

ADAPTIVE PERSUASIVE MESSAGES IN AN E-COMMERCE SETTING: THE USE OF PERSUASION PROFILES

Kaptein, Maurits, Eindhoven University of Technology / Philips Research, Den Dolech 2, HG 3.45, 5600 MB Eindhoven, the Netherlands, m.c.kaptein@tue.nl

Abstract

Technologies that are intentionally designed to change a person's attitude or behaviors are emergent. Designers of these technologies frequently use implementations of influence strategies to increase the effectiveness of their systems. In this paper we argue that there are large individual differences in users responses to implementations of influence strategies, and as such persuasive technologies would benefit from adapting to these differences. We detail how designers of persuasive technologies could use persuasion profiles to adapt to differences in individual susceptibility to different influence strategies.

To evaluate the notion of persuasion profiles we build an adaptive persuasive system in an e-commerce setting and test this system against a non-adaptive counterpart. We describe how — using Bayesian learning — designers can implement adaptive persuasive technologies. To our knowledge this is the first implementation and evaluation of a working persuasive system that utilizes dynamically created persuasion profiles to increase its effectiveness. We show increased revenues for the system that utilizes persuasion profiles. Finally, we discuss the limitations of the presented evaluation and give suggestions for future research efforts.

Keywords: Persuasion, Persuasive Technology, e-commerce

1 Introduction

Persuasive technologies — technologies that are intentionally designed to change a person's attitude or behavior (Fogg, 1999) — are emergent. Roughly a decade after the groundbreaking work of Fogg (2003), the field of persuasive technologies has taken a big flight and the literature is increasingly populated by records of persuasive applications (IJsselsteijn et al., 2006). Technologies have been designed which influence people to smoke less (Räsänen et al., 2008), assist people in losing weight (Maheshwari et al., 2008), or help people maintain a healthy workout regime (Lacroix et al., 2009). In recent years several persuasive applications have made their way out of the research labs and currently a broad range of persuasive applications is commercially available.

Persuasive technologies employ a multitude of methods to attain their intended attitude or behavioral change goals. Social science findings such as the power of *goal setting* (Locke et al., 1981), the effects of *framing* (Fogg, 2003), and classical work on *conditioning* (Rescorla, 1988) are often used to enhance the effectiveness of persuasive systems. A prime example of implementations of these techniques can be found in an e-commerce setting: Designers of online stores employ similar principles and combine knowledge of design, technology, and social science to increase consumer satisfaction and, eventually, revenues.

In this paper we present the implementation of a persuasive system which adapts its methods for attitude or behavioral change to responses of its individual users. To our knowledge this is the first implementation of a system that uses *persuasion profiles*, which were described in detail by Kaptein & Eckles (2010) during the 2010 persuasive technology conference. We first review the social science findings on influence strategies and explain how these findings lead to the idea of persuasion profiling. Next, we describe the implementation of a system that uses persuasion profiles in an online commerce setting and evaluate its performance. Finally we discuss the limitations of the current evaluation and give suggestions for further design and research efforts.

1.1 Influence strategies

The methods employed by designers of persuasive technologies to increase the effectiveness of their systems are often referred to as *influence strategies*. The array of influence strategies available for changing attitudes and behaviors can be overwhelming. Theorists have varied largely in how they individuate influence strategies: Cialdini (2001) develops six principles at length, Fogg (2003) describes 40 strategies, and others have listed over a 100 (Rhoads, 2007). While there is variation in their individuation, influence strategies are a useful level of analysis that helps to group and distinguish specific influence tactics. A large body of research focuses around the *six principles of persuasion* as identified by Cialdini (2001). These principles are:

1. *Reciprocity*: People feel obligated to return a favor, thus when a persuasive request is made by a person the receiver is in debt to, the receiver is more inclined to adhere to the request R. Cialdini (2004).
2. *Scarcity*: When something is scarce, people will value it more. Announcing that a product or service is scarce will favor the evaluation and increase the chance of purchase (West, 1975).
3. *Authority*: When a request or statement is made by a legitimate authority, people are more inclined to comply or find the information credible (Milgram, 1974).
4. *Commitment and consistency*: People do as they said they would. People try to be consistent with previous or reported behavior, resolving cognitive dissonance by changing their attitudes or behaviors to achieve consistency. If a persuasive request aligns with previous behavior people are more inclined to comply (R. Cialdini, 2001)

5. *Consensus*: People do as other people do. When a persuasive request is made people are more inclined to comply when they are aware that others have complied as well R. Cialdini (2004); Ajzen & Fishbein (1980).
6. *Liking*: We say “yes” to people we like. When a request is made by someone we like, we are more inclined to act accordingly (R. Cialdini, 2001).

These influence principles (or strategies) are related to *how* a certain attitude or behavioral change request is made, and not necessarily tied to the goal of the actual request (Kaptein & Eckles, 2010). This enables us to distinguish the *end* of a request (e.g. a persuasive application urges you to work out more) from the *means* in which the request is made (e.g. by showing you how your friends are working out, or by giving you expert advice). This property — *end-independence* — makes influence strategies useful not just for typifying specific influence attempts, but more broadly to function as a level of analysis to describe and possibly predict the effects of different implementations of the same strategy at later points in time or in different contexts.

1.2 Individual Differences in Responses to Influence Strategies

While generally effective in increasing the chance of success of an attitude or behavioral change request, studies into the effect of influence strategies have not always painted a clear picture. Because of seemingly conflicting results, research efforts have been devoted to studying individual differences in responses to influence strategies.

Several personality traits have been consistently linked to differences in processing persuasive messages: A person’s *need for cognition* is their degree of motivation to think (Cacioppo & Petty, 1982; Petty & Wegener, 1999). An increase in need for cognition makes cognitive processing of messages more likely (Cacioppo et al., 1986). Hence, those individuals scoring high in need for cognition are less influenced by implementations of influence strategies than those scoring low. It has also been shown that refraining from usage of influence strategies for those low susceptible to these strategies increases compliance (Kaptein et al., 2010).

Next to this individual difference in *overall* responses to influence strategies, research has shown systematic variation in responses to *specific* influence strategies: The *preference for consistency* scale predicts responses to the commitment strategy, such that for participants low on this scale these strategies are generally ineffective (Cialdini et al., 1995; Guadagno et al., 2001). Similar strategy specific individual differences have also been found for the *consensus* strategy (Kaptein et al., 2009).

1.3 Persuasion Profiles

Given the large individual differences in responses to influence strategies it seems worthwhile to create persuasive systems that adapt to their users’ susceptibility to individual strategies. Such systems could use a persuasion profile — *a collection of expected effects of different influence strategies for a specific individual* (Kaptein & Eckles, 2010) — to determine the way in which their advocacy is presented. Given strong indications of backfiring in cases of “wrong” influence strategy usage (Kaptein et al., 2010), systems that use persuasion profiles to tailor their influence attempts are expected to gain higher levels of compliance than systems that do not adapt to their users’ susceptibilities to different strategies.

1.4 Outline

In this paper we describe an implementation of persuasion profiles in an e-commerce setting. We describe how one can track, and adapt to, the effectiveness of multiple implementations of influence strategies for specific users. We compare the performance of an e-commerce site prior to the implementation of persuasion profiles to the period in which the site incorporated persuasion profiles.

After showing an increase on several performance indicators we discuss the limitation of this method of evaluation and provide pointers for future work.

2 Implementing Persuasion Profiles

Any product offered on an e-commerce site can be displayed differently to different users, and thus can be accompanied by one or more implementations of influence strategies. While a product can be both “*a special offer*” and an “*all times bestseller*” the work on individual differences in responses to influence strategies dictates that presenting a product with both of these strategies is not optimal for all users. Some users will be more inclined to buy the product when presented only with the scarcity strategy (*special offer*), while others will be more swayed only by the consensus strategy (*bestseller*). Finally, for those high in the need for cognition, *not* presenting any these types influence strategies would increase conversion.

2.1 Implementations of Influence Strategies

In the e-commerce domain one can currently find a multitude of examples of implementations of different influence strategies. Implementations of each of the six influence strategies identified by Cialdini (2001) are already in use:

1. *Reciprocity*: Online stores hand out a small gift to potential consumers before presenting their products
2. *Scarcity*: Online stores promote special discounts or limited time offers.
3. *Authority*: Online stores use expert endorsements or (positive) critics reviews to promote products.
4. *Commitment and consistency*: Online stores try to make consumers more committed by acknowledging product preference — e.g. add to “Wish list”.
5. *Consensus*: Online stores often show how many people already bought a specific product.
6. *Liking*: Trough social media online stores use friend recommendations to increase conversion.

With obvious constraints imposed by honesty and reality, these implementations can be implemented for each product on an e-commerce site interchangeably. Using these influence strategies is thus distinct from (e.g.) recommender systems which adapt the *end* proposition (the offered product) instead of the *means*: the way in which the advocacy is presented.

2.2 Technological Implementation

To create a system that implements persuasion profiles — a systems that adapts its usage of influence strategies based user’s responses to specific influence strategies — we believe there are three prerequisites that need to be met: First, one should be able to *manipulate* content presentation such that multiple influence strategies can be implemented. Second, one needs to be able to *identify* individual consumers. Third, it needs to be possible to *measure* the outcome of the persuasive attempt and update subsequent product presentations accordingly.

To enable this functionality we created a visitor tracker which assigns a unique id to each new visitor to an online e-commerce site and stores this as a cookie on the user’s machine¹. Next, the tracker retrieves the preferred influence strategy from a database – which contains the users persuasion profile

¹ This method of identifying individuals is not 100% accurate. However, it’s the best we can currently do without requiring any user action.

— and adapts the site accordingly. Finally, the tracker measures whether a target is clicked and then stores the outcome of the influence attempt. When a user views multiple products the implementations of influence strategies that are presented are changed dynamically between views to optimize the user’s response.

2.3 Adaptivity Algorithms

To adapt the chosen influence strategy to user’s responses to the influence strategy we used a simple Bayesian approach. The probability of a single user clicking the target can be represented as a binomial random variable, $B(n,p)$, where n denotes the number of product views accompanied by that particular influence strategy and p denotes the probability of success. Given M different influence strategies one can compute for each individual, for each strategy, probability $p_{mi} = k_{mi}/n_{mi}$, where k_{mi} is the number of observed successes after representation of strategy m , n_{mi} times to a user i . It makes intuitive sense to present a user with the strategy with the highest probability of success, thus the message where p_{mi} is highest.

If we have a large number of observations, N , for one user, this would make perfect sense. However, this decision rule will not inform a decision for a newly observed user. For a *new* user of the e-commerce store one would likely present the message m for which p_m is maximized over previously observed users². Actually — given Stein’s result (Stein, 1955; Efron & Morris, 1975) — for every user a weighted average of the p_{mi} for the individual user and those of other users — one where the estimated p_{mi} for an individual is “shrunk” toward the population mean — will provide a better estimate than an estimate based on observations of a single user alone. E.g., if the authority strategy is effective 20% of the time over all users and only 10% of the time for the specific user under consideration, the best estimate of the (real) effectiveness of the authority message p_{A1} for this user is a weighted average of these two.

A common way of modeling such a binomial random process is to use the Beta-Binomial model (Wilcox, 1981). The Beta, $Beta(\alpha,\beta)$, distribution functions as a conjugate prior to the binomial and thus allows for dynamic updates of the estimates of strategy effectiveness which are easily analytically traceable. If we re-parameterize the beta distribution as follows:

$$\pi(\theta | \mu, M) = Beta(\mu, M)$$

where $\mu = \frac{\alpha}{\alpha + \beta}$ and $M = \alpha + \beta$, then the expected value of the distribution is given by:

$E(\theta | \mu, M) = \mu_m$. In our system μ_m represents the expected probability of a successful influence attempt by a specific strateg. After specifying the probability of success μ_m of strategy m and the certainty surrounding this estimate, σ_m^2 , we can use this as our prior expectancy about the effectiveness of a specific message and update this expectancy by multiplying it by the likelihood of the observations, k , to obtain the distribution of our posterior expectation:

$$p(\theta | k) \propto l(k | \theta)\pi(\mu, M)$$

$$p(\theta | k) = Beta(k + M\mu, n - k + M(1 - \mu))$$

The newly obtained Beta distribution, $B(\mu, M)$, functions as our probability distribution with a new point-estimate of the effectiveness of the presented message given by:

² This is assuming the error costs—the effects of presenting the wrong strategy—are equal for each strategy.

$$E(\theta | k) = \frac{k + M\mu}{n + M}.$$

We maintain both a population wide (between-user) estimate of the effectiveness of different strategies based on all available observations and an individual level estimate based only on a specific user's response. After a users session the individual level estimate p_{mi} for a specific strategy m is shrunk towards the grand-mean (p_m) using the following shrinkage factor:

$$\hat{p}_{mi} = p_m + \frac{\sigma_m}{\sigma_i + \sigma_m} (p_{mi} - p_m)$$

where σ_m denotes the variance over users and σ_i denotes the variance within a user. The process described here enables estimation of strategy effectiveness of each influence strategy for each user.

3 Evaluation

After implementing the system which enables estimation of the strategy effectiveness and dynamic adaptation of the presentation of different products we set out to evaluate our use of persuasion profiles. In this section we describe how the test was setup and detail our initial results after just over four weeks of testing.

To evaluate our adaptive persuasive system it was implemented on the affiliate online store kinderkleertjes.com (*children's-clothing* in Dutch). Store performance prior to, and post implementation of, the adaptive persuasive system on a number of performance indicators is compared³.

3.1 Method

Kinder-kleertjes.com offers a selection of over 1.200 products. Products are offered by two affiliate programs and the website aims at attracting traffic through search engines and establishing click-through to the two final vendors. The site has been running from the beginning of July 2010. The website is rather limited in size, attracting 387.5 visitors a month on average during an almost 6 month period (till the 22nd November 2010). Our evaluation of the adaptive persuasive system consists of a comparison of several indicators of site performance from early July 2010 till the 21st of October 2010 when the system was not in place (baseline period), with the subsequent month when the adaptive persuasive system was implemented.

The homepage of the online store presents a random collection of 40 of the offered products together with pictures and a single sentence description. Once a visitor clicks on one of the products (or enters the site using a search term directly pointing at a product page) a product is displayed which shows a large image and a textual description of the product. The persuasive strategies are implemented on this page.

We choose to examine only a limited number of strategies and only implemented two strategies: Scarcity and Consensus. For the scarcity implementation an image that said "Special offer" and the text "*This clothing item is available today at a special discount rate*" was displayed. For the Consensus strategy an image that stated "Bestseller" and the message "*This is one of our bestselling clothing items. This product is popular with many*" were added. By keeping the original "no strategy"

³ This store is a minor player in the Dutch affiliate market. We are currently looking for partners to perform large scale tests of our system.

product representation as used from early July till the 21st of October we ended up with three product display versions that could be used for every product on the site.

To start the adaptive persuasive system an initial estimate was made of the conversion effectiveness for each of the product presentations (e.g. the probability that a single user clicked either the “Buy now” button or the image and was redirected to the vendor’s site). The estimate of the effectiveness of the *No Strategy* version was manually set to 20% ($p_m = 0.2$), that of the *Scarcity* strategy to 30% and that of the *Consensus* strategy to 24%. By setting the variances of these estimates high (at 0.1) these initial estimates are variable to new data. Practically, this implementation ensures that the first user after deployment of the adaptive persuasive system on the 21st of October did not receive the *No Strategy* product presentation but rather received the *Scarcity* version. Based on the behavior of this first user — either clicking or not clicking on the product — a new estimate of the effectiveness of the scarcity strategy was made both for this user as well as for subsequent users.

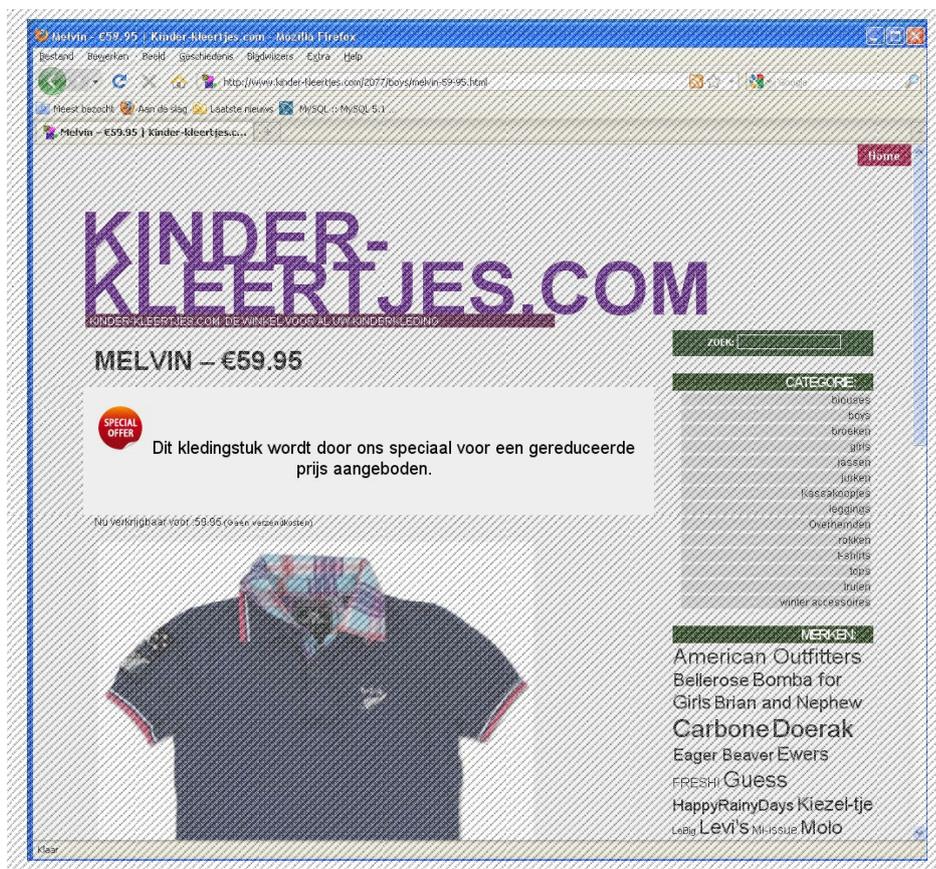


Figure 1. Overview of the product presentation page on kinder-kleertjes.com. The non-shaded area indicates the location of the implementations of the influence strategies *Consensus* or *Scarcity* (as described in the text).

3.2 Measures

To evaluate the adaptive persuasive we examine a number of performance indicators. First of all we examine how the effectiveness estimates of the different influence strategies change over time. This is done both at a between-user level and at a within-user level. Next, we compare the period in which the system was implemented with the period in which only the *No Strategy* product representation was used. We compare *click through rates* and *revenues* per unique visitor over the two periods.

3.3 Results

In the baseline period from the 8th of July till the 21st of October 1339 unique visitors visited kinderleertjes.com. Due to organic visitors and advertisements via Google Adwords, 831 unique visitors visited the site during the employment of the adaptive persuasive system.

3.3.1 Strategy Effectiveness Estimates.

Figure 2 shows the progression of the estimated (between-user) effectiveness of the three product presentations over time. It is clear that initially the estimates are variable and they are haphazardly updated based on new observations. However, after a few days the estimates stabilize and the certainty increases. Both product representations that implement a persuasive strategy significantly outperform the No Strategy implementation within 2 weeks after deployment.

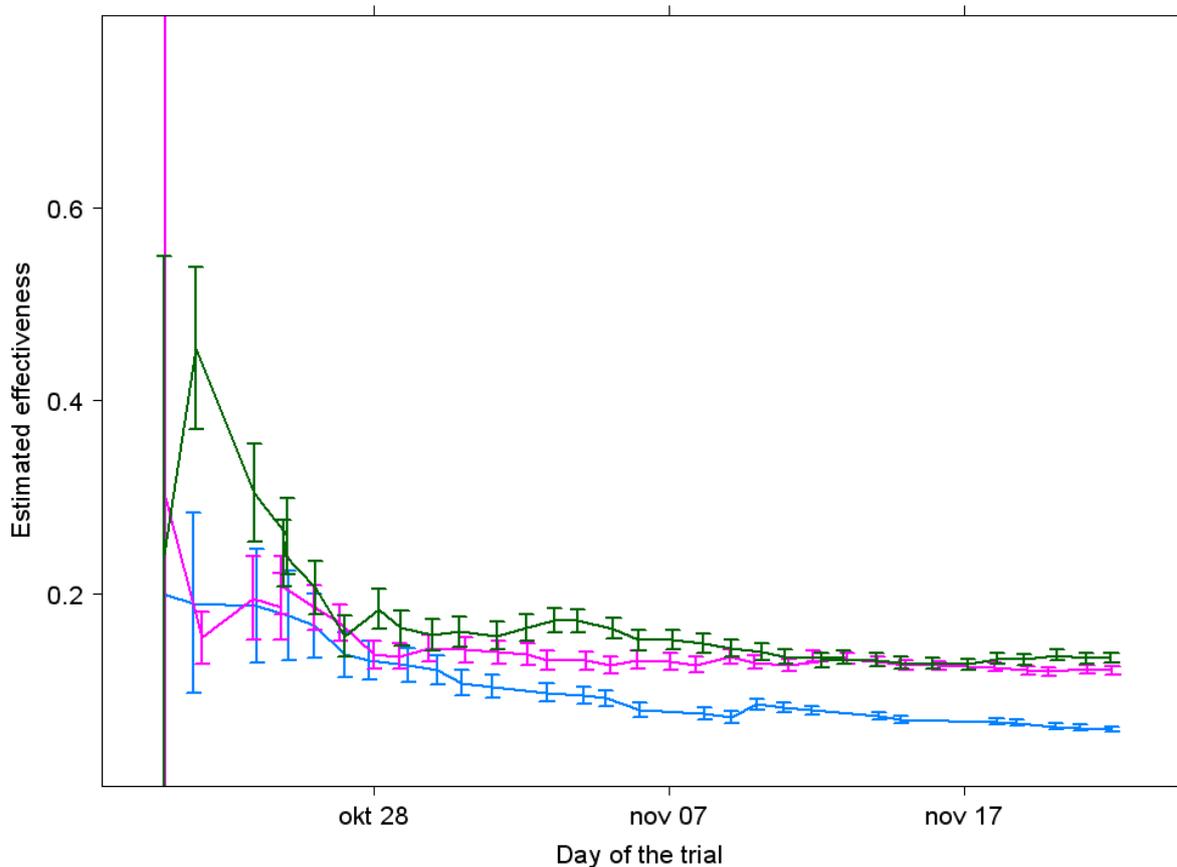


Figure 2. Progression of between-user (overall average) effectiveness estimates of the three strategies and indications of the certainty of the estimates. The blue line represents the effectiveness estimates over time of the 'Neutral' presentation. The green and pink line are the estimates for the Consensus and Scarcity implementations respectively.

While Figure 2 indicates that using persuasive strategies can aid the performance of the affiliate store by increasing click through rates, this graph alone does not show that different strategies should be presented to different users. To examine whether adaptivity is actually necessary (e.g. for some users the Consensus strategy — which is most effective overall — is not the preferred strategy) Figure 3 presents the individual level estimates of three randomly selected users who viewed at least 8 products on the website.

Figure 3 shows that for the users in the middle panel the most effective strategy is not the *Consensus* strategy but rather the *Scarcity* strategy. Hence, when adapting to users only on a *between-user* level — thus displaying the *Consensus* strategy to every user — one would obtain suboptimal results for this user. This observed individual difference is in-line with previous social science findings.

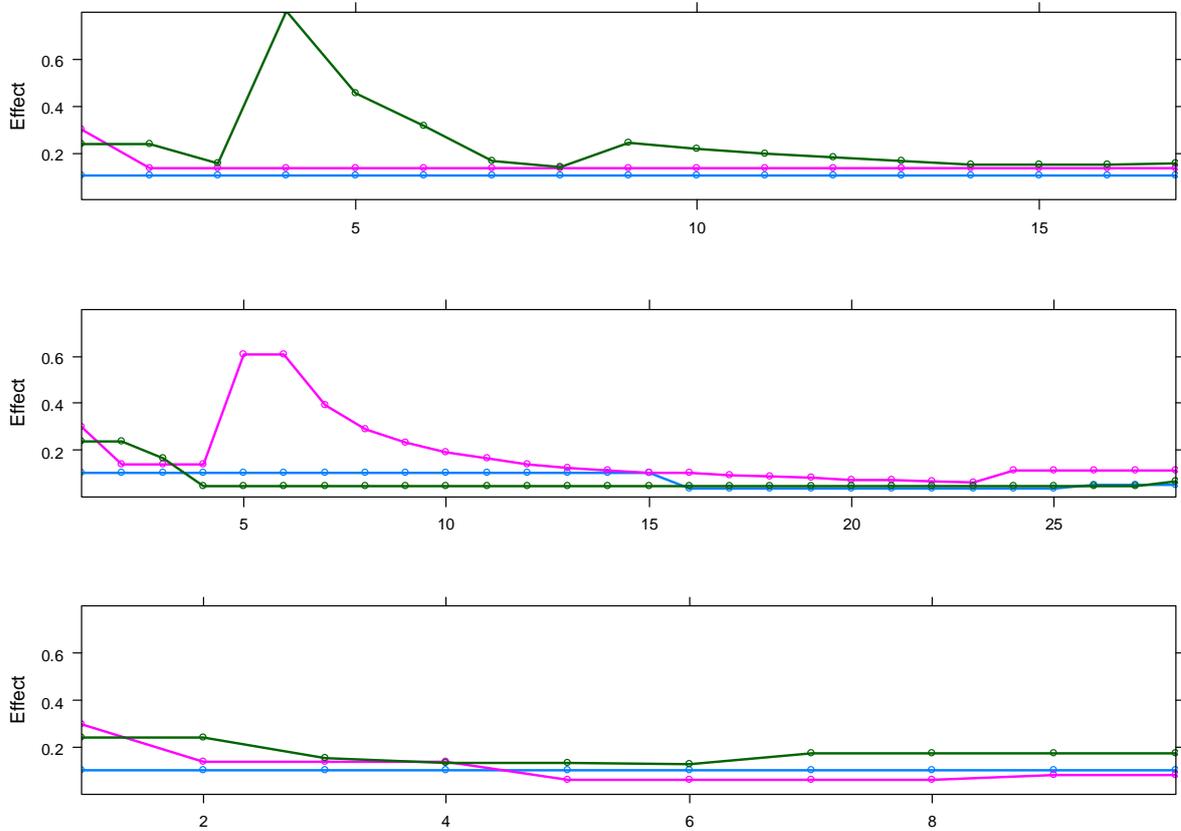


Figure 3. *Estimated effectiveness of the different product presentations for three randomly selected users. The top panel indicates that this user was first presented with the scarcity implementation which did not lead to conversion. Next, the consensus implementation was shown which was also unsuccessful. Eventually however the consensus strategy was effective twice and was estimated most effective for this user.*

3.3.2 Performance

Next to examining the predicted strategy effectiveness we examine how the online store performs before and after usage of persuasion profiles. The first performance indicator is the proportion of clicks on viewed products per user. In the baseline period 14.4% of the users of the site eventually clicked on one of the products and was taken to the vendor’s homepage. In the adaptive persuasion period this number was higher: 18.3%, showing increase in click through in the adaptive persuasion period. This increase in click through rates is statistically significant ($\chi^2=5.766$, $p=0.016$).

The average revenue generated per unique visitor in the baseline period was 0.037 Euro’s. In the the adaptive persuasion period this figure increased to 0.046 Euro’s per visitor. The average revenue per visitor is thus higher in the persuasive system period than in the baseline period, however — due to large variation — this difference is not statistically significant ($p>0.05$).

To further inspect the comparison between the baseline period and the adaptive persuasive system implementation period, Figure 4 shows an overview of the average proportion of clicks per visitor for

each day from the start of the measurement period. The figure shows that the daily proportion of users who navigate to the vendor’s page fluctuates heavily. However, a increase in clicks is visible since the start of the adaptive persuasive system period, a difference indicated by the difference between the solid and dotted horizontal lines which indicate the mean performance of each period. The vertical dotted line presents the separation between the baseline and the adaptive system periods.

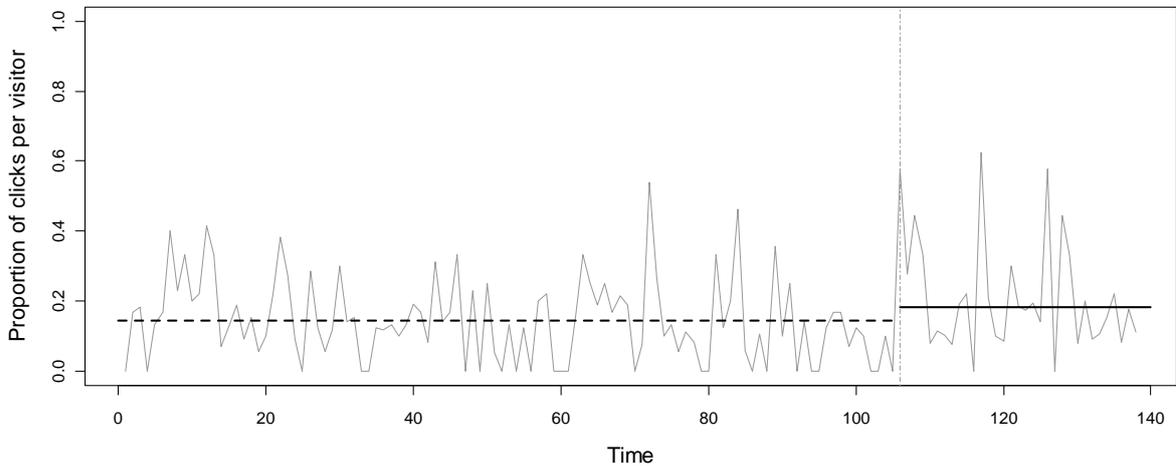


Figure 4. *Proportion of clicks per visitor since measurement start. The dashed horizontal line is the mean conversion rate prior to deployment of the adaptive system; the solid horizontal line is the mean conversion rate after deployment. The gray line is the daily conversion rate, and the vertical gray dotted line shows the cut-off between baseline and deployment.*

4 Discussion

In this paper we detailed the first implementation — as far as we know — of an adaptive persuasive system that uses *persuasion profiles* to adapt to individual differences in peoples responses to influence strategies. We tried to introduce the concept of *influence strategies* and we provided support for the argument of large individual differences in responses to these strategies. Next, we described an implementation of a system that uses these individual differences to adapt product presentations. We presented an initial evaluation of this system by (a) inspecting the adaptive process over time and (b) comparing the performance of the adaptive persuasive system with the “baseline” period in which the system was not yet in use on the affiliate online store used for the evaluation. Results show that influence strategy usage increases conversion and that there are indeed individual differences in responses to implementations of influence strategies. The adaptive persuasive system increased overall site performance on two key-performance indicators.

In our evaluation we managed to demonstrate the benefits of the use of persuasion profiles based on two simple implementations of influence strategies in a real e-commerce setting. A significant increase in click-through was obtained, and a direct increase in revenue resulted. These results however should be interpreted with caution: As Figure 4 shows there is large variability in click-through rates, and there is also a large variability in per visitor revenue over multiple days. Hence, while we believe that the identified trend of increased conversion by the use of persuasion profiles will hold, actual effect-size estimates should be based on larger scale evaluations. It would be feasible to evaluate the adaptive persuasive system over a longer period of time on a website that attracts more visitors.

With the current evaluation — comparing an implementation period with a prior established baseline period — a number of questions remains unanswered. First of all, changes in time or season (e.g. people buy more in autumn than in summer) most probably influence the results. An evaluation in

which half of new visitors is assigned to the original version of the online store, and half to the adaptive persuasive version — and where thus time is the same for both experimental groups — would prevent this confound. Please note that we have only reported relative indicators (e.g. estimates per visitor) to minimize effects originating from an overall increase in visitor rates during the implementation period. However, it is still the case that during November individual consumers could be more inclined to buy products they browse online than during previous months.

Second, since the current comparison is between the original version of the e-commerce site, without any implementation of influence strategies, and a full adaptive system it is unclear exactly which difference between the two systems increases conversion rates. It is worthwhile to setup an evaluation which would also compare (e.g.) usage of the “best” (between-user) influence strategy with the adaptive system, or usage of a randomly selected influence strategy versus the adaptive system. With these evaluations it is possible to reliably rule out mere effects of (a) influence strategy optimization at an average (between-user) level and effects of changing product representations within a session.

The system design as detailed in this paper presents a rudimentary approach of an adaptive persuasive system. Not only does the system implement only two of the six influence strategies, it also uses only one specific implementation of these strategies which is identical for each product presentation. Systems that employ a broader range of strategies and subsequent implementations probably support a larger visitor base. Besides this limitation the adaptive algorithms presented here are basic. While the individual beta-binomial models, in combination with shrinkage towards overall means after a shopping session has ended, provide a convenient and easy to compute starting point there are several feasible properties lacking in this algorithm. First of all, the algorithm does not specifically model relationships between different strategies (e.g. visitors that are more swayed by consensus arguments are also more inclined to respond to implementation of the liking strategy). Second, the presented model is not flexible enough to allow for additional predictors to explain the effectiveness of different influence strategies. The latter would be necessary to utilize (e.g.) demographic information users. In future work we hope to implement multilevel logistic regression models to predict clicking behavior so we can both include a flexible covariance matrix of strategy x subject effects, and we can easily integrate background characteristics that are known about the user based on previous purchases.

4.1 General Conclusions

In this paper we detailed the implementation of an adaptive persuasive system that used persuasion profiles to adapt product presentations on an e-commerce site to the responses of its users. The resulting profile gives, for each visitor, an estimate of their “persuadability” to different influence strategies. Comparison of the performance of an e-commerce site which uses persuasion profiles with one without such profiles shows increased conversion. While this might be commercially attractive, and while the e-commerce setting provides a great environment for testing adaptive persuasive systems, we feel designers should be cautious: A visitor’s persuasion profile represents private data similar to a personality profile and should thus be handled with care. Regulation and proper ethical consideration of possible sharing and re-use of persuasion profiles are feasible.

References

- Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, 42 (1), 116-131.
- Cacioppo, J. T., Petty, R. E., Kao, C. F., & Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology*, 51(5), 1032 - 1043.
- Cialdini, R. (2001). *Influence, science and practice*. Boston: Allyn & Bacon.

- Cialdini, R. (2004). The science of persuasion. *Scientific American Mind*, 284, 76-84.
- Cialdini, R. B., Trost, M. R., & Newsom, J. T. (1995). Preference for consistency: The development of a valid measure and the discovery of surprising behavioural implications. *Journal of Personality and Social Psychology*, 69, 318–328.
- Efron, B., & Morris, C. (1975, Jun). Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350), 311–319.
- Fogg, B. J. (1999). Persuasive technologies: Introduction. *Communications of the ACM*, 42 (5), 26-29.
- Fogg, B. J. (2003). *Persuasive technology: Using computers to change what we think and do*. Morgan Kaufmann.
- Guadagno, R. E., Asher, T., Demaine, L. J., & Cialdini, R. B. (2001). When Saying Yes Leads to Saying No: Preference for Consistency and the Reverse Foot-in-the-Door Effect. *Personality and Social Psychology Bulletin*, 27(7), 859–867.
- IJsselsteijn, W., Kort, Y. de, Midden, C. J. H., Eggen, B., & Hoven, E. van den. (2006). Persuasive technology for human well-being: Setting the scene. In W. IJsselsteijn, Y. de Kort, C. J. H. Midden, B. Eggen, & E. van den Hoven (Eds.), *Persuasive* (Vol. 3962, pp. 1–5). Springer.
- Kaptein, M., & Eckles, D. (2010). Selecting effective means to any end: Futures and ethics of persuasion profiling. In T. Ploug, P. Hasle, & H. Oinas-Kukkonen (Eds.), *Persuasive technology* (pp. 82–93). Springer Berlin / Heidelberg.
- Kaptein, M., Lacroix, J., & Saini, P. (2010). Individual differences in persuadability in the health promotion domain. In T. Ploug, P. Hasle, & H. Oinas-Kukkonen (Eds.), *Persuasive technology* (pp. 94–105). Springer Berlin / Heidelberg.
- Kaptein, M., Markopoulos, P., Ruyter, B. de, & Aarts, E. (2009). Can you be persuaded? individual differences in susceptibility to persuasion. *Proceedings of Interact 2009*. LNCS 5726, 115–118.
- Lacroix, J., Saini, P., & Goris, A. (2009). Understanding user cognitions to guide the tailoring of persuasive technology-based physical activity interventions. In S. Chatterjee & P. Dev (Eds.), *Persuasive* (Vol. 350, p. 9). ACM.
- Locke, E., Shaw, K., Saari, L., & Latham, G. (1981). Goal setting and task performance: 1969-1980. *Psychological Bulletin*, 90(1), 125-152.
- Maheshwari, M., Chatterjee, S., & Drew, D. (2008). Exploring the persuasiveness of "just-in-time" motivational messages for obesity management. In H. Oinas-Kukkonen, P. F. V. Hasle, M. Harjumaa, K. Segerstahl, & P. Øhrstrøm (Eds.), *Persuasive* (Vol. 5033, pp. 258–261). Springer.
- Milgram, S. (1974). *Obedience to authority*. London: Tavistock.
- Petty, R. E., & Wegener, D. T. (1999). The elaboration likelihood model: Current status and controversies. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (p. 416-72). New York: Guilford Press.
- Räsänen, T., Oinas-Kukkonen, H., & Pahlila, S. (2008). Finding kairos in quitting smoking: Smokers' perceptions of warning pictures. In H. Oinas-Kukkonen, P. F. V. Hasle, M. Harjumaa, K. Segerstahl, & P. Øhrstrøm (Eds.), *Persuasive* (Vol. 5033, pp. 254–257). Springer.
- Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43(3), 151 - 160.
- Rhoads, K. (2007). *How many influence, persuasion, compliance tactics & strategies are there?* <http://www.workingpsychology.com/numbertactics.html>
- Stein, C. (1955). Inadmissibility of the usual estimator of the mean of a multivariate normal distribution. In *Proceedings of the third Berkeley symposium on mathematical statistics and probability*. (pp. 197–206). Berkeley: Univ. Cal. Press.
- West, S. G. (1975). Increasing the attractiveness of college cafeteria food: A reactance theory perspective. *Journal of Applied Psychology*, 60, 656-658.
- Wilcox, R. R. (1981). A review of the beta-binomial model and its extensions. *Journal of Educational Statistics*, 6, 3–32.