

Personalizing persuasive technologies : explicit and implicit personalization using persuasion profiles

Citation for published version (APA):

Markopoulos, P., Kaptein, M. C., Ruyter, de, B. E. R., & Aarts, E. H. L. (2015). Personalizing persuasive technologies : explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies*, 77, 38-51. <https://doi.org/10.1016/j.ijhcs.2015.01.004>

DOI:

[10.1016/j.ijhcs.2015.01.004](https://doi.org/10.1016/j.ijhcs.2015.01.004)

Document status and date:

Published: 01/01/2015

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles[☆]



Maurits Kaptein^{a,*}, Panos Markopoulos^b, Boris de Ruyter^c, Emile Aarts^b

^a Radboud University, Nijmegen Archipelstraat 13, 6524 Nijmegen, The Netherlands

^b Eindhoven University of Technology, The Netherlands

^c Philips Research, The Netherlands

ARTICLE INFO

Article history:

Received 8 May 2014

Received in revised form

2 December 2014

Accepted 16 January 2015

Communicated by E. Motta

Available online 24 January 2015

Keywords:

Persuasive technology

Persuasion profiling

Adaptive persuasive systems

ABSTRACT

This paper discusses how persuasive technologies can be made adaptive to users. We present persuasion profiling as a method to personalize the persuasive messages used by a system to influence its users. This type of personalization can be based on *explicit* measures of users' tendencies to comply to distinct persuasive strategies: measures based on standardized questionnaire scores of users. However, persuasion profiling can also be implemented using *implicit*, behavioral measures of user traits. We present three case studies involving the design, implementation, and field deployment of personalized persuasive technologies, and we detail four design requirements. In each case study we show how these design requirements are implemented. In the discussion we highlight avenues for future research in the field of adaptive persuasive technologies.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

We have entered an era of persuasive technology of interactive computing systems intentionally designed to change people's attitude or behavior (cf. Fogg, 2002). A substantial body of research has demonstrated the feasibility of these technologies in a variety of contexts and for different ends, e.g., advertising (Kaptein and Eckles, 2012), promoting healthy or pro-social behaviors (Lambert, 2001; Morris and Guilak, 2009; Consolvo et al., 2008, 2009), and reducing energy consumption (see, e.g., Svane, 2007; Midden et al., 2008; Bang et al., 2006; Dillahunt et al., 2008). Still, *reliably* affecting an individual's attitude or behavior remains an elusive goal (Oinas-Kukkonen and Harjumaa, 2008). This is true despite the argument made by Fogg and Eckles (2007), in their book *Mobile Persuasion*, that persuasive systems *could* be more persuasive than their human counterparts. Their arguments are based on a number of empirical investigations showing that humans respond similar to computers as they do to humans (e.g., Nass et al., 1996; Fogg and Nass, 1997a,b) and that, compared to humans, computers could be more persistent and “always on” (Fogg, 2009; Preece, 2010).

To be effective persuasive systems should deliver the right message, at the right time, in the right way. This very general maxim (and truism) emphasizes three key elements for successful attitude and

behavior change: First, the target of the persuasive attempt needs to be receptive to the end goal of the attempt. Here, with the term “end goal” we refer to the target attitude or behavior that the technology was intentionally designed to promote (see Fogg, 1998, for a discussion on the intentionality of persuasive systems). Second, the message needs to be delivered at a time that enables the recipient to attend to it, and, if immediate action is required, one that provides the opportunity for the action (Faber et al., 2011). Finally, large variation can exist in the way in which a persuasive request is framed: a message aiming to persuade users to work out more could read “80% of users runs at least once a week” or “Fitness experts recommend that you run at least once a week”. In both cases the end goal is the same, but the argument differs substantially.

Unfortunately, the *right* time, the *right* message, and the *right* way for a persuasive request are hard to determine at design time, without knowing the specific situation and person concerned. A solution to this is to create *adaptive* persuasive systems; systems that adapt the message, the timing, and the persuasive approach to the situation at hand. The notion of *ambient* persuasion has been proposed as a (partial) answer to this challenge (cf. Aarts et al., 2007; Kaptein et al., 2009). Ambient persuasion combines the notion of *ambient intelligent systems*—systems that build on the large scale integration of electronic devices and the ubiquitous availability of digital information—and persuasive technologies; systems aimed at changing users' attitudes or behaviors (Kaptein et al., 2009). In an ambient intelligent world, massively distributed devices operate collectively while embedded in the environment using information and intelligence that is hidden in the interconnection network. Context sensing in this setting could help determine appropriate

[☆]This paper has been recommended for acceptance by E. Motta.

* Corresponding author. Tel.: +31 621262211.

E-mail addresses: maurits@mauritskaptein.com (M. Kaptein), p.markopoulos@tue.nl (P. Markopoulos), boris.de.ruyter@philips.com (B. de Ruyter), e.aarts@tue.nl (E. Aarts).

persuasive ends, fitting the user context and activities. Embedded and ubiquitous computing devices can help present the message at a location that will be noticed by users, fitting their activity and context. Till now however less is published about how to effectively personalize the persuasive approach. Our research addresses this challenge by using *persuasion profiles* to enable the personalization of the framing of persuasive attempts.

Adapting the persuasive approach to the persuadee has long been advocated throughout many fields that study persuasion. For example, marketers advocate adapting sales tactics to consumers (McFarland et al., 2006), and health-care professionals promote tailoring of the persuasive principles used to gain medication compliance (e.g., Strecher et al., 1994; Kreuter and Strecher, 1996; Dijkstra, 2005). In a similar vein, and borrowing from the marketing literature, Churchill (2013) has recently advocated the need to distinguish between *process* and *outcome* personalization; readers are referred to Churchill (2013) for an extensive explanation of these concepts which concludes with a call to think more, and more imaginatively, about them. Notably, many of these discussions focus on the “way” rather than the end-goal of a persuasive request and argue that the method itself should be personalized (Kaptein et al., 2011).

Recently, health-care professionals and researchers, most noticeably in the domain of nutrition education, are examining computer-tailored interventions. Here, tailored interventions are often created to mimic to a certain extent person-to-person counseling (de Vries and Brug, 1999; Brug et al., 2003). Both target group segmentation—which also initially emerged within marketing (Tynan and Drayton, 1987; Plummer, 1974)—and personalization based on psychological characteristics such as people’s stage-of-change (Brug et al., 1997; Prochaska and Velicer, 1997) are starting to be used. Initial evaluations show an increased effectiveness of these types of computer-tailored interventions over more traditional, “one size fits all” health education efforts (Brug et al., 1998; Brugg, 1990a,b). Noar et al. (2007) conducted a meta-analysis of the effects of tailoring on the success of health interventions based on over 50 published comparisons and derived the same conclusion: tailored interventions are more successful than generic ones.

Currently, however, most persuasive technologies described in the research literature or implemented commercially are not personalizing their “ways”. This is striking since personalization of the end-goal is common place in commercial applications. Examples of the latter can be found in the rich literature on recommender systems (Kantor et al., 2011; Gretzel and Fesenmaier, 2006), or in the attempts to serving personalized ads through behavioral targeting (Stallworth, 2010). Notable exceptions do exist: Hauser et al. (2009) discuss how persuasive attempts in e-commerce can be made more successful by tailoring to customers cognitive style. Some of these approaches have likely made their way into commercial applications but the outcomes of these attempts are hardly shared with the research community. It remains that in most current persuasive technologies outside the marketing domain, the way in which the end goal is presented, alias the approach taken to influence users, is not adapted to the individual.

In this paper we detail how persuasive technologies can adapt the ways in which their users are persuaded—irrespective of the end goal—with the aim to increase the effectiveness of the technological interventions. In the current paper we focus specifically on the content of these interventions (see Davidson et al., 2003, for a taxonomy of intervention types). Possible taxonomies of *content* are provided in several fields, most noticeably by Michie et al. (2013) in behavioral medicine: our focus here is on persuasive user feedback (see also DiClemente et al., 2001). First, we discuss briefly some of the social psychology findings which motivate that designers of persuasive technologies should use the so-called *influence principles* to persuade their users. The effectiveness of these different means to influence the behavior of users has been shown convincingly by those studying persuasion and social influence. Next, we introduce explicit and

implicit methods of personalization, and we propose four practical design requirements for the design of personalized persuasive systems. Finally, we describe three instances of adaptive persuasive systems to illustrate the challenges facing designers of such systems.

2. Persuasion and persuasive technology

In looking for a scientific foundation for designing persuasive technologies, designers and researchers often turn to social sciences that study persuasion, most notably psychology (e.g., Bless et al., 1990; Crano and Prislin, 2006). Within this large field several theories of attitude and behavior change, such as the transtheoretical model of behavior change (e.g., Prochaska and Velicer, 1997; Long and Stevens, 2004), and the theories of reasoned action and its follow up, the theory of planned behavior (see, e.g., Madden et al., 1992; Fishbein and Ajzen, 2011), have gained large support and are used actively by designers (see for example Consolvo et al., 2009). Also, classical psychological work on operant conditioning (Skinner, 1976) has made a mark on the design of persuasive technologies, most notably in efforts of gamification (Deterding, 2012). Fogg on his website on the Fogg behavioral model¹ describes a large list of influential models and theories for the design of persuasive systems such as social cognitive theory (Bandura, 1991), the heuristic-systematic model (Chaiken, 1980; Chaiken et al., 1989), the elaboration likelihood model (Petty and Cacioppo, 1986), work on resistance and persuasion (Knowles and Linn, 2004), cognitive dissonance theory (Festinger, 1957), and a number of others (e.g., Maslow and Herzberg, 1954; Heider, 1944; Deci and Ryan, 2010). Finally, work in which (heuristic) decision making of individuals is studied, under the heading of behavioral economics (Kahneman and Tversky, 1979; Kukar-Kinney and Close, 2009), has been incorporated in attempts to design effective persuasive technologies.

Psychologists often describe different influence principles that can be used to change attitudes or behaviors. Similar descriptions of fixed principles (or strategies) to change attitudes or behaviors can be found in the marketing literature under the heading of sales influence tactics (McFarland et al., 2006). In our attempt to describe adaptive persuasive systems we focus on the literature regarding influence principles as pioneered by Cialdini and Trost (1998) and later followed up on by Cialdini (2001) and Guadagno and Cialdini (2005). These principles describe distinct psychological means that designers of persuasive technologies can use to increase the effectiveness of their persuasive applications.

2.1. Influence principles

The array of influence principles that can be used to change the attitudes and behaviors of users can be overwhelming. Both researchers and practitioners have made extensive use of the categorization of persuasive messages as implementing more general influence principles. Theorists have varied in how they individuate persuasive strategies: Cialdini (2001, 2004) develops six principles at length, Fogg (2002) describes 40 “strategies” under a more general definition of persuasion, Kellermann and Cole (1994) gather 64 groups from several existing taxonomies, and others have listed over 100 distinct tactics (Rhoads, 2007). These different counts result from differing levels of exhaustiveness, exclusivity, emphasis, and granularity (Kellermann and Cole, 1994). Influence principles are however a useful *level of analysis* that helps us to group and distinguish specific influence tactics or implementations of these principles (Kellermann and Cole, 1994; O’Keefe, 1994). In this paper we focus on the six influence principles as discussed extensively by Cialdini (2001). The effectiveness of each of

¹ See <http://www.behaviormodel.org>

these principles is supported both by social psychology and by marketing literature. We detail each in turn:

- **Authority:** People are inclined to follow recommendations and suggestions originating from authorities (Milgram, 1974; Blass, 1991). Authority is considered a form of social influence (Kelman and Hamilton, 1989; Martin and Hewstone, 2003) that is effective because some levels of responsibility and obedience to authority are essential for the existence of every social community (Modigliani and Rochat, 1995; Cialdini, 2001). However, not all psychological theories predict a positive effect of authority endorsements: Fiegen and Brehm (2004) use reactance theory to explain how authority endorsements can lead to negative effects when people's perception of freedom of choice is threatened.
- **Consensus (or Social Proof):** When individuals observe multiple others manifesting the same belief or behavior, they are more likely to believe and behave similarly (Ajzen and Fishbein, 1980; Cialdini, 2004; Goldstein et al., 2008; Zhu and Zhang, 2010). Multiple processes have been posited to explain the effectiveness of the consensus principle: Asch (1956) ascribes the observed effects to mere conformity, while others postulate that implementations of the consensus principle constitute informational influence, by serving as "social proof" (Hardin and Higgins, 1996; Cialdini, 2001). Recently, Salganik and Watts (2009) showed the effects of social proof in cultural markets: in their experiments "false" initial popularity lead to actual popularity in cultural markets.
- **Consistency and Commitment:** The Consistency and Commitment principle refers to people's strive to maintain consistent beliefs and act accordingly (Cialdini, 2001). This strive has been well researched under the heading of reducing cognitive dissonance (Festinger, 1957) and can be also used to explain both attitudes and behaviors. If a person is asked to write down that he or she will stop taking the elevator and take the stairs instead, she/he will be more inclined to do so even if they did not agree on writing it down in the first place (e.g., Deutsch and Gerard, 1955).
- **Scarcity:** (Assumed) scarcity increases the perceived value of products and opportunities (Cialdini, 2001). Therefore, advertisers and sales people often use phrases like "limited release", and "while supplies last" (Lynn, 1991). There is overwhelming evidence that identifying a product or a service as scarce affects people's attitudes favorably and increases the chance of purchase (West, 1975; Inman et al., 1997; Eisend, 2008; Lynn, 1989). Multiple psychological processes have been proposed to explain the effects of scarcity, the most prominent of which is based on commodity theory (Brock, 1968) according to which consumers desire scarce products more because the possession of such products results in feelings of personal distinctiveness or uniqueness.
- **Liking:** We say "yes" to people we like. When a request is made by someone we like, we are more inclined to act accordingly (Cialdini, 2001). Overwhelming evidence of this principle is presented by studies that exploit increased liking due to increased interpersonal similarity (Garner, 2005). As a striking example: people seem more inclined to return a wallet to the lost and found when the name listed in the wallet is similar to their own—and is thus liked—than when the name of the (ostensible) owner is dissimilar (Hornstein et al., 1968).
- **Reciprocity:** People are inclined—or actually, people go through a great deal of effort—to pay back a favor (Cialdini, 2004). This influence principle—when implemented properly—is exceptionally strong, and seems to work even when it is truly not beneficial for the persuadee. When a persuadee is in debt to the source, he or she will comply with persuasive requests to even out this discrepancy (Greenberg, 1980). The principle of reciprocation is the foundation in the *tit-for-tat* principle in social dilemma games (Komorita et al., 1991). It has been shown that people even reciprocate to favors they have never asked for (James and Bolstein, 1990).

2.2. Individual differences

While each of the influence principles described above has been shown to be effective and most have been used in the design of persuasive systems, the responses to these influence principles are not always clear-cut. For example, Johnson and Eagly (1989) discuss the difficulties many experimentalists have had to replicate classical findings within the social influence field. These, and other, conflicting results are suggestive of individual differences in responses to influence principles.

2.2.1. Trait differences in overall responses to Persuasion

Much of the work on individual differences in persuasion has directly drawn on dual-process models (for example the Elaboration Likelihood Model (ELM), see Petty and Cacioppo, 1986; Petty and Wegener, 1999) to work out how new or established traits could moderate persuasion. These dual processing models explain the effects of persuasive messages by postulating two routes of information processing: the *central* route through which arguments are elaborately processed, and the *peripheral* more heuristic route of information processing. It is hypothesized that many persuasion attempts are primarily effective through peripheral processing.

Many of the studies concerning dual processing models have examined trait differences in motivations, such as need for cognition (NfC), that are associated with structural differences in peripheral and central processing of persuasive messages (Cacioppo et al., 1986). NfC refers to the tendency of people to vary in the extent to which they engage in and enjoy effortful cognitive activities (Cacioppo and Petty, 1982). NfC predicts differences in the effects of argument strength on attitudes, the degree to which individuals rely on product characteristics versus source liking (e.g., Haugtvedt et al., 1992), attitude strength resulting from processing a persuasive message (e.g., Haugtvedt and Petty, 1992), and metacognition in persuasion (e.g., Tormala and DeSensi, 2009). More generally, for many choice settings in which personal relevance is neither very low nor very high, elaborative processing of stimuli varies with NfC, such that NfC measures an individual propensity to scrutinize and elaborate on arguments via the central route (Cacioppo et al., 1996). For example, people high in NfC are more likely to scrutinize whether someone endorsing a health related product in some advertisement is actually a doctor (or an actor playing a doctor) and how this might be informative about the product. High elaboration or personal involvement both lead to increased usage of the central route to persuasion and thus less persuasion through social influence principles.

While NfC is the most widely used trait that operationalizes stable motivational heterogeneity in dual-process models, several related traits have been identified and studied (Haugtvedt et al., 2008). Measures of the need for closure (Leone et al., 1999), an aversion to ambiguity and uncertainty as well as a preference towards firm, definitive answers to questions, the need to evaluate (Jarvis and Petty, 1996), the extent to which people spontaneously evaluate objects or experiences as either good or bad, and the need for affect (Maio and Esses, 2001), the tendency to approach or avoid emotion-inducing situations and activities, have all received attention in the persuasion literature.

Recently, scholars in the persuasive technology and HCI fields have examined individual differences in the personality of users for more general personalized persuasive applications. For example Halko and Kientz (2010) explored the relationships between the Big Five personality scale (see, e.g., Gosling and Rentfrow, 2003) and a preference for distinct persuasive strategies. The authors find a number of relationships between the personality of users and the preferred types of persuasive messages such as competitive, authoritative, or reinforcement messages. In similar vein, Nov and Arazy (2013) recently explored relationships between personality and interface design. The

authors (amongst other findings) show that user's conscientiousness levels relate to their reactions to the use of social proof messages.

2.2.2. Trait differences in responses to distinct influence principles

Investigators have also drawn on the categorization of messages as implementing distinct influence principles to identify and study personality constructs that are plausibly associated with the posited processes by which particular influence principles function. For example, the commitment principle, including a range of implementations, such as in the “foot-in-the-door” principle, functions through the application of motivations for consistency. A personality scale that measures these motivations—the preference for consistency scale—predicts responses to the commitment principle, such that for participants low on this trait these principles are ineffective (Cialdini et al., 1995; Guadagno et al., 2001). This prior research has helped explain the difficulties investigators have in replicating results regarding the effectiveness of the Consistency principle. However, successful use of trait differences in studying the effects of influence principles requires a theory about the psychological processes that make the principle effective and how these might vary between individuals in the population. Such theory is not always available; even in the case of preference for consistency, there has been considerable controversy about the mechanism(s) by which foot-in-the-door is effective (Burger, 1999).

2.2.3. Measurement of individual differences

Prior research on measuring individual differences in responses to persuasion can be described as relying on *meta-judgmental measures* of personality traits. In the context of attitude strength, Bassili (1996) distinguishes between meta-judgmental measures and operative measures of attitude strength. A similar distinction applies in the context of individual differences in persuasion. *Meta-judgmental measures* of personality trait ask individuals to report judgments about the consistent, structural properties of their broadly applicable attitudes, preferences, beliefs, and behaviors. In these measures, individual's psychological processes serve as objects of their consideration. On the other hand, for *operative measures* individuals' psychological processes are in use—they are operating. Operative measures are measures that are directly linked to the cognitive processes that are responsible for the response (Vanharrevel, 2004). An operative measure of the effect of an influence principle could for example be the actual behavioral response to a message that implements the influence principle.

One could imagine using both types of measures for assessing differences in responses to persuasion to be of use in personalizing persuasive technology. A personalized persuasive system could require users to fill out a number of questionnaires to obtain meta-judgmental measures. Furthermore, the system could collect observations of actual behavioral responses to influence attempts as operative measures. Using both of these measures the system could estimate which of the above influence principles a user is most susceptible to and tailor the influence principle deployed for a specific individual.

3. Personalizing persuasive systems: explicit and implicit methods

The work reviewed in the previous sections implies that persuasive systems could adapt their choice of influence principles for distinct individuals. However, for proper personalization we need methods by which we can estimate, for each individual user, which influence principles will be the most effective. As described briefly in Section 2.2.3, two distinct methods to measure individual differences exist. We coin the use of these two methods for personalization, analogues to a previous discussion by Garde-Perik (2009) in HCI, *explicit* or *implicit* means of profiling.

3.1. Explicit profiling

Meta-judgmental measures are often obtained using questionnaires in which users are asked to reflect upon their own traits. Such an explicit approach could be used to tailor persuasive applications: if we have a questionnaire that elicits the tendencies of individual users to comply to distinct influence principles we would be able to measure these tendencies a priori, and adapt the interaction with the user according to the obtained estimates. For the tendency to comply with different influence principles such a questionnaire was recently developed; Kaptein et al. (2012) introduced the susceptibility to persuasive strategies (STPS) scale.² However, measures such as the NfC scale or the preference for consistency scale are also likely candidates for use in explicit profiling. When using explicit profiling the user will be aware that such a measure is established, consents to fill-out the questionnaire, and will often be able to know that the questionnaire scores influence his or her interactions with the system.

3.2. Implicit profiling

Next to explicit profiling, we can also use *implicit* profiling to enable personalization: in persuasive technologies which use implicit personalization, operative measures are used to estimate the individual susceptibility of users to distinct influence principles. Here the actual responses to persuasive attempts are used to personalize future interactions. For example, if an application aimed at increasing exercise levels amongst its users links a specific user to her social network for comparison, and thus uses the principle of social proof, but fails to be effective (which is easily measured using actimetry), the application could lower the estimated success of this principle and use another principle in future interactions.

Using implicit personalization, the influence principles are adapted based on interactions with the user. In this case, users might not be made aware of the profiling and resulting adaptations. Thus, implicit profiling for personalization brings separate design as well as ethical challenges to designers of personalized persuasive systems. However, with the risk of additional ethical concerns also comes a benefit of an undisturbed user experience: for implicit profiling the user merely has to use the system for it to adapt to his or her personal needs. No additional questionnaires or other types of actively user-generated data are necessary. Thus, if implemented well, implicit profiling could potentially increase the usability of adaptive persuasive systems.

3.3. Persuasion profiles

*Persuasion profiles*³ are collections of estimates of the expected effects of different influence principles for a specific individual. Hence, an individual's persuasion profile indicates which influence principles are expected to be most effective. Persuasion profiles can be based on both meta-judgmental and operative measures of persuasive susceptibility. Relying primarily on behavioral data has recently become a realistic option for interactive technologies, since vast amounts of data about individuals' behavior in response to attempts at persuasion can easily be collected.

Fig. 1 shows an example of a persuasion profile. The profile consists of the estimates of the effects of different influence principles, and the certainty around these estimates. Thus, for this user, the implementations of the consensus principle are the most effective. Implementations of the authority principle are the least effective, however the

² The author refers to influence strategies as opposed to principles, but its use is similar in meaning.

³ The idea of creating a profile of users at the level of the persuasive techniques was for the first time publicly discussed by Fogg in a statement to the Federal Trade Commission.

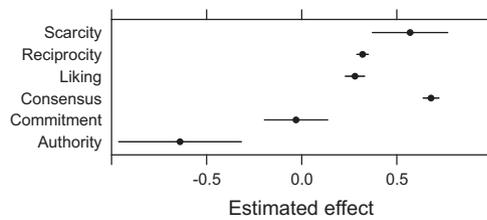


Fig. 1. Fictional example of a persuasion profile. Dots represent the estimated effect of the respective influence principles, while the bars represent the certainty around the estimates.

estimates of the effect of this principle are relatively uncertain. A persuasion profile of this kind that is consulted at runtime can ensure that the system can attend to individual differences and make an informed selection of social influence principles.⁴

3.4. Design requirements

We argue that all systems that use persuasion profiles to personalize persuasion should address four key requirements which we coin *Identification*, *Representation*, *Measurement*, and *Single inheritance*. We detail each in turn:

3.4.1. Identification: the ability to identify individual users

To be able to adapt to individual differences in responses to social influence principles, a system should be able to identify individuals (be it just by a unique key rather than linking to personal identity). More generally, *personalization* will only be possible if the user can be uniquely identified and the information needed for the personalization effort (such as the persuasion profile) can be retrieved. Thus, the technological ability to identify users both within a usage session and over multiple sessions of usage, and perhaps even over multiple devices, is key to developing personalized persuasive systems. Only once a user has been identified it is possible to personalize persuasive messages. The identification allows for the retrieval of the persuasion profile, which can subsequently be used for message selection (see Section 4 for implementation examples).

If the personalized persuasive messages are delivered through mobile devices this requirement can be addressed trivially by assuming that a device is personal to a single individual. In such a situation the phone number serves as a unique identifier. In cases of email communication similarly identification is trivial and one can assume a unique email address to belong to a unique user. However, in an ambient computing scenario the possibilities of identification are less straightforward: if our aims are to personalize communication delivered via multiple devices, to multiple users in, e.g., public spaces, then other identification methods will become a technological necessity. While more challenging, such opportunities have recently emerged: for example, designers can use the unique bluetooth key that is used by mobile devices (Consolvo et al., 2008) for identification of users in public spaces. Designers have also used face recognition or fingerprints (Cowie et al., 2001), gait analysis, or RFID badges (Schmidt et al., 2000) to identify individual users. Each of these technologies offers different advantages and disadvantages, that need to be considered by designers. In our case studies (Section 4) we use the (unique) phone number of users for identification, the email address, or a web browser cookie to identify users.

⁴ In Section 5.3 we discuss the trade-offs designer face when making selections of content based on estimated traits of individuals in more detail.

3.4.2. Representation: the ability to present the influence principles to users

Adaptive persuasive systems which personalize the influence principles used to motivate users should be able to implement various influence principles. While this may sound trivial at first, the technical ability to tailor content at the “granularity” of the use of influence principles is not something designers are used to design for. In the cases of personalization of influence principles the “end-goal” of the persuasive attempt and the means need to be separated intrinsically in the design of the systems such that for a specific goals multiple influence principles can be implemented and tailored.

For example, a digital exercise coach can influence users to exercise by having users set targets (e.g., commitment), coupling users to others (e.g., consensus), or by providing advice from a fitness instructor (e.g., authority). To enable usage of persuasion profiles, systems should have the flexibility to present their end goal (e.g., work out more) in these different ways to users. Thus, a system that manages the content of the messages that might be used to communicate to users should separate the end-goals specifically from the means.

Additionally, in the system architecture designers should distinguish the higher level social influence principles, and their respective (often textual or visual) implementations. Thus, if a persuasive system uses the authority principle then still different expert sources could be used, via different communication channels, to influence users. In each case, the authority principle is represented by a different implementation. To enable representation of *different* influence principles for different users, the system should be endowed with the ability to represent multiple principles, and multiple implementations for each.

3.4.3. Measurement: the ability to measure user traits

Using explicit measures, designers would rely on standardized (often existing) scales or measurement devices to estimate a persuasion profile. In these cases these measurements need to be available a priori, and they need to be uniquely identified for individual users. Also, likely, system designers have to create fall-backs in cases in which the data of individual users is not available for the system—for example when a user failed to fill out a questionnaire. We detail the usage of explicit measures for the design of personalized persuasive systems in our first case study below.

When designers use implicit measures to create systems that do not use a persuasion profile based on a priori measurements of the effectiveness, but rather adapt to user's responses *dynamically*, it is necessary to measure the outcome of influence attempts: measurement of the behavioral outcomes is key to implicit profiling. In this case the behavioral outcomes will be used directly to inform the users' profile (examples are provided in our second and third case study below).

While measuring the behavioral outcome of an influence attempt sounds straightforward it is not always easy to measure whether an appeal was successful, or even to determine what a measure of success of a message would entail. For example, in a digital exercise coach a prompt by a fitness instructor to run for 30 min that is followed by the user running for 20 min, 14 h after the prompt, might constitute a partial success—indicating the success of the authority principle—but might also be due to external causes. Thus, the coupling of observed behavior to the intervention content might not always be one-to-one. Furthermore, technologically not all behavioral responses are measured easily or reliably. When direct behavioral measurements are used, the behavior needs to be compatible with the influence principle (Fishbein and Ajzen, 2011).

3.4.4. Single inheritance: the ability to link behavioral observations uniquely to influence principles

When using messages or other stimuli that implement different influence principles for persuasion profiling it is important that

individual implementations (e.g., the actual messages) implement a single principle. Thus designers should design their interaction such that while an influence principle might be implemented in many ways, each implementation is distinct and only represents a single influence principle. Only if a one-to-one mapping of influence principle to the presented message is present one can reliably measure the observed effect of the influence principle. If an implementation confounds multiple influence principles then one cannot attribute the success or the failure of an influence attempt to a distinct influence principle and hence one cannot validly update the persuasion profile.

Persuasive messages implementing multiple principles would confound the effects of the principle. This might compromise the effectiveness of the most appropriate principle and make it difficult to attribute effective implementations to any single principle. It is not said that designers should never try to use multiple influence principles jointly (but, see Kaptein and Duplinsky, 2013), however implementations should not combine multiple principles and thus personalization should be enabled at the level of the *principle* not at the level of the implementation. If multiple principles are used in a specific message, then this message should be labelled as two (or more) influence attempts, each identifying the right principle.

The single inheritance requirement sounds more abstract than the first three requirements. However, it is methodologically key for the design of personalized persuasive systems.

4. Case studies: personalizing persuasive technologies

In this section we present three case studies in which persuasive technologies are personalized to their users using persuasion profiles. The first system we describe makes use of explicit profiling and thus relies on meta-judgmental measures. The next two design cases rely on implicit methods.⁵

4.1. Tailoring short text messages to prevent snacking

This first case study presents the use of explicit measures to personalize a pervasive technology (for an elaborate discussion of this design case see Kaptein et al., 2012). To evaluate the use of explicit measures in personalizing persuasive technologies we developed a system that used persuasive short text messages as a prompt to reduce snacking behavior. For the purposes of this system, we defined snacking behavior to be the act of eating (unhealthy) snacks in between meals. Earlier attempts to use text messaging to this end have achieved mixed results (McGraw, 2010; Gerber et al., 2009). In this study we tried to evaluate whether messages that are personalized are more effective than messages that are not personalized. We first used the STPS (see Kaptein et al., 2012) to measure the susceptibility of users to influence principles, and then conducted a field experiment to see if personalizing text-messages increased their effectiveness.

In this design case *identification* was accomplished using the phone number of the individual. Using mobile phones makes the identification requirement relatively straightforward as long as one is willing to assume that a mobile phone is uniquely used by a single user, which is likely not far from the truth. *Representation* was done by means of presenting persuasive text messages that were each designed to implement a single principle (see Table 1 for the exact implementations). Text messages restrict the representation of the influence principles to textual implementations. In this study we choose to implement, for each principle, multiple

messages to ensure that principles could be stimulus sampled. By ensuring that each message implemented only one principle the requirement of *single-inheritance* was also met. Finally, *Measurement* of the effects of the messages was based on self-report of users in this design case. In this design case these latter measurements were not used to inform the profiles but rather to evaluate the effect the system.

To evaluate the performance of personalized persuasion we compared three versions of the text messaging system:

1. A *personalized version* (PV): In this version of the system messages the persuasive messages contained those influence principles that users indicated to be most susceptible to.
2. A *contra-personalized version* (CPV): In this version of the system the messages contained implementations of principles that the users indicated to be least susceptible to.
3. A *random message version* (RMV): In this version of the system users receive a random selection out of all the created text messages.

4.1.1. System design and evaluation methodology

A 2-week long evaluation of the short text messaging application to decrease snacking was set up. Since snacking behavior varies substantially between people, we chose to include a 1 week baseline assessment of individual snacking behavior before introducing the three versions of our application.

Participants in the evaluation were recruited via a professional recruitment agency. A call for participation was sent out via email to potential Dutch participants between 18 and 65 years of age, with fluent understanding of English, and in possession of a mobile phone. The call for participation detailed that the evaluation of the short text messaging application would take two full weeks and would entail filling out several questionnaires and receiving daily text messages on their mobile phone. In total 162 participants completed the sign up process in full and started their participation. Participants received text messages for a period of 2 weeks (2 × 5 days, workdays only). Participants were instructed every evening to go to a designated website to fill out a short diary. The first week was used to establish a baseline snacking frequency for each user, while the intervention was employed in the second week. We included for our final analysis only those participants that filled in at least one diary during each of the 2 weeks (e.g., during the baseline measurement and during the intervention). Our final sample was composed of 73 users. The average age of the participants was 34.9 years (SD = 11.1) and 32 (43.8%) were females.

After browsing to the designated website for the first time all users filled out a small questionnaire regarding their snacking behavior. Next, participants filled out the STPS and provided their mobile phone number. Participants then received one text message a day (on workdays) for a period of 2 weeks, and subsequently filled out a small online diary every day. The diary consisted of the following question:

- How many snacks did you have today? (Open ended)⁶

For the first week participants received one text message a day which asked them to fill in their diary. In week two participants received the persuasive messages according to their version of the application. This setup enabled us to study the change(s) in snacking behavior over the course of 2 weeks between the three different versions of the

⁵ The first two design cases presented here are presented in more detail in Kaptein et al. (2012) and Kaptein and van Halteren (2013). We discuss these systems here with a focus on the design challenges and we motivate how the four design requirements are met in each case.

⁶ While it is known that self-report measures as these may suffer from social desirability bias, this is not necessarily the case and they are thought of as valid instruments for modeling behavior change (e.g., Armitage and Conner, 1999).

Table 1

The messages used in the text messaging application: For each of the four influence principles used in this trial three implementations were used. Note the mapping of many implementations to a single principle in line with requirement four.

Principle	Message
Authority	Try not to snack today. According to the College of Physicians this is an easy way to lead a healthier life.
Authority	Dietitians advise to have 3 meals a day without snacking. Try to reduce snacking.
Authority	The World Health Organization advises not to snack. Snacking is not good for you.
Consensus	90% of people benefit from reducing snacking between meals. It will boost your energy and you will live a healthier life.
Consensus	Everybody agrees: not snacking between meals helps you to stay healthy.
Consensus	Reduce snacking. You are not on your own: 95% of participants have already reduced snacking.
Commitment	The aim of this study is to live healthier. Reducing snacking is a way to achieve that.
Commitment	Try to obtain your goal for living a healthier life by not snacking. You are committed!
Commitment	You have to continue what you have started: you are participating in this test to lead a healthier life. Reducing snacking.
Scarcity	There is only one chance a day to reduce snacking. Take that chance today!
Scarcity	This test lasts only 3 weeks: you have the unique opportunity to enhance your health by reducing snacking.
Scarcity	Today is a unique opportunity to lead a healthy life. Reduce snacking.

application. Our implementations of the influence principles that were used in the text messaging application are presented in [Table 1](#).

4.1.2. Results of the evaluation

The primary test to see whether personalized messages can be effective in reducing snacking behavior is provided by a comparison of the progression of the snacking behavior during the experiment between the three experimental conditions. To statistically test the effects of our different experimental conditions we fit a multilevel model to the data. This allows us to estimate the effects of time on the number of snacks consumed, in each of the versions of the application. We started by fitting a “null” model (denoted “P” in the Tables, [Snijders and Bosker, 1999](#)) which can formally be written as $y_{ij} \sim \mathcal{N}(\mu_j, \sigma_{err}^2)$ where $\mu_j \sim \mathcal{N}(\bar{\mu}, \sigma_{\mu}^2)$ for $j = 1, \dots, N = 73$ users. From the “null” model, we build a model that includes both time and the experimental condition to explain the snacking behavior of our users (Models A–C in [Table 2](#)). Finally, to test whether the different versions of the text messaging application significantly influence the snacking behavior of our users we fit a model in which time during phase two interacts with version—essentially fitting separate time effects for the different versions during phase two.

[Table 3](#) shows the fixed effects of Model *D* to support interpretation of the effect sizes found in this study. The fixed effects show that during the baseline phase the number of snacks consumed by the users of the system does not decrease significantly. Hence, during the first week the number of consumed snacks hardly changes (if anything, it is estimated to slightly increase). During the treatment phase however the number of snacks consumed by our participants decreased significantly for *RMV* and *PV* participants. The fixed effects table indicates that the decrease in snack consumption is higher for those using the *PV* version than those in the *RMV* version. In the *PV* condition the number of consumed snacks is estimated to decrease by 0.3 every day: this implies that after a week in this condition about 2 snacks less were consumed each day in the personalized version than during the baseline period. Besides being “statistically different from zero”, we believe that this unstandardized effect is sizable enough to conclude that the personalization of persuasive messages directly impacted the eating patterns of the users in a positive and meaningful way.

4.1.3. Discussion

Our implementation and evaluation of the personalized persuasive system to reduce snacking showed that personalizing influence principles, using explicit measures, can influence the effect of these messages favorably. While the evaluation period of 2 weeks is limited, the results provide a first proof of the effectiveness of personalized persuasion. Whether these effects hold in the long run remains to be

Table 2

Comparing the null model with models including a time effect and different time effects for each condition.

Model	df	AIC	logLik	χ^2	$\Pr(> \chi^2)$
A: Null model	3	1857.70	−925.85		
B: + Time	4	1849.64	−920.82	10.07	< 0.01
C: + Time and phase	5	1843.39	−916.70	8.24	< 0.01
D: + Time P2 × Condition	7	1837.79	−911.89	9.60	< 0.01

Table 3

Overview of the fixed effects of the Model *D* including an interaction between time and condition to predict snacking behavior. Empirical *p*-values are based on MCMC simulations.

Parameter	Estimate	Std. error	<i>t</i> -value	<i>p</i>
Intercept	2.16	0.19	11.33	< 0.001
Time Phase 1	0.05	0.04	1.07	0.28
P2: Time:CPV (contra tailored)	−0.01	0.08	−0.16	0.85
P2: Time:RMV (random)	−0.22	0.07	−3.30	< 0.001
P2: Time:PV (tailored)	−0.30	0.07	−4.29	< 0.001

seen, but at least we can conclude that an initial influence attempt can be made more successful using personalized persuasion.

4.2. The PMS system

The second system that implements personalized persuasion was created to increase user engagement in a health and lifestyle service. In this design case we focussed on implicit profiling. The health service combines a 3d accelerometer optimized to detect physical activity patterns with active human and technology initiated coaching to help users gain a more active lifestyle. Within the health service user engagement is key: coaching is done via a web service, and the activity data is only analyzed after it has been uploaded to a web service. Users often fail to upload their data. To encourage *docking*—the uploading of the activity data to the web service—docking reminders are sent via email. In this design case we examined whether we could use implicit personalization to improve the effectiveness of these reminder emails. The system is called the Persuasive Messaging System (PMS) (for a more elaborate discussion of the PMS see [Kaptein and van Halteren, 2013](#)).

4.2.1. System design and evaluation methodology

The PMS was implemented on a server that was external to the health service's own system. The PMS used the unique key provided by each accelerometer to identify individual users. When a user docks,

Table 4
Influence principles and their implementations in the PMS system.

Principle	Implementation
<i>Neutral</i>	...[no insert] ...
<i>Scarcity</i>	1. We would like to remind you to connect it to your PC soon and stay in touch with [X]. Today is a great day to stay fit so make sure you do not miss out on your participation in [X]!
<i>Authority</i>	2. Any chance to connect your Activity Monitor is a chance to learn about your own activities. Take the opportunity to learn about your activities right now. 3. Experienced [X] coaches recommend frequent uploads of your activity data. This will help you to gain more insight and be more active!
<i>Consensus</i>	4. Activity experts recommend moderate to high activity on a daily basis and connecting to the [X] platform will help you to reach this target! 5. People like you who connect their Activity Monitor frequently with their PC are more likely to benefit from the program and obtain a healthy lifestyle! 6. Thousands of people are participating actively in the [X] program and they stay connected at least once a week. Join the group!

the activity data and a unique identifier are sent to the persuasive messaging service. The unique key of the accelerometer, combined with the individual email address of the users in this study, jointly satisfied the requirement of *identification*. The influence principles were represented in the body text of the email reminders that were sent to users after 3 or 6 days of inactivity. These reminders were initiated by the health service and fulfilled the requirement of *representation*. The service's server sends a request to the PMS server for the next influence principle to be used for the current user. The PMS server, upon receipt of the request, looked up the persuasion profile for that user and returned the text snippet of one of the persuasive email messages that were created for the PMS (see Table 4).

Dear (*firstname*),

How are you doing? We hope all is well.

It is 3 days since the last time you connected your ActivityMonitor

[Message]

We would like to remind you to connect it to your PC soon and stay

in touch with [X]⁷.

Sincerely,

The [X] Team

During a brainstorm session, five researchers working in the field of persuasive technology generated a large number of persuasive messages. Messages were created that implemented the *scarcity*, *authority*, and *consensus* principles. After the brainstorm a card-sorting test was used to classify messages according to their principles, and for each principle two messages were selected for use in the trial. The persuasive messages consisted of text snippets containing social influence principles that were added to the standard email reminder. This standard reminder mail is presented in the text box above. The influence principle was inserted at the [message] location of the email reminder email. Table 4 gives the implementations of the influence principles that were used. The process of designing the messages using experts in the field enabled us, in this case study, to create multiple implementations of influence principles while still adhering to the *single-inheritance* requirement (see below for the implementation of the *measurement* requirement in this case study)

To enable a priori estimation of the effect of the messages each of the messages was presented to $N=80$ participants in a pre-test. Participants were instructed to answer the question "This message would motivate me" on a seven-point ("Totally Disagree" to "Totally Agree") scale. Scores were averaged over the two implementations of each principle. The neutral message had the lowest evaluation: $\bar{X} = 3.46$, $SD = 1.44$. The messages implementing influence principles scored higher, with authority scoring highest, $\bar{X} = 4.21$, $SD = 1.59$, before consensus, $\bar{X} = 3.96$, $SD = 1.54$, and scarcity, $\bar{X} = 3.81$, $SD = 1.52$.

To evaluate the PMS system an evaluation was set up in which the system was deployed for all *new* users of the persuasive systems from the 1st of January 2011 until the 1st of July 2011.⁸ To measure the effectiveness of the reminder emails a dynamic image was inserted into the email message body, which allowed the PMS to log the fact that a user had opened an email. If, and only if, within 24 h after opening the email the same user would dock her activity monitor, then the message would be considered a success. To assess the effects of personalized persuasive messaging as opposed to the original reminder message, or messages using influence principles that were not personalized, we assigned users to one of four conditions randomly:

1. *Baseline*: Users assigned to this condition received the standard (no influence principle) docking reminder message.
2. *Best pre-tested*: Users assigned to this condition received one of the two messages implementing the authority advice.
3. *Random*: Users assigned to this condition received randomly one out of the seven messages (with probabilities equal for each of the principles).
4. *Adaptive*: Users assigned to this condition received messages suggested by the PMS personalization algorithm (see below).

In the adaptive condition we estimated the effect of the influence principles after each interaction. We assumed the principles to be independent, and we were thus looking for an estimate \hat{p}_{us} of the effectiveness of a distinct principle for an individual user. To do so we used a Bayesian approach and used a Beta(α, β) prior for each of the estimated probabilities. The beta distribution can be re-parametrized as follows:

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

where $\mu = \alpha/(\alpha + \beta)$ and $M = \alpha + \beta$, then the expected value of the distribution is given by $E(\theta|\mu, M) = \mu_m$. In our specific scenario, μ_s represents the expected probability of a successful influence attempt by a specific influence principle given the previous data. The distribution of \hat{p}_{us} can be updated using

$$p(\theta|k) \propto l(k|\theta)\pi(\theta|\mu, M) = \text{Beta}(k + M\mu, n - k + M(1 - \mu)),$$

in which k , 0, 1, is the outcome of the new observation. The PMS server ran a cron-job every 24 h to match all opened emails with the recent docking behavior and update the individual level estimates. Hence, the docking behavior after receiving the persuasive message was used to *measure* the effect of the influence principles. This dynamic updating allowed us to implicitly profile users.

To determine which message to present in the next interaction we used Thompson sampling (Scott, 2010): we obtained a single draw of each of the Beta distributions for each principle and selected the principle with the highest draw. Scott (2010) showed

⁷ The company name cannot be disclosed in this publication.

⁸ Note that the PMS trial ran on the live service of the product and was thus conducted in the field.

Table 5
Overview of the data from the persuasive messaging system evaluation.

Condition	Users	% Success [S.E.]
Baseline	271	28.49 [1.7]
Best pre-test	289	24.01 [1.5]
Random	289	25.41 [1.6]
Adaptive	280	26.49 [1.6]

Table 6
Model comparisons used for the analysis of the persuasive messaging system evaluation.

Model	BIC	logLik	χ^2	df	<i>p</i>
Model A	10,903.18	−5449.59			
Model B	10,486.91	−5240.46	418.26	1	< 0.001
Model C	10,480.00	−5228.00	24.92	9	< 0.01
Model D	10,474.54	−5222.27	11.46	3	< 0.01

that this provides an asymptotically optimal solution⁹ to the exploration–exploitation trade-off inherent in dynamic learning.¹⁰

4.2.2. Evaluation results

All users who (a), participated in the service for at least 30 days, and (b), received at least 3 email reminders during the evaluation, were included in the final dataset. For the period of the evaluation this led to a dataset describing the upload frequency and responses to reminders of 1129 users. Table 5 gives an overview of the number of users, and the success percentages in each condition. It is clear that users are relatively equally distributed over the conditions.

To analyze the data obtained in the PMS evaluation we fit a series of multilevel models. We start again by fitting a null-model, this time with a logit link. Adding average effects for the influence principles to the null-model shows no significant effect of the influence principle on success of the emails, $\chi^2 = 4.75$, $df=3$, $p=0.19$. We find a large main effect of Frequency—the number of the reminder that is sent (see Table 6, Model A and Model B). Addition of average effects of the condition to Model B shows that there is no significant average effect of the conditions, $\chi^2 = 3.19$, $df=3$, $p=0.36$. There is however a significant increase in model fit when adding varying influence principle effects by user (Table 6 Model B and Model C). Finally, condition interacting with frequency is added to the model. This interaction significantly improves model fit (see Table 6 Model D) and indicates that the effect of frequency differs between the four conditions.

Table 7 shows the fixed effects of Model D. The negative coefficient for Frequency indicates that the probability of success of a reminder message decreases over time: the first reminder is successful around 27.7% of the time (for users in the baseline condition, which is used as a reference) while the fifth reminder is successful only 17.9% of the time. The interactions of the version of the application and frequency can be interpreted in the same way: for the random version—compared to the baseline version—the drop in effectiveness of the reminders over time is lower, while that of the best-pretested version is higher (although both are not significantly different from 0). The drop in effectiveness of messages in the adaptive version is significantly lower than the baseline version: the predicted effectiveness of the

⁹ See Section 5.3 for a more detailed discussion of the exploration–exploitation trade-off.

¹⁰ The prior for the neutral (no social influence) message was set to $\bar{X}=0.39$, $Var=0.1$. In line with the pre-test of the messages the authority principle prior was set the highest, $\bar{X}=0.52$, $Var=0.1$, before consensus, $\bar{X}=0.50$, $Var=0.1$ and scarcity, $\bar{X}=0.47$, $Var=0.1$. Given the relatively large dataset in this study the prior is quickly “swamped” by the data and hence has very little influence on the study.

Table 7
Coefficients of the fixed effects of Model D. The table shows that over time (Frequency) the success of the email reminders decreases. This decrease over time is significantly smaller in the adaptive version than in the baseline version.

Parameter	Estimate	SD	<i>z</i>	<i>p</i>
(Intercept)	−0.83	0.06	−12.72	< 0.001
Frequency (Freq)	−0.14	0.01	−11.44	< 0.001
Freq. × Random	0.02	0.02	1.10	0.27
Freq. × Pre-tested	−0.01	0.02	−0.40	0.69
Freq. × Adaptive	0.04	0.02	2.74	< 0.01

fifth message in the adaptive condition is 21.5%, which is 3.6% higher than the estimated effectiveness of the fifth message in the baseline version. For the tenth message this difference is even larger: 4.8%. The adaptive version also significantly outperforms the pre-tested version, $t=3.74$, $p < 0.01$ showing that personalized messages remain effective longer than non-personalized messages.

4.2.3. Discussion

The implementation and evaluation of the PMS system highlights a number of challenges for designers who intend to use implicit measures to personalize persuasion. The systems provides solutions to each of the four design requirements: identification is based on the user's unique email address, and measurement is implemented using interactive tracking of the behavior of users. Email allows for dynamically altering the messages at the individual level, thus allowing us to represent the distinct strategies. As noted in Table 4 each message implements a single principle thus adhering to the single inheritance requirement. The evaluation shows that personalized messaging is beneficial, especially compared to the best pretested message and the status-quo.

4.3. Persuasion profiling in e-commerce

In the third evaluation we again study implicit persuasion profiling in email communication. We found an online travel agency willing to participate in an evaluation of personalized persuasion for their weekly email campaign. Because this design case was set up within a commercial email campaign our choices of experimental setup were limited: we only included a baseline and an adaptive condition. The baseline condition consisted of showing a status-quo message to users as in the original system.

In the period of 22 of July to the 16 of September 2012 we ran an email field evaluation of personalized persuasion with a total of $N=133,538$ customers. In total we collected 454,452 observations during the trial. On average each customer read 3.4 emails. Each of the participating customers was randomly allocated to either the baseline ($N=24,984$) or the adaptive ($N=108,554$) condition using draws from a Uniform(0,1) in which those with a draw < 0.2 were allocated to the baseline condition. The analysis concerns responses to 7 email batches in a period of 8 weeks, and of primary concern were the click-through rates that were obtained in the emails.

4.3.1. System design and evaluation methodology

We implemented three influence principles, authority, scarcity and social proof, in the email text. The email consisted of a recommended product, a description of the product, and (if applicable), an implementation of an influence principle in text. The emails changed weekly, and the only manipulation was the textual addition of the implementation of the influence principle.¹¹ Thus, in this design case representation was done in the emails, and measurement was achieved

¹¹ In an effort to not disclose the agency we do not publish the actual implementations.

by monitoring the clicks on the emails. *Identification* was achieved using the email address, and we ensure that the implementations had *single inheritance*. We logged the email open rate, and the click through rates. Click through rates were used to estimate p_{US} and select messages, similar to the previous study.¹²

The implicit personalization was set up as follows: before each email agency would request, by sending a list of hashed email addresses, the principles to send to distinct customers. From our experiment servers, based on the previous behavior of that customer, that of other customers, and the outcome of the random process of Thompson sampling, we recommended a principle for each customer which would be neutral (for those in the baseline condition) or a choice of neutral, scarcity, authority, or social proof (in the adaptive condition). Subsequently, the agency would send the email which contained a dynamic image to log whether the email was opened. This information was sent back to the experiment servers: if the user clicked one of the links provided in the email this indicated a success of the email (and hence of the recommended principle).

4.3.2. Results of the evaluation

Similar to the previous evaluations of personalized persuasion we fit a series of multi-level models to examine the effects of the emails. We start by fitting a null-model (Model A) and subsequently add a main effect of condition (Model B), of batch (Model C), and an interaction between batch and condition (Model D). Table 8 presents the model comparisons between these models, while Table 9 presents the fixed effects of model D. These can be interpreted as follows: the emails in both conditions become more effective over time, $\beta_{batch} = 0.32$. The average click through in the adaptive condition is higher than that in the baseline condition $\beta_{cond.} = 0.05$. Finally, there is a significant interaction between the timing of the email and the condition, $\beta_{batch \times cond.} = 0.05$, which indicates that in the adaptive condition the emails are more and more effective over time. This last finding shows that implicit personalization over time increases the click through on the emails. Note that the intercept of the model in Table 8 is omitted to prevent disclosure of absolute click-through rates. However, the estimated click through rate of the 5th batch of emails is about 15% higher in the adaptive version than in the baseline version. This difference is not only statistically significant but also (commercially) meaningful in this context.

4.3.3. Discussion

The third design case, which presented the use of personalized persuasion in commercial email campaigns shows that the baseline (status-quo) message is outperformed by the adaptive message. This difference itself can have a number of causes. However, the finding that the effectiveness of the personalized messages increase as more knowledge about users becomes available indicates the benefits of personalized persuasion.

5. General discussion and conclusions

In this paper we motivated the design of personalized persuasive systems: systems that tailor their use of psychological influence principles to the effectiveness of these principles for individual users. Next, we presented two distinct methods by which application can personalize their use of influence principles to their users. We introduced a distinction between explicit personalization of persuasion—based on meta-judgemental measures—and implicit personalization based on behavioral responses. We also introduced four design requirements for

Table 8
Model comparisons used for the analysis of evaluation 3.

Model	BIC	logLik	χ^2	df	p
Model A	332,924	−166,460			
Model B	332,924	−166,459	1.87	1	0.17
Model C	332,654	−166,323	272.39	1	< 0.01
Model D	332,604	−166,297	51.63	1	< 0.01

Table 9

Coefficients of the fixed effects of Model D (evaluation 3). The table shows that over time (*Batch*) the success of the emails increases. This increase over time is significantly larger in the adaptive version than in the baseline version.

Parameter	Estimate	SD	z	p
(Intercept)				
Batch	0.32	0.06	5.51	< 0.001
Condition	0.05	0.003	17.34	< 0.001
Condition × Batch	0.05	0.007	6.88	< 0.001

the design of personalized persuasive systems: identification, representation, measurement, and single inheritance. Three cases studies showed how these requirements can be implemented in persuasive applications and highlight that both methods of personalization are potentially of use. Each of the three designs used *persuasion profiles*—collections of estimates of the effect of distinct social influence principles for individual users—to improve its effectiveness. In the first design the profile was built using the STPS scale, while in the next two applications the profile was built dynamically by tracking the effect of distinct influence principles over multiple interactions.

The persuasion profiles presented here were based on the influence principles as listed by Cialdini (2001). This list provides a starting point for the classification of different influence attempts, as well as for the creation of different implementations of influence principles. Cialdini (2001)'s list however is not the only one that is of possible use to design personalized persuasive systems: likely, it will be beneficial for the design of persuasive technologies to start with a sufficiently informed list, such as the one adopted here, while enabling profiles to contain more (or less) principles based on distinctions in implementations that prove effective during the deployment of the system. However, it is important to note that the collection of influence principles used in persuasive technologies should be selected based on theoretical and empirical foundations instead of ad hoc.

5.1. Limitations of the current focus and evaluations

By presenting three case studies we have tried to illustrate how our four design requirements can be implemented in the actual designs of adaptive persuasive systems. The case studies are presented to illustrate the design process, and briefly demonstrate the possible effectiveness of persuasion profiling using both implicit and explicit measures. It has to be noted that the durations of the evaluations of the cases studies are limited, and hence cannot be interpreted as a proof of the long-term effectiveness of the personalization of influence principles. Likely, the effects of the use of distinct influence principles will change over time, and, at this moment, these possible time dynamics are not well understood. Persuasive technologies however provide a methodological tool to study such dynamics, and thus we encourage further research using long-term deployments of persuasion profiling, possibly with the addition of explicit time dynamics of the “ordering” of persuasive messages: e.g., it might be the case that a scarcity principle should follow a social proof principle if the effect of the social proof principle is diminishing over time.

¹² In this implementation we used a slightly more elaborate learning algorithm than in the previous design case. The algorithm itself can be found in Kaptein (2011) but is not the main aim of this paper.

The current case studies are limited not only in their duration, but also in their focus. In all three case studies the requirements were relatively easily met: identification for example is fairly trivial on mobile phones, in email, or on the web. Also the measurement of behaviors and the technological ability to change content are straightforward in these applications. The applications were actively sought out to demonstrate the implementation of the design requirements and provide a strong methodological framework for the evaluations. However, the application of the design requirements in ambient persuasive technologies—cross device, in context, etc.—needs further scrutiny.

Finally, the general implications of our design requirements need to be further developed. We have shown the applicability of the requirements for the design of systems that use persuasion profiling. However, this is only one specific instance of personalized persuasion: the applicability of the requirements for other types of personalization of persuasive systems needs to be explored further. For now we contend that *identification* will remain key in any personalized system, and that also the requirement of *measurement* (e.g., being able to determine personal “traits”) will generalize to other personalized persuasive systems. The *representation* requirement likely also extends beyond the use of persuasion profiling: if the personalized content cannot be displayed to individual users dynamically then personalization attempts seem futile. Finally, we believe that also the *single inheritance* principle will remain appropriate in more general designs of personalized persuasive systems that use implicit measures: if implicit measures are used to determine traits, then methodologically it should be possible to, without confounds, estimate the effect of specific persuasive content. If the content confounds multiple principles, strategies, or otherwise persuasive attempts, then one cannot attribute the observed attitudinal or behavior change one-to-one to the selected content. Hence, one cannot estimate the underlying trait reliably. Methodologically the observed change in attitude or behavior should be caused by a single, and known, property of the displayed content: this is exactly what the single-inheritance principle states.

5.2. Combining multiple methods of personalization in persuasive technologies

Persuasion profiles concern the means—the ways in which—people are influenced to comply to a request. Most notable personalization efforts up to now have however focused primarily on the end goals: recommender systems, like those used by Amazon.com and heavily researched by computer science researchers, select the appropriate product to endorse to individuals (end) without systematically varying or adapting the way in which a product is presented (means) (cf. Gretzel and Fesenmaier, 2006; Ochi et al., 2010; Zanker et al., 2009). It is likely that both methods of personalization will be combined in the future. Recommender systems determine which target behavior or product to offer, while persuasion profiles play a role in presenting that goal to people.

Persuasion profiles could benefit from combinations with other profiles. Target group profiling, as common in marketing practice, has the distinct benefit of being able to generalize knowledge gained over one set of individuals to other, unknown, sets. For example, if no woman has ever bought product A, one could decide to refrain from offering product A to new female clients even if no other knowledge about their previous decision making is available. In a similar fashion susceptibilities to influence principles are likely correlated with gender, age, occupation, context, etc. It is thus most likely beneficial for the effectiveness of influence attempts to combine target group profiles with persuasion profiles to obtain more accurate estimates, especially of new users of a persuasive system.

Finally, persuasion profiles can also be compared to other efforts of tailoring persuasion. Computer-tailored health education (see, e.g., de Vries and Brug, 1999; Brug et al., 2003) is an example of another approach to personalizing persuasion. In this approach

often both ends—e.g., what is a realistic health goal for the current individual—and means—e.g., in what way should the information be presented—are tailored. This tailoring is largely done in the following fashion: psychological theory is explored to determine the theoretical constructs that might be useful for tailoring (such as people's stage of change, or people's NfC). Next, experts create rules for selecting different content based on different values of the theoretical constructs of interest (Dijkstra and De Vries, 1999; Kreuter, 2000). Persuasion profiles, while more limited in scope than full breath computer-tailored interventions, allow for a selection of influence principles based on the measurement of user susceptibility to persuasion. Both approaches could be combined: expert determined rules could influence the probability of content selection a priori, while both explicit and implicit means of profiling can be used to build a persuasion profile and select implementations of influence principles.

5.3. Exploration versus exploitation in personalization

In the last two case studies we presented persuasion profiling based on implicit measures of user traits. Here, we briefly mentioned the use of Thompson sampling (Scott, 2010) to select messages given the estimated success of each message. Thompson sampling provides a strategy (or *policy*) to address the so-called exploration–exploitation trade-off that one faces in this context: there is a trade-off between selecting the message with the highest estimated click-through (exploitation) or selecting uncertain, but potentially more effective, messages to learn their effectiveness (exploration). The exploration–exploitation problem is inherent in many personalization efforts: the estimated traits (either based on implicit or explicit) contain uncertainty, and hence there is no deterministic choice for the best content.

The exploration–exploitation problem is broadly studied under the heading of the multi-armed bandit problem (see, e.g., Whittle, 1980; Berry and Fristedt, 1985; Sutton and Barto, 1998). Slot machines are colloquially known as one-armed bandits. When faced with the choice of playing on multiple slot machines, which arm should one pull? Based on sampling variability, one may have relatively more information about some slot machines than about others, hence one must trade-off pulling the arm of the slot machine that appears to have the best probability of a pay-off versus selecting alternative slot machines that appear to have inferior payoffs, but with little certainty.

Formally, the personalization problem can be regarded a *contextual bandit problem* (Yue et al., 2012; Ortega and Braun, 2013). The context X is given by the individual identifier and presents itself to the system. At each point in time, $t = 1, \dots, t = T$ the system has to select an action $A(t)$ (the persuasive content in this case) and observes a reward R_t (for example a click on the message). It then becomes of interest to study policies which prescribe actions that maximize the reward over all interactions (e.g., $\max \sum_{t=1}^T R_t | A_t, X_t$). There is a broad literature on the topic (see, e.g., Berry and Fristedt, 1985; Audibert et al., 2009; Scott, 2010), and within the marketing literature researchers are already approaching the personalization problem as a contextual bandit problem (Hauser et al., 2009, 2014; Schwartz et al., 2013). Appreciation of the inherent uncertainty in the coupling between user, message content, and observed behavior by exploring different *policies* is, in our view, a key next step for the development of personalized persuasive systems.

5.4. Current and future work

The current article introduced persuasion profiling using either explicit or implicit measurements of user susceptibility to influence principles. It is worthwhile to note the current state of this research, and contemporary attempts in similar directions. Persuasion profiles have by now been used not only to reduce snacking and to improve the effectiveness of emails as described above, but have also been used

to (a) motivate people to lead a healthier lifestyle (Sakai et al., 2011), and (b) increase the impact of online commerce (Kaptein, 2011). This latter work is preceded by several personalization attempts in this domain, most noticeably by Ansari and Mela (2005) and by Hauser et al. (2009). In the marketing literature personalization using implicit measures is becoming mainstream, and novel attempts aim to personalize promotional appeals (e.g., the persuasive principle) rather than the actual product that is recommended are emergent.

The application of personalized persuasion within typical HCI problem domains still needs to be strengthened. The above-described attempts to change users health behavior provide only anecdotal evidence of the use of persuasion profiles within HCI. There is a need to examine the use of persuasion profiles in other popular HCI problem domains such as the reduction of energy consumption (Bang et al., 2006), medication adherence (as suggested in Oinas-kukkonen and Harjumaa, 2009), and for the improvement of sleep quality (Scherini and Melo, 2010). In these fields our four design requirements are often less easily met as compared to online applications. However, the societal benefit of such applications is possibly larger and thus these domains warrant further study.

5.5. Conclusions

In this paper we described the design of personalized persuasive systems using persuasion profiles. Persuasion profiles can, as demonstrated, be built both using explicit and implicit measures of individual persuasion susceptibility. We believe that persuasion profiles, and their use in adaptive persuasive systems, will be common place in years to come. Already the technologies presented here are starting to be used in e-commerce and marketing because of their positive impact on revenue. We hope to have presented the minimal requirements to build personalized persuasive systems. Furthermore, we have demonstrated how personalized persuasive systems can be designed and evaluated in practice. We have shown how explicit measures of persuasion susceptibility can be used to personalize text messages to reduce snacking behavior, and we have detailed how implicit measures can be used to inform a persuasion profile which increases the effectiveness of email reminders to be more physically active. These design cases highlight the design challenges present when designing personalized persuasive technologies. However, there are obvious questions regarding ethics and privacy that need to be addressed if personalized persuasion is to be picked up widely by designers of persuasive systems.

In this paper we have presented evaluations of personalized persuasive technologies in field trials. Some of the presented field trials were run on the web, and in industry such trials are more prominent. We hope to encourage practitioners to disclose, to the scientific community, findings on personalization carried out in a commercial context. We also believe that it is time, within the persuasive technology community, to strengthen the discussion on the ethical side of personalized persuasive technologies (following up on the work of, e.g., Berdichevsky and Neuenschwander, 1999; Kaptein et al., 2011) and discuss explicitly the relationship of academia and industry as well as the feasibility of conducting user studies in industry.

As illustrated with the case studies, personalized persuasive technologies can be used to create effective persuasive applications. However, we need to further study the use of personalized persuasion in different domains: personalized persuasive systems can possibly be of use for reducing energy consumption, encouraging healthy lifestyles, and other behavioral change applications. These systems have been emergent in the last decade, and we feel that designers of such systems should go through great lengths to make their applications as successful as possible for individual users. Persuasion profiles are but one step in that direction. We believe that for these types of applications personalization is the ethical thing to do: if designers are aware of possible negative outcomes of their applications through the use of

“wrong” motivators—such as erroneously selected influence principles—designers of such technology would be “at fault” (e.g., Berdichevsky and Neuenschwander, 1999) when they fail to make the persuasion approach adaptive to the individual (see Kaptein and Eckles, 2010, for an initial discussion on the ethics of personalized persuasion).

References

- Aarts, E.H.L., Markopoulos, P., Ruyter, B.E.R., 2007. The persuasiveness of ambient intelligence. In: Petkovic, M., Jonker, W. (Eds.), *Security, Privacy and Trust in Modern Data Management*. Springer, Berlin, Heidelberg, pp. 367–381.
- Ajzen, I., Fishbein, M., 1980. *Understanding Attitudes and Predicting Social Behavior*. Prentice-Hall, Englewood Cliffs, NJ.
- Ansari, A., Mela, C.F., 2005. E-Customization. *J. Mark. Res.* 40 (2), 131–145.
- Armitage, C.J., Conner, M., 1999. Predictive validity of the theory of planned behaviour: the role of questionnaire format and social desirability. *J. Community Appl. Soc. Psychol.* 9 (4), 261–272.
- Asch, S.E., 1956. Studies of independence and conformity: a minority of one against a unanimous majority. *Psychol. Monogr.* 70 (9), 1–70.
- Audibert, J.-Y., Munos, R., Szepesvári, C., 2009. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theor. Comput. Sci.* 410 (19), 1876–1902.
- Bandura, A., 1991. Social cognitive theory of self-regulation. *Organ. Behav. Hum. Decis. Process.* 50 (2), 248–287.
- Bang, M., Torstensson, C., Katzeff, C., 2006. The PowerHouse: a persuasive computer game designed to raise awareness of domestic energy consumption. *Persuas. Technol.* 3962 (3962), 123–132.
- Bassili, J.N., 1996. Meta-judgmental versus operative indexes of psychological attributes: the case of measures of attitude strength. *J. Personal. Soc. Psychol.* 71 (4), 637–653.
- Berdichevsky, D., Neuenschwander, E., 1999. Toward an ethics of persuasive technology. *Commun. ACM* 42 (5), 51–58.
- Berry, D.A., Fristedt, B., 1985. *Bandit Problems: Sequential Allocation of Experiments* (Monographs on Statistics and Applied Probability). Chapman & Hall, London.
- Blass, T., 1991. Understanding behavior in the Milgram obedience experiment: the role of personality, situations, and their interactions. *J. Personal. Soc. Psychol.* 60 (3), 398–413.
- Bless, H., Bohner, G., Schwarz, N., Strack, F., 1990. Mood and persuasion. *Personal. Soc. Psychol. Bull.* 16 (2), 331–345.
- Brock, T.C., 1968. Implications of commodity theory for value change. In: Greenwald, A.G., Brock, T.C., Ostrom, T.M. (Eds.), *Psychological Foundations of Attitudes*. Academic Press, New York, pp. 243–275.
- Brug, J., 1999a. Computer-tailored nutrition education: differences between two interventions. *Health Educ. Res.* 14 (2), 249–256.
- Brug, J., Campbell, M., van Assema, P., 1999b. The application and impact of computer-generated personalized nutrition education: a review of the literature. *Patient Educ. Couns.* 36 (2), 145–156.
- Brug, J., Hospers, H.J., Kok, G., 1997. Differences in psychosocial factors and fat consumption between stages of change for fat reduction. *Psychol. Health* 12 (5), 719–727.
- Brug, J., Glanz, K., Van Assema, P., Kok, G., van Breukelen, G.J.P., 1998. The impact of computer-tailored feedback and iterative feedback on fat, fruit, and vegetable intake. *Health Educ. Behav.* 25 (4), 517–531.
- Brug, J., Oenema, A., Campbell, M., 2003. Past, present, and future of computer-tailored nutrition education. *Am. J. Clin. Nutr.* 77 (4 Suppl), 1028–1034.
- Burger, J.M., 1999. The foot-in-the-door compliance procedure: a multiple-process analysis and review. *Personal. Soc. Psychol. Rev.* 3 (4), 303–325.
- Cacioppo, J.T., Petty, R.E., 1982. The need for cognition. *J. Personal. Soc. Psychol.* 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. *J. Personal. Soc. Psychol.* 51 (5), 1032–1043.
- Cacioppo, J.T., Petty, R.E., Feinstein, J.A., Jarvis, W.B.G., 1996. Dispositional differences in cognitive motivation: the life and times of individuals varying in need for cognition. *Psychol. Bull.* 119 (2), 197–253.
- Chaiken, S., 1980. Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *J. Personal. Soc. Psychol.* 39 (5), 752.
- Chaiken, S., Liberman, A., Eagly, A.H., 1989. Heuristic and Systematic Information Processing Within and Beyond the Persuasion Context. *Unintended Thought* 212.
- Churchill, E.F., 2013. Putting the person back into personalization. *Interactions* 20 (5), 12–15.
- Cialdini, R., 2001. *Influence, Science and Practice*. Allyn & Bacon, Boston.
- Cialdini, R., 2004. The science of persuasion. *Sci. Am. Mind* 284, 76–84.
- Cialdini, R.B., Trost, M.R., 1998. Social influence: social norms, conformity, and compliance. In: Gilbert, D.T., Fiske, S.T., Lindzey, G. (Eds.), *The Handbook of Social Psychology*, vol. 2. McGraw-Hill, New York, pp. 151–192 (Chapter 21).
- Cialdini, R.B., Trost, M.R., Newsom, J.T., 1995. Preference for consistency: the development of a valid measure and the discovery of surprising behavioral implications. *J. Personal. Soc. Psychol.* 69 (2), 318–328.
- Consolvo, S., Libby, R., Smith, I., Landay, J.A., McDonald, D.W., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., 2008. Activity sensing

- in the wild. In: *Proceeding of the 26th Annual CHI Conference on Human Factors in Computing Systems—CHI '08*. ACM Press, New York, NY, USA, p. 1797.
- Consolvo, S., McDonald, D.W., Landay, J.A., 2009. Theory-driven design strategies for technologies that support behavior change in everyday life. In: *Proceedings of the 27th International Conference on Human Factors in Computing Systems—CHI 09*. ACM Press, New York, NY, USA, p. 405.
- Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., Taylor, J., 2001. Emotion recognition in human-computer interaction. *IEEE Signal Process. Mag.* 18 (1), 32–80.
- Crano, W.D., Prislin, R., 2006. Attitudes and persuasion. *Annu. Rev. Psychol.* 57, 345–374.
- Davidson, K.W., Goldstein, M., Kaplan, R.M., Kaufmann, P.G., Knatterud, G.L., Orleans, C.T., Spring, B., Trudeau, K.J., Whitlock, E.P., 2003. Evidence-based behavioral medicine: what is it and how do we achieve it?. *Ann. Behav. Med.* 26 (3), 161–171.
- Deci, E.L., Ryan, R.M., 2010. *Self-Determination*. Wiley Online Library Hoboken (John Wiley & Sons).
- Deterding, S., 2012. Gamification: designing for motivation. *Interactions* 19 (4), 14.
- Deutsch, M., Gerard, H.B., 1955. A study of normative and informational social influences upon individual judgment. *J. Abnorm. Soc. Psychol.* 51 (3), 629–636.
- de Vries, H., Brug, J., 1999. Computer-tailored interventions motivating people to adopt health promoting behaviours: introduction to a new approach. *Patient Educ. Couns.* 36 (2), 99–105.
- DiClemente, C.C., Marinilli, A.S., Singh, M., Bellino, L.E., 2001. The role of feedback in the process of health behavior change. *Am. J. Health Behav.* 25 (3), 217–227.
- Dijkstra, A., 2005. Working mechanisms of computer-tailored health education: evidence from smoking cessation. *Health Educ. Res.* 20 (5), 527–539.
- Dijkstra, A., De Vries, H., 1999. The development of computer-generated tailored interventions. *Patient Educ. Couns.* 36 (2), 193–203.
- Dillahunt, T., Becker, G., Mankoff, J., Kraut, R., 2008. Motivating environmentally sustainable behavior changes with a virtual polar bear. In: *Pervasive 2008 Workshop Proceedings*, pp. 58–62.
- Eisend, M., 2008. Explaining the impact of scarcity appeals in advertising: the mediating role of perceptions of susceptibility. *J. Advert.* 37 (3), 33–40.
- Faber, J.P., Markopoulos, P., Dadlani, P., van Halteren, A., 2011. AULURA: engaging users with ambient persuasive technology. In: Keyson, D.V., Maher, M.L., Streitz, N., Cheok, A., Augusto, J.C., Wichert, R., Englebienne, G., Aghajan, H., Kröse, B.J.A. (Eds.), *Ambient Intelligence. Lecture Notes in Computer Science*, vol. 7040. Springer, Berlin, Heidelberg, pp. 215–221.
- Festinger, L., 1957. *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford.
- Fishbein, M., Ajzen, I., 2011. *Predicting and Changing Behavior: The Reasoned Action Approach*. Taylor & Francis, London.
- Fogg, B.J., 1998. Persuasive computers: perspectives and research directions. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press/Addison-Wesley Publishing Co., Boston, pp. 225–232.
- Fogg, B.J., 2002. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann, Burlington, Massachusetts.
- Fogg, B.J., 2009. *Creating Persuasive Technologies: An Eight-Step Design Process*. Fogg, B.J., Eckles, D., 2007. Mobile persuasion: 20 perspectives on the future of behavior change. In: Fogg, B.J., Eckles, D. (Eds.), *Mobile Persuasion*. Stanford Captology Media, Stanford, pp. 1–166.
- Fogg, B., Nass, C., 1997a. Silicon sycophants: the effects of computers that flatter. *Int. J. Hum.-Comput. Stud.* 46 (5), 551–561.
- Fogg, B.J., Nass, C., 1997b. How users reciprocate to computers: an experiment that demonstrates behavior change. In: *CHI'97 Extended Abstracts on Human Factors in Computing Systems: Looking to the Future*. March. ACM, New York, pp. 331–332.
- Fuegen, K., Brehm, J.W., 2004. The Intensity of Affect and Resistance to Social Influence. Lawrence Erlbaum, Hillsdale, pp. 39–63.
- Garde-Perik, E.M., 2009. *Ambient Intelligence & Personalization: People's Perspectives on Information Privacy* (Ph.D. thesis). Technical University of Eindhoven.
- Garner, R., 2005. What's in a name? Persuasion perhaps. *J. Consum. Psychol.* 15 (2), 108–116.
- Gerber, B.S., Stolley, M.R., Thompson, A.L., Sharp, L.K., Fitzgibbon, M.L., 2009. Mobile phone text messaging to promote healthy behaviors and weight loss maintenance: a feasibility study. *Health Informatics J.* 15 (1), 17–25.
- Goldstein, N.J., Cialdini, R.B., Griskevicius, V., 2008. A room with a viewpoint: using social norms to motivate environmental conservation in hotels. *J. Consum. Res.* 35 (3), 472–482.
- Gosling, S.D., Rentfrow, P.J., Swann, W.B.J., 2003. A very brief measure of the big-five personality domains. *J. Res. Personal.* 37 (6), 504–528.
- Greenberg, M.S., 1980. A theory of indebtedness. In: *Social exchange*. Springer, Berlin, Heidelberg, pp. 3–26.
- Gretzel, U., Fesenmaier, D., 2006. Persuasion in recommender systems. *Int. J. Electron. Commer.* 11 (2), 81–100.
- Guadagno, R., Cialdini, R., 2005. Online persuasion and compliance: social influence on the Internet and beyond. In: Amichai-Hamburger, Y. (Ed.), *The Social Net*. Oxford University Press, Oxford, New York, p. 296.
- 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. *Personal. Soc. Psychol. Bull.* 27 (7), 859–867.
- Halko, S., Kientz, J.A., 2010. Personality and persuasive technology: an exploratory study on health-promoting mobile applications. In: *Persuasive technology*. Springer, Berlin, Heidelberg, pp. 150–161.
- Hardin, C.D., Higgins, E.T., 1996. *Shared Reality: How Social Verification Makes the Subjective Objective*, vol. 3. Guilford, New York, pp. 28–84.
- Haugtvedt, C.P., Petty, R.E., 1992. Personality and persuasion: need for cognition moderates the persistence and resistance of attitude changes. *J. Personal. Soc. Psychol.* 63 (2), 308–319.
- Haugtvedt, C.P., Petty, R.E., Cacioppo, J.T., 1992. Need for cognition and advertising: understanding the role of personality variables in consumer behavior. *J. Consum. Psychol.* 1 (3), 239–260.
- Haugtvedt, C.P., Liu, K., Min, K.S., 2008. Individual differences: tools for theory testing and understanding in consumer psychology research. In: *Handbook of Consumer Psychology*. CRC Press, Boca Raton, pp. 1161–1176.
- Hauser, J.R., Urban, G.L., Liberali, G., Braun, M., 2009. Website morphing. *Mark. Sci.* 28 (2), 202–223.
- Hauser, J.R., Liberali, G., Urban, G.L., 2014. Website morphing 2.0: Switching costs, partial exposure, random exit, and when to morph. *Manag. Sci.*
- Heider, F., 1944. Social perception and phenomenal causality. *Psychol. Rev.* 51 (6), 358.
- Hornstein, H.A., Fisch, E., Holmes, M., 1968. Influence of a model's feeling about his behavior and his relevance as a comparison other on observers' helping behavior. *J. Personal. Soc. Psychol.*
- Inman, J., Peter, A., Raghuraj, P., 1997. Framing the deal: the role of restrictions in accentuating deal value. *J. Consum. Res.* 24 (1), 68–79.
- James, J.M., Bolstein, R., 1990. The effect of monetary incentives and follow-up mailings on the response rate and response quality in mail surveys. *Public Opin. Q.* 54 (3), 346.
- Jarvis, W.B.G., Petty, R.E., 1996. The need to evaluate. *J. Personal. Soc. Psychol.* 70 (1), 172–194.
- Johnson, B.T., Eagly, A.H., 1989. Effects of Involvement on Persuasion: A Meta-Analysis.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47 (2), 263–291.
- Kantor, P.B., Rokach, L., Ricci, F., Shapira, B., 2011. *Recommender Systems Handbook*, vol. 54, 11. Springer, Berlin, Heidelberg, pp. 217–253.
- Kaptein, M.C., 2011. Adaptive Persuasive Messages in an E-commerce Setting: the use of Persuasion Profiles. In: *Proceedings of ECIS 2011*, Helsinki.
- Kaptein, M.C., Duplinsky, S., 2013. Combining multiple influence strategies to increase consumer compliance. *Int. J. Internet Mark. Advert.* 8 (1), 32–53.
- Kaptein, M.C., Eckles, D., 2010. Selecting Effective Means to Any End: Futures and Ethics of Persuasion Profiling. In: *Persuasive Technology* (pp. 82–93). Springer, Berlin, Heidelberg.
- Kaptein, M.C., Eckles, D., 2012. Heterogeneity in the effects of online persuasion. *J. Interact. Mark.* 26 (3), 176–188.
- Kaptein, M.C., van Halteren, A., 2013. Adaptive persuasive messaging to increase service retention. *J. Personal Ubiquitous Comput.* 17 (6), 1173–1185.
- Kaptein, M.C., Markopoulos, P., Ruyter, B., Aarts, E., 2009. Persuasion in ambient intelligence. *J. Ambient Intell. Humaniz. Comput.* 1 (1), 43–56.
- Kaptein, M.C., Eckles, D., Davis, J., 2011. Envisioning persuasion profiles: challenges for public policy and ethical practice. *ACM Interact.* 18 (5), 66–69.
- Kaptein, M.C., de Ruyter, B., Markopoulos, P., Aarts, E., 2012. Tailored persuasive text messages to reduce snacking. *Trans. Interact. Intell. Syst.* 2 (2), 10–35.
- Kellermann, K., Cole, T., 1994. Classifying compliance gaining messages: taxonomic disorder and strategic confusion. *Commun. Theory* 4 (1), 3–60.
- Kelman, C., Hamilton, V.L., 1989. *Crimes of Obedience: Toward a Social Psychology of Authority and Responsibility*. Yale University Press, New Haven.
- Knowles, E.S., Linn, J.A., 2004. *Resistance and Persuasion*. Psychology Press, Abingdon.
- Komorita, S.S., Hilty, J.A., Parks, C.D., 1991. Reciprocity and cooperation in social dilemmas. *J. Confl. Resolut.* 35 (3), 494–518.
- Kreuter, M., 2000. *Tailoring Health Messages: Customizing Communication with Computer Technology*. Routledge.
- Kreuter, M.W., Strecher, V.J., 1996. Do tailored behavior change messages enhance the effectiveness of health risk appraisal? Results from a randomized trial. *Health Educ. Res.* 11 (1), 97–105.
- Kukar-Kinney, M., Close, A.G., 2009. The determinants of consumers' online shopping cart abandonment. *J. Acad. Mark. Sci.* 38 (2), 240–250.
- Lambert, C.L., 2001. Baby think it over. *Pediatrics* 107 (4), 806.
- Leone, C., Wallace, H., Modglin, K., 1999. The need for closure and the need for structure: interrelationships, correlates, and outcomes. *J. Psychol.* 133 (5), 553–562.
- Long, J.D., Stevens, K.R., 2004. *Using Technology to Promote Self-Efficacy for Healthy Eating in Adolescents*. Technical Report 2, Lubbock Christian University, 5601 W. 19th Street, Lubbock, TX 79407, USA. www.joann.long@lcu.edu.
- Lynn, M., 1989. Scarcity effects on desirability: mediated by assumed expensive-ness?. *J. Econ. Psychol.* 10 (2), 257–274.
- Lynn, M., 1991. Scarcity effects on value: a quantitative review of the commodity theory literature. *Psychol. Mark.* 8 (1), 43–57.
- Madden, T.J., Ellen, P.S., Ajzen, I., 1992. A comparison of the theory of planned behavior and the theory of reasoned action. *Personal. Soc. Psychol. Bull.* 18 (1), 3–9.
- Maio, G.R., Esses, V.M., 2001. The need for affect: individual differences in the motivation to approach or avoid emotions. *J. Personal.* 69 (4), 583–615.
- Martin, R., Hewstone, M., 2003. Social influence processes of control and change: conformity, obedience to authority, and innovation. In: Hogg, M.A., Cooper, J. (Eds.), *Psychological Foundations of Attitudes*. Sage, London, pp. 347–366.
- Maslow, A., Herzberg, A., 1954. Hierarchy of needs. In: Maslow, A.H. (Ed.), *Motivation and Personality*. Harper, New York.
- McFarland, R.G., Challagalla, G.N., Shervani, T.A., 2006. Influence tactics for effective adaptive selling. *J. Mark.* 70 (4), 103–117.
- McGraw, K.L., 2010. *The Effects of Persuasive Motivational Text Messaging on Adherence to Diet and Exercise Programs Across Different Personality Traits*. (Ph.D. thesis). Fielding Graduate University.

- Michie, S., Richardson, M., Johnston, M., Abraham, C., Francis, J., Hardeman, W., Eccles, M.P., Cane, J., Wood, C.E., 2013. The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions. *Ann. Behav. Med.* 46 (1), 81–95.
- Midden, C., Mccalley, T., Ham, J., Zaalberg, R., 2008. Using persuasive technology to encourage sustainable behavior. In: Workshop paper at Sixth International Conference on Pervasive Computing, no. 1. Eindhoven University of Technology, pp. 83–86.
- Milgram, S., 1974. *Obedience to Authority*. Tavistock, London.
- Modigliani, A., Rochat, F., 1995. The role of interaction sequences and the timing of resistance in shaping obedience and defiance to authority. *J. Soc. Issues* 51 (3), 107–123.
- Morris, M., Guilak, F., 2009. Mobile heart health: project highlight. *IEEE Pervasive Comput.* 8 (2), 57–61.
- Nass, C., Fogg, B.J., Moon, Y., 1996. Can computers be teammates? *Int. J. Hum.-Comput. Stud.* 45 (6), 669–678.
- Noar, S.M., Benac, C.N., Harris, M.S., 2007. Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. *Psychol. Bull.* 133 (4), 673–693.
- Nov, O., Arazy, O., 2013. Personality-targeted design: theory, experimental procedure, and preliminary results. In: Proceedings of the 2013 Conference on Computer Supported Cooperative Work. ACM, New York, pp. 977–984.
- O'Keefe, D.J., 1994. From strategy-based to feature-based analyses of compliance gaining message classification and production. *Commun. Theory* 4 (1), 61–69.
- Ochi, P., Rao, S., Takayama, L., Nass, C., 2010. Predictors of user perceptions of web recommender systems: how the basis for generating experience and search product recommendations affects user responses. *Int. J. Hum.-Comput. Stud.* 68 (8), 472–482.
- Oinas-Kukkonen, H., Harjumaa, M., 2008. A systematic framework for designing and evaluating persuasive systems. In: Oinas-Kukkonen, H., Hasle, P. F.V., Harjumaa, M., Segerstahl, K., Ohrström, P. (Eds.), *Persuasive Technology, Proceedings of the Third International Conference (PERSUASIVE) 2008*, Oulu, Finland, June 4–6, 2008. Lecture Notes in Computer Science, vol. 5033. Springer, pp. 164–176.
- Oinas-kukkonen, H., Harjumaa, M., 2009. Persuasive systems design: key issues, process model, and system features. *Commun. Assoc. Inf. Syst., Berlin, Heidelberg* 24 (1), 28.
- Ortega, P., Braun, D., 2013. Generalized Thompson Sampling for Sequential Decision-Making and Causal Inference. arXiv preprint: arXiv:1303.4431.
- Petty, R.E., Cacioppo, J.T., 1986. The elaboration likelihood model of persuasion. *Adv. Exp. Soc. Psychol.* 19 (1), 123–205.
- Petty, R.E., Wegener, D.T., 1999. The elaboration likelihood model: current status and controversies. In: Chaiken, S., Trope, Y. (Eds.), *Dual-process Theories in Social Psychology*. Guilford Press, New York, pp. 41–72.
- Plummer, J.T., 1974. The concept and application of life style segmentation. *J. Mark.* 38 (1), 33–37.
- Preece, J.J., 2010. I Persuade, They Persuade, It Persuades!. In: *Persuasive Technology, 2010*, pp. 2–3.
- Prochaska, J.O., Velicer, W.F., 1997. The transtheoretical model of health behavior change. *Am. J. Health Promot.* 12 (1), 38–48.
- Rhoads, K., 2007. How Many Influence, Persuasion, Compliance Tactics & Strategies Are There? (<http://www.workingspsychology.com/numbertactics.html>).
- Sakai, R., van Peteghem, S., van de Sande, L., Banach, P., Kaptein, M.C., 2011. Personalized persuasion in ambient intelligence: the APStairs system. In: *Proceedings of Ambient Intelligence (AmI) 2011*, Amsterdam.
- Salganik, M.J., Watts, D.J., 2009. Web-based experiments for the study of collective social dynamics in cultural markets. *Top. Cogn. Sci.* 1 (3), 439–468.
- Scherini, T., Melo, P., van Craenendonck, T., Zou, W., Kaptein, M.C., August 2010. Enhancing the sleeping quality of partners living apart. In: *Proceedings of the Eighth ACM Conference on Designing Interactive Systems—DIS '10*. ACM Press, New York, NY, USA, pp. 171–179.
- Schmidt, A., Gellersen, H.-W., Merz, C., 2000. Enabling implicit human computer interaction: a wearable RFID-tag reader. In: *Digest of Papers, Fourth International Symposium on Wearable Computers*. IEEE Computer Society, Piscataway, pp. 193–194.
- Schwartz, E.M., Bradlow, E., Fader, P., 2013. Working Paper Customer Acquisition via Display Advertising Using Multi-Armed Bandit Experiments (1217).
- Scott, S.L., 2010. A modern Bayesian look at the multi-armed bandit. In: *Applied Stochastic Models in Business and Industry*, vol. 26, no. 6, pp. 639–658.
- Skinner, B.F., 1976. *About Behaviorism*. Vintage Books, New York.
- Snijders, T.A.B., Bosker, R.J., 1999. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. Sage Publications Ltd, Thousand Oaks.
- Stallworth, B., 2010. Future imperfect: googling for principles in online behavioral advertising. *Fed. Commun. Law J.* 62, 465–491.
- Strecher, V.J., Kreuter, M., Den Boer, D.-J., Koberin, S., Hospers, H.J., Skinner, C.S., 1994. The effects of computer-tailored smoking cessation messages in family practice settings. *J. Fam. Pract.* 39 (3), 262–270.
- Sutton, R.S., Barto, A.G., 1998. *Introduction to Reinforcement Learning*. MIT Press, Cambridge, Massachusetts.
- Svane, O., 2007. Helping, informing or coaxing the consumer? Exploring persuasive technology as applied to households energy use. *Policy* 1 (11), 1–11.
- Tormala, Z.L., DeSensi, V.L., 2009. The effects of minority/majority source status on attitude certainty: a matching perspective. *Personal. Soc. Psychol. Bull.* 35 (1), 114–125.
- Tynan, A.C., Drayton, J., 1987. Market segmentation. *J. Mark. Manag.* 2 (3), 301–335.
- Vanharreveld, F., 2004. Attitudes as stable and transparent constructions. *J. Exp. Soc. Psychol.* 40 (5), 666–674.
- West, S.G., 1975. Increasing the attractiveness of college cafeteria food: a reactance theory perspective. *J. Appl. Psychol.* 60 (5), 656–658.
- Whittle, P., 1980. Multi-armed bandits and the Gittins index. *J. R. Stat. Soc. Ser. B Methodol.* 42 (2), 143–149.
- Yue, Y., Hong, S., Guestrin, C., 2012. Hierarchical Exploration for Accelerating Contextual Bandits. arXiv preprint: arXiv:1206.6454.
- Zanker, M., Bricman, M., Jessenitschnig, M., Damiani, E., Jeong, J., Howlett, R., Jain, L., 2009. *New Directions in Intelligent Interactive Multimedia Systems and Services—2. Studies in Computational Intelligence*, vol. 226. Springer, Berlin, Heidelberg.
- Zhu, F., Zhang, X.M., 2010. Impact of online consumer reviews on sales: the moderating role of product and consumer characteristics. *J. Mark.* 74 (2), 133–148.