration trade-off in the *strategy formation* stage of our algorithm (see Section 4 above). Following this, we discuss our three field experiments, which are used to test our hypotheses and verify the effectiveness of the implementation of our automated adaptive selling algorithm. The results for the simulation studies and for each experiment are discussed within their respective subsections below. In addition, we discuss the pooled results across all three field studies in Section 6.

5.1 Simulation studies and results

We implement a heuristic approximation to Thompson sampling to solve the exploitation-exploration problem (Thompson, 1933; Scott, 2010). We do so in the strategy formulation stage to determine whether it is more beneficial to continue to experiment with new influence tactics with a customer or to use an influence tactic that has already been proven effective. Because there are other methods for optimizing these payoffs, we tested the viability of our method as compared to plausible alternatives in Supplements S3 and S4. Both of our simulation studies show that our simple-to-compute approach is competitive in performance (in terms of expected rewards) as compared to other current state-of-the-art methods that are more computationally intensive.

5.2 Field experiments: procedures

We conducted three field experiments to evaluate the effectiveness of the automated adaptive selling algorithm and to test our hypotheses. Given confidentiality and privacy concerns, the participating companies did not allow us to track customers through the check-out process; thus, we only report on increases or decreases in clicking behavior rather than actual revenues to comply with the conditions placed on us in return for participation.

In each field experiment, we tested our hypotheses by assigning visitors to one of the following three conditions:

- (1) In the *adaptive* condition, customers were exposed to sales influence tactics that were automatically adapted using our novel method, as described in Section 4.
- (2) In the *status-quo* condition, which represented the current practice of participating companies, customers were exposed to product presentation without influence tactics.
- (3) In the *random* condition, customers were exposed to product presentation with a randomly selected sales influence tactic. That is, for example, for each product displayed, we randomly chose one of the influence tactics to display.

Thus, we used a between-subject design and subsequently compared performance in each of the three conditions with the different methods of selecting influence tactics. In each of our field experiments, we tested *H1* by comparing the performance of the *status-quo* condition to the *random* condition. The *status-quo* condition in each experiment comprised a product presentation *without* an influence tactic. Effectively, the product presentations in the *status-quo* conditions were equal to those in the alternative conditions when the “neutral” tactic was selected. We tested *H2* by comparing the *status-quo* condition to the *adaptive* condition, and finally tested the more conservative *H3* by comparing the *random* condition to the *adaptive* condition.

Both the *status-quo* and the *random* condition can be considered as comparison (or *baseline*) conditions, whereas the adaptive condition is the treatment condition of interest, as it implements automated adaptive selling using influence tactics. A comparison of click-through rates between the *adaptive* and *status-quo* conditions most directly assesses the effect of automating adaptive selling. However, in directly comparing the adaptive condition to the *status-quo* condition, we are comparing our adaptive method to one in which no influence tactic is presented. A more conservative test is provided by comparing our adaptive method with the (static) use of influence tactics, thereby motivating the inclusion of the *random* condition.

Next, we describe the three field experiments in which we use different operationalizations of automated adaptive selling to test our hypotheses. We use a stimulus sampling approach to show that the effectiveness of our method does not depend on one specific implementation. Table I presents an overview of the descriptive statistics (number of participants and number of page views) for each of the three field trials. Finally, we pool the results of the three trials and test our hypotheses using a single meta-analysis.

5.3 Field Experiment 1: performance with a large-scale online retailer

5.3.1 Experiment 1 design. The first field experiment was run in collaboration with a large-scale online retailer that sells bath products. Automated adaptive selling was implemented on both the product-overview page and the product-detail pages[3]. To implement the different influence tactics, we changed the labels displayed on the products. Here, we either presented the neutral influence tactic (no label), the authority influence tactic (“our choice”), the social proof influence tactic (“bestseller”) or the scarcity influence tactic (“limited availability”). In this implementation, we used the authority of the vending platform to implement the authority influence tactic. We showed labels on a maximum of three products at a time on the product-overview page. Each time a product was shown with a label (or with none in the neutral version), a click on that product (on the overview page) or its addition to the shopping cart (on the product-detail page) was considered a success.

During a 12-month period, a total of 563,776 unique visitors (number of distinct individuals requesting pages from the website during the trial period) viewed 2,085,996 products, resulting in an average of 3.70 views per visitor. In this study, we allocated 19.5 per cent of the visitors to the *status-quo* condition, 80 per cent to the *adaptive* condition and 0.5 per cent to the *random* condition[4]; these latter 2,819 visitors received a random selection of influence tactics, as they browsed the online store. To evaluate the performance of automating adaptive selling during the trial, we compared the proportion of successes (clicks on the product or its addition to the shopping cart) for products supported with an influence tactic (including the neutral message) between the three conditions. We prevent the disclosure of absolute click-through rates of the participating Web company, which is

Trial	Type	Visitors	Page views	Adaptive (%)	Random (%)	Status-quo (%)
1	Bathing products	563,776	2,085,996	80.0	19.5	0.5
2	Lingerie	44,740	167,426	70.0	20.0	10.0
3	Flash sales	375,013	5,403,056	60.0	20.0	20.0

Table I.

Descriptive statistics of the field trials

Notes: Overview of the number of customers, number of page views and the distribution over conditions for each of the field trials

considered confidential, by normalizing the performance of the *status-quo* condition to 1. This allows us to test our hypotheses, without disclosing competitive information.

5.3.2 Experiment 1 results. The (normalized) performance of the random condition in this field trial was 1.061, indicating a 6.1 per cent increase in the click-through rates as compared to the *status-quo* condition. The performance of the *adaptive* condition was 1.083, indicating an increase of 8.3 per cent as compared to the *status-quo*. Pairwise comparisons show that the *adaptive* condition significantly outperformed the *status-quo* condition ($H2$; $p < 0.001$). The *random* condition did not significantly outperform the *status-quo* condition ($H1$; $p > 0.05$), nor did the *adaptive* condition significantly outperform the *random* condition ($H3$; $p > 0.05$).

These latter results are primarily attributable to the small number of observations within the random condition in combination with the overall low click-through rates. All reported p -values were computed using asymptotic z -tests. In each case, the number of observations was corrected for repeated numbers of observations within customers (to correct for any dependency between observations). Furthermore, multiple pairwise comparisons were adjusted using Bonferroni corrections. In sum, this first field study confirmed $H2$ but yielded insufficient evidence to confirm the other two hypotheses.

5.4 Field Experiment 2: performance with a small-scale online retailer

5.4.1 Experiment 2 design. The second field study was conducted in collaboration with a smaller e-commerce store selling lingerie, specifically to examine the performance of adaptive selling in smaller-scale stores. The implementation of automated adaptive selling in this context was very similar to that in Study 1: We used a set of labels to implement the influence tactics both on the product overview page as well as on the subsequent product-detail pages. Again, either a click on the product or its addition to the shopping cart was counted as a success, otherwise as a failure; as discussed previously, we were not granted access to the actual sales figures.

This study ran for six months. In total, 167,426 products were shown to 44,740 unique visitors for an average 3.74 product views per visitor. In this field trial, we allocated 20 per cent of the visitors to the *status-quo* condition, 10 per cent to the *random* condition and 70 per cent to the *adaptive* condition. We specifically chose to increase the number of customers assigned to the *random* condition, as compared to the previous trial, to allow for more precise estimates. The participating company, whose primary interest was in the comparison between the adaptive and the *status-quo* conditions, was not willing to raise the proportion of customers allocated to the random condition beyond these 10 per cent.

5.4.2 Experiment 2 results. In reporting our results, we again normalize performance in the *status-quo* condition and compare the relative performance of the *random* and *adaptive* conditions (as in Study 1). The (normalized) performance of the *random* condition was 1.012, indicating a 1.2 per cent increase in the click-through rates ($H1$). The performance of the *adaptive* condition as compared to the *status-quo* was 1.036 ($H2$). Only the difference between the *adaptive* condition and the *status-quo* condition was statistically significant ($H2$; $p < 0.05$); all other pairwise comparisons were insignificant ($p > 0.05$). Thus, again, $H2$ was confirmed.

5.5 Field Experiment 3: performance in a complex implementation

5.5.1 Experiment 3 design. The third field experiment was conducted with an online retailer offering “flash sales,” or temporary offers. In this study, our implementation of adaptive selling was more involved than in the previous field studies. We chose this more elaborate implementation to demonstrate that our method could be extended beyond simple labels

placed over product pictures to influence the full customer experience. The red boxes indicate the locations on the screen in which we implemented adaptive selling. First, on the top-right part of the page, automated adaptive selling was used to determine a dynamic menu item. These menu items state either “popular products,” “our recommendations,” or “almost sold out” to implement the social proof, authority and scarcity influence tactics, respectively. Second, a much more elaborate implementation of the influence tactics is shown in the large box at the top of the page: the “opinion of an expert” (authority) is displayed, and similar text boxes were created for the scarcity and social proof tactics. Third, small labels were placed next to the actual products on display, reflecting a similar implementation of sales influence tactics, as in Studies 1 and 2. Finally, a short summary of the information in the large box at the top of the page was displayed on the left bar, again to implement an influence tactic. Thus, in the adaptive condition, a large proportion of the user experience was influenced by our method of automating adaptive selling.

Study 3 ran for a period of six months. In total, 5,403,056 products were presented to 375,013 unique visitors, giving an average of 14.4 product views per unique visitor. In this field trial, we allocated 20 per cent of the visitors to the *status-quo* condition, 20 per cent to the *random* condition and 60 per cent to the *adaptive* condition.

5.5.2 Experiment 3 results. As before, we compared the performance of the *random* and *adaptive* conditions to that of the *status-quo* condition. The performance of the *random* condition was 1.034, indicating a rise in the click-through rates of 3.4 per cent, which was statistically significant ($H1$; $p < 0.001$). The performance of the *adaptive* condition was 1.076. Using pairwise comparisons, we demonstrate that the *adaptive* condition outperformed both the *status-quo* ($H2$; $p < 0.001$) and the *random* ($H3$; $p < 0.001$) conditions. Thus, all three hypotheses were confirmed in this third field experiment.

6. Pooled results: a meta-analysis across the three field studies

In this section, we will discuss our meta-analysis. We first pooled the data from each of the field experiments to conduct a meta-analysis. The results of each of the individual trials are summarized in Table II. The results of the meta-analysis show that the *adaptive* condition has a normalized performance of 1.035 (estimated 95 per cent confidence interval (CI): 1.032-1.037, Cohen’s $d = 0.02$) as compared to the *random* condition, and 1.077 (estimated 95 per cent CI: 1.070-1.084, Cohen’s $d = 0.08$) as compared to the *status-quo* condition across the studies. Using the pooled results renders all the pairwise comparisons statistically significant ($p < 0.001$), indicating that automated adaptive selling significantly outperforms both the *status-quo* and the *random* conditions across the different implementations in the three studies. In line with the results obtained in the third field trial, this meta-analysis thus confirms all three hypotheses.

Trial	Observations	Status-quo (A)	Random (B)	Adaptive (C)
1	2,085,996	1 ^C	1.061	1.083 ^A
2	167,426	1 ^C	1.012	1.036 ^A
3	5,403,056	1 ^{BC}	1.034 ^{AC}	1.076 ^{AB}

Table II.
Overview of the
results of the field
trials

Notes: Results of the comparisons between the different experimental conditions for each of the three field trials. Note that also in computer-to-human selling adaptation consistently outperforms non- adaptive selling. Superscript capitals denote statistically significant differences when comparing proportions pairwise using z-tests. The numbers of observations are corrected for repeated observations (dependency between observations). Pairwise comparisons are corrected for using Bonferroni corrections

Despite a small Cohen's d , the improvement in the click-through rates of about 7 per cent is not only statistically significant but also carries a high value for the participating companies. The long-term, continuous increase in customer interaction generated by our automated adaptive selling approach in each case translated into higher revenues in the adaptive group and continued use of the suggested algorithm by the participating companies after the experiments concluded (a strong managerial endorsement of the method).

7. General discussion

This paper develops, describes and evaluates a novel method of automated adaptive selling using sales influence tactics. Although there have been many attempts to automate, customize and personalize online and electronic selling in recent years, we are not aware of any that have exploited the extensive literature on adaptive (offline) selling as a theoretical driver for the development of novel, online or automated approaches. We have introduced a new method in which the stages of impression formation, strategy formulation, transmission, evaluation and adjustment in the ISTEAD adaptive selling model can be implemented automatically in online retailing. We targeted the adjustment of sales influence tactics, which are derived directly from the literature on adaptive selling. We further evaluated the effectiveness of automated adaptive selling in three field trials. Adaptive selling consistently outperformed the *status-quo* condition, suggesting the importance and potential benefits of using such an approach. Notably, the *status-quo* in each of our studies consisted of the established best practices of these sophisticated, for-profit retailers.

7.1 Managerial implications

Our demonstration of the effectiveness of automated adaptive selling using influence tactics in e-commerce has several practical implications. First, our current work highlights the prime importance of influence tactics in the (online) selling process. Clearly, influence tactics such as social proof, scarcity and authority affect consumer behavior and should be actively considered by online vendors. We have demonstrated the benefits of applying theoretical models developed in face-to-face selling to e-commerce channels, namely, by translating each stage of the ISTEAD adaptive selling process model to online selling. This analysis, and the evaluation of our method of adaptive selling, highlights the importance of “closing the loop” in online selling: just as face-to-face selling benefits from adaptive selling and e-commerce benefits from tailoring sales influence tactics to individual consumers. For online sellers, this means that once they implement influence tactics, and measure consumer responses to these tactics, they can start adapting this content. For businesses, our implementation process should be seen as just one example of what can be done with the presented algorithm-based methodology, as it can be used to implement other e-selling mechanisms and/or used in other digital settings.

Second, e-commerce is rapidly permeating digital channels that are more personal and interpersonal than browser-based Web apps, such as instant messengers and mobile apps, and social media platforms. This implies a significant shift from e-commerce to e-selling. For companies, managing this transition is vital. All of our field trial companies adopted the approach used in this research and continued to use it after the experiments concluded, something they would only do if our adaptive selling method proved profitable. One of the field trial companies, a leader in its field, commented that during the trial period, using e-selling improved its system's internal click-through rate to the same extent that the company's entire Web optimization team had been able to achieve the prior year.

Finally, the current work highlights a change in online sales practice. Currently, most companies decide what content to show to the majority of customers. Our work highlights that moving away from selecting the “best” messages on average toward dynamic adjustment based on measured responses of individual customers is more effective. This finding has major ramifications for how the persuasive content of marketing and sales should be designed and created in the era of interpersonal digital interactions. Namely, we should no longer be looking for a killer slogan or a unique selling proposition, but rather should seek out a versatile toolkit of options that can be optimized for each customer using the adaptive selling method. This is also a key message for advertising agencies and media companies to absorb.

7.2 Theoretical implications and future research directions

Our field trials highlight several future directions for research on automated adaptive selling, which also makes a theoretical contribution to the literature on (offline) adaptive selling. First, given that our approach highlights the choice of influence tactics as content to be adapted, research could examine the possible adjustment of other types of content, such as customers’ inferred product preferences, cognitive style and price sensitivity. Automated adaptive selling provides an inherent opportunity to compare different kinds of content in the sales attempts being adapted and to examine their interactions.

We compared automated adaptive selling to the *status-quo* and to a random selection of influence tactics, but other ways of selecting messages could also be contrasted. For example, it would be worth investigating the conditions that alter the selection of influence tactics at different times of the day or for different contexts. In addition, research on the interaction among multiple influence tactics (or other types of adapted content) would directly address theoretical questions regarding the human processing of influence tactics (e.g. sequentially or in parallel). Future studies should examine a larger variety of influence tactics to enhance understanding of how and to what extent they benefit from dynamic adjustments and profiling. Investigations into the use of multiple influence tactics separately, in bundles, and/or sequentially would provide relevant information about their different roles and interdependencies (Barry and Shapiro, 1992; Payan and McFarland, 2005).

We used Thompson sampling (see S3 for details) to assess the value of gaining new knowledge about customers against the value of using existing knowledge. We ran rigorous simulations that demonstrated the equivalent effectiveness of our method as compared to more computationally taxing methods. Although there are other acknowledged policies for the trading of exploration and exploitation in the computer science literature (Scott, 2010), as far as we know, this fundamental process has not been studied with salespeople: How do they decide, during the adaptive selling process, between eliciting preferences and learning about customers as opposed to trying to close the deal? Automated adaptive selling offers a unique opportunity to systematically change the policies used and to compare the behavior of such automated systems to that of salespeople. A similar approach has already shed light on the decision-making of humans in (gambling) games (Steyvers *et al.*, 2009).

Thus, our automated selling method could actively inform the human-to-human adaptive selling process or be implemented among multichannel retailers (cf. Konuş *et al.*, 2008). Although various scholars have focused on how presenting online product information can contribute to offline sales (Pauwels *et al.*, 2011), the idea of using individual-level estimates obtained online in offline retail settings is underdeveloped. Big data systems are currently built to give “next best actions” or “next based offers” that can be used online or offline

(Chitalia and Berg, 2017). This emerging practice is based on the assumption that individual-level customer behaviors obtained online are beneficial for salespeople in subsequent interactions, regardless of the channel. Restaurant or in-store personnel, for example, may benefit from knowing which influence tactic struck a chord when a specific customer chose to book a table online or to reserve an in-store pick-up after shopping online (Pöyry *et al.*, 2017).

Finally, our study highlights an avenue of research that could profitably be applied to offline adaptive selling: it is theoretically possible to allow computers to generate the exact sales messages to be used. Instead of selecting messages from an existing set (as happened in our studies), we could potentially create computer programs that create individualized sales messages. This represents a new arena of persuasive message content generation.

7.3 Limitations

We focused in our study on the click-through rates to evaluate the success of the selected influence tactics. This is an elementary measure of the success of an online sales pitch. However, the click-through rates should be related to actual sales, browsing behaviors and the time spent looking at products. In offline selling, salespersons are likely to use subtler verbal and nonverbal cues to make their adjustments. Automated adaptive selling potentially provides a unique opportunity to contrast different evaluation methods and to examine the validity and reliability of each one. Likely, in the near future, it will even be possible to make adjustments to sales pitches by examining facial expressions and specific nonverbal behaviors as captured through webcam, sensors or virtual-reality goggles. Thus, future research should consider alternative measures of success, while balancing the measures' costs and benefits.

For practical considerations, we used three out of the six influence tactics developed by Cialdini (1993). While this represents a more conservative test of our implementation of automated adaptive selling, examining all six tactics in future research would be useful.

7.4 Conclusions

The existing literature on adaptive selling opens novel avenues for the customization and personalization of online retail by offering a strong conceptual framework and highlighting many ways in which automated adaptive selling can be carried out. However, and conversely, automated adaptive selling could, in the future, enhance our theoretical understanding of the adaptive selling process. Once the ISTEAM process model is implemented as a technical infrastructure, it will be possible to conduct systematic experiments within each of its stages. On the practical side, our automated adaptive selling method proved effective in significantly improving the click-through rates. Importantly, all three companies that participated in our field trials continued to implement our method after the studies were concluded.

Notes

1. Using the terminology in the subsequent section, *H1* tests the status quo condition, *H2* tests the random condition and *H3* tests the adaptive condition.
2. We include this “neutral” influence tactic, in which a product is not accompanied by any influence tactic in our adaptive selling algorithm, because prior work has demonstrated that some customers respond adversely to the use of influence tactics (Kaptein and Eckles, 2012).

3. In this and the subsequent trials, we have removed references to the brand/company names. We reached an agreement that allows us to publish the results of these trials only if the results were anonymized in this way.
4. The proportion of visitors assigned to the random condition was kept small in all the field studies, as requested by the collaborating companies, whose main interest was the comparison to the *status-quo*.

References

- Agrawal, S. and Goyal, N. (2012), "Further Optimal Regret Bounds for Thompson Sampling", arXiv preprint arXiv:1209.3353.
- Ajzen, I. and Fishbein, M. (1980), *Understanding Attitudes and Predicting Social Behavior*, Prentice Hall, Englewood Cliffs, NY.
- Ansari, A. and Mela, C.F. (2003), "E-customization", *Journal of Marketing Research*, Vol. 40 No. 2, pp. 131-145.
- Ansari, A., Essegai, S. and Kohli, R. (2000), "Internet recommendation systems", *Journal of Marketing Research*, Vol. 37 No. 3, pp. 363-375.
- Barry, B. and Shapiro, D.L. (1992), "Influence tactics in combination: the interactive effects of soft versus hard tactics and rational exchange", *Journal of Applied Social Psychology*, Vol. 22 No. 18, pp. 1429-1441.
- Bodapati, A.V. (2008), "Recommendation systems with purchase data", *Journal of Marketing Research*, Vol. 45 No. 1, pp. 77-93.
- Chitalia, P. and Berg, S. (2017), "Towards an Efficient Method of Modeling Next Best Action for Digital Buyer's Journey in B2B", *Machine Learning and Data Mining in Pattern Recognition: 13th International Conference, MLDM, Springer, New York, NY*, (15-20 July), Vol. 10358, p. 107.
- Cialdini, R.B. (1993), *Influence: Science and Practice*, Harper Collins College Publishers, New York.
- Evans, F.B. (1963), "Selling as a dyadic relationship – a new approach", *The American Behavioral Scientist*, Vol. 6 No. 9, pp. 76-79.
- Franke, G.R. and Park, J. (2006), "Salesperson adaptive selling behavior and customer orientation: a Meta-analysis", *Journal of Marketing Research*, Vol. 43 No. 4, pp. 693-702.
- Gittins, J.C. (1979), "Bandit processes and dynamic allocation indices", *Journal of the Royal Statistical Society: Series B (Methodological)*, Vol. 41 No. 2, pp. 148-177.
- Goldstein, N.J., Cialdini, R.B. and Griskevicius, V. (2008), "A room with a viewpoint: using social norms to motivate environmental conservation in hotels", *Journal of Consumer Research*, Vol. 35 No. 3, pp. 472-482.
- Gretzel, U. and Fesenmaier, D. (2006), "Persuasion in recommender systems", *International Journal of Electronic Commerce*, Vol. 11 No. 2, pp. 81-100.
- Griskevicius, V., Goldstein, N.J., Mortensen, C.R., Sundie, J.M., Cialdini, R.B. and Kenrick, D.T. (2009), "Fear and loving in Las Vegas: evolution, emotion, and persuasion", *Journal of Marketing Research*, Vol. 46 No. 3, pp. 384-395.
- Hall, Z.R., Ahearne, M. and Suhan, H. (2015), "The importance of starting right: the influence of accurate intuition on performance in salesperson–customer interactions", *Journal of Marketing*, Vol. 79 No. 3, pp. 91-109.
- Hauser, J., Liberali, G. and Braun, B. (2009), "Website morphing", *Marketing Science*, Vol. 28 No. 2, pp. 202-223.
- Kaptein, M. and Eckles, D. (2012), "Heterogeneity in the effects of online persuasion", *Journal of Interactive Marketing*, Vol. 26 No. 3, pp. 176-188.

-
- Kaptein, M., Markopoulos, P., de Ruyter, B. and Aarts, E. (2015), "Personalizing persuasive technologies: explicit and implicit personalization using persuasion profiles", *International Journal of Human-Computer Studies*, Vol. 77 No. 1, pp. 38-51.
- Kelman, H.C. and Hamilton, V.L. (1989), *Crimes of Obedience: Toward a Social Psychology of Authority and Responsibility*, Yale University Press, New Haven.
- Kidwell, B., McFarland, R.G. and Avila, R.A. (2007), "Perceiving emotion in the buyer – seller interchange: the moderated impact on performance", *Journal of Personal Selling & Sales Management*, Vol. 27 No. 2, pp. 119-132.
- Konus, U., Verhoef, P.C. and Neslin, S.A. (2008), "Multichannel shopper segments and their covariates", *Journal of Retailing*, Vol. 84 No. 4, pp. 398-413.
- Lynn, M. (1991), "Scarcity effects on value: a quantitative review of the commodity theory literature", *Psychology and Marketing*, Vol. 8 No. 1, pp. 43-57.
- McFarland, R.G., Challagalla, G.N. and Shervani, T.A. (2006), "Influence tactics for effective adaptive selling", *Journal of Marketing*, Vol. 70 No. 4, pp. 103-117.
- Martin, R. and Hewstone, M. (2003), "Social influence processes of control and change: conformity, obedience to authority, and innovation", in Hogg, M.A. and Cooper, J. (Eds), *Psychological Foundations of Attitudes*, Sage, London, pp. 347-366.
- Montgomery, A.L. and Smith, M.D. (2009), "Prospects for personalization on the internet", *Journal of Interactive Marketing*, Vol. 23 No. 2, pp. 130-137.
- Mullins, R.R., Aharne, M., Lam, S.K., Hall, Z.R. and Boichuk, J.P. (2014), "Know your customer: how salesperson perceptions of customer relationship quality form and influence account profitability", *Journal of Marketing*, Vol. 78 No. 6, pp. 38-58.
- Parker, J.R. and Lehmann, D.R. (2011), "When shelf-based scarcity impacts consumer preferences", *Journal of Retailing*, Vol. 87 No. 2, pp. 142-155.
- Pauwels, K., Leeflang, P.S., Teerling, M.L. and Huizingh, K.E. (2011), "Does online information drive offline revenues?: only for specific products and consumer segments!", *Journal of Retailing*, Vol. 87 No. 1, pp. 1-17.
- Payan, J.M. and McFarland, R.G. (2005), "Decomposing influence strategies: argument structure and dependence as determinants of the effectiveness of influence strategies in gaining channel member compliance", *Journal of Marketing*, Vol. 69 No. 3, pp. 66-79.
- Pöyry, E., Parvinen, P. and McFarland, R.G. (2017), "Generating leads with sequential persuasion: Should sales influence tactics be consistent or complementary?", *Journal of Personal Selling & Sales Management*, Vol. 37 No. 2, pp. 89-99.
- Román, S. and Iacobucci, D. (2010), "Antecedents and consequences of adaptive selling confidence and behavior: a dyadic analysis of salespeople and customers", *Journal of the Academy of Marketing Science*, Vol. 38 No. 3, pp. 363-382.
- Scott, S.L. (2010), "A modern Bayesian look at the multi-armed bandit", *Applied Stochastic Models in Business and Industry*, Vol. 26 No. 6, pp. 639-658.
- Spiro, R.L. and Weitz, B.A. (1990), "Adaptive selling: Conceptualization, measurement, and nomological validity", *Journal of Marketing Research*, Vol. 27 No. 1, pp. 61-69.
- Steyvers, M., Lee, M.D. and Wagenmakers, E.J. (2009), "A Bayesian analysis of human decision-making on bandit problems", *Journal of Mathematical Psychology*, Vol. 53 No. 3, pp. 168-179.
- Sujan, H. (1986), "Smarter versus harder: an exploratory attributional analysis of salespeople's motivation", *Journal of Marketing Research*, Vol. 23 No. 1, pp. 41-49.
- Sujan, H., Weitz, B.A. and Sujan, M. (1988), "Increasing sales productivity by getting salespeople to work smarter", *Journal of Personal Selling & Sales Management*, Vol. 8 No. 2, pp. 9-19.

- Thompson, W.R. (1933), "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples", *Biometrika*, Vol. 25 Nos 3/4, pp. 285-294.
- Urban, G.L., Liberali, G., MacDonald, E., Bordley, R. and Hauser, J.R. (2013), "Morphing banner advertising", *Marketing Science*, Vol. 33 No. 1, pp. 27-46.
- Verbeke, W., Dietz, B. and Verwaal, E.J. (2011), "Drivers of sales performance: a contemporary Meta-analysis: have salespeople become knowledge brokers?", *Journal of the Academy of Marketing Science*, Vol. 39 No. 3, pp. 407-428.
- Weitz, B.A. (1978), "Relationship between salesperson performance and understanding of customer decision making", *Journal of Marketing Research*, Vol. 15 No. 4, pp. 501-516.
- Weitz, B.A. (1981), "Effectiveness in sales interactions: a contingency framework", *Journal of Marketing*, Vol. 45 No. 4, pp. 85-103.
- Weitz, B.A., Sujan, H. and Sujan, M. (1986), "Knowledge, motivation, and adaptive behavior: a framework for improving selling effectiveness", *Journal of Marketing*, Vol. 50 No. 4, pp. 174-191.
- Ying, Y., Feinberg, F. and Wedel, M. (2006), "Leveraging missing ratings to improve online recommendation systems", *Journal of Marketing Research*, Vol. 43 No. 3, pp. 355-365.
- Zhang, J. and Wedel, M. (2009), "The effectiveness of customized promotions in online and offline stores", *Journal of Marketing Research*, Vol. 46 No. 2, pp. 190-206.

Further reading

- Hinz, O., Hann, I. and Spann, M. (2011), "Price discrimination in E-commerce? An examination of dynamic pricing in name-your-own-price markets", *MIS Quarterly*, Vol. 35 No. 1, pp. 81-98.

Corresponding author

Maurits Kaptein can be contacted at: m.c.kaptein@uvt.nl

Appendix

Supplementary materials for the article

"Automated adaptive selling"

S1. Impression formation: estimating the effects of sales influence tactics. For the impression formation, we estimate the success of distinct sales influence tactics (or more simply, influence tactics, here abbreviated as SITs) for individual customers. In our operationalization, the success, x , is a dichotomous outcome measure: either the customer clicks on the product associated with the SIT ($x = 1$) or not ($x = 0$). We are interested in the quantity p_{is} , the probability that a customer $i = 1, \dots, i = N$ clicks on a product presented with SIT $s = 1, \dots, s = S$. We assume that $Pr(X = x_{(is)} \mid \dots) \sim \text{Bernoulli}(p_{is})$.

Our estimate of p_{is} is based on both customer-level observations as well as on average-level performance. We obtain point-estimates of SIT effectiveness at each interaction:

$$\begin{aligned}
 n_{is}(t+1) &= n_{is}(t) + 1 \\
 n_{.s}(t+1) &= n_{.s}(t) + 1 \\
 p_{is}(t+1) &= p_{is}(t) + (x_{(is)}(t) - p_{is}(t)) / n(t+1)_{is} \\
 p_{.s}(t+1) &= p_{.s}(t) + (x_{(.s)}(t) - p_{.s}(t)) / n(t+1)_{.s}
 \end{aligned} \tag{1}$$

where n_{is} denotes the number of observations for customer i and SIT s , n_s denotes the overall number of responses to SIT s counted over all customers, and p_{is} and p_s , respectively, denote the estimated click-through probabilities. Furthermore, $x_{(is)}(t)$ and $x_{(s)}(t)$ denote the outcome at interaction t which is a response of a customer i to a SIT s . When observing a data point, four scalars (p_{is} , p_s , n_{is} , n_s) are updated after which the data point is discarded. We thus limit ourselves to an implementation that is fully online (or “streaming”). This restriction bounds the computational complexity of our approach.

In practice, p_{is} is often based on a small number of observations. To improve our estimates of p_{is} , we use a simple shrinkage approach:

$$\hat{p}_{is}(t) = b_{is}(t) p_s(t) + (1 - b_{is}(t)) p_{is}(t) \quad (2)$$

where $b_{is}(t) = 1 / \sqrt{n_{is}(t)}$ is a (heuristic) shrinkage factor. Thus, the estimated success of a SIT for a customer is a combination of the average success and the individual-level success. The more individual level observations n_{is} , the more the final estimate $\hat{p}_{is}(t)$ is influenced by the observations of that individual customer.

We provide as initial starting values $p_s=0.5$ and $n_s=2$ for each s . For each *new* customer, we instantiate (or “copy”) p_{is} and n_{is} as follows for each s :

$$\begin{aligned} p_{is}(t) &= p_s(t) \\ n_{is}(t) &= \frac{1}{2(1-p_{is}(t))p_{is}(t)} \end{aligned} \quad (3)$$

to form an initial impression.

S2. Strategy formulation: earning vs learning. When a customer arrives at the webpage, we need to select a SIT. If the p_{is} were known with complete certainty for a customer i , we would select SIT s with the highest p_{is} . However, the estimates are uncertain, and we are faced with the earning vs learning (or exploration-exploitation) trade-off, see also main text. To solve this trade-off, we are inspired by an approach called Thompson sampling. The basic idea is to assume a prior distribution on the parameters of the reward distribution of every SIT, and at any interaction, select a SIT according to its posterior probability of being the most effective. We can think of our estimates of $\hat{p}_{is}(t)$ and $n_{is}(t)$ as describing a (posterior) *Beta* (α_{is} , β_{is}) distribution with parameters:

$$\begin{aligned} \alpha_{is} &= \hat{p}_{is}(t) n_{is}(t) \\ \beta_{is} &= (1 - \hat{p}_{is}(t)) n_{is}(t) \end{aligned} \quad (4)$$

We implement Thompson sampling using a random draw from each *Beta* (α_{is} , β_{is}) for each s for the current customer i and selecting s with the highest draw.

S3. Simulation Study 1: estimation. We examine, using simulations, the performance of our estimation procedure for a single SIT ($S = 1$). For $T = 10,000$ and $N = 1,000$, we simulate the ($n_i = 10$) observations obtained from customers using the following model:

$$\Pr(X = 1|i, \dots) = p_{true(i)} = 1/(1 + \exp(-(-0.2 + u_i))) \quad (5)$$

where $u_i \sim N(0, 4)$ and the observed ten observations for customer i are drawn from a Bernoulli ($p_{true(i)}$) distribution. We randomize the order of occurrence. We then run through the data set row-by-row, and estimate at each time point t the mean squared error (MSE) using:

$$MSE(t) = \frac{\sum_{i=1}^{i=N(t)} (p_{true(i)} - p_i(t))^2}{N(t)} \quad (6)$$

where $N(t)$ denotes the number of unique customers i observed at t . To examine the performance of our estimation procedure, we compare our approach to three alternative estimation regimes:

- *Comparison to the grand mean:* We maintain an estimate $p_{gm}(t) = \frac{\sum_{i=1}^{i=t} x_i}{t}$ and compute the squared error as by substituting $p_i(t)$ with $p_{gm}(t)$ for each i .
- *Comparison to the individual mean:* We maintain for each customer i an estimated average value of the observations. Thus, we substitute $p_i(t)$ in [equation \(6\)](#) by $p_{im}(t) = \frac{\sum_{i=1}^{i=n_i} x_i}{n_i}$ for each i .
- *Comparison to a hierarchical Bayes logistic model:* We fit a Bayesian hierarchical logistic model using a (hardly informative) Normal Prior on the fixed effect (intercept) and an Inverse-Wishart prior on the random effects u_i . We then obtain the maximum a posteriori (MAP) estimate of $\Pr(X=1|i, \dots)$ as the $p_i(t)$.

[Figure A1](#) shows the performance of our method (red line) averaged over 100 simulation runs compared to the grand mean (pink line), individual mean (green line) and the hierarchical Bayes approach (blue line). Note that grand mean initially ($t < 20$) performs quite competitively, but is quickly outperformed by methods that allow for heterogeneity. The individual mean, after a large number of observations within each individual, performs competitively, but suffers from a slow start for obvious reasons. Our approach and the hierarchical Bayes approach have a fairly similar performance. However, our method is computationally much more feasible: for illustration purposes, a single run for a data stream of length 10,000 took about 160 s for the hierarchical Bayes approach, and only about 30 s for our approach.

S4. Simulation Study 2: the exploration-exploitation trade-off. To examine the performance of our method in addressing the exploration-exploitation trade-off, we simulated the decision problem for $S = 4$ messages using the following hierarchical model to generate our data:

$$\Pr(X = 1|i, \dots) = p_{true}(i) = 1/(1 + e^{-(F\beta + AZ)}) \quad (7)$$

where \mathbf{Z} is a matrix of random effects for individuals i for each s , and β is a vector of fixed effects (omitting the intercept) of the messages. We choose $\beta = \{0.5, 0, -0.5, -0.8\}$ and $\mathbf{Z} \sim MVN(\mathbf{0}, \mathbf{I})$, where \mathbf{I} is a 4×4 identity matrix. We generate a data stream for $T = 100,000$, $N = 10,000$ (thus with $n_i = 10$), where at each time point t , we choose an individual i , then select message s , and finally observe

the response x_t as generated from the model. Because the true data generating probabilities $Pr(X = 1 | i, s, \dots)$ are known, we can simulate an observation, $x_{(i)}^*$ from the message with the highest probability of success (the optimal choice). We compute the *regret*:

$$R(t) = \sum_{i=1}^{i=t} (x_i^* - x_i) \quad (8)$$

of our procedure. We then compare our method to the following strategies for solving the exploration vs exploitation problem:

- *Thompson sampling*: We fit a Bayesian hierarchical model to the observed data using the model specified in [equation \(7\)](#). Next, we obtain a single draw from the posterior distribution of the parameters and select s that maximizes $Pr(x=1 | i, s, \dots)$ given that draw.
- *Greedy selection based on the grand mean*: We estimate the mean success ($x=1$) of each message s given all the observations up to t and select the message with the highest mean, $p_{\cdot,s}$.
- *Greedy selection based on the individual mean*: We estimate for each individual i , for each messages s , the mean success $x=1$. Then, for customer i , arriving in the data stream, we select message s that has the highest estimated mean $p_{i,s}$.

[Figure A2](#) displays the average cumulative regret $R(t)$ over 100 simulation runs as a function of (log) t . Greedy selection based on the individual level mean (green line) performs badly, whereas greedy selection based on the grand mean performs badly initially but seems to catch up for large values of t . Our method (red line) is competitive in its performance to the state-of-the-art (blue line) while being much less computationally demanding.




Figure A1.
Comparisons of
different estimation
procedures for the
impression formation.

Note: Our method (red line) gives a performance that is competitive to alternatives while being easy to compute




Figure A2.
Comparison of the
performance of our
method (red line)
compared to greedy
(green and pink) and
a state-of-the-art
method (blue)