Personalized persuasion in ambient intelligence

Citation for published version (APA):

Kaptein, M. C. (2012). *Personalized persuasion in ambient intelligence*. [Phd Thesis 1 (Research TU/e / Graduation TU/e), Industrial Design]. Technische Universiteit Eindhoven. <https://doi.org/10.6100/IR729200>

DOI:

[10.6100/IR729200](https://doi.org/10.6100/IR729200)

Document status and date:

Published: 01/01/2012

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.

Personalized Persuasion
in
Ambient Intelligence

This research was conducted in part at the User Centered Engineering group, Department of Industrial Design, Eindhoven University of Technology (Eindhoven, The Netherlands), in part at Philips Research Laboratories (Eindhoven, The Netherlands), and in part at the Department of Communication, Stanford University (Stanford, CA). This research was commissioned and funded by Philips Research.

An electronic copy of this thesis in PDF format is available from the TU/e library website (<http://www.tue.nl/lib>).

©2011-2012, Maurits Clemens Kaptein, The Netherlands

All rights reserved. Reproduction of this publication in whole or in part is prohibited without the prior permission from the author.

A catalogue record is available from the Eindhoven University of Technology Library

ISBN: 978-90-386-3106-6

Personalized Persuasion in Ambient Intelligence

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op donderdag 29 maart 2012 om 16.00 uur

door

Maurits Clemens Kaptein

geboren te Doetinchem

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr. E.H.L. Aarts

en

prof.dr. P. Markopoulos

Copromotor:

dr. B.E.R. de Ruyter

Contents

1	Introduction	1
1.1	Background: Persuasion in Intelligent Systems	1
1.2	Research Problem	3
1.3	Outline	5
1.4	Key Contributions	6
2	Related Work	9
2.1	Persuasive Technologies	9
2.2	Attitude and Behavioral Change	14
2.3	Conclusion: Reflections	26
3	Insight Generation I:	
	<i>Heterogeneity in Responses to Social Influence Strategies</i>	29
3.1	Introduction	29
3.2	Study 1: Identifying an upper bound in heterogeneity	31
3.3	Study 2: Examining stability of heterogeneity over time	41
3.4	Study 3: Stability of heterogeneity across context	47
3.5	Conclusions	54
4	Insight Generation II:	
	<i>Choice and Repetition Effects of Influence Strategies</i>	57
4.1	Introduction	57
4.2	Study 1: Choice Effects	58
4.3	Study 2 and 3: Simultaneous Presentation	65
4.4	Study 4: Simultaneous Presentation in Practice	73
4.5	Conclusions	76
5	Insight Generation III:	
	<i>Meta-Judgemental Measures</i>	79
5.1	Introduction: Measures of Persuasion Susceptibility	79
5.2	Study 1: Initial application of the short <i>STPS</i>	80
5.3	Study 2: Development and Validation of the <i>STPS</i>	85
5.4	Conclusions	93

6	Design Requirements	97
6.1	Conclusions From the <i>Insight Generation</i> Chapters	97
6.2	Identification, Representation, and Measurement	102
6.3	Dynamic Adaptation	104
6.4	Properties of Adaptive Persuasive Systems	111
7	Case Studies I:	
	<i>Influencing Health Decisions</i>	115
7.1	Introduction: Measures of Susceptibility	115
7.2	Study 1 & 2: Promoting Lunch Walks	115
7.3	Study 3: Influencing Snacking Behavior	122
7.4	Conclusions	136
8	Case Studies II:	
	<i>Dynamically Personalizing Persuasion</i>	139
8.1	Introduction: Dynamic Persuasive Systems	139
8.2	Design 1: APStairs, Promoting Stair Usage	140
8.3	Design 2: Influencing Docking Behavior	148
8.4	Design 3: E-Selling Applications	157
8.5	Conclusions	166
9	Reflections	169
9.1	Advancing the Design of Persuasive Technologies	169
9.2	The Possible Future(s) of Persuasion Profiles	172
9.3	Ethical Considerations	176
9.4	Future Research Challenges	180
	Bibliography	185
	Articles by Maurits Kaptein	205
	Summary	209
	Acknowledgments	211
	Curriculum	213
	Index	215
	List of Figures	219
	List of Tables	222

1

Introduction

1.1 Background: Persuasion in Intelligent Systems

Persuasion has been a subject of human investigation for millennia. The ancient Greek deemed the topic of “moving or inducing someone by words to do something” sufficiently important to give Peitho (Πειθω)—the goddess of persuasion—a place next to the better known Eros and Aphrodite. The ability to persuade others to enact a certain behavior or adopt a certain opinion has appealed to those studying design, psychology, communication, and marketing. Applications of persuasion are manifold: health care professionals use persuasion to influence their clients to lead a healthy lifestyle while marketers use persuasion to drive products sales.

Psychology and communication scholars define persuasion as “...*a symbolic process in which communicators try to convince other people to change their attitudes or behaviors regarding an issue through the transmission of a message in an atmosphere of free choice.*” (Perloff, 2003) This implies that persuasion concerns a persuader who, by the act of communication, changes the attitudes or behavior of a target. Given this definition *classical conditioning, operant conditioning, the formation of attitudes*, and the more contemporary work on *perceived and preferred attitude basis* all concern persuasion (Kaptein et al., 2009b).

These different aspects of persuasion together compose one of the most widely studied topics in social science.

Persuasion is traditionally studied in a human-to-human context. However, Reeves and Nass (1996), in their book “The Media Equation: How People Treat Computers, Television and New Media Like Real People and Places”, sum up the overwhelming evidence presented in earlier papers (e.g. Nass et al., 1994; Moon and Nass, 1996; Fogg and Nass, 1997; Nass, 1997) that people treat computers and other technologies as *social actors*. Reeves and Nass (1996) replicated some of the most compelling results of social science—such as the tendency of people to form teams and the associated between-team rivalry by mere random assignment of different colored t-shirts (Tajfel, 1982; Mackie and Cooper, 1984; Mackie, 1986)—in the context of human-computer as opposed to human-human interaction. This work and the work of earlier scholars such as Dennett (1987) and Flavell et al. (1990) “opened the door for computers to apply [...] social influence strategies (Fogg, 2002)”. Thus, interactive technologies and systems have the potential to engage *in the same persuasion processes that humans do*.

Fogg (1999), was one of the first to make explicit the importance of the findings presented in “The Media Equation” (Reeves and Nass, 1996) for the design of interactive systems designed with the intent to change human attitudes or behaviors. With his book on the topic Fogg (2002) initiated a field that is now called *persuasive technology*. Persuasive technology researchers study *interactive systems that are intentionally designed to change user’s attitudes or behavior* (Oinas-Kukkonen et al., 2008; Chatterjee and Dev, 2009; Ploug et al., 2010; IJsselstein et al., 2006, e.g.). While the initial work by Fogg (2002) focused heavily on the use of computers in their traditional sense—e.g. the screen, keyboard, mouse combination commonplace in office buildings around the world—engineers are now incorporating ubiquitous sensors and mobile devices in their systems. Because of these recent developments scholars have argued that in the future persuasive technologies will be more effective than their human counterparts since they can be more persistent and “always on” (Fogg and Eckles, 2007).

Recently, Aarts et al. (2007) combined the notions of *ambient intelligent (AmI) systems*—systems that build on the large scale integration of electronic devices and the ubiquitous availability of digital information—and persuasive technologies. In an AmI world, distributed devices operate collectively while embedded in the environment using information and intelligence that is hidden in the interconnection network

(Aarts and Ruyter, 2009). The AmI scenario extends earlier ideas about ubiquitous computing (Weiser, 1991) and provides numerous opportunities for persuasion which are out of reach for traditional computing systems.

1.2 Research Problem

Even though persuasive technologies are already commercially available (see also 2.1), a number of aspects are poorly understood. Persuasive technologies have mainly made their way into marketing applications where positive average effects—effects over groups of people—are a reason to adopt new technologies. For example, the *average* sales of an e-commerce website increases by the use of persuasion. A more recent focus of persuasive technologies to support *individual attitude and behavior change*¹ however benefits less from obtaining these average effects. Rather, these technologies should attend to *individual level* effects: the attitudinal or behavioral change of a single user. More and more technologies are marketed with the *promise* of changing the behavior of an individual user and thus *designers of these systems are required to design for individual level behavior change instead of average effects*.

The problem of designing for individual behavior change has proven notoriously hard. While the first measurable persuasion effects *on average* were surprisingly large—as in the infamous Milgram experiment (Milgram, 1974)—establishing predictable behavior or attitude change *at an individual level* has failed repeatedly. To illustrate, consider the persuasion principle that is responsible for the effects observed in the Milgram (1974) study called *Authority*. Experiments show that authority appeals (e.g. “A fitness instructor recommends you to have a 30 minute run today”) are on average more effective than similar appeals without an authority argument (e.g. “You should have a 30 minute run today”). However, this finding does not imply that *everyone* will comply or even benefit from usage of the authority strategy (see Chapter 3). Actually, even in the well-known experiments performed by Milgram (1974), about one third of the participants failed to comply to the authority argument.

Reliably affecting an individual’s attitudes or behavior through persuasion (rather than coercion) is a goal that is still out of reach of today’s persuasive technologies. Scholars like Fogg and Eckles (2007)

¹For example Philips DirectLife (<http://directlife.philips.com>) or the Lark (<http://www.lark.com>).

have hypothesized that to do so persuasive technologies should deliver (a) the right message, at (b) the right time, in (c) the right way to be able to be effective. This truism identifies three important aspects of successful attitude and behavioral change: First, the target of the persuasive attempt needs to be receptive to the end goal of the attempt. Second, the message needs to be delivered at a time that enables the user to attend to it, and if immediate action is required one that provides the opportunity for the action. Finally, large variation can exist in the *way in which the request is framed*: a 30 minute run can be proposed by a fitness expert, but can also be supported by a statement that 80% of healthy people frequently run for 30 minutes.

While large numbers of theorists and experimentalists in the field of persuasive communication have struggled over the *number of ways* that can be distinguished to identify a persuasive request, or even what defines such a way, *adapting* the way to the current user or persuadee is advocated throughout all fields that study persuasion: marketeers advocate to adapt the sales tactic used to endorse a product to a consumer (McFarland et al., 2006), while health-care professionals promote tailoring of the persuasive strategies used to gain medication compliance (Gerber et al., 2009; Ribeiro et al., 2011). Within the health-care field a large literature on *computer-tailored health education*, in which personalization that was traditionally carried out by nutrition counselors is now (partially) replaced by interactive expert systems (Brug et al., 2003), has also recently emerged (see, e.g. de Vries and Brug, 1999; Kroeze et al., 2006). In most of these cases it is advocated that the “way” instead of the end goal of a persuasive request is personalized to an individual persuadee. Kaptein and Eckles (2010) describe this adaptation of the “way” by focusing on different *means* to an end. These *means* are the core topic of this thesis.

The question “*how can the means in which a persuasive request is made by a persuasive systems be personalized to increase its impact?*” is the main research question addressed in this thesis. The question is split into several parts which are addressed separately. The first part of this thesis—the *insight generation* section—focusses on the questions concerning human behavior that need to be answered:

1. How large is the heterogeneity—the difference between individuals—in responses to different ways in which persuasive requests are framed? (Chapter 3)
2. How do people respond to multiple influence strategies (ways)

that all support the same request? (Chapter 4)

3. Can we measure—using questionnaires—people’s susceptibility to these different ways in which requests are framed? (Chapter 5)

Based on the knowledge that is gained about human behavior by empirically answering the above questions two more questions arise:

1. How do we apply the obtained susceptibility profile in the context of a persuasive system? (Chapter 7)
2. Can a system dynamically (sequentially) “learn” a user’s susceptibility based on his or her behavioral responses instead of using questionnaire measures? (Chapter 8)

These latter two questions are addressed in the *case study* section of this thesis in which a number of persuasive interventions are designed, implemented, and evaluated.

1.3 Outline

Figure 1.1 shows a graphical overview of the structure of this thesis. Chapter 1 gives a background of the problem that is addressed and states the contributions of this thesis. Chapter 2 describes the current state of the persuasive technology field by reviewing parts of the literature relevant for the questions answered in this thesis. These two chapters together thus provide the reader with an overview of the current state of the persuasive technology field.

From this starting state the thesis reports a number of experimental studies which are presented in the chapters called *insight generation* chapters, Chapters 3, 4, and 5. These studies explore human responses to repeated exposure to different persuasion principles and establish the need to personalize influence attempts at the level of *social influence strategies*. These chapters show that (a) the heterogeneity in responses to social influence strategies is large compared to their average effect, (b) combining multiple strategies does not necessarily increase persuasion, and (c) questionnaire measures can be used to reliably assess people’s susceptibility to distinct persuasive strategies.

Chapter 6 summarizes the findings from the *insight generation* studies and presents a generalized solution for the design of adaptive persuasive systems. The requirements to incorporate the findings from the previous chapters into design to create *personalized* persuasive systems are presented. Next, the *case studies* presented in the second part of this thesis describe the implementation and evaluation of systems which directly apply the findings from previous chapters. In Chapter 7 two

Figure 1.1: Graphical representation of the structure of this thesis. The *insight generation* chapters explore several questions regarding influence strategies using an experimental approach. In the *case studies* chapters these findings are validated in several field studies.

studies are presented that show how adaptation to individual *susceptibility* to different persuasion strategies increases the success of different health interventions. Next in Chapter 8 three instances of systems that use *dynamic* adaptation based on operative measures are deployed and evaluated. Finally in Chapter 9, the implications of the current findings for the design of persuasive technologies, and social science studies into the effects of persuasive strategies are discussed.

1.4 Key Contributions

The overall aim of this thesis is to advance the design of successful persuasive technologies by introducing a method to personalize the *means* by which a persuasive message is presented to users. The thesis adds to the current literature by extending our understanding of human behavior, enhancing the methods used to understand these behaviors, and advancing technology.

1.4.1 Contributions to our understanding of Human behavior

Historically human-to-human persuasion attempts have had large *impact*—as manifest in the abilities of good salesmen or debaters to convince their audiences—but they have had limited *reach*: only a limited number of individuals could be addresses at the same time. With the emergence of mass media and technology mediated communication the *reach* of persuasive attempts has increased dramatically but the average *impact* of each individual attempt has decreased. This is evidenced for example by the fact that conversion rates in physical stores are magnitudes higher than those in e-commerce (see Moe and Fader, 2004). The main contribution of this thesis to our understanding of human behavior is to provide a description of one of the antecedents of persuasive effectiveness, in a way that facilitates usage at a large scale.

To examine the impact of persuasion, the thesis estimates the heterogeneity—individual differences—in responses to social influence strategies. The thesis shows that while different means to an end such as using authority endorsements or consensus appeals all increase compliance *on average*, the stable differences between people *are larger than these average effects*. Applied to a single individual, using an authority appeal to promote a request can have *negative effects that are far larger in size than the positive effect found when averaging over groups of people*. These individual differences prove stable both over time and across contexts.

The thesis then advances our understanding of responses to persuasive appeals by specifically comparing different means in which a single request can be framed and by exploring the use of multiple of these means for a single appeal. Thus, an appeal to work out more can be supported both by a health expert (*Authority*) and by a group of similar others (*Consensus*). The studies presented in this thesis show that combining persuasive strategies does not increase chances of success (Chapter 4). Together, these findings motivate a shift from an understanding of the average behavior of groups, to the repeated behavior of individuals. The thesis shows that personalized persuasion—persuasion adapted to individual’s previous responses—outperforms persuasion based on theories of average effects. Human behavior thus proves variable between people, but consistent within them (see Chapter 3). If we want to persuade on a large scale and with a large impact, our persuasive technologies should implement personalization at the level of social influence strategies.

1.4.2 Contributions to Methodology

Social science experiments and theories are by-and-large based on the examination of average effects: effects over groups of people. While these average effects are often of interest to researchers and policy makers, they are not necessarily a good summary of the effect of psychological processes within individuals. Only if the effects of the variables under study are homogeneous can researchers use the average treatment effect as an estimate of the individual level effect. If however effects are largely heterogeneous, the average treatment effect does not properly capture the nature of the effect under study and thus should not be used for the development of individual-level theory in social science (Cf. Hutchinson et al., 2001).

This thesis presents the use of hierarchical multi-level models with

crossed random effects² to study individual differences—instead of using these models to more accurately estimate average treatment effects. The within-subject approach used in a number of the studies presented in this thesis enables researchers to estimate individual level effects of persuasion strategies while controlling for other factors. These models can also be used to compare the heterogeneity of responses of individuals to the average effect under study. If the heterogeneity is small compared to the average effect, then it is sensible that theoretical models of persuasive communication can indeed be based on average effects. However, this thesis shows that for some persuasion processes this is not the case, and thus theoretical explanations of persuasion that focus on average effects are misplaced.

1.4.3 Contributions to Technology

As a final contribution, this thesis guides designers of adaptive personalized persuasive systems by detailing the three core requirements that need to be met to implement personalized persuasion in ambient intelligence. *Identification*, a means to identify users, *representation*, the technological ability to change representations of appeals, and *measurement*, the ability to log the success of the persuasive attempt, are necessities when creating adaptive personalized persuasive systems. Several systems which address each of these three requirements in different ways are presented in this thesis. By detailing these systems and detailing the implementations of the three design requirements in different application domains the thesis advances our understanding of the design of persuasive technologies.

This thesis presents the core requirements that algorithms used to create *adaptive* personalized persuasive systems need to address. The proposed multi-level model, in which the higher level model is computed a-synchronously, is both effective in identifying the best way to deliver an appeal to an individual users as well as computationally feasible. Several methods to optimize the explore-exploit trade-off that naturally arises in the design of adaptive systems are presented to aid designers in creating personalized persuasive systems³.

²See section 3.1.1 or Gelman (2005) for clarification.

³This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2009b; Kaptein and Eckles, 2010).

2

Related Work

2.1 Persuasive Technologies

We have entered an era of persuasive technology, of *interactive computing systems intentionally designed to change people's attitudes and behaviors* (IJsselsteijn et al., 2006). These systems emerged for the first time around the 80's with a small selection of research prototypes of computing systems that were designed to promote health or increase workplace productivity. In 1981 a description of a computing system named "Body Awareness Resource Network" (BARN) was published. The system provided health information on topics such as drugs, smoking, and exercise to adolescents with the aim of improving health related behaviors (Bosworth et al., 1983). Following the massive growth of the internet the emergence of persuasive technologies truly took off. Website developers started to think actively about attitude and behavioral change and started implementing findings from social science to make their websites more effective.

This chapter describes the *current* persuasive technology field. The chapter begins by listing a number of examples of persuasive technologies that are exemplar of the different target behaviors or attitudes than have gained interest from designers of persuasive systems. Next, the chapter specifically addresses persuasion in ambient intelligence by

making explicit the ambient intelligence vision and identifying the opportunities for persuasive systems.

Since this thesis focusses on the use of social influence strategies the chapter briefly reviews parts of the social science literature that are concerned with attitude and behavioral change and are deemed relevant for the design of persuasive systems. The literature on social influence and possible taxonomies of social influence strategies is brought forward. A special focus in this section is on previous work that identifies *individual differences in responses to social influence strategies*. The overview of persuasive systems and the review of the literature on social influence strategies are then combined to detail the current state of persuasive technologies. Two main limitations of the current state are identified which give rise to the questions that are addressed in Chapters 3, 4, and 5.

2.1.1 Examples of Persuasive Technologies

The number of persuasive systems grew rapidly in conjunction with the growth of the web. Early persuasive applications focussed on education of the public. A few years later, the landscape of persuasive technologies started to shape up, and it became clear that four application areas warranted the most attention of researchers and practitioners: Health, Safety, Environmental conservation, and marketing.

The first health related persuasive systems concerned the 5-A-Day Adventures cd-rom distributed by Dole, and Baby Think It Over (BTIO) (King and Tester, 1997; Lambert, 2001; De Anda, 2006). The first of these used animated characters, music, and small games to encourage children to eat more fruit and vegetables. BTIO is an interactive doll first released in 1998 which is still on the market today. Through simulation the doll teaches young potential parents the pitfalls of parenthood. BTIO was included in several obligatory high-school courses to prevent teenage pregnancy. More recently health related commercial products have hit the market that focus on maintaining an active lifestyle (e.g. Philips DirectLife and Nike +) and adopting a healthy sleeping regime (e.g. Lark and MyZeo). Persuasive technologies to increase general health also include text-based interventions to aid users in giving up smoking (Khaled et al., 2008; Preece, 2010) or consume less snacks (Long and Stevens, 2004; Bubb, 2007; Arteaga et al., 2009; Kaptein et al., 2011a).

The second thriving field, that of safety, was defined early on by www.kustomsignals.com. This website educates users of the advan-

tages of driving safely and reducing their overall driving speed (King and Tester, 1997). Other educative websites that train users to increase safety in many domains have emerged after this example (Yeo et al., 2009; Mintz and Aagaard, 2010; Chittaro and Zangrando, 2010). Next to web-based applications persuasive technologies that go beyond the “boxed” computer paradigm have also been developed: the Hygiene Guard system by King and Tester (1997) consisted of a series of interactive badges for restaurant and hospital employees combined with an interactive system installed in the restrooms. Employees were encouraged to wash their hands by the public display of their failure to do so on their badges, making the behavior salient to others.

A more recent focus of persuasive technology researchers and practitioners has been on the development of systems to increase environmentally friendly behaviors. The introduction of Smart (energy) Meters—or the promise of such an introduction—has generated several research projects examining the effects of different types of technology initiated feedback on user’s energy consumption (see, e.g. Midden et al., 1983; Svane, 2007; Midden et al., 2008; Ham et al., 2009; Bang et al., 2006). Besides providing feedback in the home, several research projects have focussed on influencing energy behavior of office workers (Lockton et al., 2008). Finally, applications that monitor and display driving behavior have been successful at reducing the fuel consumption of drivers (Meschtscherjakov et al., 2009). Given the emergent shortage of oil and other natural resources to satisfy humanities energy needs, more persuasive applications in this field are likely to be developed in the coming years.

Persuasive technologies used for marketing purposes have marked both the initial growth of the field and are still the most salient example. While starting off by using online implementations of themes that had worked for decades in the consumer world, persuasive technologies for marketing purposes now extend beyond traditional marketing: applications track consumers and their preferences over multiple stores and across media to efficiently target their appeals (see, e.g. Zanker et al., 2009; Amit and Zott, 2001). With this tracking and tailoring—adapting to individual preferences—persuasive marketing applications seem to again be a step ahead of their counterparts in different application domains.

Recently health-care professionals and researchers, most noticeably in the domain of nutrition education, are following the marketers in providing 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 showed the increased effectiveness of these types of computer-tailored interventions over more traditional, “one size fits all” health education efforts (Brug et al., 1998, 1999; Brug, 1999). Noar et al. (2007) several years later conducted an elaborate 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.

2.1.2 Persuasion in Ambient Intelligence

Although persuasive applications initially originated on the web and were mainly implemented on standard computing systems, currently persuasive applications excel in using additional sensing equipment to create a more persuasive experience. More and more do persuasive technologies combine a number of computing devices for measurement, signaling, and user interaction. With this shift, persuasive systems designers frequently find themselves at the forefront of the ambient intelligent vision. Ambient Intelligence (AmI) (Aarts, 2010) extends the Ubicomp scenario brought forward by Weiser (1991). Similar to ubiquitous computing AmI is concerned with the integration of computing devices in the environment. Different however is the focus on the user instead of on the technology (Koch and Schlichter, 2001). Ubiquitous computing researchers were mainly concerned with the enabling technologies (Aarts and Ruyter, 2009), but AmI researchers work on applications that suit user needs and that combine distributed computing systems and a high level of system intelligence.

The technology and vision created by AmI researchers and practitioners lends itself well to develop persuasive applications. Persuasive technologies benefit from unobtrusive measurements of their effects, multiple media and devices to reach out to their users, and a high level of system intelligence to provide effective persuasion. Services like Philips DirectLife provide a clear example of *persuasion in ambient intelligence*: A measurement device is coupled with interactive (human) coaching via the web and the smartphone. Multiple devices and a high level of intelligence—be it partly provided by humans—are combined

to serve the goal of the user to lead a more active lifestyle.

In three of the most prominent application areas of persuasive technologies, health, safety, and energy consumption, the AmI vision is almost a requirement to provide adequate feedback about the effectiveness of the persuasive application. For example, to determine the effectiveness of a persuasive application aimed to reduce energy consumption multiple measurement devices to accurately assess actual energy consumption need to be in place. Furthermore, feedback needs to be provided to users at different times, via different channels, and relevant to the current usage situation.

2.1.3 Persuasive technology design

Several scholars have put effort into the development of frameworks to support the design and or evaluation of persuasive systems. Following the work of Fogg (2002) researchers and practitioners have tried to capture the complexities that underly the design of successful persuasive technologies. The most noticeable of such design frameworks are the *8-step design process* by Fogg (2009b) and the *Persuasive Systems Design (PSD)* model brought forward by Oinas-Kukkonen and Harjumaa (2008a).

The 8-step design framework by Fogg (2009b) starts with (1) a choice of a target behavior. Next, designers are encouraged to (2) determine the target audience and (3) find out why the target behavior is currently not performed. Designers should then (4) find an appropriate channel to reach the target behavior and (5) find prior examples of persuasive applications that are relevant for the current problem. Imitation of successful others (6), quick testing and iteration (7), and expansion on success (8) then finalize the steps to build successful persuasive technologies according to Fogg (2009b).

The PSD model presented by Oinas-Kukkonen and Harjumaa (2009) also aims at aiding designers in their efforts to create persuasive applications. The PSD model separates three steps in the design of persuasive systems. First, the designer should understand the key issues behind persuasive systems. Second, the designer of a persuasive system should analyze the persuasion context by identifying the intent, the event, and the strategy that is used. Finally, the designer of a persuasive system should consider the system qualities further elaborated upon by identifying primary task support, dialogue support, credibility support, and social support.

The PSD model which was initially published in combination with a

clear example for each of the three phases (Oinas-Kukkonen and Harjumaa, 2008b, 2009; Törning et al., 2009) provides designers with guidelines for the design of persuasive systems. Throughout, the framework uses knowledge derived from several social science disciplines—in the form of persuasive principles—to recommend design choices. Oinas-Kukkonen and Harjumaa (2008a) take this even further by evaluating persuasive applications by their implementation of these principles, implying that successful persuasive applications implement as many principles as possible.

Both Oinas-Kukkonen and Harjumaa (2009) and Fogg (2009b) thus emphasize the importance of using persuasive tactics—or social influence strategies—to increase the effectiveness of persuasive systems. Fogg (2002, 2009b) identifies 42 of such tactics, and Oinas-Kukkonen and Harjumaa (2009) identifies 28 design principles, most of which can be found in the social science literature typically described as *social influence strategies*.

2.2 Attitude and Behavioral Change

The study of attitude and behavioral change is too elaborate to properly review in this thesis. Based on the work presented in (Kaptein et al., 2009b) this section gives a concise overview of the relevant findings of these associated fields for the design of persuasive technologies. Kaptein et al. (2009b) structure the attitude and behavioral change literature relevant for the design of persuasive technologies along two axis. The first axis is typified by the *attitude-behavior continuum* and defines the target of the persuasive intervention: Does the intervention try to (a) change attitude, (b) elicit a new behavior, or (c) sustain an already performed behavior? The second axis distinguishes between *the number of sources* that are used for a persuasive attempt and ranges from strategies that rely on a single source (1), to those relying on the effects of multiple sources (2). Figure 2.1 gives a graphical overview of these axis and provides a structure for the review presented in this section.

Attitudes (a)

Attitudes are defined in multiple ways by different scholars, but at the core is the notion of evaluation: *Attitudes are viewed as a summary of evaluations of objects (thoughts, ideas, products, behaviors, etc.) along a dimension ranging from positive to negative* (Petty et al., 1994, 1997). The strength of an object-evaluation association—e.g. the ease

Figure 2.1: Two axis model as presented in Kaptein et al. (2009b).

with which an evaluation of an object can be retrieved from memory—influences the final evaluation of the object (Fazio, 1993; Bargh et al., 1996). Research has for example shown that a failure to recall a specified number of positive evaluations relating to an object decreases the final evaluation (Fazio, 1995). Next to the accessibility of an attitude, the object evaluation is influenced by the ambivalence of the attitude. The ambivalence of an attitude is the extent to which the attitude is based on consistent information. This literature views the final positive or negative evaluation of an object as a result of several negative and positive evaluations of aspects of the object (Cacioppo and Berntson, 1994; Breckler, 1994). Interesting from this perspective is that even though a final evaluation of an object might be strongly positive, the ease with which this evaluation can be changed depends not only on the

strength—generally it is assumed that strong attitudes are less easily changed (Tesser, 1993)—but also on the ambivalence of the base of the attitude: The more pro’s and con’s the attitude was based on, the more compliance with a counter attitudinal requests increases (Leippe and Eisenstadt, 1994).

Research has shown that there is a distinction between attitudes having an affective base, and those having a cognitive base, and that this origin of attitudes influences how attitudes can be changed. Practically this distinction shows that attitudes towards some topics are more easily changed using affective—emotional—arguments, while some benefit from cognitive—fact—based arguments (Petty et al., 1994). This distinction could be of importance for the design of ambient persuasive systems since an attitude with an emotional basis is likely not easily changed by cognitive arguments and vice versa.

Initial Behavior (b)

The MOA, Motivation, Opportunity, and Ability, model was introduced by MacInnis and Jaworski (1989), and elaborated upon by Rothschild (1999), and is primarily used to explain why a specific behavior occurs amongst competing behaviors. The basic principle of the MOA model is that the likelihood for a single behavior to be performed depends on the motivation to perform the behavior, the opportunity to perform the behavior, and the ability to perform the behavior (de Heer and Poiesz, 1998). Motivation is the predisposition of the person performing the behavior towards the behavior. Motivation is often split between *intrinsic* and *extrinsic* motivation (Ryan and Deci, 2000). Intrinsic motivation refers to motivation to enact a behavior for its inherent satisfaction—its alignment with ones personal values or attitudes—and not for some separable external consequence (West et al., 1975; Deci, 1975). External motivation refers to motivation which is controlled by externalities that are not part of the activity or behavior they are influencing (Deci et al., 2001). The most common examples of external motivation are reward and punishment. Internal motivation seems to increase the likelihood of the behavior being performed, and perhaps more importantly seems to lead to sustained behavior (Deci, 1975)

Once a person is motivated, be it internally or externally, to perform a behavior the likelihood of enactment is dependent on the opportunity—the extent to which the external environment enables the behavior—and a persons ability—the extent to which a person possesses the skills and knowledge necessary to enact. For example, even though

a person is very willing to throw a brick through a window, first one needs to find the brick (opportunity) and secondly, one needs to be able to aim, control the arm muscles, and strike a hit (ability). In a brick-less situation this behavior becomes unlikely even though motivation and ability might be high (Robben and Poiesz, 1993).

Sustain behavior: Conditioning (c)

When a desired target behavior is performed it is often important to sustain that behavior. The classical psychological approach to sustaining behavior is that of conditioning. Conditioning can be separated into the field of *classical conditioning* (Gormezano et al., 1987)—relating two previously unrelated stimuli—and *operant conditioning* (For an overview see Skinner, 1976)—enforcing behavior by the use of reward and punishment. For persuasive purposes the field of operant conditioning, which enables sustaining of voluntary performed behaviors, is most relevant.

Once favorable attitudes towards a behavior have been formed, and the behavior is performed for a first time, several methods of stimulating or inhibiting the behavior can be of use. Traditionally one separates reinforcements, consequences of the behavior that make the behavior more likely, and punishments, consequences that render the behavior less likely. Since both of consequences can either be removed or added, there are four possible reinforcement schemes (Ferster, 1957):

1. *Positive reinforcers*: A behavior is followed by a positive stimulus (reward).
2. *Negative reinforcers*: A behavior is followed by the removal of a negative stimulus (punishment).
3. *Positive punishment*: A behavior is followed by a negative stimulus.
4. *Negative punishment*: A behavior is followed by the removal of a positive stimulus.

All four of these schemes can be used by ambient persuasive systems to reinforce behavior. Since the notion of operant conditioning is relatively old numerous effects have been described and researched. A lot of work has been done on extinction (e.g. Kirkpatrick et al., 1964; McNaughton, 1984), the effects of neither reinforcing nor punishing a behavior, after a behavior has been conditioned. Typically one observes an *extinction boom*: a sudden frequent outburst of the behavior. Next to extinction effects a number of different punishment and reinforcement

schedules have been researched. The literature also identifies different schedules of reinforcement and their behavioral outcomes (Ferster, 1957). One can distinguish between *time* reinforcers or *ratio* reinforcers; a reward or punishment after a number of times the behavior has been performed (ratio) or a specific time after the first occurrence of the behavior (time). Both of these can be done specifically or variable, in which fixed implies that the behavior is reinforced only on regular intervals.

Single Source processes (1)

Besides the *attitude—behavior—sustaining behavior* continuum the social science literature has focussed on important source effects—properties of the persuader—that influence compliance. The next two sections briefly summarize the effects that have been identified for a single source, and for multiple sources.

Several single source characteristics influence the effectiveness of a persuasive request. Heavily researched of these (single) source effects are perceived friendliness, perceived similarity, mimicry, and reciprocity. The conclusions are summed below:

- A greater perceived legitimate *authority* increases compliance to a persuasive request (Milgram, 1974; Miller et al., 1995; Slater et al., 2006).
- Sources that are considered *friendly*, or are liked by the receiver, increase compliance to their requests (Cialdini, 2004; Kaptein et al., 2010b).
- A greater *similarity* of the source of the request and the receiver increases compliance. (Burger et al., 2004; Festinger, 1954; Murray et al., 1987; Schultz, 1999).
- *Mimicry*—similarities in behavior—by a source relative to a receiver increases compliance (Chartrand and Bargh, 1999; Chartrand et al., 2005; Kaptein et al., 2010b).
- People are inclined to return—*reciprocate*—a favor (Cialdini, 2004; James and Bolstein, 1990; Komorita et al., 1991).
- *Repetition*—multiple exposures to the same source over time—increases compliance (Latané et al., 1995; Nowak et al., 1990).

Multiple sources (2)

Next to the influence of the perception of the receiver of a single source making a persuasive request, a number of multiple source effects exist.

These effects are distinct from single source effects. Research in the area of multiple sources acknowledges that the compliance of an individual to a persuasive request is not merely determined by the current interaction between the source and the receiver, but also by previous interactions with others, interactions with the same source, and the number of repetitions. Thus, ones attitude or behaviors are dependent upon the social environment in which one lives. The main findings are:

- *Social proof* is a powerful persuader. In uncertainty people look at the behavior and attitudes of others to determine their own (Cialdini, 2001; Latané, 1970; Latané and Nida, 1981).
- The *number of people* in reference group (for example in cases of social proof) or the absolute number of sources making a request influences compliance. More sources with a similar cue lead to higher compliance (Latané, 1996; Latané and Bourgeois, 1996).
- *Immediacy* leads to compliance: people comply with those that are close or more intimate to them (Latané, 1996; Nowak et al., 1990).
- People generally seek *consensus*. People generally tend to minimize stress and conflict arising from competing opinions or behaviors (Asch, 1955).

It is important to note that many scholars identify a positive attitude towards a behavior key for the behavior to be performed (e.g. MacInnis and Jaworski, 1989; Fogg, 2009a). As such, the formation of attitudes, and the processing of appeals to influence attitudes, is of importance for the design of persuasive systems.

2.2.1 Dual Processing Models

Investigators in psychology often explain and predict how implementations of influence strategies affect consumer attitudes using dual-process models. According to the Elaboration Likelihood Model (ELM) (Cacioppo et al., 1986), persuasive messages can affect attitudes through both central and peripheral routes. The central route is characterized by elaboration on and consideration of the merits of presented arguments. On the other hand, the peripheral route is characterized by responses to cues associated with but peripheral to the central arguments of the advocacy. The latter occurs through the application of simple, cognitively “cheap”, but fallible rules.

Frequently, the use of these cognitively “cheap” rules—such as “a product is almost out of stock”—leads to a fast and relatively accurate

appraisal of the merits of the product: If the product is almost out of stock, a large number of prior customers may have bought the product based on product merits and opportunities to buy in the future may be rare or high cost (Verhallen and Robben, 1994). Thus, without engaging in full and cognitively costly processing, a consumer can make a choice based on an accurate peripheral cue (Goldstein and Gigerenzer, 2002). This suggests why implementations of many *social influence strategies*—or compliance-gaining tactics in the marketing literature (Payan and McFarland, 2005)—are effective: via the peripheral route they provide a cognitively cheap shortcut to, on average, effective decision making about product merits. The central route and the peripheral route label the endpoints of a continuum, as any message will have its effects both via the central route and the peripheral route, and can affect the use of either route (Petty and Cacioppo, 1986; Petty and Wegener, 1999).

2.2.2 Social Influence Strategies

The array of persuasion and influence tactics that can be used to change attitudes and behaviors in consumers can be overwhelming. Both researchers and practitioners have made extensive use of the categorization of persuasive messages as implementing more general influence strategies. Theorists have varied in how they individuate influence 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 taxonomies, and others have listed over 100 (Rhoads, 2007). These different counts result from differing levels of exhaustiveness, exclusivity, emphasis, and granularity (Kellermann and Cole, 1994). Influence strategies are however a useful *level of analysis* that helps to group and distinguish specific influence tactics or implementations of these strategies (Kellermann and Cole, 1994; O’Keefe, 1994). For designers of persuasive messages, classifying “almost out of stock” as an implementation of the *scarcity* strategy (Cialdini, 2001) provides research-based expectations about that message’s effects—across products and individuals. This thesis focusses on the six influence strategies described by Cialdini (2001). Multiple implementations of each of these strategies have been shown to be effective in previous laboratory and field experiments.

(1) Authority

When an authority figure tells people something to do, they typically do it (Milgram, 1974; Blass, 1991). Consumers are therefore frequently

faced with authority endorsements of products such as “expert reviews”. 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); thus, obedience has an evolutionary advantage. 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. Thus, while frequently used to influence consumers, theory allows for authority arguments to have either a positive or a negative effect on an individual’s attitudes and behaviors.

(2) Consensus

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). This effect of consensus is used to influence people’s decision making by stating that products are bestsellers or by displaying other consumers’ positive evaluations of a product. Multiple processes have been posited to explain the effectiveness of the consensus strategy: Asch (1956) ascribes the observed effects to mere conformity, while others postulate that implementations of the consensus strategy constitute informational influence, by serving as “social proof” (Hardin and Higgins, 1996; Cialdini, 2001). But theorists have also identified processes that would lead to adverse effects of consensus. For example, Conway and Schaller (2005) use attribution theory to explain how individuals might come to question the value of a consensus opinion and may be compelled to reject it.

(3) Consistency and Commitment

The Consistency and Commitment strategy 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 they will be more inclined to do so even if they did not agree on writing it down in the first place (Deutsch and Gerard, 1955). People will try to be

consisted with their writing. The consistency principle as such also explains the power of commitment: people will act as they told or wrote they will.

(4) Scarcity

Assumed scarcity increases perceived value of products and opportunities (Cialdini, 2001), so advertisers and salespeople often use phrases like “limited release”, and “while supplies last” (Lynn, 1991). There is overwhelming evidence that identifying a product or service as scarce will favorably affect consumer attitudes 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) and states that humans desire scarce products more because the possession of such products produces feelings of personal distinctiveness or uniqueness. Other authors have expanded this idea by positing a psychological trait, the Need for Uniqueness (Fromkin, 1970; Snyder and Fromkin, 1980), which drives people to seek out products that make them feel unique. Other theoretical explanations for the effectiveness of the scarcity include reactance theory (Brehm, 1966; Clee and Wicklund, 1980), personal equity theory (Seta and Seta, 1982), and dissonance theory (Festinger, 1957). Not all of these theories predict a positive effect of scarcity in all situations; for example, reactance theory predicts scarcity effects only when the product under consideration represents an important right or freedom.

(5) 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 strategy is presented by studies that exploit increased liking due to interpersonal similarity. For example, people are more inclined to return a wallet to the lost and found when the name listed in the wallet is similar to their own than when it is dissimilar (Hornstein et al., 1968).

(6) Reciprocity

People are inclined—or actually, people go through a great deal of effort—to pay back a favor (Cialdini, 2004). This social influence strategy—when implemented properly—is exceptionally strong, and seems to work even when it is truly unbeneficial 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. The strategy of reciprocation which is also the foundation in the *tit-for-tat* strategy, the most favorable algorithm to win social dilemma games (Komorita et al., 1991), seems rational. However, it has been shown that people even reciprocate to favors they had never asked for (James and Bolstein, 1990).

2.2.3 Individual Differences

Despite the large body of work investigating social influence strategies and the theoretical models such as the ELM to explain their effectiveness researchers have had serious difficulties in replicating previous findings. For example, a thorough meta-analysis (Johnson and Eagly, 1989) of the research on the effects of argument strength on persuasion—as frequently used in ELM research to appeal to either peripheral or central processing—has found mixed results. This highlights the importance of protocols and well developed stimuli to replicate some of the field’s most compelling findings. Because of these, and other, difficulties in replication researchers have investigated properties of context, messages, and individuals to further understand persuasion processes. The following section reviews the most important advances in the studies of individual differences in responses to social influence strategies.

Need For Cognition

Much of the work on individual differences in persuasion has directly drawn on dual-process models—and the ELM in particular—to work out how new or established traits could moderate persuasion. Many of these studies have examined trait differences in motivations, such as need for cognition (*NfC*, Cacioppo et al., 1986), that effect differences in peripheral and central processing of persuasive messages. Thus, *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 difference in 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 product is actually a doctor (or an actor playing an actor) 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 strategies.

While *NfC* is the most widely used trait that operationalizes stable motivational heterogeneity in dual-process models, several relating traits have been identified and studied (Haugtvedt et al., 2008). Measures of individuals' need for closure (Webster and Kruglanski, 1994), need to evaluate (Jarvis and Petty, 1996), and need for affect (Maio and Esses, 2001) have all received attention in the persuasion literature. The approach of using the ELM to generate hypotheses about how new or established traits affect persuasion processes continues to yield insights into the many relationships between personality and persuasion. On the other hand, the ELM does not immediately suggest traits that would be associated with differences in responses to *distinct influence strategies*, but taxonomies of influence strategies are widely used by those designing influence attempts as detailed in the previous section. The flexibility of the ELM—whereby any cue can serve in many different roles—can also make it difficult to extract specific predictions about how a trait might affect persuasion in practice.

Preference for Consistency

Investigators have fruitfully drawn on the categorization of messages as implementing distinct influence strategies to identify and study personality constructs that are plausibly associated with the posited processes by which particular influence strategies function. For example, the commitment strategy, including a range of implementations such as in “foot-in-the-door”, functions through the application of motivations for consistency. A personality scale that measures these motivations—*preference for consistency*—predicts responses to the commitment strategy, such that for participants low on this trait these strategies are ineffective (Cialdini et al., 1995; Guadagno et al., 2001). This prior research has helped explain the difficulties investigators have had in replicating consistency results. Successful use of this approach in studying other influence strategies requires a theory about the psychological processes that make that strategy effective and how these might vary in the population. Such settled 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).

Measurement of Individual Differences

The prior research can be described as relying on *meta-judgmental measures* of personality traits. In the context of attitude strength, Bassili (1996) distinguished 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 traits 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. Many questionnaire-based measures of personality traits are meta-judgmental measures. On the other hand, for *operative measures* individuals' psychological processes are in use: they are operating. For example, an operative measure of *NfC* might be a summary of differences in an individuals' responses to strong and weak arguments about multiple topics.

One argument in favor of operative measures of individual differences in persuasion is that it has been difficult to construct valid meta-judgmental measures. While the meta-judgmental measure of preference for consistency is now widely used and accepted (Guadagno and Cialdini, 2010), it followed several failed attempts at creating a measure that would predict heterogeneity in responses to implementations of the commitment strategy (see Cialdini et al., 1995). There are also reasons for using operative measures that follow from issues in applied settings: It might not always be possible to administer questionnaires to users of a persuasive technology. Finally, meta-judgmental measures might be, due to their context independence and evaluative nature, less predictive of responses to a given persuasive system than operative measures obtained within a specific context.

Discussion on Individual Differences

While investigations of individual differences in responses to persuasion have proven fruitful in an experimental context, the theories often leave designers of persuasive systems uninformed about the ways in which individual differences can be addressed when designing persuasive systems. Even though some social influence strategies are theorized to be ineffective for a part of the user population it is often unclear how these users should be identified and what changes need to be made to a

system to facilitate these users. Within computer-tailored health education, an active contemporary personalization of persuasion attempt, researchers therefore often rely at least partly on the expertise of practitioners (counselors). These experts create a series of “if-then” statements that inform the presentation of different persuasive content to different individuals (Brug et al., 2003) instead of creating this content solely based on theories of attitude and behavioral change.

Given the numerous indications of individual differences in responses to social influence strategies it is also surprising to see that the magnitude of these differences has not been compared to the frequently reported upon average effects. The average effects—the effects over multiple people—defined the formulation of strategies and subsequently their way into the design frameworks for persuasive technologies. However, other than in marketing applications, persuasive technologies often focus on attitude or behavioral changes for *individual users* and as such the average effects are not the core interest: the possible range of effects of a influence strategy on a specific user will determine the success of the persuasive system. The current literature hardly informs designers about the relative importance of attending to individual differences versus implementing different social influence strategies identified in the literature.

2.3 Conclusion: Reflections

Fogg (2002) identified the core opportunities of persuasive systems: Persuasive technologies can be more efficient than their human counterparts by being more persistent, by being always on, and by *tailoring their persuasive attempts to responses by their users*. A review of the current practice however shows that while the first two objectives are met, the third one is frequently neglected. This is true even though social science investigations into the effects of social influence strategies indicate that there are individual differences that could be attended to. Differences which, once they are attended to, possibly increase the effectiveness of persuasive systems.

This neglect has two causes: First, the social science literature is generally not informative enough for designers of persuasive systems to decide on which individual differences to attend to, and thus which attributes of persuasive systems to tailor to individuals. While social influence strategies constitute a frequently used feature of persuasive technologies that have received considerable research attention, even there

the magnitude of the average effects compared to the individual effects is unclear. Thus, designers cannot readily estimate the importance of tailoring their use of social influence strategies to individual users. Second, persuasive technology designers are generally not very concerned about measuring the effects of their interventions. Many evaluations of persuasive systems that are deployed in the field concern usability evaluations rather than *effectiveness* evaluations. This point is eminent by the lack of focus on effect measurement in the design frameworks that are proposed to aid designers in their efforts to design persuasive applications (e.g., Oinas-Kukkonen and Harjumaa, 2009; Fogg, 2009b). Thus, often it is unclear what the effects of existing applications on user’s attitudes and behaviors actually are.

In the subsequent chapters the individual differences in responses to social influence strategies, the effects of the usage of multiple strategies for a single application, and possible means of measuring individual differences are further explored. These chapters focus on uncovering properties of social influence strategies, and people’s responses to these strategies, that inform designers in their quest to design effective persuasive technologies ¹.

¹This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2009b; Kaptein and Eckles, 2010).

3

Insight Generation I: *Heterogeneity in Responses to Social Influence Strategies*

3.1 Introduction

This first *insights generation* chapter studies the heterogeneity—individual differences—in responses to social influence strategies. This chapter uses a selection of the social influence strategies as the basis to explore the difference in magnitude between average treatment effects—the effects that are usually attended to in the literature on social influence strategies—and the individual level effects.

In three studies this chapter examines the size of the individual differences relative to the average effect of persuasive strategies, and the stability of these individual level estimates over time and contexts. The method brought forward here focusses on operative measures: Measures of the effect of different influence strategies at the moment the psychological processes are at play. All three studies present people with different implementations of a number of influence strategies, and measure their responses. Based on the responses an estimate of the effectiveness of each strategy can be obtained, and this estimate *is* the

operative measure of an individual's *persuadability* by a specific influence strategy.

3.1.1 Analysis Using Multi Level Models

All three studies presented in this chapter use multilevel hierarchical models with crossed-random effects to obtain the operative measures of the individual level effects of different influence strategies. This section describes how these models are build up, and why they provide a good method to estimate and study individual differences. The models extend common linear regression in the sense that some parameters (or batches of parameters)—called the “random effects”—are constrained by a prior distribution over their values.

To understand hierarchical models it is convenient to start with an explanation of the so-called “null” model (Snijders and Bosker, 1999), a model which includes only an overall intercept and varying intercepts for each individual. In this “null” model, the “scores” of an individual are modeled using the average score across individuals (present for each individual) and an individual level estimate. The latter represents the difference in score of a specific individual from the average score of others. The model is written as

$$y_{ij} = \bar{\mu} + \mu_i + \sigma_{err}^2 \quad (3.1)$$

where

$$\mu_i \sim \mathcal{N}(0, \sigma_\mu^2) \quad (3.2)$$

for $i = 1, \dots, N$ people with $j = 1, \dots, J$ observations per person. The score of an individual is estimated by the sum of the overall intercept $\bar{\mu}$ and the per-participant intercept μ_i . The latter is constrained to have a Gaussian distribution with mean 0 and variance σ_μ^2 . Thus, $\bar{\mu}$, and σ_μ^2 are the parameters that are estimated.

Let us examine how this “null” model represents individual differences: Suppose y_{1j} is participant 1's evaluation of a product j . The model described in equation 3.1 predicts this individual's evaluations based on a weighted average of the evaluations of other participants, as well as the (prior) evaluations of the same participant. If an individual—participant 1—evaluates products more positively than the average rater, her individual μ_1 will be positive. If multiple participants score consistently much higher or lower than the grand average, $\bar{\mu}$, then σ_μ^2 will be large. Thus, σ_μ^2 is a direct measure of the size of the heterogeneity in responses between people. This principle can be ex-

tended to incorporate effects of multiple influence strategies at the level of individuals. In this case σ_μ^2 is replaced by a matrix Σ that holds the variances and co-variances of the different individual level effects. All three of the studies presented in this chapter explicitly investigate Σ to draw conclusions about the individual differences in responses to social influence strategies. These models can be seen as a modern version of analysis of variance (ANOVA) for factors with many levels (Gelman, 2005).

3.2 Study 1: Identifying an upper bound in heterogeneity

Study 1 was designed to identify an *upper-bound* on heterogeneity in the effects of influence strategies in a product-evaluation context. Study 1 estimates an "upper-bound" since this study estimates individual level effects of social influence strategies at one specific point in time. Thus, the heterogeneity that is shown by inspection of Σ combines both individual level *traits* that cause differences in responses to social influence strategies, as well as temporary *states*.

The main in this first study focus was to see whether the heterogeneity in responses to social influence strategies was large in comparison to the average effect. Some heterogeneity is expected but two different scenario's can be pitted against each other. First, influence strategies could have a large positive average effect and individual level effects are close to that average. Thus, the effects of an influence strategy for each individual are positive compared to a baseline in which the strategy is not presented. Second, it could be the case that even though influence strategies have a large positive average effect, the heterogeneity is such that for a group of people the effects of the strategies are zero or even negative. In the latter case the heterogeneity warrants further attention.

3.2.1 Method Study 1

Procedure

Participants were invited by email. After signing up for the study, participants received a link to the study Web site. The study was called "Evaluating books" and asked prospective participants, "what would you like to read?" Participants were then instructed that they would be presented with 14 science fiction novels, and that they would

be asked to evaluate each of these books.¹ Participants were shown a list of reasons why the books were selected for presentation. These messages implemented three influence strategies as well as one control message (“A random selection”):

1. *A random selection*: These books are randomly selected from the book database. (*Control*)
2. *The majority view*: These books are selected because they have been sold widely or are appreciated by the majority of readers. (*Consensus*)
3. *Rare or special items*: These books are selected because they are rare or limited in their run. (*Scarcity*)
4. *The experts recommendation*: These books are selected from the book database because they are recommended by experts in the field. (*Authority*)

These messages were selected out of a larger set of messages by four persuasion researchers. The messages were selected based on their external validity (these messages are actually used on e-commerce websites) and the expert judgment that they implemented the social influence strategy. This latter criteria is somewhat arbitrary but the quality of the messages can be assessed after collecting the data: If the average effect of the messages is in the direction of, and in the order of magnitude of, previous findings, one can be confident that the messages indeed implement the social influence strategies as intended.

After these introduction messages participants then presented with fourteen Web pages in sequence. Each page contained an image of a book, a short textual description, and four questionnaire items. The books were selected from Amazon.com and all fell in the same price range (\$10 to \$12), all had approximately an average rating of three out of five stars, and none were bestsellers. Descriptions of the books were adapted to be of similar length.

Located just above the description and cover image that was presented for each book was a message that either implemented one of the influence strategies or was a control message. Two implementations of each strategy were delivered in sequence. For example, a participant would first be presented with two books that were ostensibly selected because they represented “the majority view.” The first book would be accompanied by the first implementation of the strategy (“Over a

¹Participants rated two books for each of 14 messages, only eight of which are analyzed here.

Table 3.1: Influence strategies and their respective implementations as used in Study 1.

Strategy	Implementation
<i>Control</i>	1. A random selection. 2. This book is randomly selected from our product offerings.
<i>Scarcity</i>	3. This is a limited edition signed by the author! 4. There are only 50 copies of this book left nationwide!
<i>Consensus</i>	5. Over a million copies sold! 6. Voted best fictional book by college students! (Princeton Review)
<i>Authority</i>	7. "I would recommend this book to anyone." - Stephen King. 8. "Every household should have a copy of this" - American Authors Book Review Committee.

million copies sold!") and the second accompanied by the second implementation of the same strategy ("Voted best fictional book by college students"). Each participant was exposed to both implementations of each of the strategies. Table 3.1 shows the implementations used in this study. Participants were randomly assigned to different orders of presentation of the influence strategies. Book order was kept constant to control for differences in the appeal of the books. Both order effects as well as possible interactions between book and implementation effects were controlled for at a between-subjects level.

Measures

Each of the books was evaluated by participants on four items on a ten-point scale:

1. How likely would you be to recommend this book to your friends? (*Very unlikely* - *Very likely*)
2. How much would you enjoy reading this book? (*Would not enjoy at all* - *Would enjoy very much*)
3. How would you judge the quality of this book? (*Very poor quality* - *Very good quality*)
4. How likely would you be to buy this book if you were going to buy a novel? (*Very unlikely* - *Very likely*)

After participants evaluated the books, participants completed a questionnaire addressing several meta-judgmental measures. This started with the standard 18-item measure of Need for Cognition (*NfC*) (Cacioppo and Petty, 1982) (Cronbach’s $\alpha = 0.893$). Next, participants reported their susceptibility to each of the six influence strategies identified by Cialdini (2001) by responding to 26 items specifically created for this purpose (See Chapter 5). Finally, participants were asked about their age, gender, and academic major.

Participants

Participants were 179 english speaking university and community college students enrolled in introductory research methods courses, who participated for partial course credit. Of the participants 111 (62.01%) were females. The mean age of participants was 24.3 ($SD = 7.99$).

3.2.2 Results Study 1

This study resulted in a dataset describing the evaluation of $B = 14$ books by $N = 179$ subjects. For each participant, the evaluated books are accompanied by one of $K = 8$ implementations of influence strategies (two implementations for each of $S = 4$ strategies). Finally, each book was evaluated using $Q = 4$ different items, each on a ten-point scale. Data analysis involved a series of comparisons of mixed-effects models (Baayen et al., 2008) and subsequent examination of the estimated parameters of the selected model. These models are: (*A*) a model with no heterogeneity in the effects of influence strategies, (*B*) a model in which only the effect of using *any* strategy varies from person to person, and (*C*) a model in which the effects of each strategy vary from person to person. Comparison of Models *A* and *B* corresponds to testing the null hypothesis that there are no individual differences in overall “persuadability” by the consensus, authority, and scarcity influence strategies. Comparison of Models *B* and *C* tests the null hypothesis that individual differences in the effects of influence strategies are exhausted by individual differences in overall “persuadability”.

Each of these three models can be written

$$y_{jbq} \sim \mathcal{N}(X_{jb}\beta_j + \alpha_b + \eta_q, \sigma_{err}^2) \quad (3.3)$$

with $\beta_j \sim \mathcal{N}(\bar{\beta}, \Sigma_\beta)$ for $j = 1, \dots, J = 179$ subjects, $\alpha_b \sim \mathcal{N}(0, \Sigma_\alpha)$ for $b = 1, \dots, B = 8$, books and $\eta_q \sim \mathcal{N}(0, \sigma_\eta^2)$ for $q = 1, \dots, Q = 4$ questions.

The design matrix X_{jb} is a 1432×4 matrix consisting of a column of ones and indicators for each of the three strategies. Thus, β is a

179×4 matrix of intercepts and coefficients for each strategy for each individual. Finally, $\bar{\beta}$ and Σ_β are, respectively, the vector of the means and the covariance matrix of the individual $\times$ strategy coefficients.

In this formulation, the three models differ only in their constraints on Σ_β . Model *A* only allows for between-person variation in an intercept, so it has

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (3.4)$$

where σ_I^2 is the between-person variance of the intercept. Model *B* allows for the effect of all three of the strategies to vary *together* by person:

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & \sigma_{I,P}^2 & \sigma_{I,P}^2 & \sigma_{I,P}^2 \\ \sigma_{I,P}^2 & \sigma_P^2 & \sigma_P^2 & \sigma_P^2 \\ \sigma_{I,P}^2 & \sigma_P^2 & \sigma_P^2 & \sigma_P^2 \\ \sigma_{I,P}^2 & \sigma_P^2 & \sigma_P^2 & \sigma_P^2 \end{bmatrix}. \quad (3.5)$$

Here σ_P^2 is the between-person variance of the overall effect of all three strategies—a measure of general “persuadability”—and $\sigma_{I,P}^2$ is the covariance of the intercept and this strategy effect. Finally, in Model *C* the entries of this covariance matrix are unconstrained,

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & \sigma_{I,c}^2 & \sigma_{I,a}^2 & \sigma_{I,s}^2 \\ \sigma_{I,c}^2 & \sigma_c^2 & \sigma_{c,a}^2 & \sigma_{c,s}^2 \\ \sigma_{I,a}^2 & \sigma_{c,a}^2 & \sigma_a^2 & \sigma_{a,s}^2 \\ \sigma_{I,s}^2 & \sigma_{c,s}^2 & \sigma_{a,s}^2 & \sigma_s^2 \end{bmatrix}, \quad (3.6)$$

where σ_c^2 , σ_a^2 , and σ_s^2 are the between-person variances of the effects of the consensus, authority, and scarcity strategies, respectively.

Table 3.2 presents the comparisons of these three models. Model *C* is preferred. This directly supports the hypothesis that there are individual differences in the effects of these three influence strategies. Furthermore, this heterogeneity is not exhausted by an heterogeneity only in an overall effect of the three strategies, as compared to *no*-strategy. Having rejected the two null hypotheses corresponding Models *A* and *B*, the estimated parameters of Model *C* are now further examined.

Table 3.2: Table comparing the preferred model (*C*), in which each of the strategy effects vary by person, with the model without varying effects of strategies by person (*A*), and the model which includes heterogeneity only for the control vs. strategy contrast (*B*). Presented are the degrees of freedom (*df*), Bayesian Information Criteria, *BIC*, the log-likelihood, *logLik* of each of the models, and the χ^2 and *p*-value for the model comparisons.

	df	BIC	logLik	χ^2	<i>p</i>
Model <i>A</i> :	8	24373.08	-12151.93		
Model <i>B</i> :	10	24096.83	-12005.16	293.55	<.0001
Model <i>C</i> :	17	23390.42	-11621.67	766.97	<.0001

Magnitude of Heterogeneity

The magnitude of heterogeneity in effects of influence strategies is characterized by examining $\hat{\Sigma}_\beta$ for the preferred model, Model *C*. In particular, this analysis compares the size of the diagonal entries of $\hat{\Sigma}_\beta$ to the estimated average effects $\hat{\beta}$ and the variances of other sources of variation in this study. Table 3.4 presents the estimated standard deviations of all of the random effects in Model *C* along with the correlations among the strategy $\times$ person effects; it is a summary of $\hat{\Sigma}_\beta$.

Table 3.3: Estimates of fixed effects in the preferred model. Using the control messages as the reference, each of the point estimates of the average effect of the influence strategies on book evaluations is positive. Empirical *p*-values computed with draws from the posterior using MCMC (see Baayen et al., 2008).

	Estimate	<i>S.E</i>	<i>t</i>	<i>p</i>
Intercept	4.25	0.33	12.91	0.0002
Authority	0.37	0.15	2.51	0.0064
Consensus	0.44	0.14	3.11	0.0020
Scarcity	0.06	0.14	0.43	0.6484

Consistent with the choice of established influence strategies, failing to include the strategy factor (i.e., adding the constraint that $\bar{\beta} = [\mu, 0, 0, 0]$, where μ is an overall intercept) results in significantly worse fit than Model *C*, $\chi^2(3) = 13.92$, $p = 0.003$. Table 3.3 presents the estimates of the average effects in Model *C*. For the implementations used in this study, the authority and consensus strategies had the

largest estimated effect. The average effects of the authority ($p < 0.01$) and consensus strategies ($p < 0.01$) are both significantly positive.

Figure 3.1: Comparison of heterogeneity in the effects of influence strategies with the average effects of those strategies. The solid black vertical lines are the estimated average effects of each strategy, as compared with the control message. The black curves are the estimated normal distribution of strategy effects for the population, while the gray curves are the density of the estimates of the strategy effects for this sample. Estimates are from Model *C*.

Figure 3.1 is a graphical comparison of the estimated average effect of each strategy and the estimated distributions of strategy $\times$ person effects. The solid vertical lines indicate the average effect of each strategy compared to the control message (vertical dotted line). The solid black density is the estimated Gaussian distribution of the strategy $\times$ person effects in the population.

As an illustration of the magnitude of individual differences in influence strategy effects, one can consider how common it is for the effect of the consensus strategy to be *negative*, despite its average effect being significantly positive. This analysis shows that for 41.3% (95% *CI* [35.8, 45.3]²) of the participants the estimated effect of consensus is negative.

²95% confidence intervals in brackets were computed using the Bayesian pigeon-hole bootstrap with $R = 1000$ (Owen, 2007).

Qualitatively, one can compare the different standard deviations presented in Table 3.4: The estimated standard deviation of participants' responses to books not accompanied by influence strategies (the intercept varying by person, $\hat{\sigma}_I^2$) is of similar magnitude as the standard deviation of the residuals $\hat{\sigma}_{err}^2$. The same is true for the estimated standard deviation of participants' responses to books accompanied by each of the influence strategies. Thus, in this study the effects of influence strategies differ as much between people as does their evaluation of different books.

Table 3.4: Summary of random effects in Model *C*. The square root of the diagonal elements of $\hat{\Sigma}_\beta$ are the first four rows of the standard deviation column, while the remaining entries are presented as correlations to the right.

Grouping		SD	Corr.		
Person	Intercept	1.54			
	Authority	1.76	-0.35		
	Consensus	1.69	-0.39	0.51	
	Scarcity	1.65	-0.40	0.40	0.59
Book	Intercept	0.58			
Question	Intercept	0.53			
Residual		1.62			

Structure of the Heterogeneity

Comparison of Models *A*, *B*, and *C* supported the hypothesis that the structure of heterogeneity in influence strategy effects is more complex than single dimension of overall “persuadability”. In particular, the covariance matrix for the strategy effects Σ_β in Model *C* allows for more complex relationships among the effects of each strategy. While the effects of the three strategies are moderately correlated, they also have substantial unique variation, as manifest in Figure 3.2. This figure makes clear that for large numbers of individuals one specific social influence strategy leads to higher product evaluations than *no*-strategy, while for other strategies the sign of the effect is reversed.

Demographics and Meta-Judgmental measures

To examine whether demographic and personality measures that are typically used in marketing practice would be able to sufficiently capture the individual differences identified a series of model comparisons

Figure 3.2: Estimated influence strategy effects for each participant in Model C , as compared to the control messages. Note that for some individuals, the estimated effect of one strategy is negative, while the estimated effect of another strategy is positive.

Table 3.5: Table examining the use of demographic and meta-judgmental measures. While both aid significantly in explaining the variance in responses to social influence strategies, explicitly modeling heterogeneity is still preferable according to the log-likelihood comparisons.

	df	BIC	logLik	χ^2	p
Model A :	8	24373.08	-12151.93		
Model A_d :	20	24419.88	-12123.42	57.02	<.0001
Model A_m :	44	23524.52	-11572.80	1101.25	<.0001
Model C_m :	53	22733.58	-11138.58	868.45	<.0001

was performed including these measures. Table 3.5 presents these four model comparisons. First, Model A —the baseline model without varying strategy effects—is compared to a model that includes age (split in three groups) and gender interacting with the social influence strategy, Model A_d . These measures significantly increase model fit and thus aid in explaining the observed responses. Next, both the obtained measures of NfC as well as the Big Five are added to the model, Model A_m . Despite the large increase in degrees of freedom this model is preferable (even according to conservative BIC). This shows that meta-judgmental measures can indeed be beneficial to attend to individual differences in persuasion processes. However, the final comparison shows that allowing for heterogeneity in responses to influence strategies—as done in Model C —still significantly improves model fit, Model C_m . Thus, the demographics and meta-judgmental measures included in this analysis seem *insufficient* to capture all individual level variation.

3.2.3 Discussion Study 1

Study 1 enables comparison of the heterogeneity in responses to social influence strategies to the often attended to average effect(s). The heterogeneity is very large compared to the average effects and this warrants further attention. This heterogeneity is not exhausted by differences in *overall* responses or persuadability to all of the influence strategies considered together; rather, according to model selection criteria such as the BIC, one *should* use a model that includes varying effects for each of the influence strategies. Furthermore, models including this heterogeneity were preferable even when including demographic or meta-judgmental measures as independent predictors for responses to each strategy.

3.3 Study 2: Examining stability of heterogeneity over time

While the analysis of Study 1 produced unbiased estimates of the effect of each influence strategy for each participant, each of these estimates had relatively high uncertainty and it is expected that the estimated heterogeneity is upwardly biased. Transient intra-individual variation, such as that caused by differences in mood, fatigue, and situation, are not separately identified, as participants only rated books during a single session. This problem is common but ignored (cf. Watson and Strayer, 2010) in operative measures of traits. In particular, moods—emotional states lasting from several hours to several days—can affect elaboration (Bless et al., 1990), thereby moderating the effects of influence strategies on attitudes. To address this issue, in Study 2 participants encountered implementations of the same influence strategies over three sessions.

3.3.1 Method Study 2

Study 2 was also conducted in a product-evaluation context in which participants rated books. Rather than being a study ostensibly about book preferences, Study 2 was ostensibly a user experience evaluation of a new online bookstore; during each of three sessions approximately a week apart, participants were guided through a series of tasks in which they could freely browse and search this bookstore (see Figure 3.3). In addition to rating a number of books, participants were asked to perform other tasks fitting with the cover story. Like the Amazon.com homepage—a popular book selling e-commerce website—books were organized under titles that implemented an influence strategy. Individual book description pages also featured implementations of the social influence strategies used in this study.

Procedure

Once signed up for the study, participants received an email explaining that they would receive a link to the study website once every week for three subsequent weeks. Participants received the emails on Mondays and were asked to participate in that week's session by that Wednesday. Participants who had not participated by Wednesday morning received a reminder email.

Participants began each session by clicking on a link to the study Web site in the Monday or Wednesday emails. After an instruction page, participants were taken to an online bookstore augmented with an instruction and questionnaire “bar” at the bottom of the page. Each

Figure 3.3: Appearance of the bookstore in Study 2. The authority strategy is implemented in the orange box just below the title of the book. The evaluation questions and the tasks for participants appeared in the bottom bar.

session consisted of approximately 20 tasks such as “Please find a book called [book name] and add it to your shopping cart”; six of these tasks asked participants to go to the home page of the bookstore, find a particular book, and rate that book.

The home page of the bookstore displayed books under categories implementing influence strategies. The *no*-strategy (“random selection”) message was omitted. Study 2 included three different implementations of each of the strategies used in Study 1 (see Table 3.6).

During each session participants were asked to rate six books. Participants rated two books for each of the three strategies within each session. These two books were presented with the same implementation and were rated in sequence. The order of the three strategies over each session was randomized and the specific implementations used in

Table 3.6: Influence strategies and their respective implementations as used in Study 2.

Strategy	Implementation
<i>Scarcity</i>	1. <i>Limited Edition</i> . This book is a limited edition and signed by the author. Availability is limited.
	2. <i>Almost Out of Stock</i> . This book is almost out of stock. There are only a few copies left so make your purchase now.
	3. <i>Collector's Item</i> . There are a limited number of prints of this edition and each book is signed and numbered. A pure collector's item.
<i>Consensus</i>	4. <i>Over a Million Copies Sold</i> . This book is a great bestseller. Over a million copies have been sold worldwide.
	5. <i>Voted Best Fiction by Readers</i> . This book was among the highest rated novels in recent in reader surveys. Everyone agrees: this book is a must read.
	6. <i>International Bestseller</i> . Worldwide sales of this book continue to increase. Now it is climbing best-seller charts in the United States also.
<i>Authority</i>	7. <i>Experts' Choice</i> . This book is generating buzz among industry experts. Based on the Experts Book Exchange Top 20, this book is among the most talked about in the past year.
	8. <i>Recommended by Top Authors</i> . This book is a top pick this season among other top novel authors.
	9. <i>Critics' Favorite</i> . Critics might be critical, but none of them had complaints about this book. This critic's favorite received positive reviews throughout the popular press.

that particular session were randomized. The order of books and the combination between books and strategy implementation was also randomized over participants.

Measures

Participants' demographic data, including age and gender, was collected at enrollment in the social science course. In Study 2 the same meta-judgmental measures as used in Study 1 were also collected. Due to the

longer duration of Study 2 the first of the product evaluation items was removed.

Participants

Participants were 70 community college students enrolled in social science courses participating for partial course credit. 59 (84.3%) of the participants were female. The mean age of participants was 25.1 ($SD = 9.9$).

3.3.2 Results Study 2

Figure 3.4: Summed raw book evaluations for two participants selected to illustrate variation and consistency in responses to influence strategies over the three sessions.

The analysis of the product evaluations collected in Study 2 is similar to that in Study 1. However, since each participant evaluated books accompanied by each of the three strategies during each of three sessions, one can separately examine their responses by session. To illustrate, Figure 3.4 presents two participants' product evaluations. Across the three sessions, participant *A* appears to evaluate books accompanied by implementations of the consensus strategy most favorably, while Participant *B* evaluated the books accompanied by scarcity more positively.

Model Comparisons

The dataset collected in Study 2, unlike Study 1, allows analysis of participants' responses to books accompanied by each strategy over the three sessions. On the other hand, all books evaluated by participants

in Study 2 are accompanied by authority, consensus, or scarcity; there are no “control” messages for comparison. As in Study 1, a test of the null hypothesis of no heterogeneity in responses to influence strategies corresponds to the comparison of a series of mixed-effects models. The first model, Model *D*, includes varying intercepts for sessions and participants. The second model, Model *E*, adds strategy $\times$ person effects.

Both models can be written as

$$y_{jbq} \sim \mathcal{N}(\mu + X_{jb}\beta_j + \alpha_b + \eta_q, \sigma_{err}^2) \quad (3.7)$$

where μ is the overall intercept, $\beta_j \sim \mathcal{N}(\bar{\beta}, \Sigma_\beta)$ for $j = 1, \dots, J = 70$ participants, $\alpha_b \sim \mathcal{N}(0, \sigma_\alpha^2)$ for $b = 1, \dots, B = 18$ books, and $\eta_q \sim \mathcal{N}(0, \sigma_\eta^2)$ for $q = 1, \dots, Q = 3$ questions.

In this model X_{jb} is a matrix of indicators for strategies and sessions such that β is a 70×6 matrix of coefficients of the coefficient for each strategy and each session for each participant. $\bar{\beta}_{(i1,i2,i3)} = 0$ represents the absence of “fixed” effects of strategies and sessions. The strategy and session coefficients varying by participant are modeled independently from each other. Thus, Σ_β has a block structure,

$$\Sigma_\beta = \begin{bmatrix} \Sigma_\beta^{(s)} & 0 \\ 0 & \Sigma_\beta^{(t)} \end{bmatrix}, \quad (3.8)$$

in which $\Sigma_\beta^{(s)}$ is the covariance matrix for strategy coefficients and $\Sigma_\beta^{(t)}$ is the covariance matrix for session coefficients. In Model *D*, each element of $\Sigma_\beta^{(s)}$ is set to zero, while in Model *E* it is unconstrained.

Table 3.7 shows the comparison of these two models. Allowing for heterogeneous influence strategy effects significantly improves model fit. This shows that also in this study there is variation in responses to the different implementations influence strategies. In Study 2 the variation is observed across the different sessions and is thus *not caused by transient intra-individual differences*.

Table 3.7: Table comparing Model *E* with a model without varying effects of strategies by person (Model *D*). The comparison shows stable heterogeneity in responses to influence strategies over the three sessions.

	df	BIC	logLik	χ^2	<i>p</i>
Model <i>D</i> :	10	16045.65	-7981.64		
Model <i>E</i> :	16	15967.27	-7917.73	127.81	<.0001

Table 3.8 summarizes the random effects of Model *E*. Qualitative comparison to Table 3.4 shows that the point estimates of the standard deviations of the strategy $\times$ person effects are slightly lower in Study 1 than in Study 2: this accords with the expectation that the estimates in Study 1 are upper bounds because of transient intra-individual variation. The correlations between the strategy $\times$ person effects are larger in Study 2. This is expected for two reasons: First, in Study 1, the overall intercept of each participant is separately identifiable from the strategy $\times$ person effects—which is illustrated by the negative correlations between the by person intercepts and the strategy $\times$ person effects. In Study 2 this is not the case, and thus the correlations between strategies include an overall response tendency. Second, this observation is consistent with the idea that the estimates in Study 1 include noise from transient intra-individual variation and are thus attenuated.

Table 3.8: Summary of random effects of Model *E*. The square root of the diagonal elements of $\hat{\Sigma}_{\beta}^{(s)}$ and $\hat{\Sigma}_{\beta}^{(t)}$ are rows one to three and four to six of the standard deviation column respectively, while the remaining entries of each are presented as correlations to the right.

Grouping		SD	Corr.	
Subject	Consensus	1.40		
	Authority	1.51	0.86	
	Scarcity	1.51	0.79	0.80
Subject	Session 1	1.03		
	Session 2	1.33	0.77	
	Session 3	1.13	0.70	0.83
Book	Intercept	0.36		
Question	Intercept	0.47		
Residual		1.81		

Demographics and Meta-Judgmental Measures

As in Study 1, Study 2 explore several models that use common marketing measures of consumers—such as gender, age, and personality—to explain the observed differences. The results are presented in table 3.9. Contrary to Study 1 the demographic measures of age and gender do not aid in explaining differences in responses to social influence strategies. Measures of *NfC* and personality do aid in explaining responses to social influence strategies. However, these typical measures are again

Table 3.9: Table comparing the Model D with models using demographic or meta judgmental measures of personality.

	df	BIC	logLik	χ^2	p
Model D :	10	16045.65	-7981.64		
Model D_d :	19	16117.46	-7980.48	2.32	0.9853
Model D_m :	37	14441.03	-7070.37	1820.22	<.0001
Model E_m :	42	14385.46	-7022.29	96.14	<.0001

insufficient to capture the variation that can be modeled using operative measures.

3.3.3 Discussion Study 2

The large upper bound on the magnitude of heterogeneity in the effects of influence strategies found in Study 1 motivated estimating this heterogeneity over multiple sessions. The analysis above demonstrates that the heterogeneity in effects identified in Study 1 largely cannot be attributed to transient intra-individual variation. Rather, the observed variation is the result of stable individual differences in responses to distinct influence strategies. Even after using meta-judgmental measures of traits as predictors, including strategy $\times$ person effects significantly improved model fit. This finding highlights that while meta-judgmental measures might aid in understanding responses to social influence strategies, operative measures—the estimated strategy $\times$ person coefficients—reflect additional variation. A variation that, given its magnitude, can be of key-importance for successful persuasion in ambient intelligent systems.

3.4 Study 3: Stability of heterogeneity across context

Study 3 explores whether people’s responses to persuasive strategies in one specific context are predictive for their responses in another context. This is of theoretical importance since it strengthens the assumption that the heterogeneity uncovered in Study 1 presents a stable individual trait that is consistent over different settings. This enables theorists to explain this heterogeneity at an individual level instead of at the level of individual and context. Investigation over contexts is also important in an applied setting: can people’s responses to one persuasive system predict their responses to another system?

3.4.1 Method Study 3

Study 3 consisted of two online parts that were, in the perception of the participants, unrelated studies. The first study of the set of two was of the evaluation of an online bookstore (referred to as “Bookstore Study”) and the second was the evaluation of an online music store (further referred to as “Musicstore Study”). Both studies were conducted online, and in both studies participants were asked to evaluate a new online store. The appearances of the stores were made as distinct as possible and the contact addresses on both studies were different to ensure that participants did not link the two studies together. Figure 3.3—the implementation of the bookstore used also in Study 2—and 3.5 show the different implementations of the online stores used in this study.

Figure 3.5: Appearance of the Musicstore Study. The vertical bar on the left is used to present assignments and questions to participants while they can freely browse the online store.

Procedure

Both of the studies were only open for participation at specific time-intervals. This ensured that participation in the Bookstore study was at least one week separated from participation in the Musicstore study. In both studies participants were first shown an instruction page. Next, participants were taken to an online store which was augmented with an instruction and questionnaire “bar” at the bottom—for the Bookstore (see Figure 3.3 study—or at the left side—for the Musicstore study (see Figure 3.5)—of the page. Each study consisted of approximately 20 tasks such as “*Please find a book called [book name] and add it to your shopping cart*”; twelve of these tasks asked participants to go to the home page of the (music)- bookstore, find a particular (album)-book, and evaluate that (album)- book. After locating the product, participants rated the products on the same three questions as used in Study 2.

Table 3.10: Influence strategies and their respective implementations as used in the Bookstore Study.

Strategy	Implementation
Control	1. <i>Random Selection</i> . These books are randomly selected from our book offerings.
	2. <i>Just a book...</i> Have a look at our books! This is just an arbitrary selection of the books on our website.
Consensus	3. <i>Over a Million Copies Sold</i> . This book is a great bestseller. Over a million copies have been sold worldwide.
	4. <i>International Bestseller</i> . Worldwide sales of this book continue to increase. Now it is climbing best-seller charts in the United States also.
Authority	5. <i>Experts’ Choice</i> . This book is generating buzz among industry experts. Based on the Experts Book Exchange Top 20, this book is among the most talked about in the past year.
	6. <i>Recommended by Top Authors</i> . This book is a top pick this season among other top novel authors.

Located just above the description and cover image of the items presented in both studies was a message that either implemented one of the two influence strategies or a control message. Two implementations of each strategy were delivered in sequence. For example, a

participant would be instructed to find a specific item. The item was presented on the front-page under the category “Recommended by Top Authors”. When selecting the item, participants were shown the implementation of the strategy (“Over a million copies sold!”). Participants always rated two items belonging to the same “category” in sequence. The strategy order was randomized in time and in their presentation on the front page. Furthermore, the books presented for each strategy were randomized for each participant. Each participant rated all items of a study in a single session and was thus exposed to all the implementations. Table 3.10 shows the implementations used in the Bookstore study and Table 3.11 shows those used in the Musicstore study.

Table 3.11: Influence strategies and their respective implementations as used in the Musicstore Study.

Strategy	Implementation
<i>Control</i>	1. <i>Random Selection.</i> These albums are a random selection out of our large bluegrass collection.
	2. <i>Pure Bluegrass.</i> Here are some arbitrary picks out of our bluegrass offerings.
<i>Consensus</i>	3. <i>Over a Million Copies Sold.</i> This album is a great bestseller. Over a million copies have been sold worldwide.
	4. <i>International Bestseller.</i> Worldwide sales of this album continue to increase. It is climbing the charts in the United States also.
<i>Authority</i>	5. <i>Experts’ Choice.</i> This album is generating buzz among music industry experts. Based on the Experts Album Top 20, this album is among the most talked about in the past year.
	6. <i>Recommended by Top Musicians.</i> This album is a top pick this season among top performing musicians.

For each sub-study, after rating $2 \times 2 \times 3 = 12$ items—three strategies, two implementation per strategy, and two (albums)- books per implementation—participants were redirected to another website. For the Musicstore study participant’s were shown a “thank you for participating” message and were directly awarded course credits. For the Bookstore Study participants were forwarded to a post-questionnaire.

Measures

Participants' demographic data, including age and gender, was collected at enrollment in the social science course. In Study 3 the same meta-judgmental measures were used as in Study 1 and 2.

Participants

Participants were community college students enrolled in a research methodology course. As part of the course requirements all students enrolled in the course had to participate in a number of research projects. Participants could signup online for various studies, of which our Bookstore Study—presented in the signup system as “Evaluate an On-line Bookstore”—and Musicstore Study—presented as “Kalua Music”—were two. Participants were unaware of the link between the two studies since the vast majority of studies posted in the signup system were indeed unrelated. In total 247 participants participated in the Bookstore Study, and 162 participated in the Musicstore Study. Since participation in both studies was not enforced in any way the merger of these two data files lead to the data used in further analysis which describes the Bookstore study *and* Musicstore study outcomes of 153 participants. 111 (72.55%) of the recruited participants were female. The mean age of participants was 24.3 ($SD = 8.63$). The study was conducted online and participants participated from their own computers.

3.4.2 Results Study 3

Model Comparisons

Study 3 is analyzed in a similar fashion as Study 1 and 2. The core model comparison(s)—the direct test of the significance of adding individual level effects of the persuasive strategies—is similar to those presented in Study 1: Since Study 3 contains a baseline condition it is first tested whether an *overall* difference in responses to persuasive persuasive strategies at an individual level increases model fit. Next, it is tested whether responses to *distinct* strategies increases model fit over the overall persuadability. Similar to Study 2 these individual level estimates are established over the two context thus only capturing variation that is stable over contexts.

The general model can again be written as:

$$y_{jbq} \sim \mathcal{N}(X_{jb}\beta_j + \alpha_b + \eta_q, \sigma_{err}^2) \quad (3.9)$$

with $\beta_j \sim \mathcal{N}(\bar{\beta}, \Sigma_\beta)$ for $j = 1, \dots, J = 153$ subjects, $\alpha_b \sim \mathcal{N}(0, \Sigma_\alpha)$ for $b = 1, \dots, B = 24$, products and $\eta_q \sim \mathcal{N}(0, \sigma_\eta^2)$ for $q = 1, \dots, Q = 3$ questions.

No main effect of context was included in the final model given that its *average* effect was not statistically significant.

The design matrix X_{jb} is a matrix consisting of a column of ones and indicators for each of the two strategies used in this study. Thus, β is a 153×3 matrix of intercepts and coefficients for each strategy for each individual. The three models that are compared again differ only in their constraints on Σ_β . Model *F* only allows for between-person variation in an intercept, so it has

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (3.10)$$

where σ_I^2 is the between-person variance of the intercept. Model *G* allows for the effect of all three of the strategies to vary *together* by person:

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & \sigma_{I,P}^2 & \sigma_{I,P}^2 \\ \sigma_{I,P}^2 & \sigma_P^2 & \sigma_P^2 \\ \sigma_{I,P}^2 & \sigma_P^2 & \sigma_P^2 \end{bmatrix}. \quad (3.11)$$

Here σ_P^2 is the between-person variance of the overall effect of all three strategies and $\sigma_{I,P}^2$ is the covariance of the intercept and this strategy effect. In Model *H* the entries of this covariance matrix are unconstrained,

$$\Sigma_\beta = \begin{bmatrix} \sigma_I^2 & \sigma_{I,c}^2 & \sigma_{I,a}^2 \\ \sigma_{I,c}^2 & \sigma_c^2 & \sigma_{c,a}^2 \\ \sigma_{I,a}^2 & \sigma_{c,a}^2 & \sigma_a^2 \end{bmatrix}, \quad (3.12)$$

where σ_c^2 and σ_a^2 are the between-person variances of the effects of the consensus and authority strategies.

As with Study 1, given the inclusions of a baseline condition, it is interesting to see whether the implementations of the persuasive strategies lead to an overall increase in model fit, and whether the $\hat{\beta}$'s of both Authority and Consensus as a main effect are significantly positive. Model comparisons indeed show a positive main effect of the strategies: Model fit improves significantly when comparing a simple model with only Subject level intercepts and an overall intercept with a model that includes average effects of the two strategies, $\chi^2 = 13.09$, $df_{\text{delta}} = 2$, $p < 0.01$. The estimates of $\hat{\beta}$ for Authority is 0.157, $t = 2.126$, $p < .05$, and that for Consensus is 0.153, $t = 2.270$, $p < .05$. Both are,

as expected, significantly positive.

After establishing this average effect the heterogeneity in responses to persuasive strategies is explored. Table 3.12 shows the outcomes of the model comparisons of Model F , G , and H . While modeling overall persuadability—responses to all influence strategies as opposed to no strategy—significantly improves model fit, allowing for varying effects of the strategies leads to a large increase of the fit of the model. This replicates the findings obtained in Study 1 and Study 2, but this time over the two contexts. Thus, the predictions of responses in both the bookstore as well as the musicstore improve when modeling heterogeneity in responses to persuasive strategies.

Table 3.12: This table shows a comparison of a model without any effects of strategy with models that include an overall strategy effect or a strategy specific effect. It is clear that the model with varying specific strategy effects by individual is preferred.

	Df	BIC	logLik	χ^2	p
Model F :	7	46826.15	-23380.59		
Model G :	9	46786.00	-23351.23	58.72	< .001
Model H :	12	46778.71	-23333.66	35.13	<.001

Similar to the analysis presented in Study 1, table 3.13 shows the estimates of the random effects of Model H . Again it is clear that there are sizable individual differences—differences that compare to the ones found in the previous two studies.

Table 3.13: Summary of the random effects of Model H . The square root of the diagonal elements of Σ_β are the first three rows of the SD column. The remaining entries are presented as correlations ($Corr$).

Groups	Name	Variance	SD	$Corr$	
subject	(Intercept)	1.55	1.24		
	Consensus	0.32	0.57	-0.10	
	Expertise	0.44	0.67	-0.091	0.619
item	(Intercept)	0.09	0.30		
question	(Intercept)	0.34	0.58		
Residual		4.21	2.05		

Demographics and Meta-Judgmental measures

Similar to Study 1 and 2 the effects of both demographics as well as personality measures on prediction are examined. Table 3.14 compares the model fit of Model F with models including these predictors. As in the previous studies the addition of personality data improves prediction, but does not exhaust the variation caused by the heterogeneity in responses to persuasive strategies. This last fact is eminent by the significant increase in model fit from Model F_m to Model H_m : Addition of varying Strategy $\times$ Subject effects increases model fit according to the log-likelihood comparisons.

Table 3.14: Table comparing the Model F with models using demographic or meta judgmental measures of personality.

	Df	BIC	logLik	χ^2	p
Model F :	7	46826.15	-23380.59		
Model F_d :	15	46894.36	-23377.57	6.04	0.6427
Model F_m :	33	40545.01	-20121.78	6511.57	< .001
Model H_m :	38	40545.02	-20098.96	45.66	< .001

3.4.3 Discussion Study 3

Study 3 extends the results found in Study 1 and 2 by addressing the stability of individual differences in responses to persuasive strategies over contexts. The results show that while the Study replicated the main effect that is often reported upon in the literature for both Consensus as well as Authority, the individual differences are again larger in size. This again leads to the result that the model fit is better when modeling individual level effects than when modeling average effects.

3.5 Conclusions

This chapter explored individual differences in responses to persuasive strategies using operative measures. The results of Study 1 show the potential importance of this heterogeneity by comparing the differences between individuals to the often attended to average effects of persuasive strategies. It is clear that for a number of people some strategies that are effective *on average* have negative effects. This opens the door for adaptation of persuasive attempts at the level of the persuasive principles identified by Cialdini (2001).

Study 2 extends the findings of Study 1 by identifying the stable part of the heterogeneity uncovered in Study 1. This stability shows that the heterogeneity identified in Study 1 is, by and large, identifying *traits* in stead of temporary states. Thus, people who respond in a certain way to a distinct influence strategy right now are likely to do so in the future. This makes the individual differences something that can be attended to by designers of persuasive systems. Study 3 explored stability across contexts. The model comparisons show that while the average effect(s) significantly improves prediction, the individual level effects are even more important to decrease prediction error.

In each of the three studies presented above the possibility of explaining the observed heterogeneity via the use of measurements of individual differences typically used in marketing—such as demographics and personality measures—is explored. In each case, the heterogeneity observed in responses to influence strategies is not exhausted by these commonly used measures. This does not prove that meta-judgmental measures of individual differences will never sufficiently capture the observed variation: it is likely that the inclusion of more, or different, predictors would better explain the heterogeneity. However, the analysis does show that the selected meta-judgmental measures fail to fully explain the variation observed in operative measures of responses to influence strategies.

While the three studies presented in this chapter show the importance of attending to individual differences between people the studies do not posit a mechanism through which the heterogeneity can be explained. As is clear in Chapter 2, Section 2.2.2, multiple competing mechanisms have been proposed to explain the success or failure of distinct strategies. Each of these mechanisms can be in place when the strategies are operated upon. The social science literature is however largely uninformative about the relative importance of different mechanisms, let alone their interactions. This renders these theoretical explanations hardly useful for the design of persuasive systems³.

³This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2011c; Kaptein and Eckles, 2010; Kaptein et al., 2011b; Kaptein and Eckles, 2012).

4

Insight Generation II: *Choice and Repetition Effects of Influence Strategies*

4.1 Introduction

The studies presented in the previous *insights generation* chapter (Chapter 3) demonstrate that there is large heterogeneity in responses to influence strategies which is consistent both over time as well as over contexts. However, all three of these studies concerned instances of influence strategy usage in which (a) the influence strategy used to endorse an appeal was *selected by the experimenter* and (b) only *a single implementation* of an influence strategy was used per persuasive attempt. Both of these are however not self-evident: (e.g.) Frequently in e-commerce products are presented with a number of different influence strategies at the same time (e.g. the product is both a bestseller *and* almost out of stock) and users frequently are offered a choice about the strategy they “prefer”: users can choose to look at the bestseller list, or in persuasive applications for health and wellbeing can choose to contact an authority figure for active coaching. There is however no clear guidance provided by the social science literature in these respects: the

ELM (e.g.) could lead one to believe that all influence strategies—when regarded peripheral cues—have their effect through the peripheral route and thus their effects would add up. It could however also lead one to hypothesize that the use of multiple influence strategies increases central processing and subsequently decreases the effects of the strategies. In this chapter we test these conflicting hypotheses.

This chapter examines the effects of both a choice between influence strategies as well as the simultaneous usage of these strategies. Study 1 examines the effects of choice on the effectiveness of social influence strategies. Studies 2 to 4 examine the effects of using *multiple* influence strategies to support a *single* appeal as opposed to the selection of one specific strategy. The problem of the simultaneous presentation of multiple influence strategies to support a single appeal has been understudied but is valuable to designers of persuasive systems. Only within the marketing literature serious attempts have been made to tackle this questions and the results are mixed: Barry and Shapiro (1992) find that using multiple social influence strategies—or *sales tactics* in that branch of literature—can be detrimental for compliance. Thus, when a single influence strategy is used (e.g. this product is almost out of stock) this seems more effective then using multiple strategies (e.g. the product is out of stock and a bestseller). Falbe and Yukl (2008) however derive a different conclusion from observing multiple human to human influence attempts within a company setting: they show that managers who are flexible and use multiple influence tactics on the same target to support a single appeal are generally more successful than those sticking to a single influence tactic. Both of these existing studies are however correlational and thus they do not provide causal evidence. Studies 2 to 4 extend this existing work by experimentally testing the effects of multiple influence strategies versus a single influence strategy.

4.2 Study 1: Choice Effects

Study 1 aimed at manipulating a persuasive scenario to examine the effects on compliance of two main factors: *disclosure*, explicitly stating to the user that an influence strategy is being used in an attempt to change their opinion, and *choice*, providing users with a choice between multiple persuasive strategies. Arguably, free choice can be a powerful tool for designers to make persuasive technologies more transparent. However, designers must be cognizant of the actual effects of disclosure as it influences user compliance (proven in subsequent findings).

Study 1 used a decision task in which participants were asked to rank a number of items (Nass et al., 1996). Participants first ranked the items in the order they believed was most important before they were ostensibly given advice from one of two groups: a group of individuals who previously completed the task successfully (implementing the consensus strategy) or by an expert (implementing the authority strategy). This “source of advice” was either randomly assigned to the participant (no choice), assigned ostensibly based on knowledge (no choice, knowledge), or chosen freely by the participant (free-choice). Next, the use of this strategy for persuasive purposes was either disclosed or not. This results in a 2×3 between subject design with three levels of choice and two levels of disclosure. Behavioral compliance levels as well as participants’ subjective interpretation of the influence strategies were assessed.

4.2.1 Method Study 1

Participants

Participants were recruited from a list of university students registered for an introductory research methodology course. A total of 112 participants were recruited. Fifty three of the participants were female (47.3%) and 59 were male (52.7%). The mean age of participants was 22.2 years ($SD = 3.3$).

Procedure

The study was conducted entirely online. Participants received a link to the online study via email and were asked to complete the study which would take approximately 30 minutes to complete. The first task of the study was an item-ranking task in which participants were asked to rank 12 items in order of their importance for survival in the arctic. Participants were introduced to the scenario in the following way:

“You have just survived the crash of a small plane. Both the pilot and co-pilot were killed in the crash.

It is mid-January, and you are in Northern Canada. The daily temperature is 25 below zero, and the night time temperature is 40 below zero. There is snow on the ground, and the countryside is wooded with several creeks criss-crossing the area. The nearest town is 20 miles away. You are dressed in city clothes appropriate for a business meeting. You manage to salvage twelve items that you can use to try to survive.”

After the introduction, participants were shown the following list of items and given the opportunity to rank the items from 1 (most important to survival) to 12 (least important to survival):

1. A 20'x 20' piece of heavy-duty canvas.	7. A hand ax.
2. A cigarette lighter.	8. Iodine water purification tablets (50 tablets).
3. A compass.	9. A loaded .45-caliber pistol.
4. Dehydrated milk (8 pounds).	10. A loud signal whistle.
5. Duct tape (25' roll).	11. One box of matches.
6. An extra shirt and pair of pants for each survivor.	12. A sectional air map made of plastic.

After ranking the items participants were told—after an ostensible 6 seconds analysis of their ranking—that “...*some of your rankings were correct, but some could use improvement. You will now get the chance to revise your answers.*”

Participants then saw a screen which stated that they would have the opportunity to revise their answers either based on: “*The advice of successful students*” or “*the advice of an arctic expert*”. Both sources of advice were presented with a picture of the respective source. These descriptions of the two sources of advice were pre-tested with 145 subjects to determine overall strategy preference in this survival scenario context. In 144 of the 145 cases, our pilot participants—again college students—chose the expert advice. Hence, it was clear that this implementation of the expert advice is the preferred social influence strategy for this specific task.

Next, one third of our participants were randomly assigned to receive the message that they would be “*randomly assigned to one of the two sources of advice*”. In reality, everyone was assigned to receive the expert advice—the preferred advice as based on the pre-test. Another third of our participants received a message that they were assigned to receive the expert advice because “...*this is the advice you will like best*”. The ostensible analysis of participants rankings, which was shown to all participants, was implemented to increase the realism of this condition. The final third of our participants had the option to select the advice they would prefer most; these participants had a free-choice between the two strategies. Consistent with our pre-test results all of the participants in this condition ended up choosing the expert advice.

After the choice manipulation, participants received short feedback about 6 of their 12 ranked items. For each participant, irrespective of their ranking, it was suggested to: move their most important item (ranked 1) down to rank 7, move the item ranked at 3 to 5, move 5 to 12, 7 to 1, 10 to 3, and 12 to 10. These suggestions were presented one by one, and, during the presentation of these suggestions, the disclosure conditions were implemented. Half of the participants received the message “You ranked [the item] at number [ranking by the participant]. The expert ranked [the item] as [more/less] important at number: [suggested rank].” The other half of the participants were shown the same advice *and* a message box with the following message: “*Please note that research shows that people tend to be persuaded by experts*”.

After reading each of the six recommendations, participants were shown a screen displaying their initial ranking of the twelve items compared side by side with the ranking of the expert. Participants could then re-rank their items as desired based on the expert’s advice. For participants in the disclosure condition, the aforementioned message box was again displayed. After re-ranking the items, participants were asked to evaluate the advice that was given. The study ended with a short set of demographic questions.

Measures

As a primary measure of the effectiveness of the persuasive attempt, (the advice given by the expert) a compliance score was computed. The compliance score is the sum of the number of ranks changed between initial rank and suggested ranks for the items for which a re-rank was suggested. Hence, if the item that was initially ranked at 1 was moved to position 5 (suggested position was 7) in the final rankings, participants received a score of $(|7 - 1| - |7 - 5|) = 4$, (the maximal suggested change minus the actual distance between the suggested rank and the final rank) for that item. The maximum compliance score was $(6 + 2 + 7 + 6 + 7 + 2 =) 28$. Negative scores could be obtained when items were moved in greater rank-distance to the suggestions than the initial rank, however, this did not occur in the study.

In addition to the actual compliance score, the perceived usefulness of the expert advice and participants confidence in their final ranking were measured. Perceived usefulness was measured using the following 5 ten-point items (Cronbach’s $\alpha = 0.934$):

1. *How useful was the advice provided by the expert?*
2. *How much did you respect the opinion of the expert?*

3. *Did the advice from the expert change your opinion?*
4. *How helpful was the advice from the expert for your ranking in the arctic survival task?*
5. *How satisfied are you with the help from the expert?*

Participants' confidence in the final ranking was measured using the following 2 ten-point items (Cronbach's $\alpha = 0.859$):

1. *How confident are you in your final ranking?*
2. *How satisfied are you with your final ranking?*

4.2.2 Results Study 1

Compliance

To examine the effects of disclosure and choice on compliance, a compliance score was computed for each participant and these scores were analyzed using a 2×3 between-participants ANOVA. There was a significant main effect of choice, $M_{no,random} = 18.2$, $M_{no,knowledge} = 21.0$, $M_{free-choice} = 22.2$, $F(2,106) = 3.38$, $p = 0.038$. Furthermore, there was a statistically significant main effect of disclosure on the compliance scores, $M_{no} = 22.1$, $M_{yes} = 18.8$, $F(1,106) = 6.55$, $p = 0.012$. No interaction between choice and disclosure was found, $F(2,106) = 0.10$, $p = 0.905$. Figure 4.1 shows an overview of the results.

It is clear that compliance to the expert's advice is higher when the persuasive intent of this implementation was not disclosed. Using post-hoc tests, the main effect of choice was further examined. The free-choice condition differed significantly from the random-assignment condition, $p = 0.011$. None of the other conditions differed significantly from each other.

Usefulness of the Advice

A mean usefulness of the advice was computed for each participant. The usefulness score was analyzed using a 2×3 between-participants ANOVA. There was a significant effect of choice, $M_{no,random} = 7.1$, $M_{no,knowledge} = 7.0$, $M_{free-choice} = 8.0$, $F(2,106) = 4.19$, $p = 0.018$. Also, there was a significant main effect of disclosure on the usefulness scores, $M_{no} = 7.7$, $M_{yes} = 7.0$, $F(1,106) = 4.74$, $p = 0.032$. No interaction between choice and disclosure was found, $F(2,106) = 0.24$, $p = 0.787$. Figure 4.2 shows an overview of the results.

It is clear that the advice from the expert is perceived as more useful when it is not disclosed to participants the fact that advice from an expert tends to influence opinion. Thus, disclosure diminishes the

Figure 4.1: The effects of choice and disclosure on compliance. Shown are the estimated marginal means and standard errors of the compliance scores. On the x-axis, are the three choice levels, and the two separate lines represent the two disclosure conditions.

perceived usefulness of this implementation of the authority strategy. Furthermore, it is clear that free-choice increases the perceived usefulness of the advice over conditions in which participants are not free to choose for this advice source. Post-hoc tests show that the free choice condition differed significantly in its perceived usefulness from both the *no-choice, random condition*, $p = 0.031$, and the *no-choice, knowledge condition*, $p = 0.028$.

Confidence

Analysis of the confidence scores using a 2×3 ANOVA showed a significant main effect of choice, $M_{no,random} = 6.1$, $M_{no,knowledge} = 6.7$, $M_{free-choice} = 7.7$, $F(2,106) = 6.20$, $p = 0.003$. No significant main effect of disclosure, $F(2,106) = 3.23$, $p = 0.076$, and no significant interaction were found, $F(2,106) = 1.22$, $p = 0.300$. Using post-hoc tests, it was clear that the free-choice condition differed significantly from both the no choice, random condition, $p = 0.001$, and the no choice, knowledge condition, $p = 0.031$.

Figure 4.2: The effects of choice and disclosure on the perceived usefulness of the advice. Shown are the estimated marginal means and standard errors of the usefulness scores.

4.2.3 Discussion Study 1

Study 1 shows that both choice—the ability to select the persuasive strategy that is used for the specific influence attempts—and disclosure—the revelation of the effect of this strategy—affect compliance. Free-choice led to higher compliance to the request. Disclosure of the fact that the expert advice generally has an influence on others, led to decreased compliance.

Free-choice by users of different persuasive strategies leads to higher compliance, usefulness, and user confidence ratings than system assigned strategies. This implies that designers of adaptive persuasive systems should try to involve users in the selection of different influence strategies. The effectiveness of free-choice is likely a result of the consistency principle: Once people make a certain choice, they will go to great lengths to stick to this choice (Cialdini, 2001). Hence, once a user chooses to comply to an expert, they will try to adapt their behavior to be consistent with this choice.

Disclosing the general effect of the use of the expert strategy reduced its effectiveness. This was probably due to the fact that disclosing the

use of a persuasive strategy leads to a higher elaboration state, reduces peripheral processing, and thus lessens the impact of the influence strategy through the peripheral route (see Cacioppo et al., 1986). This finding might imply that secrecy about the use of influence strategies can be beneficial for persuasive systems. The implications of this study for the design of persuasive systems are further detailed in Chapter 6.

4.3 Study 2 and 3: Simultaneous Presentation

While Study 1 focused on the effects of both disclosure and choice on the effectiveness of using influence strategies, Study 2 to 4 focus on the applied question whether adaptation is actually necessary: *Is it not just possible to use multiple strategies simultaneously?* In these studies we investigate the main effects of the usage of multiple strategies—and thus not the individual differences in responses to multiple strategies.

4.3.1 Method Study 2

Participants

Forty-four undergraduate students volunteered to participate by accepting an email invitation with a link to the study website. This email invitation went out to a total of 136 possible participants, giving the study a 32.4 percent response rate. Of the final sample 25 (56.8 percent) were females. The average age of the sample was 23.8 (SD=7.6).

Procedure

The first part of the study consisted of the same arctic survival item-ranking task as used in Study 1. After participants were finished ranking these items, a period of 6 seconds was spent ostensibly analyzing their ranking before participants were told *"...some of your rankings were correct, but some could use improvement. You will now get the chance to revise your answers."* After this message the experimental manipulations were implemented and participants had the opportunity to revise their rankings.

Manipulations

In this experiment four conditions were used on two dimensions: (a) the number of strategies used (*one* or *multiple*) and (b) whether, in the multiple strategy condition, the implementations were congruent (*yes* or *no*). In the single strategy condition participants received “advice” on how to change their rankings either from an expert or based on the consensus of a group of similar others. Participants were told: *"You*

will have the chance to revise your answers based on advice from..." and then were shown a picture and a brief textual description of their advice source. Despite the different source labels, all participants were exposed to the same advice.

The authority strategy was implemented as follows: Participants were told that they would receive advice originating from a "survival expert". This was supported with the notion that: "*You will get tips on how to better rank your items based on the knowledge of an arctic expert.*" The consensus strategy was implemented by stating that participants would receive advice from "Other students" and was further elaborated on by stating: "*You will get tips on how to improve your ranking of items based on the consensus of other students who have generally done well on these types of problems.*"

Congruence of sources was nested within the multiple strategy condition. In the congruent condition, the advice of both sources agreed. The message read: "*The group of successful students agreed with the expert*" for five out of the six suggestions. To increase realism, one of the six suggestions stated that the sources *disagreed* with each other. These two numbers were reversed in the incongruent condition so that the two sources disagreed on five out of six suggestions and agreed on one.

These are the four experimental groups in the experiment:

1. *Single strategy-Authority*. Advice from the authority source only. $N=10$
2. *Single strategy-Consensus*. Advice from the consensus source only. $N=12$
3. *Multiple strategy-Congruent*. Agreeing advice from the authority source and consensus source. $N=10$
4. *Multiple strategy-Incongruent*. Disagreeing advice from the authority source and consensus source. $N=12$

Measures

Measures in Study 2 were the same as those used in Study 1.

4.3.2 Results Study 2

For each of the dependent variables the main-effect of the number of strategies was first examined. Next, separate analyses were performed to test the effects of the specific strategy that was used, the number of strategies used, and the congruency of the messages under the multiple strategy condition.

Compliance

For the actual compliance to the advice given in the four experimental conditions, we found no significant main-effect of the number of sources: The average compliance score for the single source conditions, $\bar{X} = 15.8$, $S.E. = 1.96$, was similar to that of the multiple strategy condition, $\bar{X} = 17.2$, $S.E. = 1.62$ $t(42) = 0.55$, $p = .582$. Within the single strategy condition a strong effect—as expected based on the pre-test for Study 1—of the actual strategy that was used was found: Participants in the authority condition, $\bar{X} = 23.6$, $S.E. = 1.87$, complied much more to the advice than participants in the consensus condition, $\bar{X} = 9.3$, $S.E. = 1.61$, $t(20) = 5.80$, $p < .001$. Within the multiple strategy condition both the incongruent group, $\bar{X} = 17.6$, $S.E. = 2.34$, and the congruent group, $\bar{X} = 16.8$, $S.E. = 2.33$ had approximately similar mean compliance scores, $t(20) = 0.235$, $p = .816$. In all of the conditions the compliance scores were significantly higher than 0, indicating that the advices in all of the conditions influenced the ratings of participants.

Figure 4.3: The effects of the use of single or multiple strategies, either authority or consensus and congruent or incongruent, on compliance.

Figure 4.3 shows the means and standard errors for each of the experimental groups. When conducting a one-way four level ANOVA on this data, there is a significant main-effect of condition, $F(3,40) = 8.097$, $p < .001$. Bonferroni corrected pairwise comparisons show that the single strategy consensus condition scores significantly lower than

all other conditions, while the single strategy authority condition scores significantly higher. Table 4.1 shows the mean differences between each of the four conditions, their standard errors, and the p -value for each possible pairwise comparison. The results indicate that when choosing the optimal single strategy for a specific context, adding other strategies can have a detrimental effect on compliance. It further shows that when one of the sources is preferable, the effect of congruency is small.

Table 4.1: Post-hoc comparisons of the experimental conditions in Study 2. SS=Single strategy, MS=Multiple strategies, Auth=Authority, Cons=consensus, Congr=congruent, Incongr=incongruent.

(A) condition	(B) condition	$(\bar{A} - \bar{B})$	$S.E.$	p -value
MS Congr	SS Cons	8.25	2.79	.005
MS Congr	MS Incongr	0.78	2.93	.790
MS Congr	SS Auth	-6.02	2.93	.046
SS Cons	MS Congr	-8.25	2.79	.005
SS Cons	MS Incongr	-7.47	2.93	.015
SS Cons	SS Auth	-14.27	2.93	.001
MS Incongr	MS Congr	-0.78	2.93	.790
MS Incongr	SS Cons	7.47	2.93	.015
MS Incongr	SS Auth	-6.8	3.06	.032
SS Auth	MS Congr	6.02	2.93	.046
SS Auth	SS Cons	14.27	2.93	.001
SS Auth	MS Incongr	6.8	3.06	.032

Confidence

Analysis of the confidence scores (Cronbach's $\alpha = 0.783$) shows that the average confidence score for the single source conditions, $\bar{X} = 7.7$, $S.E. = .29$, was similar to that of the multiple strategy condition, $\bar{X} = 8.1$, $S.E. = .25$ $t(42) = 0.673$, $p = .504$. Within the single strategy condition no significant effect of strategy was found: Participants in the authority condition, $\bar{X} = 7.8$, $S.E. = .52$, were as confident as participants in the consensus condition, $\bar{X} = 7.7$, $S.E. = .61$, $t(20) = 0.061$, $p < .952$.

Different from the previous results on compliance, a significant effect of congruency was found: Within the multiple strategy condition the confidence in the final rating based on advice from incongruent sources, $\bar{X} = 7.5$, $S.E. = 0.17$, was lower than that based on advice from congruent sources, $\bar{X} = 8.8$, $S.E. = .43$, $t(20) = 2.992$, $p = .007$. Hence, while

incongruent advice did not lead to lowered compliance, it did lead to a lowered confidence in the final rankings.

4.3.3 Discussion Study 2

Study 2 quantitatively shows a predisposition of participants to respond to one preferred social influence strategy, and solely that strategy. Compliance was greatest in the condition in which the advice came from *only* the preferential source (authority). Having multiple sources of advice agree on the recommendation had not only no positive impact on compliance levels but actually had a slightly negative effect when compared to the preferred strategy (Table 4.1, bold). This finding could potentially be a result of increasing cognition and elaboration moving from a strictly peripheral processing approach to a, higher elaboration, central route. While peripheral processing and central processing should be regarded as end-points on a continuum (Petty and Cacioppo, 1986), a move towards more central processing is likely to lower the effects of influence strategies that are hypothesized to be effective primarily via the peripheral route.

The authority and consensus strategies employed in this study are likely to be primarily effective via the peripheral route—although a *legitimate* authority strategy or credible consensus strategy can also be effective via central processing (see also Petty et al., 1997)—and thus smaller effects of the influence strategies due to an increase in elaboration could be expected. The added advice could also have introduced a sense skepticism or lack of trust with intention of the application (observed in the variance in the usefulness-index across conditions). However, the result is clear that in some situations using multiple strategies can be detrimental as compared to the presentation of a single, correct, strategy.

Another plausible explanation for the difference in effectiveness of different strategies (as witnessed in the difference in compliance between the consensus and the authority strategy), as well as the lack of an additive effect of using both strategies, can be found in the work of McGuire (see, e.g. 1981, 1995) on the effectiveness of communication campaigns. McGuire (1981) describes how different response “steps”, ranging from the initial exposure to persuasive communication, via attention, comprehension, behavior change, to structural change may require different inputs: different types of messages, via different channels, and from different sources. It is thus likely that while both of the strategies used in Study 2 are effective for persuasion, they are effective for different

response steps. Combining strategies that are effective inputs for different response steps could be ineffective for a specific step, or at least appear ineffective when only one of the response-steps is measured in an experiment. The possible combination(s) of different persuasion inputs (such as the different influence strategies used in this study) and their effectiveness to influence different response steps is an interesting avenue for further study.

4.3.4 Method Study 3

Study 2 raised an important question that still needs to be answered to have a clear picture of the correct implementation approach: If there is no clearly preferred strategy (i.e. the implementations of both sources of advice are equally influential), is a single strategy implementation still optimal? In study 3 we address this question by repeating the setup of Study 2 but this time with equally preferred sources of advice to see if there is any context in which multiple strategies are significantly more effective than a single strategy.

Creating Equal Implementations of Influence Strategies

To create implementations of both the Consensus and Authority strategy that were equally preferred sources of advice in the item-ranking scenario we pre-tested a number of different implementations of both the Authority strategy and the Consensus strategy in a similar way as way pre-tested participants preference towards the strategy in study 2: A group of participants was invited via email to, after being introduced to the item-ranking task, choose which sources of advice they would like advice from.

After several changes to our wordings and several small pilots we choose to reword the implementation of the Authority strategy to make it less preferential given the item-ranking scenario. Participants were told that they would receive advice originating from a “doctor”. This was supported with the notion that: “You will get tips on how to better rank your items based on the knowledge of a doctor.” The consensus strategy was implemented like in Study 1. A pre-test of these two implementations by 69 participants showed that 32 participants wanted to receive advice from the expert, while 37 participants choose the consensus advice. Thus, these implementations in Study 3 enable us to examine the effect of using a single or multiple strategies when the strategies are (almost) equally preferred.

Participants

48 undergraduate students volunteered to participate by accepting an email invitation with a link to the study website. The email invitation was sent out to a total of 113 possible participants, giving the study a 42.5 percent response rate. The final sample consisted of 28 (58.3 percent) females. The average age of the sample was 21.3 (SD=2.19). None of the participants in Study 3 had previously participated in Study 2.

Procedure

The procedure used in Study 3 was identical to that of Study 2, with the exception that the authority source and the consensus source were designed to be equally preferential.

4.3.5 Results Study 3

Compliance

For compliance to the advice given in the four experimental conditions, we found no significant main-effect of the number of sources: The average compliance score for the single source conditions, $\bar{X} = 19.0$, $S.E. = 1.49$, was not significantly different from that of the multiple strategy condition, $\bar{X} = 16.7$, $S.E. = 1.70$ $t(46) = 1.029$, $p = .309$. Within the single strategy condition no effect—as expected based on the pre-test—of the actual strategy that was used was found: Participants in the authority condition, $\bar{X} = 19.5$, $S.E. = 2.34$, complied equally to the advice as participants in the consensus condition, $\bar{X} = 18.5$, $S.E. = 1.97$, $t(22) = 0.300$, $p = .767$. Within the multiple strategy condition participants in the incongruent group, $\bar{X} = 12.3$, $S.E. = 2.52$, complied less to the advices than those in the congruent group, $\bar{X} = 21.1$, $S.E. = 1.51$, $t(22) = 2.98$, $p < .01$.

Figure 4.4 shows the means and standard errors for each of the experimental groups. When conducting a one-way four level ANOVA on this data, there is a significant main-effect of condition, $F(3, 44) = 3.267$, $p < .05$. Bonferroni corrected pairwise comparisons (See Table 4.2) show that the Multiple Source Incongruent condition scores significantly lower than all of the other conditions.

Confidence

Analysis of the confidence scores (Cronbach's $\alpha = 0.908$) shows that the average confidence score for the single source conditions, $\bar{X} = 7.2$, $S.E. = .24$, was similar to that of the multiple strategy condition,,

Figure 4.4: The effects of the use of single or multiple strategies, either authority or consensus and congruent or incongruent advice on compliance in Study 3.

Table 4.2: Post-hoc comparisons of the experimental conditions in Study 3. SS=Single strategy, MS=Multiple strategies, Auth=Authority, Cons=consensus, Congr=congruent, Incongr=incongruent.

(A) condition	(B) condition	$(\bar{A} - \bar{B})$	<i>S.E.</i>	<i>p</i> -value
MS Congr	SS Cons	2.5	2.99	.409
MS Congr	MS Incongr	8.8	2.99	.006
MS Congr	SS Auth	1.6	2.99	.600
SS Cons	MS Congr	-2.5	2.99	.409
SS Cons	MS Incongr	6.3	2.99	.043
SS Cons	SS Auth	-.92	2.99	.761
MS Incongr	MS Congr	-8.8	2.99	.006
MS Incongr	SS Cons	-6.3	2.99	.043
MS Incongr	SS Auth	-7.2	2.99	.021
SS Auth	MS Congr	-1.6	2.99	.600
SS Auth	SS Cons	.92	2.99	.761
SS Auth	MS Incongr	7.2	2.99	.021

$\bar{X} = 7.15$, $S.E. = .38$ $t(46) = 0.837$, $p = .407$. Within the single strategy condition no significant effect of strategy was found: Participants in the authority condition, $\bar{X} = 7.3$, $S.E. = .36$, were as confident as participants in the consensus condition, $\bar{X} = 7.1$, $S.E. = .34$, $t(22) = 0.340$, $p < .737$. No significant effect of congruency was found: Within the multiple strategy condition the confidence in the final rating based on advice from incongruent sources, $\bar{X} = 7.5$, $S.E. = 0.65$, was the same as that based on advice from congruent sources, $\bar{X} = 7.5$, $S.E. = .41$, $t(22) = 0.001$, $p = 0.99$.

4.3.6 Discussion Study 3

The results of Study 3 extend the findings of Study 2. In Study 2 it became clear that using multiple social influence strategies to support a single goal does not necessarily lead to increased persuasion. Especially when one of the social influence strategies used is less effective than other strategies the overall persuasion is lowered. Study 3 examined whether, when social influence strategy implementations are equally preferred, a similar result holds. The results show that (a) the results in study 2 are indeed at least partially explained by aversion of one non-preferential strategy but, (b) even if two strategies are equally effective *the persuasion does not always add up*. This latter finding indicates that combining multiple strategies for a single appeal can be an unattractive option for designers: Persuasion might not be increased if both strategies are successful, and it can be *decreased* if one of the social influence strategies selected by the designers is suboptimal. The results of both of these studies thus support the idea that designers should be cautious about combining multiple influence strategies to support a single request.

4.4 Study 4: Simultaneous Presentation in Practice

Study 2 and 3 demonstrated the importance of a careful implementation approach for persuasive arguments. These studies used controlled laboratory manipulations to prove that not only is it the type of persuasive strategy that matters, but more importantly, how and with what other strategies an argument is implemented. Study 4 takes the previous findings and implements the above strategies in a traditional consumer-facing advertisement setting.

4.4.1 Methods Study 4

Participants

North American Google Search engine users, $N = 197102$, between the ages of 18 and 55 were exposed to one out of six advertisement belonging to one of our experimental conditions.

Procedure

Six Google search advertisements were created to solicit users to participate in a study. All six advertisements were titled "Participate in a Study!" with varying 140 character descriptions that fell into two conditions: advertisements using only a single social influence strategy versus those that use multiple social influence strategies. Three social influence strategies were used in this study: a) consensus, b) authority, and c) scarcity. Our main aim was to compare the performance of an advertisement that implements all of these strategies to one that implements only a single of these strategies.

In the Single Strategy condition an implementation of one of the social influence strategies was shown to participants in the textual add to encourage them to participate in our online study:

1. *"100s of others have taken this study before."*
2. *"Professor Ford recommends taking this study."*
3. *"There are only 18 hours left to participate in this study."*

Where the first implements the influence strategy consensus, the second authority, and the third scarcity.

In the Multiple Strategy condition implementations of multiple strategies in a single advertisements were shown to participants. To control for implementation order we showed one of the following advertisements to participants in this condition:

1. *"100s participated, & Professor Ford recommends it. Only 18 hours left."*
2. *"Prof. Ford recommends it, 100s participated, only 18 hours left."*
3. *"Only 18 hours left, & Professor Ford recommends it. 100s took it."*

The study ran for 21 days and each advertisement was systematically alternated over time to ensure an even distribution over the allocated time period. Once a participant clicked on the advertisement the success of that advertisement was logged. After clicking participants were taken to a landing page that asked if they would like to participate in a study (informed consent). If they selected "Yes", they were taken to

the "study" page where they were told to rank pictures one at a time based on a scale of one to five of how the picture made them feel. This latter task was unrelated to the study reported here.

Measures

The primary measure of compliance to the conditions was the click-through rate (clicks / views) for the advertisements in each condition. This score measures the direct effectiveness of the multiple versus single strategy conditions.

4.4.2 Results Study 4

Click-through

Table 4.3 shows the number of views and clicks on each of the advertisements. Aggregated over the different stimuli in the two conditions there were 87356 views and 316 clicks in the single strategy condition. The click-through rate in this condition was .36 percent. In the multiple strategy condition the number of views was 109746 and the number of clicks was 195. This is an average click-through of .18 percent. The multiple strategy condition thus performed significantly worse than the single strategy condition, $\chi^2 = 63.1$, $p < .001$.

Table 4.3: The number of views and the click-through rate of the advertisements used in Study 4.

Cond.	Add text	Views	Clicks	%
Single	1. "100s of others..."	25825	123	.47
	2. "Professor Ford..."	24509	61	.25
	3. "There are only..."	37022	132	.36
Multiple	1. "100s participated..."	23546	52	.22
	2. "Prof. Ford..."	21390	51	.24
	3. "Only 18 hours left..."	64810	92	.14

To illustrate the effectiveness of each of the advertisements the beta-binomial distribution was used to model the success of each of the advertisements independently (gray lines) and aggregated over conditions (black and gray solid lines). Figure 4.5 plots each of these probability density distributions. The vertical lines are the 97.5% and 2.5 percentiles of the distribution. It is clear that each of the advertisements that implements a single strategy scores higher than those implementing multiple strategies. Thus, the effects of the social influence strategies do not to add up when used simultaneously in a single advertisement.

Figure 4.5: Model based click-through behavior based on the different advertisement versions. Solid lines represent average of the click-through for single strategy advertisements (black) and multiple strategy advertisements (gray). Dotted lines represent individual advertisements.

4.4.3 Discussion Study 4

Study 4 replicated the findings of Study 2 and 3 in an externally valid setting: Usage of multiple influence strategies for a single appeal does not always increase persuasion. Study 4 used several single influence strategy and multiple influence strategy implementations to test this hypothesis. It is clear that both averaged over all tested advertisements, as well as for each individual advertisement, a single strategy was more persuasive.

4.5 Conclusions

The studies in this chapter explored a number of questions that originate from a desire to build personalized adaptive systems—a desire that logically follows the findings presented in the previous chapter. The studies explored free user choice of usage of different influence strategies and the disclosure of the persuasive intent of implementations of influence strategies. Furthermore, they explored the implications of using multiple distinct influence strategies to advocate a single end. From Study 1 it is clear that designers of personalized persuasive systems can use

free-choice as a means of making their applications more persuasive. Users can choose the type of argument or social influence strategy that appeals to them most. However, when the persuasive intent of the social influence strategy is made salient persuasion is decreased. This latter finding highlights the importance for designers to be cautious about their presentation of social influence strategies.

Studies 2 to 4 show the importance of choosing the “correct” social influence strategy. The studies in Chapter 3 uncovered large heterogeneity at an individual level. This chapter adds to this finding by showing that using multiple social influence strategies—possibly appealing to different types of people—does not necessarily increase persuasion. Rather, a single influence strategy can be more persuasive. This finding was consistent even when two different social influence attempts had similar average effects: combining the strategies does not significantly increase persuasion. Study 4 showed that these results also hold in an applied context. In the presented studies, the combination of influence strategies converges in its persuasion to the least successful influence attempt. While these results do not show that every possible combination of influence strategies will always be less effective than the use of a single strategy, the results do make clear that designers of persuasive systems should not always expect additive effects of different influence strategies. The implications of these findings for the design of persuasive systems are further discussed in Chapter 6¹.

¹This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2011b, 2012).

5

Insight Generation III: *Meta-Judgemental Measures*

5.1 Introduction: Measures of Persuasion Susceptibility

The previous chapters have shown that (a) large differences between individuals exists in their responses to social influence strategies, and (b) that using multiple social influence strategies to support a single goal is not necessarily beneficial to increase compliance. These findings support the idea that social influence strategies as identified by Cialdini (2001) provide a useful taxonomy to distinguish different persuasive messages. This chapter explores whether the individual differences observed in responses to social influence strategies can be measured *a priori* using meta-judgmental measures.

Designers of persuasive technologies can use the results presented in the previous chapters in two ways: First, designers can aim to measure possible susceptibility to strategies before deploying the technology. Thus, based on questionnaire measures a profile of an individual user can be created and used to adapt message presentation of a persuasive system. Second, persuasive technologies can be designed that dynamically update the representation of persuasive messages based on a user's response to implementations of social influence strategies. This chapter

explores the first of these two options and presents the development of the *STPS: The Susceptibility to Persuasive Strategies Scale*.

The *STPS* is not the first instrument designed to assess *a-priori* people's possible responses to subsequent exposure to social influence strategies. Initially researchers working on dual-processing models of attitude change created measures to classify the tendency of people to either elaborate a lot (take the central route to persuasion) or elaborate a little (using the peripheral route). Those with a high tendency to elaborate—those high in the “Need for Cognition” (*NfC*)—are less susceptible to peripheral cues. Elaborate development of the scale shows that the 18-item version presented by (Cacioppo and Petty, 1982) is both internally consistent (reliable) and externally valid. Next to measures of the *overall* tendency to be persuaded by implementations of social influence strategies researchers have also developed strategy-specific measures: Cialdini et al. (1995) present the “Preference for Consistency (PFC)” scale and show that it adequately predicts responses to implementations of the consistency and commitment strategy (Nail et al., 2001).

In this chapter this previous work is extended by developing a scale, called the “Susceptibility to Persuasive Strategies scale (*STPS*)”, which addresses susceptibility to each of the six social influence strategies identified by Cialdini (2001) separately. Study 1 presents an initial attempt to develop such a scale. The 12-item scale that is developed is analyzed to determine its factor structure. Next, participants scores are related to behavioral responses to implementations of the consensus and reciprocity strategy. Study 2 further develops the scale by increasing the number of items and performing a more formal reliability and validity analysis. The 26-item scale that is presented in Study 2 has a clear six factor solution representing each of the social influence strategies. The variance explained by this solution is reasonable, and the overall low correlation to the *NfC* scale shows that the *STPS* partly measures a distinct trait. Participants scores on the *STPS* are finally related to responses to social influence strategies (as obtained in Study 1 of Chapter 3). The *STPS* measure significantly aids in explaining responses to implementations of social influence strategies.

5.2 Study 1: Initial application of the short *STPS*

As an initial test of the *STPS* an experiment was setup in which (a) susceptibility to each influence strategy was measured using two items,

and (b) participants behavioral response to a persuasive request that was supported by implementations of social influence strategies was observed. In this study a request is made to participants and a system delivers implementations of multiple influence strategies to try and increase compliance. The main aim of this first study was to show that people who are more susceptible to distinct influence strategies will comply to a greater extent with a request that is supported by an implementation of that influence strategy.

5.2.1 Methods Study 1

In Study 1 participants were asked to fill in an online questionnaire on their experience of their social relationships. The questionnaire consisted of 42-items and was constructed using items from the connectedness questionnaire as presented by (van Bell and Smulders, 2009) as well as 12 items measuring respondent's susceptibility to persuasion. After filling in the questionnaire participants were asked to provide the experimenter with email addresses of friends that might be willing to participate in the same study. This request was either supported (*S*) or not supported (*NS*) by implementations of social influence strategies.

Development of the Short STPS

To measure participants susceptibility to social influence strategies two items asking how likely participants would be to respond to an implementation of that strategy were developed. For each social influence strategy, reciprocity, scarcity, authority, commitment, consensus, and liking, two applied instances of implementations were created. Table 5.1 presents all of the 12 items used in this study grouped by social influence strategy. Participants were asked to specify the extent to which they agreed or disagreed with the statement. Response categories ranged from (1), *Completely agree* to (7) *Completely disagree*.

Conditions

To operationalize the supported (*S*) and not supported (*NS*) conditions the request to provide several email addresses was either *not* supported by any additional message, or was supported by two persuasive statements. The first implementation relied on the principle of consensus: *"All of the other participants provided several email addresses to us"*. The second implementation relied on the principle of reciprocation: *"In return for providing us with your friend's addresses, we will send you a copy of the results of our study"*. Both implementations were added

Table 5.1: The 12 susceptibility items—presented per influence strategy—that compose together the short-STPS.

Strategy	Item Name	Susceptibility item
Reciprocation	Recip_1	When a family member does me a favor, I am very inclined to return this favor.
	Recip_2	I always pay back a favor.
Scarcity	Scarce_1	I believe rare products (scarce) are more valuable than mass products.
	Scarce_2	When my favorite shop is about to close, I would visit it since it is my last chance.
Authority	Auth_1	I always follow advice from my general practitioner.
	Auth_2	When a professor tells me something I tend to believe it is true.
Commitment	Commit_1	Whenever I commit to an appointment I do as I told.
	Commit_2	I try to do everything I have promised to do.
Consensus	Cons_1	If someone from my social network notifies me about a good book, I tend to read it.
	Cons_2	When I am in a new situation I look at others to see what I should do.
Liking	Like_1	I accept advice from my social network.
	Like_2	When I like someone, I am more inclined to believe him or her.

in a clearly visible block right before the presentation of twenty empty text fields which could be used to supply the email addresses.

Participants

454 respondents were invited by email to participate in the study. Email addresses were taken from lists of respondents to previous research projects in the HCI domain. Participants' were asked for their age and gender at the end of the questionnaire. Out of the 454 initially invited participants 82 took part in the study leading to a response rate of 8.1%.

The email addresses collected during the study—those provided by the participants—were not used to invite subsequent participants. Of those who completed the study, 47 (57.3%) were male and 35 (42.7%) were female. The average age of the respondents was 37 years ($SD = 13.3$).

5.2.2 Results Study 1

Scale Reliability

The internal consistency of the responses to the 12-item susceptibility to persuasive strategies scale was examined using both reliability analysis as well as a principal component analysis. Cronbach's Alpha of the full scale was 0.609 which is relatively low. A Principal Component Analysis on the 12 items shows that, when rotating the solution and allowing correlations between the components, 5 components have Eigenvalues higher than one. The items created to address susceptibility to implementations of the *Liking* strategy and those created to address *Consensus* strategy are highly correlated and thus grouped into one component. The last *Liking* item loads correlates both to the *Consensus* as well as the *Scarcity* items. The rotated component loadings are presented in Table 5.2.

Cumulative, the five components explained 71.6% of the total item variance. An average score of all susceptibility items was computed for each respondent to indicate *overall* susceptibility. Furthermore, scores were computed for the two relevant sub-scales (susceptibility to consensus and susceptibility to reciprocity) by averaging the score on the two items relating to a specific strategy. In both cases the two respective items correlated highly ($r = .620, p < .001$, Cronbach's $\alpha = 0.77$ for the *Reciprocity* items and $r = .672, p < .001$, Cronbach's $\alpha = 0.80$ for the *Consensus* items).

Effects on Compliance

The scores on the compliance measure—the number of provided email addresses—were examined for normality. The Kolmogorov-Smirnov statistic and the Shapiro-Wilk statistic clearly showed that deviation from the normal was significant ($KS = .319, p < 0.001$; $KW = .651, p < 0.001$). Based on this preliminary analysis non-parametric statistics were used to test the effects of the *S* and *NS* conditions on compliance.

A Mann-Whitney U test showed that the number of email addresses provided in the *S* condition ($Meanrank = 50.35$) was significantly higher than that in the *NS* condition ($Meanrank = 31.74, p < 0.001$).

Table 5.2: Component loadings based on a principal components analysis with oblimin rotation for the 12 items of the short-*STPS*. Loadings smaller than .3 are suppressed.

	C. 1	C. 2	C. 3	C. 4	C. 5
Recip_1			.854		
Recip_2			.880		
Scarce_1	.722				
Scarce_2	.819				
Auth_1				.717	
Auth_2				.775	
Commit_1					.850
Commit_2					.850
Consens_1		.895			
Consens_2		.886			
Like_1		.402			
Like_2	.657	.403			

Furthermore, there was a significant positive relationship between individual's overall susceptibility to persuasive strategies and the number of email addresses provided ($\rho = .227$, $p < .05$).

To test whether the scores on the sub-scales of the *short-STPS* were predictive of the response of participants to a request supported by implementations of the two strategies used in this experiment the relationship between participants score on the sub-scale and their response to the request was further examined. Participants score on the susceptibility to consensus strategy sub-scale correlated highly with their response to the persuasive request when the request was supported by this strategy: $\rho = .672$, $p < .001$. This correlation was much lower for participants in the *NS* condition, $\rho = .336$, $p < 0.05$. This latter low correlation is expected since the consensus strategy was not shown in the *NS* condition. This interaction between the conditions and the scores on the *STPS* is statistically significant, $p < .05$ ¹.

For the reciprocity strategy, both in the *S* as well as the *NS* conditions, no significant correlation between participants score on this sub-scale of the *STPS* and their subsequent compliance was found. The fact that, for the reciprocity strategy, the scores on the short-*STPS* do not relate to the responses to this strategy can be explained by the way

¹This difference in ρ 's was tested using the *V*-statistic as presented by Hays (1973).

the reciprocity strategy was implemented: a promise to reciprocate often has a much weaker effect than providing an actual favor in advance (Cialdini, 2001).

5.2.3 Discussion Study 1

Study 1 provides initial evidence for the usefulness of the *STPS* by showing that compliance to a persuasive request supported by an implementation of the consensus strategy is highly correlated to participants self-reported susceptibility to this scale. However, the study also highlights that the two items per strategy as presented here might be too limited to properly differentiate between each of the six social influence strategies: the correlations between the Liking and the Consensus items were very high². Also, for the second social influence strategy that was implemented to increase compliance—reciprocity—no relationship between self-reported susceptibility and behavioral response was found. While the results obtained in this initial study are thus mixed, the high correlation between the behavioral measure and the meta-judgmental score on the consensus items inspired further investigation of the opportunity to measure the susceptibility of respondents to different social influence strategies using a questionnaire.

5.3 Study 2: Development and Validation of the *STPS*

Based on the results of Study 1 the *STPS* was further developed. Via structured brainstorms a more elaborate item set was developed and tested. By creating four to six items measuring participants susceptibility to each of the social influence strategies identified by Cialdini (2001), the internal validity of each of the sub-scales is increased. Next, the scores on each of the sub-scales are used to explain responses to three of the six social influence strategies separately.

5.3.1 Item Construction

Starting from the 12 item measure used in study 1 additional items were created for each of the six latent variables of interest. In a session with a group of five persuasive technology researchers eight or more items per variable were composed. Items were constructed to fit the underlying latent variable as much as possible and to appeal both to specific instances of the influence strategy (e.g. “I always follow advice

²The alternative explanation that the correlation between *Liking* and *Consensus* is so strong in actuality that one will not be able to design a scale that measures them separately is proven wrong in Study 2.

from my general practitioner”) as well as to broad statements of the latent variable of interest (e.g. “I am very inclined to listen to authority figures”). In a pre-test with $N = 9$ participants the understandability and clarity of each of the items was evaluated and per variable 5 – 6 items were selected for further testing. The complete item set used for evaluation is presented in table 5.3.

Table 5.3: Items used in the development of the *STPS*. Items marked with * are selected after the principal component analysis.

Principle	Abbreviation	Susceptibility item
Reciprocity	Recip_1*	When a family member does me a favor, I am very inclined to return this favor.
	Recip_2*	I always pay back a favor.
	Recip_3*	If someone does something for me, I try to do something of similar value to repay the favor.
	Recip_4*	When I receive a gift, I feel obliged to return a gift.
	Recip_5*	When someone helps me with my work, I try to pay them back.
Scarcity	Scarce_6*	I believe rare products (scarce) are more valuable than mass products.
	Scarce_7*	When my favorite shop is about to close, I would visit it since it is my last chance.
	Scarce_8*	I would feel good if I was the last person to be able to buy something.
	Scarce_9*	When my favorite shampoo is almost out of stock I buy two bottles.
Authority	Scarce_10*	Products that are hard to get represent a special value.
	Auth_11	I always follow advice from my general practitioner.
	Auth_12	When a professor tells me something I tend to believe it is true.
	Auth_13*	I am very inclined to listen to authority figures.

... continues on next page

Table 5.3: Items used in the *STPS* (...continued)

Strategy	Abbreviation	Susceptibility item
Commitment	Auth_14*	I always obey directions from my superiors
	Auth_15*	I am more inclined to listen to an authority figure than a peer.
	Auth_16*	I am more likely to do something if told, than when asked.
	Commit_17*	Whenever I commit to an appointment I always follow through.
	Commit_18*	I try to do everything I have promised to do.
	Commit_19*	When I make plans I commit to them by writing them down
	Commit_20	Telling friends about my future plans helps me to carry them out
	Commit_21*	Once I have committed to do something I will surely do it.
Consensus	Commit_22*	If I miss an appointment, I always make it up.
	Consens_23*	If someone from my social network notifies me about a good book, I tend to read it.
	Consens_24*	When I am in a new situation I look at others to see what I should do.
	Consens_25	I will do something as long as I know there are others doing it too.
	Consens_26*	I often rely on other people to know what I should do.
Liking	Consens_27*	It is important to me to fit in.
	Like_28	I accept advice from my social network.
	Like_28	When I like someone, I am more inclined to believe him or her.
	Like_28*	I will do a favor for people that I like
	Like_28*	The opinions of friends are more important than the opinions of others.

... continues on next page

Table 5.3: Items used in the *STPS* (...continued)

Strategy	Abbreviation	Susceptibility item
	Like_28*	If I am unsure, I will usually side with someone I like.

5.3.2 Scale Validation

To determine the internal validity of the *STPS* it was administered to $N = 215$ participants. Participants were under-graduates enrolled for a research methodology course, and the *STPS* was administered as part of the intake-questionnaire required for students who intended to sign up for the course. All participants filled out both the 32 items of the susceptibility scale as well as the 18 item Need for Cognition scale Cacioppo et al. (1986). Need For Cognition (*NfC*) was included to also assess the external validity of the overall scale by comparison to a known construct.

Method

The items were administered using a 7-point scale. The end-points of the scale were labeled with “Completely Disagree” and “Completely Agree” where a score of 7 marked “Completely Agree”. The mid-points of the scale were not labeled. Participants were provided with a “Don’t know” option for each item. The average age of the participants was 23.0 years ($SD = 7.4$). Of the participants 167 were female (75.9%). Participants filled out the questionnaire online using their own computer after receiving an email with a link to the study.

Results

Similar to Study 1, a principal component analysis was conducted to examine the internal consistency of the scale. Figure 5.1 shows the eigenvalues of all the components with a line superimposed to indicate the clearest cut-off. It is clear that, contrary to Study 1, a six component solution provides an appropriate fit for these data.

The total cumulative variance explained by all six components was 52%, which is in a common range for multidimensional constructs. The six component solution was further examined using Oblimin rotation. Table 5.4 gives an overview of the loadings of each of the items on the components. Because of relatively low loadings, or high cross-loadings, it was decided to remove six items from the scale: *Auth_11*,

Figure 5.1: Scree plot showing the eigenvalues (y-axis) of each of the extracted components (x-axis). A clear increase is visible from component 6 upwards indicating a proper fit of a six factor solution.

Auth_12, Commit_20, Consens_25, Like_28, and Like_32. Refitting the six component solution to the 26-item scale led to a cumulative variance explained of 56%, and no cross-loadings of factors over .3. In table 5.3 the 26 items used to compute final scores on the six latent variables of the *STPS* are marked with an *.

For each of the six components a composite score was computed by averaging over the 3–5 items in each sub-scale. Table 5.5 presents each of the sub-scales with their appropriate descriptives. Overall, this analysis indicates that the six factors underlying the *STPS* are moderately internally consistent. The correlations between the scores on the sub-scales range from .2 to .4. As a check of the external validity of the scale we examined the correlation between the total composite score of participants on the *STPS* and their score on *NfC* (Cronbach's $\alpha = 0.89$): The correlation between these two scales was -0.14 , $p < 0.05$.

The sub-scales of the 26-item *STPS* are moderately internally reliable, and the correlations between the separate sub-scales are relatively low, indicating that the *STPS* indeed measures people's susceptibility to 6 distinct strategies. The internal reliability of a number of the sub-scales is questionable—particularly that of the *Scarcity*, *Consensus* and *Liking* strategies. This is probably due to the nature of the items,

Table 5.4: Component loadings based on a principal components analysis with oblimin rotation for 32 items of the *STPS*. Loadings smaller than .3 are suppressed.

	C. 2	C. 3	C. 1	C. 4	C. 6	C. 5
Recip_1			0.34			
Recip_2			0.72			
Recip_3			0.73			
Recip_4			0.67			
Recip_5			0.67			
Scarce_6					0.77	
Scarce_7					0.33	
Scarce_8					0.44	
Scarce_9					0.34	
Scarce_10					0.84	
Auth_11	0.37	0.45				
Auth_12		0.44	0.40			
Auth_13		0.73				
Auth_14		0.71				
Auth_15		0.75				
Auth_16		0.59				
Commit_17	0.84					
Commit_18	0.63					
Commit_19	0.51					
Commit_20						
Commit_21	0.80					
Commit_22	0.76					
Consens_23				0.54		
Consens_24				0.53		
Consens_25		0.40		0.42		
Consens_26				0.63		
Consens_27				0.53		
Like_28				0.64		
Like_29						0.58
Like_30						0.70
Like_31						0.51
Like_32				0.53		

which address both specific persuasion attempts as well as more general tendencies to comply to the different strategies. An examination of this

difference in the specificity of the items, and especially the ability of participants to adequately judge their own responses to distinct influence strategies in different situations, should be a subject for further investigation.

Table 5.5: Overview of the composite scores of the *STPS*. Presented are the mean, standard deviation, and Cronbach's α of each of its sub-scales.

Sub-scale	# items	Mean (SD)	Cronbach's α
Reciprocity	5	5.3 (0.83)	0.75
Scarcity	5	4.7 (0.98)	0.63
Authority	4	4.3 (1.10)	0.75
Commitment	5	5.1 (0.97)	0.81
Consensus	4	4.1 (0.98)	0.60
Liking	3	5.1 (0.91)	0.61
Total Composite	26	4.8 (0.59)	0.85

The negative correlation of the composite *STPS* with *NfC* shows that those high in overall susceptibility to persuasion score low on *NfC*. This is expected since a higher tendency to elaborate would lead to a lowered susceptibility to strategies that function via the peripheral route which is likely at least partly the case for the strategies addressed by the *STPS*.

Overall, the above analysis shows that the *STPS* presents a moderately internally reliable scale to measure participants susceptibility to *distinct* influence strategies.

Usage of the *STPS* to Predict Responses

To examine the external validity of the scores obtained using the *STPS* the results presented in Study 1 of Chapter 3 are further analyzed by using the scores obtained on the *STPS* by participants in this study. The *STPS* was administered two weeks prior to participation and was part of a inclusion questionnaire which was one of the requirements of the research methods course from which the participants were recruited. Table 5.6 shows the model comparison of the preferred model presented in the analysis of Study 1, chapter 3 (Model *C*) to a larger model that includes participants scores on the *STPS* on the sub-scales Authority, Consensus, and Scarcity interacting with the strategy that was used (Model *C_{Suscept}*).

Table 5.6: Model comparisons of the preferred model of Study 1, Chapter 3, and a model with interactions of *STPS* scores on the appropriate strategies and the strategy in use.

	df	BIC	logLik	χ^2	p
Model C :	17	23390.42	-11621.67		
Model $C_{Suscept}$:	29	23462.78	-11605.94	31.46	< 0.001

A further examination of model $C_{Suscept}$ shows the effects of each of the *STPS* measures: The coefficient of the *STPS* measure of Authority interacting with the Authority strategy is positive, $\beta = 0.37$, $t = 2.89$, $p < .01$. The same is true for the Scarcity coefficient, $\beta = 0.30$, $t = 2.63$, $p < .05$. These results show that—as expected—people’s responses on the *STPS* are indeed positively related to their responses to the respective strategies. Thus, participants who scored higher on their self reported susceptibility to (e.g.) the authority strategy, provided a more positive evaluation of products supported by this strategy. The Consensus coefficient is also estimated positive, $\beta = 0.14$, but is not significantly different from zero, $t = 0.88$, $p > .05$. This latter relationship might be non significant due to the low internal reliability of the Consensus sub-scale: this creates a potentially noisy estimate of participant’s true susceptibility to the consensus strategy.

5.3.3 Discussion on the development of the *STPS*

Study 2 extended the initial development of the *STPS* presented in Study 1. Finally, a 26-item scale (as presented in Table 5.3) was developed to measure participants susceptibility to distinct social influence strategies. The scale brought forward has sufficient internal validity, as shown by its clear 6-component solution and the obtained reliability coefficients. Furthermore, scores on sub-scales of the *STPS* significantly improve prediction of the responses to persuasive appeals that are supported by implementations of social influence strategies. Thus, those who score high on (e.g.) susceptibility to Authority—as measured using this four item sub-scale of the *STPS*—comply more to persuasive requests that are supported by implementations of the Authority strategy.

5.4 Conclusions

Two studies explored the possibility of using meta-judgmental measures of susceptibility to influence strategies to explain subsequent behavioral responses to requests supported by social influence strategies. In study 1, susceptibility to the consensus strategy as indicated on the short-*STPS* related positively to compliance to an implementation of this strategy. In Study 2 the *STPS* was further developed by adding extra items to each sub-scale. In a subsequent evaluation, the full *STPS* proved useful in explaining responses to influence strategies at a later point in time.

While of some success, the results of the studies also show that predicting responses to social influence strategies based on meta-judgmental measures is not straightforward. In Study 1 no relationship between self-reported susceptibility to the reciprocity strategy and the actual behavioral response to this strategy was found. In this study this can be explained by the findings presented in Chapter 4: only one of the social influence strategies that was implemented appealed to participants and led to the increase in compliance. However, other explanations are also plausible: Reciprocity was not implemented properly (see e.g. Cialdini, 2001), or the short version of the *STPS* used in this study fails to properly measure participant's susceptibility to this strategy.

During the advanced scale development in Study 2 it was clear that for a number of sub-scales of the *STPS* the internal reliability was relatively low ($0.6 < \alpha < 0.7$). This can have several origins, two of which are expected based on previous work.

First, there is a relatively large body of literature that shows that people are generally not very capable of judging their own responses to persuasion (e.g. Bassili, 1996). This would imply that estimates based on self-report are noisy, and thus have a low internal reliability. This first is hard to address within the questionnaire itself—although methods for improving ones own meta-judgments by aiding recall of specific situations do exist (Bowling, 2005)—and operative measures of susceptibility could prove more externally valid than meta-judgmental ones (See also Chapter 8).

Second, the *STPS* queries susceptibility based on both very general statements, as well as specific persuasion attempts. These differences in specificity between the items might cause different judgments within the same sub-scale and thus might lower the internal reliability (e.g.

Froggatt, 1969; Aleamoni and Thomas, 1980). This latter issue could be addressed in further development of the *STPS* scale by explicitly examining items of general nature and those addressing specific persuasion attempts. This could improve the reliability of the scale at both levels separately.

The negative relationship between the *STPS* and *NfC* shows that, partially, the findings presented above might be explained by differences in elaboration between people that are addressed by the *NfC* scale. However, the correlation of the *STPS* to *NfC* is low, indicating that at least partially a different mechanism is addressed. This is further strengthened by the fact that different sub-scales of the *STPS* correlate to the responses of people to distinct influence strategies. Thus, the *STPS*—while capturing partly a predisposition to elaborate—also proved be useful in predicting the responses to *distinct* influence strategies. Given that the responses to the *STPS* are meta-judgments, and thus likely occur under central processing, the partial functioning of the influence strategies via the central processing route is probably also captured. As noted in 4.5 influence strategies can be effective via *both* routes, and thus the focus on influence strategies, as opposed to the processing mechanism itself, can be valuable to predict the responses to influence strategies in practice.

While the *STPS* significantly increased model fit for the data collected in Study 1 of Chapter 3 and thus aids in explaining the responses to different persuasion attempts, the effects are only moderate. Thus, people’s judgments of their own responses to influence strategies did not correspond one-to-one with their *actual* responses to the influence strategies that were used. This could be caused—as discussed above—by the fact that people are not always able to accurately assess their own susceptibilities. However, it could also be due to the different levels of processing between the actual persuasion attempts and the act of filling out the questionnaire. It is likely that the central processing route (Petty and Cacioppo, 1986; Petty and Wegener, 1999) is most prominent when filling out a questionnaire, while the influence strategies that are used function at least partially via more peripheral, and less conscious, processing.

The development of the *STPS* as presented in this chapter strengthens the overall idea that social influence strategies are useful as a level of analysis to identify differences between individuals. The *STPS* can aid designers of persuasive systems to tailor their influence attempts to their users. By measuring user susceptibilities to distinct influence

strategies and adapting strategy selection accordingly the effectiveness of persuasive systems is probably increased. However, the *STPS* should be regarded a starting-point rather than an end-point in the process of personalizing persuasive attempts³.

³This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2009a, 2011a).

6

Design Requirements

6.1 Conclusions From the *Insight Generation* Chapters

The three *insight generation* chapters, Chapter 3, 4, and 5, examined questions that are of key importance for designers of persuasive systems but have largely been neglected by social science researchers:

1. How large is the heterogeneity in responses to different ways in which persuasive requests are framed¹?
2. How do people respond to multiple strategies (ways) that support the same request?
3. Can we measure—using questionnaires—people’s susceptibility to different ways in which requests are framed?

In the current chapter the answers to these questions are used to draw three conclusions that inform the design of ambient persuasive systems. This chapter argues that designers of adaptive persuasive systems should *attend to individual differences* in responses to social influence strategies, should *actively choose strategies*, and should create *persuasion profiles* to manage and use the results derived in the previous chapters.

¹Only if this heterogeneity is large it makes sense for designers to attend to it in their design of persuasive systems.

After motivating these three conclusions three key implications for the design of ambient persuasive systems are derived: If designers want to attend (dynamically) to the individual differences that are uncovered in the previous chapters then their systems should have the ability to *identify users, represent different social influence strategies, and measure their effectiveness*. Finally, this chapter details the conceptual implementation of an adaptive persuasive systems.

6.1.1 Conclusion 1: Designers of Persuasive Systems should attend to Individual Differences in Responses to Persuasive Strategies

The three studies presented in Chapter 3 demonstrated the difference in magnitude of the average effect of influence strategies and the individual level (or conditional) effects. While the average effects have been of importance for the optimization of persuasive systems, persuasion in ambient intelligence would clearly benefit from a personalized approach: one in which the conditional effects are attended to.

The average treatment effect (ATE) that is estimated in the most noticeable studies on persuasion, such as Milgram (1974)'s early studies on authority and Cialdini (2005)'s studies on the effects of consensus, can be large: in Cialdini (2005)'s study on the use of paper signs to encourage towel re-usage implementations of social influence strategies led to an increase in re-usage of 26.3%: A strong and important effect that inspired both researchers and practitioners. In Study 1 of Chapter 3 a similar result is found: usage of both authority and consensus strategies on average significantly increase people's evaluation of products. However, Study 1 of Chapter 3 also enabled estimation of the effect of different strategies *conditional* on individuals. For the first time the ATE could be contrasted to the conditional effects, as done in figure 3.1. This direct contrast shows the relative importance of the individual level effects as opposed to the ATE's: While the ATE's might be statistically significant, the size of these effects is small compared to the differences in responses between individuals. The ATE as such is a bad predictor of responses of individuals to influence attempts: Even for the strategy with the largest positive ATE usage of this ATE to predict an individual level response would have the wrong sign (positive vs. negative) for 41% of the users of a persuasive system.

It is important to consider the different effects that can be observed when looking at the effects of social influence strategies. Researchers have classically attended to the main effect of *Strategy*—the ATE. The

studies in the first Insights Generation chapter, Chapter 3, focussed mainly on a *Strategy* $\times$ *Person* interaction: the effect conditional on the individual. However, there will probably also be *Strategy* $\times$ *Context* and *Strategy* $\times$ *Time* interactions: For some persuasive technology application areas one social influence strategy might be more important than another and this might differ over time. However, as long as the three-way *Strategy* $\times$ *Person* $\times$ *Context, or Time, or...* interaction, is smaller than the *Strategy* $\times$ *Person* interaction designers can meaningfully attend to the conditional effects across contexts and time points.

The findings in Chapter 3 indicate that while designers of persuasive systems should consider usage of social influence strategies—both due to their ATE’s as well as their individual level effects—they should not expect a homogeneous effect on their users. While averaged over user groups systems that use social influence strategies are more effective, which explains their prominence on e-commerce and other marketing related systems, they are not necessarily more effective for *individual* users and might even have a negative effect. Since the persuasive technology field is moving more and more towards *individual* lifestyle or *individual* energy consumption change it is evident that the individual differences should be attended to. If designers of ambient persuasive systems want to deliver on their promise to change an individual behavior they should mind the conditional effects. An increase in this conditional effect will eventually also increase the ATE’s, however increasing the ATE itself should not be the main aim of designers.

6.1.2 Conclusion 2: Designers of Persuasive Systems should Select Persuasive Strategies

Influence professionals have all through history attended to individual differences both in the ends—the goals—of their requests as well as the means. Persuasive technologies however have mainly been focussed on adapting the end goal of their persuasive request to individuals while using the same or similar means. The results presented in Chapter 3 clearly indicate that the individual level effects of the ways in which request are made are important and should be attended to. From this conclusion it is straightforward to ask whether persuasive systems should combine multiple social influence strategies to appeal to each user, or whether a selection should be made for individual users.

The studies presented in Chapter 4 directly address these issues. Study 2 to 4 examined different combinations of social influence strate-

gies used to support a single appeal and measure compliance. In different contexts, and with different *Strategy* $\times$ *Context* effects, it is clear in each of these studies that a “well chosen” social influence strategy ensures higher compliance than a combination of social influence strategies. Even when two strategies have similar expected ATE’s (Study 3) their combination does not increase compliance. However, when one strategy has a small ATE (Study 2) this decreases the persuasiveness of other social influence strategies leading to lower overall compliance both for individuals as well as on average. This was true not only in an experimental setting, but also in a real-life application (Study 4).

Two plausible explanations for the observed effects of the simultaneous usage of multiple influence strategies exist. First of all, usage of multiple strategies probably increases elaboration and as such leads to more central information processing. The effect of peripheral cues—such as the use of influence strategies—will likely be smaller in high elaboration than in low elaboration. A second explanation is provided by the work on the negativity bias (Rozin and Royzman, 2001): People pay more attention to negative events than positive events. Thus, if one of the influence strategies does not appeal to a user and as such has a negative effect, that negative effect will overshadow any positive persuasion gained from well chosen social influence strategies.

Combining the results of the three studies on simultaneous usage of influence strategies (2 to 4) with the results obtained in Chapter 3 makes clear the need for designers of persuasive technologies to select social influence strategies specifically for individual users. The failure to do so will lead to lowered compliance, and in the case of applications designed for *individual* lifestyle or *individual* energy consumption change can lead to aversive effects. Study 1 presented in Chapter 4 aids designers partially in how to choose persuasive strategies and highlights the design challenges involved: While users are able to identify their preferential strategy by choosing a “source of advice”, disclosure of the persuasive attempt of such a source seems to reduce compliance. Thus, designers can actively involve users in the choice of persuasive strategies but caution needs to be given to the implementation of such a choice scheme.

6.1.3 Conclusion 3: Designers should Create and Use Persuasion Profiles

When systems represent individual differences as variation in responses to influence strategies and adapt to these differences, they are engaging

in *persuasion profiling*. Persuasion profiles are collections of estimates of the expected effects of different influence strategies for a specific individual. Hence, an individual’s persuasion profile indicates which influence strategies—one way of individuating means of attitude and behavior change—are expected to be most effective. Persuasion profiles can potentially be constructed based on several metrics, including demographics, personality measures (meta-judgmental measures), and behavioral data (operative measures) or combinations of these. 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 are currently collected. These data describe how people have responded to presentations of certain products (e.g. e-commerce) or have complied to requests by persuasive technologies (e.g. the DirectLife Activity Monitor (Lacroix et al., 2009)).

Figure 6.1: Example of a persuasion profile. Dots represent the estimated effect of the respective persuasive strategies, while the bars represent the certainty around this estimate.

Figure 6.1 shows an example of a persuasion profile. The profile shows the estimates of the effects of different social influence strategies on compliance and the certainty around these estimates. Thus, for this user the implementations of the consensus strategy are most efficient. Implementations of the authority strategy are least efficient for this user however the estimate of the effect of this strategy is relatively uncertain. An *accurate* persuasion profile ensures that designers can attend to individual differences and can choose social influence strategies.

The *STPS* presented in Chapter 5 presents a validated 26-item scale

to determine people susceptibility to different social influence strategies *a priori* using meta-judgmental measures. The scores on the *STPS* indicate people’s susceptibility to each of the six social influence strategies identified by Cialdini (2001) and as such can be directly used by designers of persuasive systems (or by the systems themselves) to attend to individual differences and choose social influence strategies. In the following *case study* chapter (Chapter 7) the applied value of profiles based on measurements obtained using the *STPS* for health related interventions is assessed.

Next to using meta-judgmental measures to build a persuasion profile, the profile can also be build, or updated, by observing behavioral responses of users to different social influence strategies and thus obtaining operational measures. This approach allows designers to create adaptive persuasive systems that dynamically adapt their selection of influence strategies to responses to persuasion over time (See Chapter 8). Persuasive Technologies likely benefit from an approach in which both sources of information about users are combined to obtain accurate conditional estimates.

Persuasion Profiles, employed by interactive systems, can be of use to create computer-tailored interventions as described in the nutrition education literature. However, the focus of persuasion profiles is narrower: where much of the computer-tailored intervention literature focuses on personalization, feedback, and adaptation (Dijkstra, 2005) — and sometimes *feedback* includes the use of influence strategies (e.g. Campbell et al., 1999) — persuasion profiles focus *only* on the use of influence strategies. This narrower focus makes that the “rules” behind the selection of strategies do not have to be governed by an expert system using “if-else” statements — as is frequently true for more general computer-tailored interventions (Brug et al., 2003) — and can be learned dynamically based on observations of user behavior (see 6.3). The choice for influence strategies as a level of personalization is motivated by the large heterogeneity demonstrated in Chapter 3.

6.2 Identification, Representation, and Measurement

When designers attend to the conclusions presented above then they will design adaptive persuasive systems: “*systems that select the appropriate influence strategy to use for a specific user based on the estimated success of this strategy.*” To be able to build such adaptive persuasive systems (i.e., systems that use persuasion profiles to adapt their influ-

ence strategy usage to individual users), designers should create systems that are capable of identifying their users, representing different social influence strategies, and measuring their effectiveness (Kaptein et al., 2011a; Kaptein, 2011b; Sakai et al., 2011).

6.2.1 Identification

To be able to adapt to individual differences in responses to social influence strategies, a system must be able to identify individuals (be it just by a unique key). Only once a user has been identified can the social influence strategy that is used to support a persuasive appeal be adapted to this user. Currently, many means of identification exist: In an online marketing contexts cookies and IP-addresses are already frequently used to tailor appeals and this usage can easily be extended to include persuasion profiles. However, in an ambient intelligence scenario the possibilities of identification are more diverse: Designers have used the unique bluetooth key that is used by mobile devices (Kostakos, 2008), face recognition (Hazen et al., 2003) or fingerprints (Caplan, 1990) to identify individual users. When such identification mechanisms are combined individual users can be tailored to, both offline as well as online, and persuasion profiles can be used over a multitude of persuasive applications.

6.2.2 Representation

Adaptive persuasive systems need to be able to implement various social influence strategies. 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 different ways to users. In the system architecture designers should distinguish social influence strategies, and their respective implementations. Thus, if a persuasive system uses the authority strategy then still different expert sources could be used, via different communication channels, to influence users. In each case, the authority strategy is represented by a different implementation. Previous persuasive technologies largely utilized a single (or a limited set) of implementations of social influence strategies to change user attitude or behavior. Mostly, the social influence strategies employed used to be selected based on (large) reported ATE's in the social science literature and were thus static. To enable representation of *different* social

influence strategies for different users the system infrastructure should change to facilitate such representations.

6.2.3 Measurement of success

When designers create systems that do not merely used a persuasion profile based on *a priori* measurements of the effectiveness of social influence strategies as obtained using the *STPS* but rather adapt to user’s responses *dynamically*—e.g. during the usage of the product—it needs to be possible to measure the effect of an influence attempt. While this sounds straightforward it is not always easy to measure whether an appeal was successful, or even what a measure of success would entail. For example, in a digital exercise coach a prompt by a fitness instructor to run for 30 minutes that is followed by the user running for 20 minutes 14 hours after the prompt might constitute a partial success—indicating the success of the authority strategy—but might also be due to external causes. Furthermore, technologically not all behavioral responses are easily or reliably measured.

Historically designers of persuasive technologies have not always quantified the outcomes of their interventions. Including *measurement* as a core requirement of persuasive systems will have a large impact on the design of, and our ability to evaluate, persuasive technologies.

6.3 Dynamic Adaptation

Once the three prerequisites identified above are met, and thus a persuasive system is able to identify its users, represent different social influence strategies, and measure the effect of the implementations of different strategies, the system can be made to adapt to user responses. While different machine learning algorithms could be used for such a goal, this section presents a relatively simple self-learning system² capable of using persuasion profiles by considering an example in which identification, representation, and measurement are relatively easy.

Consider an ambient persuasive system designed to encourage users to save energy by using a revolving door (which keeps the heat in) instead of a sliding door that is next to it (See Figure 6.2). This setup is common in hotels and office buildings and often one can find a paper sign motivating visitors to indeed take the revolving door. Designers could use several technologies (e.g. face recognition or Bluetooth scanning,

²The system presented here is called “relatively simple” since the algorithm does not include (e.g.) correlations between strategies or possible uneven loss functions.

see Chapter 8) to identify visitors when they enter the hotel lobby. The same technology can also be used to assess whether a visitor took the revolving door or the sliding door. Based on this knowledge about the visitor and records of earlier decisions a message implementing the *right* influence strategy can be selected and displayed on a screen instead of the paper note.

Figure 6.2: Graphical representation of a user entering the building (from below). The user is identified, and a message is displayed on the screen. Next, the user chooses between the sliding door (left) and the revolving door (right).

Suppose there are only two messages to display, one implementing the authority strategy—“*The general manager of this hotel urges you to...*”—and one implementing the consensus strategy—“*80% of our visitors always use...etc.*”. The system then needs a mechanism to choose the message that is most likely to be effective for the current visitor. It is intuitive that for a *new* visitor the system should present the message which has lead to the highest compliance for other, previously observed, visitors. Hence, this would be the message with the highest average effect. If this message is successful then there is no need to try different messages on subsequent visits. However, when the selected message is not effective, it might become attractive to present another message on a subsequent visit. This decision depends on the initial success probabilities of the messages under consideration, the variance of effectiveness of messages between visitors, and the number of success or failures observed for the current visitor.

The probability of a single visitor taking the revolving door on multiple occasions can be regarded a binomial random variable $B(n, p)$ where n denotes the number of approaches the visitor has made to the doors and p denotes the probability of success: the probability of taking the revolving door. Given M different messages one can compute for each

individual, for each message, probability $p_m = k_m/n_m$ where k_m is the number of observed successes after representation of message m , n_m times to a specific visitor. It is reasonable to present an approaching visitor with the message with the highest probability of success, thus the message where p_m is highest.

For a large number of observations N of one visitor this would make perfect sense. However, this will not inform a decision for a newly observed visitor. For a *new* visitor one would present the message m for which p_m is maximized for previously observed visitors³. Actually—given Stein’s result (Stein, 1955; Efron and Morris, 1975)—for every user a weighted average of the p_m for an individual user and those of other users—one where the estimated $\hat{p}_m$ for an individual is “shrunk” toward the population mean—will provide a better estimate than an estimate based on observations of a single visitor alone. E.g., if the authority message is effective 70% of the time over all visitors and only 30% percent of the time for the specific visitor under consideration, the best estimate of the (real) effectiveness of the authority message $\hat{p}_A$ for this visitor is a weighted average of these two.

6.3.1 Adapting to individual behavior using Stein Estimation

To include both the known effectiveness of a message for others, and a specific visitors previous responses to that same message, into a new estimate of message effectiveness, p_m , designers can use a Bayesian approach. A common way of including prior information in a binomial random process is to use the Beta-Binomial model (Wilcox, 1981). The $Beta(\alpha, \beta)$ distribution functions as a conjugate prior to the binomial. The beta distribution can be re-parametrized as follows

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

where $\mu = \frac{\alpha}{\alpha+\beta}$ and $M = \alpha + \beta$. The expected value of the distribution is then given by: $E(\theta|\mu, M) = \mu_m$. In our specific scenario, μ_m represents the expected probability of a successful influence attempt by a specific message. The variance of this estimated success probability is given by:

$$Var(\theta|\mu, M) = \sigma^2 = \frac{\mu(1-\mu)}{M+1}.$$

After specifying the probability of success μ_m of message m and the certainty about this estimate σ_m^2 as the prior expectancy about the ef-

³This is assuming the error costs—the effects of presenting the wrong message—are equal for each message.

fectiveness of a specific message and updating this expectancy by multiplying it by the likelihood of the observations one obtains the posterior expectation

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

in which $k \in \{0, 1\}$, is the outcome of the new observation. The newly obtained Beta distribution, $B(\mu, M)$, functions as our probability distribution of the estimated success with a new point-estimate of the effectiveness of the presented message given by

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

In this way a persuasive system can keep estimates of the success of influence attempts for different social influence strategies on individual users—a persuasion profile—and update these estimates based on subsequent observations.

6.3.2 Decision Rule to Choose a persuasive Strategy

The Beta-Binomial model described above allows estimation of the effectiveness of message m , include prior knowledge, and updating these estimates based on new observations. As such one can maintain a record of both the point estimate, μ_m , and its certainty, σ_m^2 , for each specific visitor. To determine which message to present next, one could pick the message which has the highest μ_m . However, if σ_m^2 is large this decision rule might not be feasible given that—from a frequentist perspective—the difference between effectiveness estimates might not be statistically significant. To address this we can choose to show the message with the highest estimate when this estimate is “certain enough”—in the binomial case only once sufficient observations have occurred. In uncertain situations we can randomly present one of the H messages which have the highest estimates out of the total set of point estimates of M messages whose confidence intervals overlap. This decision rule would avoid presenting each new visitor with only the single most effective message when responses to messages are very variant.

Because the Beta distribution is not necessarily symmetrical the variance σ_m^2 provides an inadequate starting point to compute confidence intervals. This problem can be solved using simulation methods: By generating a number of draws (say 10,000) from the specified Beta distribution and computing (e.g.) the 20th and 80th percentiles one can

compute a more appropriate confidence interval. The above described decision rule for $M = 2$ would can be described as:

$$M_{selected} = \begin{cases} 1 & \mu_1 > Perc(80)_2 \\ 2 & \mu_2 > Perc(80)_1 \\ Rand(1,2) & \text{otherwise.} \end{cases}$$

Thus, if the estimated effectiveness of a message 1, $\hat{p}_1 = \mu_1$, is higher than the 80th percentile of message 2, $Perc(80)_2$, the system presents message one⁴. If the confidence level of one of the two messages overlaps with the highest scoring point estimate the system randomly presents one of the two messages and observes the response.

Another, more recent, solution to this selection problem—which is analogous to the multi-armed bandit problem in mathematics—is presented by Scott (2010). His proposed *randomized probability matching* depends on obtaining a single draw of each of the Beta distributions for each strategy and comparing theses draws. At a specific occasion the strategy representing the highest draw will be shown. Scott (2010) shows through simulation that this strategy of selecting from competing random variables with differing levels of uncertainty provides an consistent and efficient solution to the explore-exploit problem (Macready and Wolpert, 1998).

6.3.3 Simulating the effectiveness of the Persuasive Strategy Algorithm

To test the presented Beta-Binomial approach in the $M = 2$ scenario a simulated dataset presenting different visitors observed at multiple points in time was created. The simulated data describes the message success of two different messages for four different groups of visitors with 20 visitors each on 50 approaches to the doors. The four groups represent (1) *general insusceptible visitors*—those that respond favorable to only 10% of the message which implements strategy A and 50% to strategy B , (2) *susceptible visitors*, $A = 40\%$, $B = 90\%$, (3) *visitors susceptible to message B* , $A = 10\%$, $B = 90\%$, and (4) *visitors susceptible to message A* , $A = 90\%$, $B = 10\%$. Table 6.1 shows an excerpt of the simulated data.

Based on these simulated data the population estimates of message

⁴The 80th percentile is arbitrarily chosen and should depend on the effectiveness estimates given the problem at hand.

Table 6.1: Overview of the simulated data for the 4 different user groups.

	Types	User	Occasion	Strat A	Strat B
1	1	1	1	0	0
2	1	1	2	0	0
3	1	1	3	0	1
..	..	..	..	..	..
..	..	..	..	..	..
1000	4	20	50	1	0

effectiveness for each message are: $\hat{p}_A = 0.38$, $\hat{p}_B = 0.58$. Thus, overall message *B*—the consensus message—was most effective. Next, for each visitor, for each occurrence at the doors, the effectiveness is modeled. Strategies are selected as specified by the decision rule and the (simulated) outcome is recorded. Next, the algorithm updates its expectancy for the selected message and iterates through all occurrences. To ensure a flexible starting point for each user the prior variance of each estimate at the first encounter were set to be relatively high: $\sigma_A^2 = \sigma_B^2 = 0.05$.⁵ Figure 6.3 shows that the proposed algorithm identifies the four simulated visitor groups. In the upper left panel convergence to strategy *B* is slow—it takes about 40 observations before *B* is consistently presented as best strategy. This is due to the low overall compliance of visitors in this group. With higher compliance and/or larger differences in effectiveness of the two strategies convergence is faster: The bottom right panel of figure 6.3 shows how within ± 10 observations strategy *A* is identified as the most successful strategy.

6.3.4 Limitations of the proposed solution

There are a number of drawbacks of the proposed Beta-Binomial solution to create adaptive persuasive systems. Besides the fact that when the number of strategies grows the number of necessary occasions for convergence will increase, there are three more fundamental issues which are not addressed by this algorithm. First the algorithm described here does not use a shrunk estimate on each occasion: After including the initial knowledge of the behavior of other visitors the model is specific for an individual visitor. While this provides quick adaptation there is no opportunity to adapt estimates based on changing population wise

⁵One could estimate this variance based on the between visitor variance of the point estimates.

Figure 6.3: Progression of point estimates and certainty of 2 messages for four users, one from each group.

trends. Population wise trends will be useful to reduce error in the estimation of individual level estimates and possibly capture (e.g.) seasonal effects of distinct strategies on compliance. Second, since the estimates for the effectiveness of the strategies are treated independently there is no way to of “borrowing strength” (Gelman and Hill, 2007) based on correlations with other strategies. Both of these concerns could be addressed using a multilevel approach. Finally, the proposed model provides no method of including prior beliefs about the distribution of visitor profiles over a population: Multilevel models with flexible priors on the individual level estimates—as opposed to the commonly used Gaussian distribution—could address this difficulty.

6.4 Properties of Adaptive Persuasive Systems

In this chapter the results of the three *insight generation* chapters were summarized. These results motivate the need for designers to build adaptive persuasive systems: Systems that use persuasion profiles to adapt to responses of individual users to increase their effectiveness. The next two Chapters (7 and 8) explore the effects of usage of both static (Chapter 7) and dynamic (Chapter 8) persuasion profiles in persuasive systems. Before presenting these case studies it is important to make clear what makes persuasion profiles—and thus profiles that describe people’s susceptibility to different means to an end—distinct from other profiles that identify people preferred ends—such as those used by recommender systems.

6.4.1 End-independence

Means-adaptive persuasive technologies—technologies that use persuasion profiles—are distinctive from other adaptive systems such as recommender systems in their end-independence: a persuasion profile created in one context can be applied to bringing about other ends in that same context or to behavior or attitude change in a quite different context. This feature of persuasion profiling is best illustrated by contrast with end adaptation. Any adaptation that selects the particular *end* (or goal) of a persuasive attempt is inherently context-specific. Though there may be associations between individual differences across context (e.g., between book preferences and political attitudes) these associations are themselves specific to pairs of contexts. On the other hand, persuasion profiles are designed and expected to be independent of particular ends and contexts.

It is important to clarify exactly what is required for end-independence to be obtained. If a persuasion profile is end-independent then this does not imply that the effectiveness of all social influence strategies for a specific individual is constant across all contexts. Consistent with the results reviewed in Chapter 2, it is clear that influence strategy effectiveness depends on, e.g., the type of behavior change or the involvement of users. That is, the most effective influence strategy for a system to employ, even given the user’s persuasion profile, would depend on both context and target behavior. However, end-independence requires only that the *difference between the average effect of a strategy for the population and the effect of that strategy for a specific individual is relatively consistent across contexts and ends*.

6.4.2 Non-disclosure

Designers of means-adaptive persuasive systems can be tempted not to disclose the adaptation of their system. This can be contrasted with end adaptation, in which it is often advantageous for the persuader to disclose the adaptation and the adaptation is (potentially) easy to detect. For example, when Amazon.com recommends books for an individual it makes clear that these are personalized recommendations—thus benefiting from effects of apparent personalization and enabling presenting reasons why these books were recommended (E.g. Ochi et al., 2010). In contrast, with means adaptation, not only may the results of the adaptation be less visible to users (e.g. emphasizing either “*Pulitzer Prize winning*” or “*International bestseller*”), but disclosure of the adaptation may reduce the target’s attitude or behavior change.

It is hypothesized that the effectiveness of social influence strategies is, at least partly, caused by automatic processes. According to dual-process models (Cacioppo et al., 1986), under low elaboration message variables manipulated in the selection of influence strategies lead to compliance without much thought. Disclosure of means adaptation can increase elaboration. This in turn might decrease the effectiveness of influence strategies which operate primarily via the peripheral route. More generally, disclosure of means adaptation is a disclosure of persuasive intent, which can increase elaboration and resistance to persuasion (Cf. Tormala and Petty, 2004).

6.4.3 Conclusions

Based on the empirical results presented in the *insight generation* chapters it is clear that persuasive technologies would benefit from personalization of the selection of influence strategies. The ambient intelligence scenario provides designers with the opportunity to identify, present, and measure the effects of different influence strategies and is thus particularly suited for the design of such systems. This chapter presented *persuasion profiles* as a means for designers to create systems that personalize their persuasive attempts.

The next two chapters will explore the use of persuasion profiles in actual persuasive interventions. These *case study* chapters each present multiple persuasive systems that are personalized—or adapted—to individuals. The first case study chapter, Chapter 7, focuses on the use of static persuasion profiles obtained using meta-judgmental measures to optimize several health interventions. The second case study chapter, Chapter 7, presents three persuasive systems that use dynamic persua-

sion profiles based on operative measures to increase exercise, and drive sales. Thus, the next chapters explore the use of persuasion profiles in a number of the core application areas of persuasive technologies⁶.

⁶This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2011a; Kaptein, 2011b; Sakai et al., 2011).

7

Case Studies I: *Influencing Health Decisions*

7.1 Introduction: Measures of Susceptibility

The previous chapters set the stage to create persuasive systems that adapt to the individual susceptibility of users to different social influence strategies. Two distinct approaches can be feasible for designers: (a) designers could use meta-judgmental measures of personality or susceptibility to adapt persuasive interventions, or (b) designers could use operative measures—people’s responses to persuasive strategies—to dynamically adapt their persuasive systems. In this chapter the first option is explored. Study 1 and 2 presented in this chapter use a short version of the *STPS* to adapt influence strategy selection to individuals. Study 3 of this chapter uses the full *STPS* to adapt a persuasive system to individuals at the level of social influence strategies. All three studies presented here focus on persuasive interventions that encourage users to lead a healthier lifestyle.

7.2 Study 1 & 2: Promoting Lunch Walks

The first two studies presented in this chapter explore only a distinction between people that score high and low on their susceptibility to per-

sualion. These studies thus focus only on overall susceptibility to persuasion instead of distinctions between different strategies. The analysis presented in Study 1 and 3 of Chapter 3 showed that this already is an important distinction and thus the first applied exploration concerns only this dimension.

7.2.1 High and low Susceptibility to Persuasion

Given the applied context, Study 1 and 2 in this chapter used a shortened version of the *STPS* to determine *overall* susceptibility of participants to persuasive strategies. Constraints of the company in which the studies were performed led to the selection of a maximum of *seven* questionnaire items to determine participant's susceptibility. To select the most distinctive seven items of the *STPS* an overall susceptibility to persuasion score—the average over all items of the *STPS*—was computed for each of the participants in the evaluation of the *STPS* presented in Chapter 5. Next, a linear regression of all 26 items of the *STPS* on the final susceptibility score was used to determine which of the individual items were most powerful in explaining the overall susceptibility to persuasion score. This led to the selection of the following seven items:

1. *Products that are "hard to get" represent a special value.*
2. *I would feel good if I was the last person to be able to buy something.*
3. *I believe rare products (scarce) are more valuable than mass products.*
4. *I always follow advice from my general practitioner.*
5. *I am very inclined to listen to authority figures.*
6. *I always obey directions from my superiors.*
7. *I am more inclined to listen to an authority figure than to a peer.*

1933 Knowledge workers located in one single office park were invited by email to participate in this 7-item questionnaire to assess their overall susceptibility to persuasion. 516 Participants completed the online questionnaire. The items were scored on a 7-point scale ranging from totally disagree to totally agree (Cronbach's $\alpha = 0.646$). For each participant an overall score was computed: the average of the 7 susceptibility to persuasion items. Based on the susceptibility to persuasion scores, three profiles were determined: the low susceptible's, the moderately susceptible's, and the high susceptible's. Since the studies focus on providing an example of the use of persuasion profiles participants

were selected with a considerable difference in susceptibility to persuasion score. Thus, in this study only participants with either low or high susceptibility to persuasion scores were selected to participate. The *low susceptible* ($N = 136$)—the lowest scoring quartile – had scores ranging from 1.00 to 3.29. The *highly susceptible* ($N = 140$)—the highest scoring quartile—had scores ranging from 4.57 to 6.14.

Both Study 1 and 2 use a so-called *Extreme Group Approach* (EGA) in which participants are selected based on extreme scores on the short version of the *STPS*. This approach has implications for the accuracy of the effects that are estimated based on the collected data and the power of the experiment. However, “*The use of EGA may be a matter of necessity in situations when a researcher has limited resources and wishes to maximize the power for detecting the presence of an effect. [...] In such situations, and with proper considerations, EGA may be a useful tool to improve the odds of detecting an effect, if it truly exists*” (Preacher et al., 2005). This was exactly the aim of the two initial studies: they provide *initial* evidence for differences in the effects of influence strategies on people who indicate different susceptibilities. Given limited resources, and a motivation to explore this question in a field experiment with high external validity but inherent lower internal validity, the EGA approach was considered appropriate for this goal.

The two studies presented here examined the applicability of persuasion profiles for promoting health-related behaviors. Study 1 focused on physical activity and study 2 focused on fruit intake. Below, the common methodology used in both studies is presented.

7.2.2 Common methodology

In both studies the low and the high susceptible to persuasion groups were invited by email to participate in a health related activity. Study 1 focused on physical activity by inviting participants by e-mail to join for a lunch walk. Study 2 focused on fruit intake by inviting participants by e-mail to express their opinion about an initiative to provide a daily fruit snack. In both studies half of the participants were randomly assigned to the *persuasive implementation(s)*, (*PI*) condition and half of the participants were assigned to the *no persuasive implementation(s)*, (*NPI*) condition. In the PI condition, the invitation e-mail was supported by a number of persuasive messages while in the NPI condition no persuasive messages were included. The studies thus employed a 2 (PI vs NPI) $\times$ 2 (Highly vs Low susceptible) between subjects design.

After receiving the invitation email participants were asked to sign

up using an online form. To gain insight into the degree of compliance to the invitations, three measures of compliance were distinguished:

1. *Interest*: Participants' click on the email.
2. *Intention*: Participants' response to the main question in the online form (e.g. the sign up for a lunch walk).
3. *Behavior*: Participants' subsequent behavior.

The first of these measures quantifies the immediate (attitudinal) response to the persuasive message. The second measure quantifies the behavioral intention after the message, and the third measure quantifies the actual behavior.

7.2.3 Method Study 1: Lunch walks

In Study 1 participants were invited to join a lunch walk. Participants received an email with an invite and a link to sign up for one of two possible time slots during lunch. After clicking on the link participants could sign up for one of the two time slots. After signing up participants were asked to print a form with their name on it and bring it to the lunch walk enabling monitoring of the actual behavioral response.

Participants in the NPI condition received an email stating: "*We would like to invite you for the [Company] lunch walk. The [Company] fun4health committee was founded 2 months ago to promote general health of [Company] employees and affiliates.*", the time of the lunch walk and the link to sign up. Participants in the PI condition received the same email with an addition of the following three messages:

1. *Both physicians and general practitioners recommend at least 30 minutes of moderate activity, such as walking, during a day. The lunch walks are a great place to start!* [Authority]
2. *We expect a lot of people so please sign up before all available slots are filled.* [Scarcity]
3. *In other companies 1000's of people are already joining in on similar initiatives.* [Consensus]

Study 1 was conducted at two points in time—referred to as *Study 1a* and *Study 1b*—with the same groups of highly and low susceptible participants, because it was expected that unpredictable weather conditions could be experienced as a barrier for behavioral compliance for the outdoor activity. Each of our participants worked in the same industrial area with a common dining facility which was the starting point for the lunch walks.

7.2.4 Results study 1

In total, 276 respondents were invited to participate in study 1a. Of these 136 belonged to the low susceptible group, and 140 belonged to the highly susceptible group. About half of the participants received an email without the persuasive cues and about half received an email with persuasive cues. Table 7.1 gives an overview of the results of the study 1a. It is clear that in the PI condition participants overall showed significantly more *interest* (PI = 23.4%, NPI = 15.6%, $\chi^2 = 2.700$, $p = 0.050$), and have a significantly higher behavioral *intention* (PI = 8.5%, NPI = 3.0%, $\chi^2 = 3.887$, $p = 0.024$) than in the NPI condition. No significant effect of the use of persuasive messages was found on actual *behavior*.

Table 7.1: Results study 1a: Percentage of respondents responding favorably.

		NPI	PI	χ^2	p (one-sided)
<i>Main effect</i>					
	Interest	15.6%	23.4%	2.700	0.050
	Intention	3.0%	8.5%	3.887	0.024
	Behavior	1.5%	3.5%	1.189	0.138
<i>Interaction</i>					
Low susceptible	Interest	17.2%	16.7%	0.007	0.468
	Intention	3.1%	6.9%	1.012	0.157
	Behavior	.	1.4%	0.895	0.172
Highly susceptible	Interest	14.1%	30.4%	5.426	0.010
	Intention	2.8%	10.1%	3.124	0.039
	Behavior	2.8%	5.8%	0.758	0.174

The observed *main* effect of the use of social influence strategies can be explained by the high compliance of the *high susceptible* participants (e.g. *interest*: PI = 30.4%, NPI = 14.1%, $\chi^2 = 5.426$, $p = 0.010$). For *low susceptible* participants no main effect of the persuasive message is observed (e.g. *interest*: PI = 16.7%, NPI = 17.2%, $\chi^2 = 0.007$, $p = 0.468$). This interaction between the condition and the susceptibility score is statistically significant, $\chi^2 = 10.463$, $p < .01$ ¹. Also for the behavioral intent, $\chi^2 = 14.728$, $p < .001$, and the actual behavior, $\chi^2 = 4.669$, $p < .05$, a significant interaction between the condition and

¹A Cochran-Mantel-Haenszel χ^2 test of the null hypothesis that two nominal variables are conditionally independent was used to test these interactions.

participant's susceptibility score was found. Thus, while overall the use of persuasive messages increased the participation in health related behavior, the actual cause of this effect is a very high compliance by high susceptible's while there is no statistically significant difference between the NPI and the PI conditions for the low susceptible's.

In study 1b, a new invitation was send out to 268 people—eight people signed out for any follow up mails after the invite for study 1a and were not invited again. Table 7.2 shows the results of this second trial. There was a significant main effect on the *interest* measure (PI = 10.9%, NPI = 4.6%, $\chi^2 = 3.761$, $p = 0.026$). As in study 1a, this main effect of persuasive implementation disappeared when looking only at low susceptible's (e.g. *interest* PI = 5.4%, NPI = 8.8%, $\chi^2 = 0.570$, $p = 0.251$). In this second trial it was found that for the *intention* measure the low susceptible's complied significantly less when persuasive implementations were used in the invitation message (PI = 0.0%, NPI = 7.0%, $\chi^2 = 5.357$, $p = 0.017$). Testing specifically for the interactions between the conditions and the susceptibility score showed that, contrary to Study 1a, only the interaction on the Interest measure was statistically significant, $\chi^2 = 13.412$, $p < .001$.

Table 7.2: Results study 1b: Percentage of respondents responding favorably.

		NPI	PI	χ^2	p (one-sided)
<i>Main effect</i>					
	Interest	4.6%	10.9%	3.761	0.026
	Intention	3.1%	2.2%	0.196	0.329
	Behavior	0.8%	1.5%	0.293	0.294
<i>Interaction</i>					
Low susceptible	Interest	8.8%	5.4%	0.570	0.251
	Intention	7.0%	.	5.357	0.017
	Behavior	1.8	.	1.308	0.127
Highly susceptible	Interest	1.4%	17.5%	11.049	0.001
	Intention	.	4.8%	3.603	0.029
	Behavior	.	3.2%	2.384	0.062

7.2.5 Methods study 2

Study 2 was similar to study 1: Again both high and low susceptible's were invited to take part in a health related activity. This time an email was sent to 267 participants which explained that plans were

being made to start a fruit distribution service at the main building of the office campus. It was explained that participants would be able to pick up a piece of fruit every day. The alleged goal of the email was to inquire about possible interest for such a project. Participants could click on a link in the email to state their interest in such a service. In the *PI* condition the following lines were added to the email: “*Eating two pieces of fruit a day is recommended by the World Health organization. Our service would make it easier to reach that target*” [Authority]. And: “*Other companies have picked up similar ideas by providing fruit during lunchtime for reduced prices for employees. If we all join in, we could make this service happen!*” [Consensus].

Contrary to Study 1, only two types of compliance were measured: *interest* (did the participant click on the email link), and *intention* (did the participant respond to the subsequent survey).

7.2.6 Results study 2

Table 7.3 shows that the results slightly differ from those obtained in study 1: There is no significant main effect of the persuasive implementations (e.g. *interest* PI = 23.4%, NPI = 21.8%, $\chi^2 = 0.149$, $p = 0.350$). When looking at the low susceptible’s and the high susceptible’s separately it is clear that the absence of a main effect is probably best explained by an interaction: Low susceptible’s seem to comply less to a message with persuasive implementations (e.g. *interest* PI = 18.8%, NPI = 25.8%, $\chi^2 = 0.919$, $p = 0.196$) while high susceptible’s seem to comply more (e.g. *interest* PI = 27.8%, NPI = 17.2%, $\chi^2 = 2.159$, $p = 0.071$). However, both of the interactions are not statistically significant at a five percent level (two sided test) in Study 2. Thus, contrary to Study 1a and 1b, no decisive evidence the effect of the persuasive messages differed between the low and high susceptible participants was collected.

7.2.7 Discussion Study 1 and 2

The results presented in these two studies suggest that individuals indeed differ in their compliance to health-related messages supported by social influence strategies. When analyzing these differences between highly susceptible people and low susceptible people it is evident that a positive effect of persuasive messages is obtained *only* for high susceptible’s and is absent or even *negative* for low susceptible’s. While the effects identified in these studies are likely over-estimated because of the use of the EGA, the two studies do show that at least for a small

Table 7.3: Results study 2: Percentage of respondents responding favorably.

		NPI	PI	χ^2	p (one-sided)
<i>Main effect</i>					
	Interest	21.4%	23.4%	0.149	0.350
	Intention	15.9%	15.6%	0.004	0.476
<i>Interaction</i>					
Low susceptible	Interest	25.8%	18.8%	0.919	0.196
	Intention	21.0%	13.0	1.468	0.082
Highly susceptible	Interest	17.2%	27.8%	2.159	0.071
	Intention	10.9%	18.1%	1.369	0.089

portion of people *not* using an influence strategy to support a request can be more effective than using influence strategies.

The applied value of the distinction between high and low susceptible's in Study 1a and 1b demonstrates that a first step in tailoring persuasive systems can be provided by focusing not on distinct influence strategies, but rather on the usage of *any* strategy. This is inline with predictions that could be derived from dual processing models of persuasion: susceptibility to persuasive strategies describes a tendency to process messages peripherally. As such, susceptibility to persuasion directly relates to *NfC* as an earlier attempt to measure people's tendency for peripheral or central processing (See also 5.4 for a discussion of the relationship between the *STPS* and *NfC*).

7.3 Study 3: Influencing Snacking Behavior

Study 1 and 2 demonstrated the use of differences in responses to persuasive strategies as measured using meta-judgmental measures: Differences in *overall* susceptibility lead to different responses to requests that are either supported by, or not supported by, implementations of influence strategies. Study 3 extends this finding by using the full *STPS* as brought forward in Chapter 5 to determine susceptibility not just at an overall level but rather at the level of distinct social influence strategies. Study 3 describes the development of a persuasive system that uses short text messaging to reduce people's snacking behavior to examine the effects of personalization on compliance to the system. The first section describes the development of the SMS based system and the development of persuasive messages that implement different social

influence strategies. Next, the empirical evaluation of a system that implements either *personalized* messages, *counter-personalized messages*, or *random* messages is presented.

7.3.1 Designing an SMS Intervention to Reduce Snacking

To evaluate whether the *STPS* could aid in message selection and whether personalization of influence strategy usage has the desired effects, a two week text-messaging intervention was setup. In the trial, short text messages (SMS) on a mobile phone were used as a prompt to reduce snacking behavior. Mobile phones have frequently been used as a platform to employ persuasive technologies because of their pervasiveness in everyday life, and their ability to be “at the right place, at the right time” (Kass, 2007). Mobile applications have been developed and tested in all area’s of the persuasive technology field: From applications that help chronic patients manage their disease as those presented by Franklin et al. (2008) to the persuasive service that promotes the sexual health of teenagers presented by Parkes et al. (2005). In all instances both researchers as well as users recognize the power of mobile devices for pervasive persuasion attempts.

The current SMS intervention focussed on reducing snacking behavior by sending SMS messages that prompted reduced snacking. Mobile interventions for a healthy diet have previously not only been part of the design research space, but have also made their way into randomized controlled clinical trials. Patrick et al. (2009) show a larger effect on weight loss of an SMS based intervention program versus a paper based one over a four month time period. They conclude that “text messages might prove to be a productive channel of communication to promote behaviors that support weight loss in overweight adults”. In another paper by the same team (Consolvo et al., 2009) they argue in favor of the mobile phone as a general carrier for health related interventions and underline the impact such technologies can have. The current study evaluates whether messages that are *personalized* to the user—as measured using the *STPS*—are more effective than messages that are not tailored towards this goal.

Design of Persuasive Text Messages

To design the actual messages that would be sent during the trial two persuasive technology researchers independently tried to generate as many messages as possible for each of the six strategies identified by Cialdini (2001) that would be usable given the context. In total 42

messages were created. After combining the lists it became eminent that both the *Liking* as well as the *Reciprocity* strategies were hard to implement in the context of mobile text messaging to reduce snacking behavior. For the *Liking* strategy to be successfully implemented there is need for a “bond” between the persuader and the receiver of the message. Given that there is no clear social actor in play which receivers of the message could relate to this strategy was hard to implement. For the *Reciprocity* strategy to be most effective, a favor has to be done to the persuadee prior to the persuasive request. These strategy thus both proved hard to implement and the two created implementations for these strategies were omitted.

To be able to evaluate the messages that were created and select the messages to be used in the final intervention an evaluation was setup with ten researchers with expertise in the field of Human-Computer Interaction. Each researcher received a small description of each of the four social influence strategies that we tried to implement to become familiar with these terms. Next, researchers were shown the 40 messages one by one, and were asked to categorize them into the following categories: (a) Implements the Authority strategy, (b) Implements the Consensus strategy, (c) Implements the Scarcity strategy, (d) Implements the Commitment strategy, or (e) Other / Don’t know.

After the ten researchers rated each of the messages the ratings were analyzed by looking at agreement matrices between the researcher ascribed categories. For the *Authority* strategy there was a general high agreement: for all of the messages that intended to implement this strategy at least 70% of the researchers ascribed the message to this strategy. Three of the implementations designed to implement the Authority strategy were selected: one of these was ascribed to this strategy by all raters, the other two by 90% of the raters. Similarly three implementations of the Consensus strategy, three of the Scarcity strategy, and three of the Commitment strategy were selected for inclusion in the trial. It has to be noted that the messages that implemented the Scarcity strategy were the least identifiable, with one of the three selected messages only ascribed to this strategy by 60% of the raters. Table 7.4 gives an overview of the messages that were selected, the strategies they aim to implement, and the agreement between raters.

Table 7.4: The messages used in the intervention. For each of the four social influence strategies used in this trial three implementations are used.

Strategy	Message	Agree
Authority	Try not to snack today. According to the College of Physicians this is an easy way to lead a healthier life.	100%
Authority	Dietitians advise to have 3 meals a day without snacking. Try to reduce snacking.	90%
Authority	The World Health Organization advises not to snack. Snacking is not good for you.	90%
Consensus	90% of people benefit from reducing snacking between meals. It will boost your energy and you will live a healthier life	90%
Consensus	Everybody agrees: not snacking between meals helps you to stay healthy	90%
Consensus	Reduce snacking. You are not on your own: 95% of participants have already reduced snacking	90%
Commitment	The aim of this study is to live healthier. Reducing snacking is a way to achieve that.	100%
Commitment	Try to obtain your goal for living a healthier life by not snacking. You are committed!	90%
Commitment	You have to continue what you've started: you are participating in this test to lead a healthier life. Reducing snacking.	90%
Scarcity	There is only one chance a day to reduce snacking. Take that chance today!	90%
Scarcity	This test lasts only 3 weeks: you have the unique opportunity to enhance your health by reducing snacking	70%
Scarcity	Today is a unique opportunity to lead a healthy life. Reduce snacking	60%

7.3.2 Method Study 3

A two week long trial to evaluate the different message conditions between subjects was setup. Since snacking behavior is very variable between people, a one week baseline assessment of individual snacking behavior was included before introducing the different messaging conditions and establishing their effects on the snacking behavior within a single user.

Participants

Participants in this study 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 study would take two full weeks and would entail filling out a several questionnaires and receiving daily text messages on a mobile phone. In total 334 potential participants clicked on the link in the call that took them to a designated website and were presented with the introduction questionnaire. At the end of the introduction questionnaire participants were asked to provide their mobile phone number. After providing their phone number a text message with an activation code to login at the designated website was sent to participants. In total 162 participants fully completed the sign up process and activated their study participation.

After signing up participants received text messages for a period of two weeks (2×5 days, workdays only). Participants were instructed—both prior to the study as well as via the text messages—to go to the designated website every evening to fill out a short diary. The first week was used to establish a baseline snacking frequency for each participant, while the intervention was employed in the second week. In the analysis only those participants that filled in at least one diary during each of the two weeks (e.g. during the baseline measurement and during the intervention) were included. The final sample was composed of 73 participants. The average age of the participants was 34.9 years ($SD = 11.1$). Of our final sample 32 (43.8%) were females. Upon completion participants were awarded research participation credits with a monetary value of two Euro's (default amount provided by the research agency).

Procedure & Measures

After arriving at the designated website all participants in the study first filled out a small questionnaire regarding their snacking behavior, their shopping behavior, and their motivation to decrease snacking. Next, participants were administered the *STPS* and provided their mobile phone number to sign up for the text messaging part of the study. Participants then received one text message a day (on workdays) for a period of two weeks, and subsequently filled out a small online diary every day.

The introduction questionnaire included all 32-items of the initial *STPS* scale, although only the 26 items that scored consistently on the factors of interest were used to allocate participants to conditions. Next to the *STPS* the questionnaire included the following questions:

1. How often a week do you usually visit a supermarket to buy ingredients to prepare a dinner?
2. How often (a week) do you prepare your own meals?
3. Would you like to eat healthier? (Scored *Yes*, *No*)
4. Do you feel you generally eat healthy dinners? (Scored *Yes*, *No*)

Finally, participants were asked for their age and gender and proceeded to the sign up procedure.

During the sign up procedure participants provided their mobile phone number and received a text message with an eight digit authorization code. Participants filled in their authorization code on the study website to create a personal profile and supplied a user name and password for subsequent logins. After logging in with their user name and password participants were asked to fill in the first diary. This diary—in the first week—consisted of the following questions:

1. How many snacks did you have today? (Open ended)
2. How many *unhealthy* snacks did you have today? (Open ended)
3. How healthy was your nutrition today? (Five-point scale, *Very unhealthy* to *Very healthy*)

For the first week participants received one text message a day which asked them to fill in their diary for that day. This phase of the research was the baseline period (Phase 1).

After Phase 1 participants entered the second week (Phase 2) in which they received the persuasive messages according to the experimental condition they were allocated to (See 7.3.2 for details). The persuasive message contained an implementation of either the Authority, Consensus, Commitment, or Scarcity strategy as described in 7.3.1.

After receiving the persuasive message participants were again asked to fill out their online diary. In Phase 2 one additional question was added to the diary: “How useful was the text message your received?” scored (1) *Not at all useful* to (6) *Very useful*. After receiving five persuasive messages during Phase 2 and filling out their last diary, participants received a “thank you for your participation” message on their screen.

Conditions

Based on participants answers to the *STPS* a mean score on each of the six variables of interest—their self-reported susceptibility to each of the six strategies—they were automatically and randomly allocated to one of the following three conditions:

1. *The personalized condition* (PC): Participants assigned to this condition received, during Phase 2, messages that were randomly selected implementations of the two strategies (of the four implemented in this study) that they had the highest susceptibility scores on. Hence, the messages were adapted to their personal profile to be most effective.
2. *The contra-personalized condition* (CPC): Participants assigned to this condition received random implementations of the two strategies they had the lowest mean scores on as judged from the *STPS*. Hence, the messages were adapted to be the least effective.
3. *The random condition* (RC): Participants in this condition received randomly selected messages out of the full set of messages presented in table 7.4. This represents a common usage of persuasive messages in persuasive systems.

7.3.3 Results Study 3

Overview

In total 506 diaries were filled out by the 73 participants included in the analysis. Of the included participants a large majority (80.8%) indicated to be motivated to eat healthier—this even though most participants claimed to eat healthy already (89.0%). All of the participants indicated to prepare (cook) their own meals at least once a week, and 45.2% indicated to prepare their own meals more than five times a week. 95.9% of participants visited a supermarket to purchase food at least once a week. These figures indicate that participants were largely individually responsible for their own food consumption and shopping

habits and thus personal text-messages could possibly influence their behavior. Based on the scores on the *STPS*, 22.8% of our respondents indicated to be most susceptible to the scarcity strategy, 14.9% to the authority strategy, 53.1% to commitment strategy, and 9.1% to the consensus strategy.

Snacking Behavior

The primary test to see whether messages personalized to participant's scores on the *STPS* can be effective in reducing snacking behavior is provided by an examination of the progression of the (self reported) snacking behavior over time between the three experimental groups. The daily diary contained three questions that are indicative for the effect of adaptive or contra-adaptive messages on snaking behavior and we examined each separately. The primary measurement—the self-reported number of snacks eaten by participants each day—is graphically represented in Figure 7.1. It is clear from this figure that—while variable over days—the snacking consumption decreased over time for both the *RC* and the *PC* conditions while it did not decrease in the *CPC*. It is eminent that the decrease is largest from timepoint 6 onwards: this corresponds to the entry into phase 2 of the experiment and thus the actual separation of messages between the three conditions.

Figure 7.1: Overview of the average number of snacks eaten each day by our participants separated for the three experimental conditions.

To statistically test the effects of our conditions over time a multilevel model with varying intercepts for participants was fitted to the data. From the “null” model

$$y_{ij} \sim \mathcal{N}(\mu_j, \sigma_{err}^2) \quad (7.1)$$

where $\mu_j \sim \mathcal{N}(0, \sigma_{\mu_j})$ a model is built that includes both time as well as condition to explain the snacking behavior of our participants. Adding time—centered around the phase shift—as a fixed factor to the model (Thus $y_{ij} \sim \mathcal{N}(\mu_j + \beta_{time}X_{time}, \sigma_{err}^2)$ significantly increases model fit (see table 7.5, model *B*). In this model the coefficient of time is significantly negative, $\beta_{time} = -0.06$, $t = -3.18$, $p < .01$, indicating a small but significant overall decrease in the number of snacks consumed by our participants during our trial.

After examining the effects of time, a phase $\times$ time interaction was added to the model to allow for different effects in the baseline and treatment phase. This significantly increases model fit, see table 7.5 model *C*, which indicates that the effects of the messages (over time) differs for the two phases. Finally, to test whether the conditions significantly influence the snacking behavior of our respondents a model is fitted in which time during phase two interacts with condition—essentially fitting separate time effects for the different conditions during phase two. Here again, model comparisons show a significant increase of model fit, see table 7.5 model *D*.

Table 7.6 shows the fixed effects of model *D*—the model allowing for different slopes for the different conditions during phase 2 of the trial. These fixed effects show that, even though averaged over all conditions snack intake decreased significantly during our intervention, reality is more granular: During the baseline phase the number of snacks does not decrease significantly. During the treatment phase the number of snacks consumed by our participants decreases significantly only for *RC* and *PC* participants. The fixed effects table also shows that the estimated decrease in snack consumption is higher for those in the *PC* condition than those in the *RC* condition. This latter difference between the slope of the *RC* condition over time and that of the *PC* condition over time is however not statistically significant. The difference between the slope of the *CPC* condition and that of the *RC* condition is statistically significant, $t = 2.89$, $p < .05$. The same is true for the difference in slope between the *CPC* condition and the *PC* condition, $t = 3.28$, $p < 0.01$.

Figure 7.2 shows the number of snacks per time-point for each par-

participant (both the time-point as well as the number of snacks are jittered to prevent overlap). Imposed on these raw data points are the overall time trend (light gray), the trend in during the first phase of the trial (black solid line), and the three time trajectories during the treatment phase. Of these latter three the black dashed line shows the *CPC* fit, the dotted line shows the *RC* fit, and the dashed and dotted line shows the *PC* fit. Participants who received personalized messages decreased their snacking intake during phase two of the experiment. Participants who received random messages also decreased their snack intake but did so to a slightly smaller extent. Finally, there is no evidence that participants who received messages that were contra-personalized decreased their snack intake.

Table 7.5: Comparing the null model with models including a time effect and different time effects for each condition.

	Df	AIC	logLik	χ	p
Model A:	3	1857.70	-925.85		
Model B:	4	1849.64	-920.82	10.07	< .01
Model C:	5	1843.39	-916.70	8.24	< .01
Model D:	7	1837.79	-911.89	9.60	< 0.01

Figure 7.2: Overview of the number of snacks eaten each day by our participants with superimposed the regression lines of the effects over time of the different conditions.

Results were very similar for the number of *unhealthy* snacks that participants ate. Again the model including an interaction between condition and time was favored by likelihood ratio tests. Snacking significantly decreased in the random condition, $\beta_{Time:RC} = -0.13$, $t = -2.49$, $p < .01$. The effect was even larger in the personalized condition, $\beta_{Time:PC} = -0.21$, $t = -3.34$, $p < .01$. The effect of time in phase 2 for the CPC was not significantly different from zero, $\beta_{Time:CPC} = -0.02$, $t = -0.31$, $p = .68$. Direct comparisons between the effects of the conditions over time showed that—as for the number of snacks that respondents ate—the effect of the *PC* condition, $\beta_{Time:PC}$, was significantly different from that of the *CPC* condition, $t = 2.96$, $p < .05$. The *PC* condition was however not significantly different from the *RC* condition.

Table 7.6: Overview of the fixed effects of the model including an interaction between time and condition to predict snacking behavior (Model *D*). Empirical p-values are based on MCMC simulations.

	Estimate	Std. Error	t-value	p
Intercept	2.16	0.19	11.33	< .001
Time Phase 1	0.05	0.04	1.07	.28
Phase 2: Time:CPC	-0.01	0.08	-0.16	.85
Phase 2: Time:RC	-0.22	0.07	-3.30	< .001
Phase 2: Time:PC	-0.30	0.07	-4.29	< .001

Next to the self-reported number of consumed snacks the perceived healthiness of the participants diet as a function of the text-messages (and the progression over time) is an interesting indicator of the success of the intervention. Analyzing the progression over time of participant's responses to the question "*How healthy was your nutrition today*" shows that the perceived healthiness of participants nutrition is relatively stable over time (See table 7.7). Only for the *PC* participants there is a significant increase during the treatment phase. Figure 7.3 shows the ratings of our participants at each time-point (again jittered to prevent overlap). It is clear that the healthiness ratings are, overall, relatively stable over time. However, the *CPC* participants perceived their diet during the treatment phase as *less* healthy than in the baseline phase, while for *CT* participants the opposite was the case. Similarly to the previous comparisons, the effect over time of the *PC* condition on the perceived healthiness of the diet is significantly higher than that of the *CPC* condition, $t = 2.37$, $p < .05$, but not from the *RC* condition.

Table 7.7: Overview of the fixed effects of the model including an interaction between time and condition predicting perceived healthiness of nutrition.

	Estimate	Std. Error	t value	p
Intercept	3.44	0.08	41.43	< .001
Time Phase 1	0.01	0.03	0.51	.61
Phase 2: Time:CPC	-0.03	0.05	-0.58	.52
Phase 2: Time:RC	0.05	0.04	1.35	.13
Phase 2: Time:PC	0.08	0.04	1.88	< .05

Figure 7.3: Overview of the perceived healthiness of daily nutrition score for participants with superimposed the regression lines of the effects over time of the different conditions.

Additional analysis

During phase 2 of the experiment (e.g. time-point 6 to 10) participants were asked to evaluate the usefulness of the messages they received. Figure 7.4 shows the usefulness evaluations per time-point with superimposed the fixed effect line of the effect of time for each of the three experimental conditions. The fixed effects of the interaction between time and condition on the perceived usefulness ratings are presented in table 7.8. It is clear that the text messages are perceived as more useful in the personalized condition than in the contra-personalized condition,

with the random condition scoring in between. Here again, the *PC* differed significantly from the *CPC* condition, $t = 4.52$, $p < .01$, but not from the *RC* condition.

Figure 7.4: Overview of the usefulness ratings over time. Added are the estimated regression lines of the effect of time.

Table 7.8: Overview of the fixed effects of the model including an interaction between time and condition predicting perceived usefulness of the received message.

	Estimate	Std. Error	t value	p
Intercept	2.58	0.31	8.32	< .001
Time:CPC	-0.06	0.05	-1.36	.12
Time:RC	0.09	0.04	2.17	< 0.01
Time:PC	0.15	0.04	3.50	< 0.001

A number of respondents dropped out of the study and the drop-out rates were examined in more detail. Since only from the start of Phase 2 the messages that are received by participants differ for the three different experimental conditions it was decided to look closely at the dropout rates in the second phase of the study as a function of the experimental condition. Participants who started phase 2 by filling out their diary at either time-point 6 or 7 but who did not finish by filling out diary number 10 were marked as drop-outs. In total 61.1% of our

participants dropped out during phase 2. When looking at the different conditions it is clear that the drop-out rate was higher amongst *CPC* participants (63.8%) than amongst *PC* participants (56.6%). The effect of condition on dropout rate in phase 2 is however only marginally statistically significant, $\chi^2(2, 72) = 4.83$, $p = .08$. No significant differences with the *RC* participants were observed.

7.3.4 Discussion Study 3

The SMS trial presented in Study 3, which tested the external validity of measurements obtained using the *STPS* in a *in situ* intervention, showed that while persuasive text-messages can be effective in changing people's behavior (snacking) and attitude (perceived healthiness and usefulness) these changes depend on the right choice of influence strategy for the right participant. Using *contra*-personalized persuasive messages, and hence messages participants thought they would *not* be susceptible to, did not decrease snacking behavior significantly, and led to a low estimation of the usefulness of the text messaging intervention. Conversely, the personalized persuasive messages, those adapted to fit with the profile derived using the *STPS*, lead to a decrease in snacking behavior and an increase in the perceived usefulness of the intervention over time. The decrease is estimated to be around 0.3 snacks a day during the course of the experiment: Hence, after about three days personalized persuasive messages decrease the daily snack consumption by about one snack. This is an effect large enough to be of importance for the physical health of participants.

The personalized messages used in this study also seemed to outperform a random selection of messages: consistently the estimated effects of the personalized messages were larger than those of the random messages. However, these differences were not statistically significant. Future researchers should however try to reduce noise in the measurement of snacking behavior when assessed in the wild. The inclusion of additional covariates that explain snacking behavior might increase the precision of the estimated effects of the different conditions. Also, replication of Study 3 with a larger sample size, or over a longer period of time, could aid in more consistently estimating the effects of personalized persuasion.

7.4 Conclusions

This first *case study* chapter examined via three separate persuasive interventions the applied benefits of using personalization of persuasion based on meta-judgmental measures in persuasive technologies. First, Study 1 and 2 demonstrated that meta-judgmental measures of overall susceptibility to persuasion—people’s tendency to comply to social influence strategies—were useful to improve the effectiveness of different health interventions: using influence strategies was more successful for those high in susceptibility to overall persuasion than for those low on this trait. These two studies also show that implementing social influence strategies for those low in susceptibility to these strategies can have a negative effect on compliance. In Study 3 these findings are advanced by showing the use of personalization of social influence strategy usage on the level of distinct strategies. Using implementations of the *wrong* strategies—the contra-personalized condition—led to *lowered* compliance in this latter study. This finding emphasizes the importance of selecting strategies that fit with the individual users of a persuasive technology.

The case studies presented in this chapter highlight two of the difficulties that accompany the use of persuasion profiles that are based only on meta-judgmental measures. First of all, it is *not* always possible to elicit elaborate questionnaire measures from all of the potential users of a persuasive system. The first two studies had to rely on a very short measure of susceptibility to persuasion due to constraints enforced by the organization in which the study was carried out. This is likely to happen more often in applied settings and makes accurate assessment of a persuasion profile via meta-judgmental measures less likely. Second, each of the studies presented here used a static presentation of social influence strategies: irrespective of responses to the persuasive request participants in a certain condition received a measure implementing, or not implementing, the social influence strategies that were associated with their assigned condition in the experiment. This even though for some people in some of the conditions the social influence strategies employed had an adverse effects. The use of meta-judgmental measures alone for personalized persuasive systems will inherently be relatively static and this might lead to suboptimal results. Since people might not be able to correctly judge their own susceptibility, or susceptibility is measured with error, it could be beneficial to update the chosen social influence strategy after observing behavioral responses of the user.

Theoretically interesting is the finding that contra-personalization has a negative effect on compliance, which might sometimes even be larger than the positive effect of personalization of influence strategies (See e.g. Study 1b). This large effect of employing the “wrong” influence strategy might be caused by psychological reactance (Brehm, 1966): wrongly selected influence strategies cause some users to contradict the requests because they perceive the influence attempt as a threat to their behavioral freedom or autonomy. Roubroeks et al. (2011) have already observed reactance in people’s responses to persuasive technologies .

The results presented in this chapter provide a first glimpse at the applied benefits of the use of persuasion profiles. Personalized persuasion consistently lead to the highest compliance and, in the studies presented here, to a more active and healthier lifestyle. Even more, contra-personalization led to adverse effects: for some users of the service the persuasive messages led to lower physical activity or an increase in unhealthy food consumption.

In the next case study chapter these findings are extended by moving from “interventions” to the design of interactive persuasive systems. Three systems that use persuasion profiles based on operative measures are presented. These systems adapt to the susceptibility of users to social influence strategies by observing user’s responses to request accompanied by implementations of distinct strategies².

²This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2010a, 2011a).

8

Case Studies II: *Dynamically Personalizing Persuasion*

8.1 Introduction: Dynamic Persuasive Systems

In Chapter 7 persuasion profiles based on meta-judgmental measures of susceptibility to social influence strategies were used to manipulate the selection of persuasive messages in several health related interventions. The chapter made clear that personalized messages lead to higher compliance than those that are not personalized. However, in a large number of applied situations in which persuasive technologies are used it is hardly possible to obtain meta-judgmental measures from all possible users of the system.

The current chapter presents three case studies in which *dynamically* adapted persuasive systems are tested: systems that create a persuasion profile dynamically by observing the behavior response of users to implementations of distinct social influence strategies. The three systems presented here focus respectively on increasing physical activity, increasing user engagement with a health service, and influencing consumer decisions. In each of the three designs different means

of *identification*, *representation*, and *effect measurement* as identified in Chapter 6 are implemented. The effectiveness of *dynamic adaptation* is each time assessed by comparing the performance of the system when using persuasion profiles to a similar system that offers persuasive messages that are not personalized to individual users.

8.2 Design 1: APStairs, Promoting Stair Usage

The first design focusses on increasing the daily physical activity of users. Due to the high levels of sedentary work subsequent low levels of physical activity are a problem for the general health and wellbeing of people (Ades, 2001). The system presented here—called APStairs—focusses on increasing stair climbing at work since benefits of this type of physical activity are clear and well-accepted. It has been documented that being more physically active improves both mental and physical health, which decreases the occurrences of obesity, cancer, cardiovascular diseases, and depression (Ades, 2001; Grediagin et al., 1995; Penedo et al., 2004).

To meet the definition of “physically active”, people are encouraged to accumulate a minimum of 30 minutes of brisk walking, which corresponds to 4 METs—Metabolic Equivalent of Task, a rate of energy consumption (Ainsworth et al., 2000)), on at least five days of the week. One possible alternative to brisk walking is climbing stairs. According to Teh and Aziz (2002) the gross energy cost of stair climbing is 9.6 METs and the gross energy cost of stair descending is 3.2 METs. Since most work sites offer a choice between stairs and elevator and because of the clear health benefits of stair use compared to elevator use this was the first application area of an adaptive personalized persuasive system.

8.2.1 System Design

The APStairs system, like any *dynamic adaptive persuasive system*, needs to be able to recognize its users, present a message implementing different social influence strategies, and measure the effect of the messages on user behavior. The following sections describe how these are implemented in the APStairs system.

Identification

The APStairs system was deployed in the lobby of an office building. When users entered the office lobby, they needed to be identified with a unique key to keep track of the effects of the adaptive persuasive

system on their individual behavior during multiple visits to the office building. After exploring several options for in-situ identification, it was decided to use Bluetooth inquiry-based scanning to uniquely identify users. Most mobile phones and laptops include Bluetooth functionality, and a number of people have its discoverable mode always enabled. A previous study which also used Bluetooth signals for identification reported that 7.5% of people had their Bluetooth enabled (Kostakos and Neill, 2008), allowing for remotely scanning the device ID.

To verify this assumption, a pilot test was run in the lobby where the system would be installed. Based on the observation of over a hundred office workers entering the building, it appeared that around 8% ($N = 153$) of the people had Bluetooth enabled on their mobile phones and were detected by the Bluetooth scanners. Thus, the system could unobtrusively assign a unique key to 8% of the visitors entering the building.

Representation

The APStairs system needed to be able to present different implementations of a multitude of social influence strategies to users. This was realized by placing a large television screen in the lobby of the office building and displaying persuasive messages on the screen. The messages were aimed at making people take the stairs instead of the elevator. In the selected building, there was a small hallway leading to both the stairs and the elevators after entering the building. This enabled strategic placement of the screen to grab user’s attention and present the message, while providing them ample time to make a conscious decision to use either the stairs or the elevator after reading the persuasive message.

Part of the process of creating the representation was the formulation of different persuasive messages, each implementing a social influence strategy. A structured brainstorm session with four interaction designers resulted in 24 messages, based on Cialdini (2004)’s six principles of persuasion, intended to persuade office workers to climb the stairs. To verify whether these messages were good representations of Cialdini (2004)’s influence strategies, a closed card-sorting test was conducted. Following an explanation of the strategies, participants ($n = 20$) were asked whether they understood the messages and then requested to sort them into categories—each category corresponding to one of the social influence strategies. In case participants had difficulties assigning messages to a specific strategy, the option to rate them as “unknown”

was provided.

During the brainstorm, no valid implementations of the Reciprocity strategy came up. Since this strategy requires an ongoing interaction with a social agent, something not present in the APStairs system, it was decided not to include this strategy. Additionally, many participants did not ascribe the Scarcity and Liking messages generated in the brainstorm to the intended categories. Based on these results, it was decided to only implement messages belonging to the *Authority*, *Commitment and Consistency*, and *Consensus* strategies since these were uniquely defined in the card-sorting task. Finally, an online questionnaire was conducted to examine to what extent each of the messages were understandable, believable and effective.

All 24 messages from the brainstorm session were evaluated by 55 participants. All of the messages were generally rated as understandable. Next, based on the ratings for each message, the three messages that were deemed most effective (i.e., obtained the highest score on the question “*How effective do you think this message is?*”) and believable (“*To what extent do you believe this message?*”) were selected. Table 8.1 presents these selected messages.

Measurement of Effect

To measure the success of the messages on a user that was previously identified by the Bluetooth scanner, scanners were installed on every floor of the five-story office building. Each scanner independently scanned for nearby Bluetooth-enabled mobile devices and continuously uploaded its scanning results to the central server. These scanning results were sufficient to decide between a user taking the stairs (which led to a successful scan on each floor of the building), or a user taking the elevator (which led to a scan on the ground floor). Every two hours during deployment of the system, the scanning patterns for each of the users were rated by a *human rater* to decide between stair and elevator usage. Ambiguous scanning patterns were transcribed as “unknown” and not used in the computation of the success probability for the different messages.

Figure 8.1 schematically describes the APStairs system. Near the entrance of the office building, a Bluetooth scanner identified the device ID of a user entering the office building. This device ID was sent to a server which in turn retrieved the estimated effectiveness of the different influence strategies for that specific user, and decided which message to

Table 8.1: Influence strategies and their implementations as used in the APstairs system.

Strategy	Implementation
<i>Commitment</i>	1. Planned to become healthier? Start by taking the stairs! 2. Committed to get in shape? Start by taking the stairs! 3. Promised yourself to be more physically active? Take the stairs!
<i>Consensus</i>	4. 70% of the people in this building already take the stairs. What about you? 5. The majority of the people in this building takes the stairs. Join them now! 6. Follow many other people; take the stairs!
<i>Authority</i>	7. “You get a good exercise by taking the stairs instead of the elevator.”—Bert Clarenbeek, gym instructor. 8. Doctors recommend taking the stairs. 9. “Taking the stairs helps you shape up your buttocks.” — Jessica de Groot, zumba instructor.

show. The message was shown on the screen in the hallway, and the subsequent response of the user—taking the stairs or the elevator—was recorded. Next, the estimate of the effect of the message that was shown was updated, and the message to show on the user’s subsequent visit was determined by the algorithm as described in Chapter 6.

8.2.2 System Evaluation

To evaluate the APStairs system, it was employed for five weeks in the hall of an office building. To be able to test the effects of individual level adaptation as opposed to the mere effect presenting persuasive messages, half of the visitors of the building received dynamically adapted messages as they entered the building multiple times, and half received a randomly selected message.

Method

The system was set up in a five-story high office building. The height of this building is in line with the average flights of stairs an employee is willing to climb (Kerr et al., 2001). Furthermore, the stairs and

Figure 8.1: Schematic setup of our adaptive persuasive system. Bluetooth scanners are present at the building entrance and at each of the floors.

elevators in this building are easily accessible and centrally located in the building. To increase the number of successful Bluetooth scans, each week a €50 voucher was given out randomly to a person who had his or her mobile Bluetooth switched on (discoverable mode enabled). Posters and flyers were distributed throughout the building to make people aware of this opportunity and to increase Bluetooth activation amongst the office workers in the building.

Users—office workers entering the building whose Bluetooth key was scanned—were randomly assigned to one of two conditions: (1) the adaptive condition, in which the system chose a random message belonging to the persuasion strategy with the highest probability of success for the identified user, and (2) the non-adaptive condition, in which the system chose a random message belonging to one of the three persuasion strategies. The system was in service for 22 workdays. During this time, each identified user was shown either an adapted message or a random message, depending on their assigned condition and the stage in the deployment. Subsequently, the user’s behavioral response was recorded. During the first 15 days of the deployment of the system, the speed of adaptation was set low (by imposing the 80% confidence bounds on the estimates as in Chapter 3) as to explore the responses

of each user to each of the strategies. In the next seven workdays, the confidence bound was limited and the available information about each user in the adaptive condition was exploited: users in the adaptive condition were only shown messages with a high estimated success rate. After collecting data, the system was also evaluated through a number of short interviews. During these interviews, we asked questions such as: “*Do you think this message was addressed to you?*”, “*Do you think the message is shown at the right moment?*”, and “*Do you think messages like this will change your behavior?*”

Results

A total of 34 office workers were scanned by three or more scanners located throughout the building multiple times, allowing us to track their behavior. However, for eight of these people, the scanning results did not lead to a clear inference pattern to judge stair or elevator usage, so these users were removed from further analysis. As a result, the analyses as presented here conducted on only 26 users. Of these 26 users, 16 were assigned to the adaptive condition, and 10 were assigned to the non-adaptive condition. The number of repeated visits by each user ranged from 2 to 16.

Figure 8.2 shows an overview of the estimates of the successfulness of the persuasive strategies employed by the APStairs system over time. For each of the two users presented in the figure, the initial estimated success of each strategy—again, messages implementing a single strategy were treated equally—corresponded to the estimates based on the pretest. Thus, the estimated success of messages implementing the *Commitment and Consistency* strategy was highest for each new user. However, due to the large prior uncertainty—which was set manually—surrounding this estimate, this strategy was not necessarily the first strategy shown to every new user entering the office lobby.

For user 94—see top panel of Figure 8.2—the first message that was presented implemented the *Authority* strategy (e.g., “Doctors recommend taking the stairs.”). This message was successful—the user took the stairs after being shown this message—and thus the estimated success of this message was adjusted upwards. During the second visit of this user, a message implementing the *Commitment and Consistency* strategy was shown. This message was also successful for this user and, thus, its estimated success for future visits was adjusted upwards. On the third visit, user 94 received an implementation of the *Consensus* strategy: “*Follow many other people; take the stairs!*”. This strategy,

Figure 8.2: Overview of the estimated effectiveness of different social influence strategies for two users of the APStairs system.

however, was not successful as the user took the elevator after seeing this message, and its estimated success was adjusted downwards. On visit four, user 94 was presented with another implementation of the *Authority* strategy, which at that visit was not successful. At this point in time, the system entered the exploit stage—the last seven days of deployment. Thus, for the fifth visit, this user was shown an implementation of the *Commitment and Consistency* strategy: its estimated effect was higher than that of the other strategies and also relatively certain. As expected, this last message was succeeded by stair usage of this user. User 14 clearly converges to a preference for the *Commitment*

and *Consistency* strategy, as was also observed for user 94¹.

To see whether there was a difference in the compliance to persuasive messages between the adaptive condition and the non-adaptive condition, the proportions of stair taking were calculated for users in both conditions. In the adaptive condition, a distinction has to be made between the explore period (with low confidence bounds and, thus, a close to random selection of messages) and the exploit period (the period in which only the message with the highest estimated success was shown). Consequently, the data collected were divided into two parts. Regretfully, the majority of the data collection for both conditions occurred during the explore stage. In this stage, both systems showed a similar random pattern of messages to users.

As expected, the proportions of stair use during the explore stage are almost equal for both conditions (see Figure 8.3). The gray vertical line marks the beginning of the exploit stage. Even though the estimated success-rates of the two systems seem to diverge—with the adaptive version of the system being more successful—this trend is not statistically significant. This means that based on the limited time deployment of this adaptive persuasive system, it is not possible to distinguish its average effectiveness from a system that uses random persuasive messages.

Figure 8.3: Comparison of the success-rates in the two experimental conditions during the 5 work-weeks of deployment of the system.

¹All observations of this user occurred during the “explore” period.

Finally, a total of 12 (possible) users—i.e., office workers entering the building at the day the system was dismantled—were interviewed for a qualitative evaluation of the APStairs system. During these short, ten minute interviews, the office workers displayed diverse responses to the system. One interviewee stated to feel that “...*the messages seemed to address me, I always tried to read them*”. Another user, however, remarked that: “*In the morning, I do not pay attention to anything*”. A number of office workers responded very positively to the system; they interpreted the system as a nice way to make people aware about physical activity. This was true also after the adaptive nature of the system was disclosed to them. On the other hand, some office workers also responded negatively; they did not think that messages like this would ever be able to change people’s behavior.

Overall, users commented that the messages were delivered using a big noticeable screen that was placed in the right location: clear in sight, and at the moment of their decision to use either the stairs or the elevator was made. Moreover, the timing and duration of the messages was found to be adequate, and the content indeed triggered people’s consciousness about stair taking behavior.

8.2.3 Discussion

The APstairs system is an adaptive persuasive system created to increase stair usage amongst office workers. The system was deployed “in the wild” for several weeks. Although the evaluation tried to compare the effectiveness of the adaptive system with a random presentation of messages, it did not collect sufficient data to reach a statistically significant conclusion. However, the APStairs system showed the adaptation process in its full breadth, and for a number of individuals the system seemed to converge to their personal most successful strategy. This demonstration and the detailed description of the implementation of our adaptive persuasive system enables designers of persuasive technologies to include personalization at the level of social influence strategies to possibly increase the effectiveness of their systems.

8.3 Design 2: Influencing Docking Behavior

The second system that implements dynamic persuasion profiles was created to increase user engagement in a health and lifestyle service. The service combines an “activity monitor”—a 3d accelerometer optimized to detect physical activity patterns—with active human and

technology initiated coaching to help users gain a more active lifestyle. During a 12 week program users can set activity goals and monitor their progress on the web service that accompanies the product.

Within the service user engagement is key: Coaching happens primarily via the web service, and activity data is only stored and analyzed after it has been updated to the web service. Updating can happen only via a physical connection of the activity monitor to the user's computer. Users that fail to upload are thus deprived of feedback and coaching and as such the persuasive technology has little impact on their daily activity pattern. To encourage docking—the uploading of the activity data to the web service—docking reminders are sent via email to users that have failed to upload for a certain number of days. The application of persuasion profiles presented here focusses on optimizing the effectiveness of these reminder email messages by including implementations of social influence strategies into the messages. The system is called the Persuasive Messaging System (PMS).

8.3.1 System Design

As with the previous adaptive persuasive system presented in this chapter, *identification*, *representation*, and *measurements* are necessary requirements to build and use persuasion profiles in the PMS. The PMS only focusses on the reminder emails and thus identification, representation, and success measurement are only defined in this context.

Identification

The PMS used the unique key provided by each activity monitor to identify individual users. When a user docks—connects their activity monitor to their computer—the activity data and a unique identifier are sent to the persuasive service. The PMS was implemented on another server that was external to the persuasive service's own system. To ensure privacy and security the user ID was hashed and a timestamp and the hashed ID were sent to the PMS server. This allowed the PMS server to log each docking event for each individual user.

Representation

Representation of the persuasive messages was done in the email reminders that were sent to users that had refrained from docking for either three or six days. Reminders were initiated by the persuasive service. The service's server sent a request to the PMS server to request the next social influence strategy to be used for the current user

(identified by their hashed ID). The PMS server, upon receipt of the request looked up the persuasion profile for that user (all stored using the hashed user ID's) and returned the text snippet of one of the persuasive email messages that were created for the PMS.

To create the social influence strategy implementations five persuasive technology researchers brainstormed a large number of messages. Messages were created that implemented the Scarcity, Authority, and Consensus strategies. After the brainstorm a card-sorting test was used to classify messages according to their strategies, and for each strategy two messages were selected for use in the trial. The persuasive messages consisted of text snippets containing social influence strategies that were added to the standard docking reminder email already in use by the service. This standard reminder mail looked as follows²:

Dear (firstname),

How are you doing? We hope all is well. It is 3 days since the last time you connected your Activity Monitor.

[Message]

We would like to remind you to connect it to your PC soon and stay in touch with [X]³.

Sincerely,

The [X] Team

The social influence strategy was inserted at the *[message]* location of the email reminder email. Table 8.2 gives the implementations of the social influence strategies as used in the PMS. Since the original service reminder message was also used there are four different social influence strategies in use (one being neutral) in the PMS. By combining the hashed User ID with the message ID the service's reminder server was able to dynamically construct a message for a specific individual users of the persuasive service.

²The email reminders and social influence strategies were translated into four languages, English, Dutch, German, and Spanish, to facilitate all users of the service.

³The reference to the name of the service is omitted in this thesis.

To enable estimation of the possible effect of these messages each of the messages was presented to $N = 80$ participants in a pre-test. Participants were instructed to read each of the (full) messages and answer the question “*This message would motivate me*” on a seven-point (Totally Disagree to Totally Agree) scale. Scores over two implementations of the social influence strategies were averaged, and average scores for each strategy were subsequently used to estimate the successfulness of the different social influence strategies at an average level. The neutral message had the lowest evaluation: $\bar{X} = 3.46$, $SD = 1.44$. The messages implementing social influence strategies scored higher, with authority scoring highest, $\bar{X} = 4.21$, $SD = 1.59$, before consensus, $\bar{X} = 3.96$, $SD = 1.54$, and scarcity, $\bar{X} = 3.81$, $SD = 1.52$.

Table 8.2: Influence strategies and their implementations in persuasive messaging system.

Strategy	Implementation
<i>Neutral</i>	...
<i>Scarcity</i>	<ol style="list-style-type: none"> 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]! 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.
<i>Authority</i>	<ol style="list-style-type: none"> 3. Experienced [X] coaches recommend frequent uploads of your activity data. This will help you to gain more insight and be more active! 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!
<i>Consensus</i>	<ol style="list-style-type: none"> 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!

Measurement

The service’s email server, after consulting the PMS, sends emails to remind users to dock their activity monitor. Hence, the reminder message containing a specific social influence strategy is successful if, within a certain time period after reading the email, the Activity Monitor is indeed docked. To measure this effectiveness a small image was inserted into the email message body which allowed the PMS to log the fact that a user with a specific hashed ID opened an email. If, and only if, within 24 hours after opening the email the user with that ID docked her activity monitor the message was considered a success and thus whatever influence strategy implemented in the opened message (neutral, scarcity, authority, or consensus) was successful for that user. The PMS ran a cron-job every 24 hours to match all opened emails with the recent docking behavior and update the individual level persuasion profiles according to the responses to messages send the last 24 hours.

8.3.2 System Evaluation

To evaluate the PMS system an evaluation was setup 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. The system was implemented on a separate server external to the persuasive system.

Method

To assess the effects of adaptive persuasive messaging as opposed to the original reminder message, or messages using social influence strategies that were not adapted to the individual, the system evaluation assigned new users during the evaluation period to one of four conditions:

1. *Baseline*: Users assigned to this condition received the standard docking reminder message. This message did not contain any implementations of social influence strategies. This condition was included to be able to compare the PMS to the current implementations
2. *Best Pre-tested*: Users assigned to this condition received randomly one of the two messages implementing the authority advice—this message was judged most motivating in the pre-test evaluation of the messages. This condition was included to compare adaptive selection of social influence strategies to the “best” average strategy.
3. *Random*: Users assigned to this condition received randomly one out of the seven versions of the message (with probabilities equal

for each of the strategies). This condition was included to compare adaptive messaging with alternating messages.

4. *Adaptive*: Users assigned to this condition received messages suggested by the adaptive persuasive system algorithm as described in Chapter 3. Thus, for the first few messages the selection was random. If users displayed a clear preference for one of the social influence strategies after receiving multiple reminder emails the reminder message was personalized to include only those strategies users were susceptible to.

No interventions other than the changes in the email messages according to the conditions were implemented in the evaluation. For each user the PMS logged each dock event, the email reminders that were sent and opened, and whether or not an email—implementing a specific strategy—was a success or not. This data allows analysis of the direct effects of the strategies, but also allows analysis of “drop-outs”: people that have not been active—have not docked—in the program for 28 days. This latter measure is most important for the service since this is exactly what the docking reminder messages are trying to prevent.

In the adaptive condition—similar to the APStairs system—the prior expectancy of the success of the different social influence strategies had to be set. Before the trial no information was available about the effects of the reminder message and thus the estimates were (a) set close together, and (b) set with large uncertainty to be updated quickly by new data. 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 strategy prior was set the highest, $\bar{X} = 0.52$, $Var = 0.1$, before consensus, $\bar{X} = 0.50$, $Var = 0.1$ and scarcity, $\bar{X} = 0.47$, $Var = 0.1$. Contrary to the APStairs system, randomized probability matching was used to select messages in the adaptive condition.

Results

For this analysis all users that (a) participated in the service for at least 30 days, and (b) received at least 3 email reminders during the trial period, 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 8.3 gives an overview of the number of users, the success percentage, and the number of drop-outs in each condition.

The table gives an overview of the aggregated results of the evalu-

Table 8.3: Overview of the data from the persuasive messaging system evaluation.

Condition	Users	% Success [S.E.]	% dropouts [S.E.]
<i>Baseline</i>	271	28.49 [1.7]	25.09 [2.6]
<i>Best Pre-test</i>	289	24.01 [1.5]	25.95 [2.6]
<i>Random</i>	289	25.41 [1.6]	20.76 [2.4]
<i>Adaptive</i>	280	26.49 [1.6]	18.93 [2.3]

ation of the PMS system. It is clear that users are relatively equally distributed over the conditions. Furthermore, at first glance the effects on the average number of reminders and the overall success of the reminders of the experimental conditions is rather low. However, the number of dropouts is lower in the adaptive messaging condition, indicating that the adaptive persuasive system is somewhat more successful in keeping users engaged in the program.

To analyze the data obtained in the PMS evaluation a series of multilevel models is fit to the data describing the successes of each of the reminders send to users included in the trial. Similar to the model(s) presented in Chapter 3 a null model—although this time with a logit link given the binomial outcome of the success measure—is first fit to the data. Thus, the model estimates the success probability of each message that is send to users based on an overall intercept and per-person intercepts that are distributed normally with mean zero and a variance estimated from the data.

Adding average effects for the social influence strategies to the null-model shows no significant effect of strategy on success of the emails, $\chi^2 = 4.75$, $df = 3$, $p = 0.19$. Thus, this average effect is omitted from further model comparisons. Further model comparisons show a large main effect of Frequency—the number of the reminder that is send (See Table 8.4, Model *A* and Model *B*)—on the probability of a reminder being successful. Addition of average effects of condition to Model *B* shows that there is no significant average effect of conditions, $\chi^2 = 3.19$, $df = 3$, $p = 0.36$. There is however a significant increase in model fit when adding varying social influence strategy effects by participant (Table 8.4 Model *B* and Model *C*) which is consistent with the findings of large heterogeneity in responses to social influence strategies presented in Chapter 3. Thus, allowing for heterogeneity in the effects of the different influence strategies explains a significant portion of the

variance in responses to the PMS⁴. Finally, condition interacting with frequency is added to the model. This interaction significantly improves model fit (See Table 8.4 Model *D*)

Table 8.4: Model Comparisons used for the analysis of the persuasive messaging system evaluation.

Model	BIC	logLik	χ^2	<i>Df</i>	<i>p</i>
Model <i>A</i> :	10903.18	-5449.59			
Model <i>B</i> :	10486.91	-5240.46	418.26	1	< 0.001
Model <i>C</i> :	10480.00	-5228.00	24.92	9	< 0.01
Model <i>D</i> :	10474.54	-5222.27	11.46	3	< 0.01

Table 8.5 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 other interactions of condition and frequency can be interpreted in the same way: For the random condition—compared to the baseline condition—the drop in effectiveness of the reminders over time is lower, while that of the best-pretested condition is higher (although both are not significantly different from 0). The drop in effectiveness of messages in the adaptive condition is *significantly* lower than the baseline condition: The predicted effectiveness of the 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 condition. For the tenth message this difference is even larger: 4.8%. Similarly, the adaptive condition significantly outperforms the pre-tested condition, $t = 3.74$, $p < .01$. The adaptive condition does not significantly outperform the random condition.

Finally, a key comparison of the evaluation of the PMS system concerns the overall drop-out rate obtained in the different conditions. While table 8.2 already shows that the drop-out rate in the adaptive condition is lower than in the other conditions, a logistic regression model is fitted to the drop-out data to further examine this result. Model comparisons show that the effect of Condition on drop-out is not statistically significant, $\chi^2 = 5.55$, $df = 3$, $p = 0.14$. Table 8.6 shows the

⁴This model comparison is similar to that presented in Study 1 in Chapter 3 between Model *B* and Model *C*. See also Equation 3.6 in Chapter 3.

Table 8.5: 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 condition than in the baseline condition.

	Estimate	<i>S.D.</i>	<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. $\times$ Random	0.02	0.02	1.10	0.27
Freq. $\times$ Pre-tested	-0.01	0.02	-0.40	0.69
Freq. $\times$ Adaptive	0.04	0.02	2.74	< 0.01

estimated coefficients of the logistic regression model predicting drop-out probability in the different experimental conditions. Again, the baseline condition was used as a reference. The model fit shows that the estimated drop-out probability for users in the baseline condition is 25.1%, while that in the adaptive messaging condition is only 18.9%. This difference however is also not statistically significant.

Table 8.6: Fixed effects of the logistic regression model predicting the drop-out probability of users in the PMS trial.

	Estimate	<i>S.D.</i>	<i>z</i>	<i>p</i>
(Intercept)	-1.0937	0.1401	-7.81	0.000
Pre-tested	0.0452	0.1940	0.23	0.815
Random	-0.2457	0.2017	-1.22	0.223
Adaptive Message	-0.3610	0.2071	-1.74	0.081

8.3.3 Discussion

The second design presents the development and evaluation of the Persuasive Messaging System for the docking-reminder messages. In this adaptive persuasive system users are identified by a unique key associated with their activity monitor which is an integral part of their usage of the service. After inactivity—failure to dock—for a period of 3 or 6 days users received a reminder email. In this email the authority, consensus, and scarcity strategy were implemented to increase compliance. The social influence strategies were added to the email messages in such a way that they were interchangeable and could thus be personalized. Finally, the effect of the messages was measured by combining logging

of the opening of the email messages via a dynamic image in the content of the email and user's logged docking behavior.

Results of the evaluation of the PMS system show the benefits of using persuasion profiles: the (repeated) success of the reminder messages is higher when using personalized persuasion than when using the default message or any of the other message creation methods explored in the evaluation. More importantly, in the long run, this led to a decrease in drop-out of the program for this condition compared to the baseline condition. Even though this signal is not very strong and only marginally significant, $.05 < p < .10$, the estimated business impact is substantial: on average users that receive personalized reminders were 6.2% less likely to drop out of the program than those receiving the original reminder message. This large effect size warrants further investigation. The current estimates, and more specifically the estimated difference between the random and the adaptive condition, is possibly not statistically significant because of the limited time period of the study: a portion of users that *would* drop out in the future has not done so during the evaluation. A study with a larger sample size and over a longer period could improve the precision of these estimates.

8.4 Design 3: E-Selling Applications

The third adaptive persuasive system presented in this chapter focusses on the application of persuasion profiling in a marketing setting. Historically (online) marketing has been on the forefront of persuasive technologies. The e-commerce application presented here provides an opportunity to test the premises of persuasion profiling on large groups of people and determine the economical relevance.

This third system—called the Persuasion API—is developed on top of an existing e-commerce website. The website kinder-kleertjes.com is an affiliate website that sells children's clothing primarily to Dutch customers. Affiliate websites use product feeds—xml files describing the product collection—of other online vendors to create a customer facing website aimed at search engine optimization and user experience. Users can browse the affiliate website for products. Purchase of products however happens at the website of the vendor providing the product feed. The affiliate website receives a share of the revenue that is generated from the lead originating on the affiliate site.

Kinder-kleertjes.com offers a selection of over 1.200 products. Products are offered by two affiliate programs and the website aims at at-

tracting traffic through search engines and increasing click through to the two final vendors. The site, in its current form, has been running from the beginning July 2010. The website is rather limited in size, attracting 388 visitors a month on average during from July 2010 till April 2011 ⁵.

8.4.1 System Design

The Persuasion API system was created on top of the original site and implemented as a Application Programming Interface (API). The website *kinder-kleertjes.com* makes a HTTP call to an external server to request the appropriate social influence strategy to use for the current visitor. The remote server returns the ID of the social influence strategy that should be used and this strategy is presented to the visitor of *kinder-kleertjes.com*. Finally, after browsing the product that is accompanied by the suggested social influence strategy *kinder-kleertjes.com* sends an update to the Persuasion API server to communicate whether or not the influence attempt was a success. Also in this system *identification*, *representation*, and *measurement* are necessary requirements to build a system that adapts social influence strategy usage at the individual level.

Identification

To identify users a Javascript visitor tracker which assigned a unique id to each new visitor and stores this as a cookie on the user's machine was created. For each new user, based on their IP address and the current time, a unique User ID was created and send to the Persuasion API server. This unique ID is also stored on the machine of the user. Every time a user visits *kinder-kleertjes.com* the client side Javascript tracks whether a Persuasion API User ID exists, and if not a new User ID is created.

Representation

To be able to create the adaptive persuasive system, the presentation of different social influence strategies to support the product offerings needs to be possible. The home page of the *kinder-kleertjes.com* presents a "random" selection of the offered products together with pictures of the products and single sentence descriptions (See Figure 8.4 on the left). Once a visitor clicks on one of the products (or enters the site using a search term directly pointing at a product page)

⁵The timeframe during which the system evaluations took place.

the product is displayed using a large image and a textual description. The social influence strategies were implemented on this product display page. Only two strategies were implemented for this evaluation: Scarcity and Consensus. For the scarcity implementation a image that said “Special offer” and the text *“This clothing item is available today at a special discount rate”* was displayed. For the Consensus strategy an image that stated “Bestseller” and the message *“This is one of our bestselling clothing items. This product is popular with many”* were added. By keeping the original “no strategy” product representation the evaluation concerned three product display versions that could be used for every product on the site (an example of the consensus strategy implementation is presented on the right of Figure 8.4).

Figure 8.4: Overview of kinder-kleertjes.com: Left the home page, in the middle a product presentation page. The implementation of the persuasive strategies are shown on the right with their position on the product presentation page indicated by the arrow.

Measurement

Also for this adaptive persuasive system a measurement of effect of the social influence strategy is necessary. Since the goal of the affiliate store is to obtain high click-through rates to the vending website of the products that are offered, this click-through was used directly as a measure of success: Any click on the image or link of the product on the product page implementing a social influence strategy was regarded a success for that strategy. A failure to click—because the user browsed to a different product or left the website—was regarded as an unsuccessful influence attempt. The success of the attempt was communicated back to the external server running the Persuasion API system. When a user left one of the product pages the UserID, the social influence strategy ID, and the success of the social influence attempt, were send to the

server and the estimated success of that social influence strategy for the current user was adapted accordingly. Every 12 hours the average level estimates were also updated based on the new observations.

8.4.2 System Evaluation

To evaluate the implementation of the Persuasion API server on *kinder-kleertjes.com* two separate tests were setup. First, the system was implemented for a period of 3 months and the performance of the *kinder-kleertjes.com* website during the period in which the system was setup was compared with its performance on a number of Key Performance Indicators (KPI's) *prior to the implementation of the system*. Second, during a 4 month period *half* of the visitors of the *kinder-kleertjes.com* website were randomly assigned to the version of the website that implemented Persuasion API, while the other half of the visitors during that time period was assigned to browse the original website. This second evaluation has the methodological advantage that the timeframe of the comparison is the same between both the original *kinder-kleertjes.com* website and the version of the site that implements persuasion profiling.

In both evaluations initial estimates of the success of the different social influence strategies, and the certainty around these estimates, were needed to start the system. The estimate of the effectiveness of the *No Strategy* version was manually set to 20% ($p_m = 0.2$), that of the *Scarcity* strategy to 30% and that of the *Consensus* strategy to 24%. Thus, in line with the literature, the expectation was that the product offerings supported by implementations of social influence strategies were more successful then the original “no strategy” representation. By setting the variances of this estimate high (e.g. $\sigma_A^2 = 0.1$) these initial estimates were very variable to new data and thus had practically no impact once the system was in place for a period of time. Practically this implementation ensures that the first user after deployment of the adaptive persuasive system did not necessarily receive the *No Strategy* product presentation as had been the case in the period prior to the deployment of the system but received one of the three versions—which one was in this early stage “random” due to their overlapping confidence intervals. Based on the behavior of this first user—either clicking or not clicking on the product—a new estimate of the effectiveness of the consensus strategy was made and the product presentation was updated accordingly both for the specific user as well as for subsequent users. Strategies were selected using randomized probability matching as described by Scott (2010).

To evaluate the adaptive persuasive system a number of indicators are of interest. First of all it is interesting to see how the effectiveness estimates of the different influence strategies change over time. This can be done both at a between-user level (corresponding to A-B testing that is flexible over time) but also at a within-user level to see whether indeed strategy preferences differ between users. Next to examining the functioning of the adaptive persuasive system a comparison can be made (a) between the period in which the Persuasion API was implemented with the period in which only the *No Strategy* product representation was used, and (b) between the performance of the system for users randomly assigned to the Persuasion API system and those assigned to the no strategy version of the affiliate store. Comparisons of click through rates from the affiliate site to the vendor site per unique visitor and revenues per unique visitor are compared for each of the methods of evaluation. Evaluation *I* focussed primarily on the functioning of the adaptive persuasive system by looking into the behavior of individual users. Evaluation *II* focusses primarily on the comparison of Key-Indicators.

Results evaluation I: Comparison over time

During the baseline period—from the 8th of July till the 21st of October 2010—1339 visited *kinder-kleertjes.com*. Due to organic visitors and advertisements via Google Adwords to ensure sufficient traffic, 831 visited the site during the employment of the adaptive persuasive system, from the 21st of October until the 8th of November.

The first analysis concerns an overview of the effect of social influence strategies during the deployment phase of the system. Figure 8.5 shows the progression of the estimated (between-user) effectiveness of the three product presentation versions over time. It is clear that initially the estimates were very variable and were haphazardly updated based on new observations. However, after a few days the estimates quickly stabilized and estimate uncertainty decreases. Both product representations that implement a persuasive strategy significantly outperform the *No-strategy* implementation which was previously used on the website. This replicates the social science finding that implementations of these strategies at an average level are effective to increase compliance.

While figure 8.5 indicates that using persuasive strategies can aid the performance of the affiliate store by increasing click through rates, this graph alone does not indicate that different strategies should be

Figure 8.5: Progression of between-user (overall average) effectiveness estimates of the three strategies and indications of the certainty of the estimates. Initially the estimates are very variable but over time they stabilize. The dates on the x-axis are jittered to increase readability. Error bars are conservative and based on between user variance not on the total number of observations.

presented to different users—as is done using the Persuasion API system. To examine whether adaptation is actually necessary (e.g. for some users the *Consensus* strategy—which is most effective overall—is not the preferred strategy) figure 8.6 presents the individual level estimates of two users who viewed at least 8 products on the website. The figure shows (e.g.) that for a the first users the most effective strategy is *not* the *Consensus* strategy but rather the *Scarcity* strategy. Hence, when adapting to user preference only on a between user level—thus displaying the most effective strategy overall to every user—one would obtain suboptimal results for a number of users.

Next to examining the predicted strategy effectiveness, analysis concerns the performance of the affiliate online store *kinder-kleertjes.com* before and after usage of the Persuasion API system. The first interesting indicator of the performance of the affiliate website is the proportion of clicks on products per user. During the baseline period 14.4%

Figure 8.6: Estimated effectiveness of the different persuasive strategies for two randomly selected users (out of those who viewed more than 8 products)—x-axis denotes the number of products viewed.

of the users of the site eventually clicked on one of the products and was taken to the vendor's homepage. In the adaptive persuasion period this number was higher: 18.3%, indicating an increase in click through in the adaptive persuasion period. This increase is statistically significant ($\chi^2 = 5.766$, $p = 0.016$). Furthermore, during the baseline period each unique visitor on average created a revenue for the affiliate store of €0.037, while in the adaptive persuasion period €0.046 were generated per visitor. This latter difference is however, due to the relatively small incidence of actual purchases, not statistically significant. The average revenue per visitor is thus higher in the adaptive persuasive system period than in the baseline period.

Results evaluation II: Comparison during same timeframe

Evaluation I concerned a comparison of the performance of kinderleertjes.com prior to implementing the adaptive persuasive system and during implementation. However, since these time-frames are different this is a mayor confound in the results presented: the conversion might have been higher due to the effects of time, for example consumers might be more inclined to buy product in autumn than in summer. To control for this confound the second evaluation considered a single timeframe in which half of the consumers used the original no strategy version of the website, while the other half was shown the website implementing the

Persuasion API system. During the period between the 20th of February and the 2nd of April 1449 visitors visited *kinder-kleertjes.com*.

Analysis first again concerns the comparison of the three different product representations: The *no-strategy* representation, the Scarcity representation, and the Consensus representation. Figure 8.7 shows the performance of these during the time of the trial. Again, it is clear that the presentations accompanied by social influence strategies lead to higher conversion than those. Contrary to evaluation *I* however the persuasiveness of the scarcity implementation seems to “wear off”.

Figure 8.7: Progression of between-user (overall average) effectiveness estimates of the three strategies and indications of the certainty of the estimates. As in the first evaluation the products displayed accompanied by a social influence strategy *on average* have higher conversion rates.

The setup of the second evaluation enables comparison of the no-strategy system performance with the adaptive system performance within the same timeframe. Figure 8.8 shows the estimated conversion rate of each of the two systems during evaluation *II*. While initially estimates are unstable due to small sample sizes, the estimates stabilize quickly and already after one month of system usage the conversion rate

of the adaptive persuasive system is consistently and significantly better than the performance of the original version of the affiliate store. During the second evaluation the average conversion rate of the no strategy version of the site was 9.4% while that of the adaptive persuasive system was 13.5%. This difference in proportions is statistically significant ($\chi^2 = 6.3856$, $df = 1$, $p = 0.012$). As in evaluation *I* the adaptive persuasive system also led to an increased revenue per visitor: €0.034 vs. €0.041. However, also in this study, due to the relatively rare occurrence of purchases this latter difference is not statistically significant.

Figure 8.8: Comparison of estimated conversion rates for consumers in the holdout condition—those visiting the no strategy version of the website—and consumers in the test condition—those looking at the adaptive persuasive system.

8.4.3 Discussion

The third design presented in this chapter tested the use of adaptive persuasion in an e-commerce setting. The Persuasion API system that used persuasive strategies and dynamically adapted these strategies to individual consumers improved conversion in two separate evaluations. The results again show the use of persuasion profiles: for each individual user a profile (see for examples Figure 8.6) which estimates of the effects of the different persuasion strategies was created and used. In

both evaluations a significant increase in conversion was obtained—as compared to the original implementation of the affiliate store—when using the adaptive persuasive system.

8.5 Conclusions

This *case study* chapter presented the design and evaluation of three adaptive persuasive systems. For each of the systems different means of *identification*, *representation*, and *effect measurement* were used to enable creation and usage of persuasion profiles. In each system user's responses to distinct social influence strategies could be identified, and these differences were attended to by the systems. The description of these systems and the evaluations of the effects of adaptive persuasion should motivate designers to use persuasion profiles in the design of their persuasive systems.

For the latter two of the three case studies presented in this chapter, attending to individual differences in responses to influence strategies led to a significant increase in compliance compared to not using influence strategies at all. Both the docking frequency of the activity monitor as well as the frequency of ordering products online increased by using an adaptive persuasive system. These case studies thus (a) are exemplars of implementations of adaptive persuasive systems that use dynamic persuasion profiles and (b) strengthen the results brought forward in the previous chapters that personalized persuasion indeed increases compliance.

8.5.1 Meta-Analysis: Personalization vs. Random Selection

The result of a direct comparison between a personalized persuasive system and system that uses a random selection of implementations of influence strategies is not as clear as the comparison of personalized persuasion with not using influence strategies at all. The comparison between personalization and random selection was only powerfully made in Study 2 of this chapter and Study 3 of Chapter 7. In both of these studies the estimates of the success of the personalized condition over time were higher than those of the random condition. These differences were however not (or only marginally) statistically significant.

Given the similarity of the results obtained in these two studies that compare the effectiveness of random message presentation over time with personalized messaging over time it is however possible to combine the data obtained in both studies to perform a meta-analysis.

Figure 8.9 shows a forest plot indicating the effect size of the comparison between the random and the personalized conditions over time (using unstandardized model coefficients) and the outcomes of the estimated meta-analysis model. For this comparison the most important dependent variables of each of the two studies are used: the estimated reduction in snacking consumption is used from Study 3 of Chapter 7, and the change in estimated success rate of the email reminders is used from Study 2 of this chapter. For both studies the coefficient which indicated the change in the effectiveness of each subsequent persuasive message is used in the analysis.

The meta-analysis shows a small but statistically significant difference between the two different ways of selecting influence strategies, $z = 2.08$, $p < .05$ ⁶. Even though the method of personalizing differs between the two studies—in Study 3 of Chapter 7 the personalization was based on meta-judgmental measures while in Study 2 of this chapter operative measures were used—the results of the meta analysis show a difference in the effectiveness of personalized persuasion as compared to random message selection. This latter analysis, combined with the fact that the estimated effects of personalized persuasion are large enough to be of practical importance, justifies further studies into the effects of personalized persuasion, as well as the use of (adaptive) personalized persuasion in practice.

Figure 8.9: Forest plot of the meta-analysis using a fixed effect model of the two studies comparing adaptive and random strategy selection.

⁶The meta-analysis was done using the [R] `metafor` package. A fixed effects meta-analysis model was estimated using the model coefficients reported earlier.

8.5.2 Closing remarks

The three designs presented in this case study chapter, contrary to the studies presented in Chapter 7, used operative measures of susceptibility to persuasion to *dynamically* derive persuasion profiles. The two methods of profiling presented in Chapter 7 and Chapter 8 are however easily combined: Meta-judgmental measures of susceptibility as obtained using the *STPS* can be used as a starting point for a dynamic profile, instead of using the average response to a distinct strategy as done in the designs presented in the current chapter. Meta-judgmental measures can thus be used to (partly) overcome the cold-start problem that many learning algorithms face (Lam et al., 2008), while dynamic adaptation can overcome changes in user’s responses to social influence strategies due to time, context, or a suboptimal self-assessment ⁷.

⁷This chapter is (partly) based on earlier publication(s): (Kaptein, 2011b; Sakai et al., 2011; Kaptein, 2011a; Kaptein and van Halteren, 2012).

9

Reflections

9.1 Advancing the Design of Persuasive Technologies

This thesis advances the persuasive technology field by examining ways in which persuasive technologies can personalize their persuasive attempts to gain compliance. The persuasive technology field consists of a mixture between technology and social science: primarily findings from social science have been used to create and design technology mediated or initiated interventions. The work in this thesis first adds to our understanding of human decision making as influenced by the usage of social influence strategies, and second uses this understanding to design and evaluate several persuasive applications. In doing so the thesis opens up a new design space: that of *personalized persuasive technologies*.

The three insight generation chapters (Chapters 3, 4 & 5) in this thesis first of all replicate the effects of influence strategies on human decision making on average: on average people's evaluations of products were positively influenced by implementations of social influence strategies, participants were more likely to change their own opinions, and click-through on advertisements was improved. However, these insight generation chapters also make clear that these average effects are only a part of the story, and perhaps a relatively uninteresting part: the dif-

ferences between people in their responses to social influence strategies are far larger than the similarities, up to a level where a strategy such as authority can consistently, over time and context, have a negative effect on *over one third of the people*. Chapter 3 clearly showed that attending to these effects is more beneficial to predict behavior than attending to the average effects.

From this understanding, questions about measurement of these “susceptibilities” of individuals to distinct influence strategies, and the effects of possible combinations of multiple influence strategies, naturally arise. The answers to the latter question are presented in Chapter 4: combinations of multiple influence strategies lead to compliance (on average) that is as low as the worst strategy in the context of study. Chapter 5 shows that susceptibilities to influence strategies can be measured *a priori*: meta-judgmental measures of susceptibility obtained using questionnaires aid in our understanding and prediction of people’s responses to social influence strategies. These three *insight generation* chapters together stress the importance of attending to individual differences in responses to social influence strategies and are exemplar for the estimation of these differences: the multilevel models presented in Chapter 3 enable quantitative comparison of the average effects with the individual level effects.

The results presented in the *insight generation* chapters should inspire technologists to create personalized persuasive technologies. Influence strategies prove valuable to do so. People’s susceptibilities to distinct strategies differ, and these differences are large enough to motivate personalization. Also, the differences can be measured and can thus be attended to. These conclusions lead to the development of *persuasion profiles*: Individual level collections of estimates of the success of distinct social influence strategies (and the certainty around these estimates). Persuasion profiles are thus estimates of the Strategy $\times$ Individual effects.

Theorists and practitioners have—and rightfully so—also identified Strategy $\times$ Context and Strategy $\times$ Product effects (see, e.g. Payan and McFarland, 2005; Midden et al., 2008). Thus, the persuasive strategy that is most effective does not *only* depend on the individual, but also on the context and the behavior change goal at hand (This is also clear in Study 2, Chapter 4). However, this poses no threat for the validity and practical usefulness of *persuasion profiles*: as long as the Strategy $\times$ Individual effects are larger than three way Strategy $\times$ Individual $\times$ Context (or product or ...) effects then these average effects can easily

be controlled for. Specifically, a persuasion profile can be recomputed in each new context to take Strategy $\times$ Context effects into account. The newly obtained profile can guide strategy selection in the new context. One of the goals of persuasion research should be to estimate the sizes of these different (interaction) effects and compare these to each other.

Chapter 6 describes two distinct ways of creating a persuasion profile: one based on meta-judgmental measures and one based on operative measures. Historically meta-judgmental measures of individuals have been easier to obtain and as such have been the major tool of social scientists studying individual differences. In line with this tradition the *STPS* (See Chapter 5) was presented to obtain measurements of individual level susceptibility to social influence strategies. However, thanks to the advances in interactive technologies, measurement of behavior is also within reach of social scientists. This enables usage of operative measures: measures of the effects of stimuli based on the process being in-play instead of being reflected on. Both sources of information are valuable to start and update a persuasion profile.

The two *case study* chapters address these two sources of information about people's susceptibilities separately: Chapter 7 presents the evaluations of several health related social influence technologies that are personalized based on meta-judgmental measures, while Chapter 8 extends this work by showing examples of dynamic personalized persuasive technologies based on operative measures. A combination of the two is straightforward and should be attended to in the future. In Chapter 7 the average effects of choosing the "wrong" influence strategy for distinct individuals are also shown: Persuasive attempts can have an overall negative effect if the wrong strategies are deliberately selected. This highlights the importance of personalized persuasion in an applied setting.

The three persuasive systems presented in Chapter 8 are more advanced, and should mark the area of persuasive technologies to come. Via *identification*, *representation*, and *effect measurement* designers can build adaptive personalized persuasive technologies in ambient intelligence. These systems are not only more effective than their non-adaptive or non-personalized counterparts, they also provide opportunities to study human behavior that have hitherto been unimaginable. The effects of social influence strategies can, due to technology which unobtrusively measures behavioral outcomes, be estimated over time and within individuals. This combination of adaptation and measurement is the core of the contribution of this thesis to the persuasive

technology field.

9.2 The Possible Future(s) of Persuasion Profiles

This thesis presents *Persuasion Profiles* as a means of personalizing persuasive technologies. Persuasion profiles are collections of estimates of the effect of distinct social influence strategies (over implementations) for individuals. By also including the certainty around these estimates, persuasion profiles can be used to select social influence strategies to target individuals while optimizing the explore vs. exploit trade-off (Scott, 2010). Persuasion profiles seem relatively stable over time—they describe people’s susceptibilities as a *trait*—and are possibly also useful for selecting social influence strategies *across contexts*.

Persuasion profiles can be built using meta-judgmental measures, operative measures, or a combination of both. As suggested in Chapter 6 meta-judgmental measures can function as a starting point for a persuasion profile, while operative measures can be used to update the estimates and their associated certainty in later stages. While the profiles presented in Chapter 7 and 8 were relatively simple—e.g. social influence strategies were modeled independently of each other—it is possible to include correlations between strategies, average Strategy $\times$ Context effects, or more traditional marketing variables such as demographics and personality into the estimation process. Each of these variables is likely to improve the accuracy of the estimated effects of social influence strategies. Consequently, including these variables might improve the effectiveness of persuasive systems whether they focus on health behaviors, education, marketing or any of the other application areas of persuasive technologies.

In this thesis persuasion profiles have been built based on the social influence strategy taxonomy proposed by Cialdini (2001). This taxonomy provides a starting point for the classification of different influence attempts, as well as for the creation of different implementations of influence strategies. Cialdini (2001)’s taxonomy however is not the only one that is of possible use to design adaptive persuasive systems. Estimates of the number of social influence strategies range from four to over a hundred (see: Rhoads, 2007) and different granularities are probably useful in different domains. It is however not necessary to settle on a *definite* taxonomy to enable wide-scale usage of persuasion profiles: operative measures and their *correlations* can be used to distinguish amongst implementations and strategies. Likely, it will be beneficial

to start with a sufficiently useful taxonomy, such as the one adapted in this thesis, while enabling profiles to contain more (or less) strategies based on distinctions that prove effective. Thus, data collected to build persuasion profiles should distinguish between strategies and their implementations. Next, adaptive persuasive systems should be flexible enough to identify clusters of implementations which subsequently can lead to the discovery of new strategies¹.

Given an accurate persuasion profile it is useful to envision the future use of such a profile. Like any profile that is already created by interactive technologies persuasion profiles could potentially be shared across systems and services, be combined with other profiles, and be disclosed to their users. The next sections describe each of these in more detail.

9.2.1 Sharing

Imagine the following scenario ²:

“Eleven-year-old Joey has resisted following his doctor’s weight-loss plan—skipping his daily bike ride to play video games, trading his healthy lunches for soft drinks, and sneaking snacks at night. His dad, Martin, is frustrated. Joey seems frustrated too. Concerned about some recent bullying, Martin goes online to check the parental-control settings on his son’s social network account. He then sees something he’s never noticed before: Joey’s persuasion profile. During his extensive use of the site, Joey has been shown countless ads for products, games, and other websites. Each time the persuasion profile recorded Joey’s response: Did he click the ad? Did he play the game? Reading on, Martin learns that Joey rarely acts on recommendations from experts or celebrities but likes what his peers like and does what his peers do. Martin thinks, “I wonder if we could use that to help Joey.” He finds an online weight-loss program that helps kids motivate each other by sharing their successes.”

In this scenario, the knowledge gained about Joey in the context of his social network service is used by his dad to personalize persuasion in a health related setting. Obviously, this type of sharing of persua-

¹It has to be noted that taxonomies that are of large granularity (e.g. are composed of a large number of strategies) will in practice require more observations to create a persuasion profile.

²This scenario is taken from (Kaptein et al., 2011c).

sion profiles is not limited to human interventions: Technologies could actively leverage persuasion profiles built up in one context (e.g. people’s responses to social influence strategies in online marketing) and use these in another context, such as the emerging health related persuasive technologies. If the Individual $\times$ Strategy $\times$ Context effects are small compared to the Individual $\times$ Strategy effect, this will lead to higher success rates of influence attempts by “borrowing strength” from estimates obtained in the other context. Higher success rates of influence attempts due to sharing of persuasion profiles are of value to designers of persuasive systems, and as such persuasion profiles can become part of the marketing currency of the future (Kaptein et al., 2011c). Once a single specification for the representation of a persuasion profiles has been established—be it as simple as a *JSON* standard—sharing, and trading of profiles is within our near future.

Shared persuasion profiles can however not merely be used for good: While most persuasive technology researchers seem to agree that improving the effectiveness of a weight-loss program is a noble cause, other purposes are well imaginable. Having learned, on a social network site, that a specific individual is more swayed by consensus arguments, a local political party might decide to personalize the information presented in their advertising emails. A goal or practice largely deemed less noble. These possibly detrimental effects of sharing persuasion profiles should actively be considered.

9.2.2 Combinations with other profiles

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. In this thesis the two have been treated as unrelated, but relationships between ends and means are likely. Thus, *both* profiles would benefit from integration.

Next to combinations with end personalization, persuasion profiles

are also likely to 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 women 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 social influence strategies are likely correlated with gender, age, occupation, etc. etc. It is thus most likely beneficial for the effectiveness of influence attempts to combine these 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 — as well as 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 strategies — means — based on the measurement of user responses to persuasion as opposed to expert judgments. However, both approaches could be combined: expert determined rules could influence the probability of content selection *a priori*, while dynamic adaptation could be used to update these probabilities.

9.2.3 Disclosure

Next to sharing persuasion profiles, and combining them with other existing profiles, there is an opportunity to disclose persuasion profiles to individuals. A first attempt to examine the effects of this disclosure was made in Chapter 4, Study 1, in this thesis. The current practice of profiling and filtering, of which persuasion profiles are merely one example, that is more and more an integral part of interactive technologies is bound to have its effects on individuals and on society as a whole. Authors like Pariser (2011) persuasively describe the possible detrimen-

tal effects of personalization and as such the ways in which researchers, designers, and practitioners deal with individual level profiles should be a part of active consideration.

It is however in no way necessary for persuasion profiles or other profiles to remain undisclosed to users. Design efforts can make profiles graphically inspect-able, editable, or removable by users. These actions allow users to gain control over the filtered world that is more and more a part of their every day lives.

Persuasion profiles, contrary to many of the other profiles collected for personalization, have a distinct property that makes their disclosure even more useful for users. Most end-oriented profiles, such as those stored by recommendation engines, provide little knowledge to the user that he or she was not already aware of (Awad and Krishnan, 2006; Goy et al., 2007). Persuasion profiles however likely collect and store information that is not part of the conscious deliberation of users. This thought is supported by the lack of predictive power of meta-judgmental measures over operative measures. Apparently people respond consistently to social influence strategies, in ways they have difficulties with predicting themselves. Thus, disclosure of this type of knowledge can be meaningful for users.

The actual ways of disclosing persuasion profiles needs future attention. Study 1 of Chapter 4 showed the possible negative effects on compliance of disclosure of persuasive intent. This would inhibit designers of persuasive systems to disclose persuasion profiles. However, the active selection of social influence strategies by users had beneficial effects on compliance. This implies that designers could actively involve users in the creation and maintenance of their own persuasion profiles. This is an opportunity that should be explored further.

9.3 Ethical Considerations

According to Berdichevsky and Neuenschwander (1999) “persuaders have long stood on uneasy ethical grounds”. From the more recent beginning of the study of interactive persuasive technologies researchers and practitioners have questioned the ethics of developing and deploying of persuasive systems. Several attempts to develop frameworks or principles for the ethical evaluation of persuasive systems have been undertaken (e.g., Berdichevsky and Neuenschwander, 1999; Fogg, 2002). Obviously, personalization, or the use of persuasion profiles, is by no means an exception to the standard unease associated with persuasive

technologies or persuaders in a broad sense.

Theorists have tried to guide designers of persuasive technologies in their quest to design ethically sound systems. Berdichevsky and Neuenschwander (1999) present a decision tree for ethical evaluation of persuasive technologies and the moral responsibility of system designers. This decision tree identifies how intent, predictability of outcomes, and ethical judgment interact to determine the proposed judgement and response to the system designers. According to this decision tree, the system designer is (a) praiseworthy if the outcome is intended and good; (b) not responsible if the outcome is unintended and not reasonably predictable, or if the outcome is reasonably predictable and good but not intended; and (c) otherwise at fault and blameworthy. Berdichevsky and Neuenschwander (1999) additionally offer eight principles that they regard as heuristics that could be justified within rule-based consequentialism.

While useful for designers, this typology, including the heuristics, does not tackle properly the origin of “unease” associated with persuasion and the design of persuasive systems. The decision tree focusses primarily on the goal of the persuasive attempt but does not question the *process* of persuasion in itself. Given that personalized persuasion through persuasion profiles is an extension of the persuasion process, irrespective of its goals, the ethical discussion of persuasion profiles should be separated from the discussion of goals or predictability of unwanted outcomes. Even though these are often related, and thus all an integral part of most frameworks on the ethics of persuasive technologies, there are two distinct mechanisms through which persuasion can be deemed ethical or unethical irrespective of the end goal. These concern the impact personalization of persuasion has on people’s privacy and autonomy, according to Haworth (1991) the core concepts by which to evaluate the ethics of new technologies.

9.3.1 Impact on Privacy

A large number of scholars in different fields have theorized about privacy and threats to privacy in their respective fields (Burgoon et al., 1989; Altman, 1976; Leino-Kilpi et al., 2001). These discussions are often inspired by the initial meaning of privacy as it derives from the Latin word “*privatus*” which means “to deprive” (Leino-Kilpi et al., 2001). Hence, privacy is often regarded “the ability of a person to withdraw (parts of) their personal information from the outside world” (Rawnsley, 1980). People deem objects, or thoughts, *private* which they

do not wish to share others with openly. A persuasion profile, and thus a collection of estimates of the success of influence strategies for a distinct individual, will by many be regarded as private. Thus, public sharing of such a profile—without control or consent—would be considered *a violation to their privacy*. This property however is shared at large with other profiles created about people by their use of interactive technologies.

The privacy of digital profiles has received considerable attention in the public debate. Profiles like those kept by social networking sites are on one hand by the general public regarded private, while on the other hand much of the information in these profiles is publicly available without a means for individuals to control the distribution of this information. In this sense persuasion profiles do not provide a novel threat to privacy but rather should be subject to privacy regulations that apply to other profiles. Largely, regulators are converging to regulations which empower users to a degree in which users have to consent profiles being created, users can inspect and delete their profiles, and users have the possibility to stop or limit sharing of their profile. These kind of regulations now apply to low level profiles—such as behavioral targeting information stored in session cookies—as well as elaborate user profiles on social network sites (Cf. Stallworth, 2010).

Designers of ubiquitous computing applications have also actively considered privacy and the privacy concerns of their users (see e.g.: Iachello, 2005). These attempts mainly focus on providing guidance for designers of ubiquitous technologies. Iachello (2005) introduces “the principle of proportionality” which enables designers to make decisions concerning privacy by balancing the usefulness of an application with its impact on privacy. It seems however that such a framework is insufficient to properly evaluate persuasion profiles: while the usefulness of a persuasion profile for a health intervention might be high and thus privacy concerns can be partly discounted, that same profile can be of use in other persuasive communications which the user does not find useful at all. The latter would logically lead a designer to impose heavy constraints on the ability of a system to collect and store personal data. To properly evaluate persuasion profiles, and *other technologies that store private information which can be used across applications and technologies*, new frameworks are needed.

For the time being, the ethical impact of persuasion profiling has to be evaluated with regard to imposing distinct threats to privacy *which had not otherwise been imposed by other technologies or applications*.

Given the large similarity of persuasion profiles to other types of profiles stored by interactive systems it seems that persuasion profiles are not a *novel* threat to privacy. The existing threat(s) however should be regulated, and persuasion profiling should be made compliant with such regulations.

9.3.2 Impact on Autonomy

Autonomy refers to people's ability to make rational and un-coerced choices and decisions. While heated debate is still ongoing as to whether there exists *any* level of autonomy (see, e.g. Ryan and Deci, 2006; Walter, 2001), there is a consensus amongst the general public that by and large choices and decisions are made by free will. As a consequence, at least in the western world, personal responsibility is attributed to such decisions with the accompanying praise or punishment.

The empirical studies of the effects of social influence strategies however show that decisions are often impacted by social influence strategies which seem to function on the border of conscious awareness (cf. Olson and Fazio, 2002). Thus, these types of choices might not be (fully) suspect to free will. Increasing *awareness* to social influence strategies, and thus elaboration and the ability to make a rational choice, often changes the effect of influence strategies (Petty and Cacioppo, 2001; Booth-Butterfield and Welbourne, 2002; Tormala and Petty, 2004). If the use of social influence strategies is made salient, decisions seem to come from a higher level of autonomy provided that indeed unconscious effects on decisions are largely perceived as threats to autonomy. From this perspective persuasion profiles that are used but not disclosed to users would pose a direct threat to the autonomy of individuals.

However, once a persuasion profile that has been created based on operative measures—as opposed to meta-judgmental measures which query those susceptibilities that are clearly conscious to people—is disclosed to users, the profile may very well *increase autonomy*. By educating people about their own decision processes people can evaluate whether they perceive their responses to distinct social influence strategies as rational and preferable. If this is not the case then a disclosed persuasion profile could aid as a prime to make the effects of social influence strategies on decision making more salient and as such bring them into the space of conscious elaboration. Via this route persuasion profiling could, opposed to many other types of profiles which are largely within conscious awareness, enhance autonomy.

9.3.3 Reflections on Ethics

To conclude, personalized persuasion itself, when evaluated based on privacy and autonomy concerns, is *not* necessarily unethical. However, the ends to which it is brought to bear could very well be, in which case designers are, to quote (Berdichevsky and Neuenschwander, 1999), “at fault”. Considered in isolation personalized persuasion, and the associated persuasion profiles, do not extend threats to privacy beyond current profiling practices and thus similar regulations should apply. Persuasion profiles however, opposed to most other profiles have the potential to increase human autonomy via education and insight into ones own behavior. As such when brought to bear for good ends, complying to privacy regulations, disclosed persuasion profiles are not ethically “uneasy”.

Next to threats to privacy and autonomy as primary criteria to evaluate the ethical impact of persuasion profiles one could also approach the discussion from a different stance. Persuasion itself currently is a part of everyday life: it is (e.g.) the primary tool of parents and educators. As such, the act itself seems to be acceptable. However the fact that persuasion profiles are technology initiated instead of human initiated might raise additional concerns: The technology does not have a “moral compass” to judge its own intends by. Furthermore, technologies are different from human persuaders in their ability to store, remember, and share information. These considerations are important and should be in the mindset of both researchers and designers working on persuasive technologies whether they employ persuasion profiling or not.

9.4 Future Research Challenges

This thesis presented a new view on the study of social influence strategies and their usage in persuasive technologies. The thesis showed that the heterogeneity in responses to social influence strategies is far larger than their average effect, and demonstrated several applications of this principle to influence human behavior. The technologies created to address this heterogeneity, *adaptive persuasive systems*, also inspire to new questions about human behavior and decision making as a function of social influence strategies. Questions which, through the application of multilevel models to estimate individual level effects as well as the deployment of ubiquitous sensing technologies to measure user behavior, can hopefully be answered in the future. Adaptive persuasive technolo-

gies can be a *tool* for further psychological research and should address the effects of social influence strategies at an individual level and over time. Furthermore, the focus on quantification of effects brought about in Chapter 8 warrants further research attention: Researchers both in social sciences as well as persuasive technology should attend more to the size and importance of the effects found in their evaluations as opposed to the mere existence of an effect.

9.4.1 Persuasion at an individual level

By and large researchers of influence strategies, persuasion, and persuasive technologies have focussed on average effects of the one term use of a persuasive intervention or manipulation. This thesis shows the importance of attending to the individual level effects, rather than the average level effects, which can be very different.

The difference between the average level effects and the individual level effects of social influence strategies and other psychological phenomena warrants future research for two reasons: From an applied perspective, technologies or interventions should deliver on their promise to change behavior of their users, not of *other users on average*. With its focus areas outside of marketing persuasive technologies frequently are designed to influence individuals. These individuals use the persuasive system to change their own attitudes or behavior and as such this is what the systems should be designed for. Now that technology has advanced to a level where designers can address individual users their individual behavior becomes important.

More theoretically important however is the progression and building of theories in social science. Many theories describe how individuals respond to stimuli, based on average effects of groups responding to these stimuli. Such an average can—and with sufficient sample size will—be statistically significant even though the estimated average effect might not occur for any of the individuals studied. Any individual-level theory that builds on empirical evidence of averages is thus likely at fault when large heterogeneity is present. Researchers need to understand what causes differences between people and consistency within people, rather than explaining artificial consistencies between people by postulating individual level processes.

Good examples of studies or theories about individuals and their average effects do already exist: Dynamic Social Impact theory as proposed by (Latané and Bourgeois, 1996) proposes a mechanism through which individual attitudes are formed which in turn, when modeled as

a group, leads to observable average level effects. Researchers should attend to the size of individual level effects versus average effects, define theories at the appropriate level, and through simulation and empirical studies show the link between the two.

9.4.2 Persuasion over time

Systems such as those presented in Chapter 8 enable monitoring of responses to social influence strategies by individuals over time. Up till now most studies into the effects of persuasion have concerned single points in time, or single interventions, in which pre- and post measures (or control and treatment groups) are compared. However persuasive systems are largely built not to promote a single behavior at a single instant but rather to change prolonged behavior and even lifestyle of their users. It is not very likely that influence strategies that are effective at one point in time remain effective throughout, and perhaps knowledge of conditioning and reinforcement schedules becomes more important when trying to sustain a behavior that was initially elicited using social influence strategies. Given that persuasive systems aim to elicit some kind of change in their users designers should understand how their systems influence users and how subsequently the “changed” user will respond to upcoming interactions with the system.

Future research should thus focus on the dynamics of attitude and behavior change within individuals over time. Already theorists are concerned with such questions, as evidenced by the recent work of Fogg (2010). However, without the involvement of sensing technologies it is hard to estimate effects over time. Adaptive persuasive systems, with identification, representation, and effect measurement, provide a platform to study the effects of prolonged influence efforts and to examine the dynamics of long term behavioral change. Due to technology, for the first time, human behavior can now be systematically studied over long periods of time. This type of “technology enabled research” is emerging within social psychology (see e.g. Killingsworth and Gilbert, 2010)) and is likely the area in which persuasion and social influence research can flourish in the future.

9.4.3 Effect size and Oomph

The last future direction that hopefully becomes more prominent in the future is more methodological in nature and is a change that has been suggested by scholars in the fields of psychology (Rouder et al., 2009; Siegfried, 2010), sociology, and economics (Ziliak, 2008): a focus

on effect size and real-world importance (Oomph) as opposed to within sample fit and statistical significance. The studies presented in Chapter 3 enabled estimation and comparison of sizes of effects as opposed to the mere proof of a non-null effect. In these studies the effect sizes were the reason to focus on individual level effects rather than average effects.

Effect size and Oomph reappeared several times in this thesis, noticeably in Chapter 7 and Chapter 8. While in these studies the estimated parameter were not always statistically different from 0, the importance of adaptive persuasion was mainly motivated by the impact on business: More people stayed in the DirectLife program and conversion rates in an online affiliate store rose. Here, noisy effects—which effects in social science are bound to be in a out-of-laboratory setting—can be extremely meaningful even if they do not “reach” statistical significance.

Only a study of effects, and their quantitative implications, can decide which statistically significant findings scientist should attend to and base theories on. And, on the flip side, quantitative implications and Oomph could also drive scientist to attend to non-significant effects. This critical evaluation of quantities, versus the qualitative notion of the existence of an effect, should be more prominent in all social science and human technology research—this thesis provides some examples³.

³This chapter is (partly) based on earlier publication(s): (Kaptein et al., 2011c; Kaptein and Eckles, 2010).

Bibliography

- Aarts, E. (2010). Ambient Intelligence. *The invisible future the seamless integration of technology into everyday life*, 5355(5):317–320.
- Aarts, E. and Ruyter, B. D. (2009). New research perspectives on Ambient Intelligence. *Journal of Ambient Intelligence and Smart Environments*, 1(1):5–14.
- Aarts, E. H. L., Markopoulos, P., and Ruyter, B. E. R. (2007). The Persuasiveness of Ambient Intelligence. In Petkovic, M. and Jonker, W., editors, *Security, Privacy and Trust in Modern Data Management*. Springer.
- Ades, P. A. (2001). Cardiac rehabilitation and secondary prevention of coronary heart disease. *The New England Journal of Medicine*, 345(12):892–902.
- Ainsworth, B. E., Haskell, W. L., Whitt, M. C., Irwin, M. L., Swartz, A. M., Strath, S. J., O’Brien, W. L., Bassett, D. R., Schmitz, K. H., Emplaincourt, P. O., Jacobs, D. R., and Leon, A. S. (2000). Compendium of physical activities: an update of activity codes and MET intensities. *Medicine & Science in Sports & Exercise*, 32(9 Suppl):S498–504.
- Ajzen, I. and Fishbein, M. (1980). *Understanding attitudes and predicting social behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Aleamoni, L. M. and Thomas, G. S. (1980). Differential Relationships of Student, Instructor, and Course Characteristics to General and Specific Items On a Course Evaluation Questionnaire. *Teaching of Psychology*, 7(4):233–235.
- Altman, I. (1976). Privacy: a conceptual analysis. *Environment And Behavior*, 8(1):7–29.
- Amit, R. and Zott, C. (2001). Value creation in E-business. *Strategic Management Journal*, 22(6-7):493–520.
- Arteaga, S. M., Kudeki, M., and Woodworth, A. (2009). Combating obesity trends in teenagers through persuasive mobile technology. *ACM SIGACCESS Accessibility and Computing*, 94(94):17–25.
- Asch, S. E. (1955). Opinions and Social Pressure. *Scientific American*, 193(5):31–35.

- Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs*, 70(9):1–70.
- Awad, N. F. and Krishnan, M. S. (2006). The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly*, 30(1):13–28.
- Baayen, R. H., Davidson, D. J., and Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4):390–412.
- Bang, M., Torstensson, C., and Katzeff, C. (2006). The PowerHouse : A Persuasive Computer Game Designed to Raise Awareness of Domestic Energy Consumption. *Persuasive Technology*, 3962(3962):123–132.
- Bargh, J. A., Chaiken, S., Raymond, P., and Hymes, C. (1996). The Automatic Evaluation Effect: Unconditional Automatic Attitude Activation with a Pronunciation Task. *Journal of Experimental Social Psychology*, 32(1):104–128.
- Barry, B. and Shapiro, D. L. (1992). Influence Tactics in Combination: The Interactive Effects of Soft Versus Hard Tactics and Rational Exchange. *Journal of Applied Social Psychology*, 22(18):1429–1441.
- Bassili, J. N. (1996). Meta-judgmental versus operative indexes of psychological attributes: The case of measures of attitude strength. *Journal of Personality and Social Psychology*, 71(4):637–653.
- Berdichevsky, D. and Neuenschwander, E. (1999). Toward an ethics of persuasive technology. *Communications of the ACM*, 42(5):51–58.
- Blass, T. (1991). Understanding behavior in the Milgram obedience experiment: The role of personality, situations, and their interactions. *Journal of Personality and Social Psychology*, 60(3):398–413.
- Bless, H., Bohner, G., Schwarz, N., and Strack, F. (1990). Mood and Persuasion. *Personality and Social Psychology Bulletin*, 16(2):331–345.
- Booth-Butterfield, S. and Welbourne, J. (2002). The Elaboration Likelihood Model: Its Impact on Persuasion Theory and Research. In Dillard, J. P. and Pfau, M. W., editors, *The Persuasion Handbook Developments in Theory and Practice*, pages 195–211. Sage.
- Bosworth, K., Gustafson, D. H., Hawkins, R. P., Chewning, B., and Day, T. (1983). Adolescents, health education, and computers: the Body Awareness Resource Network (BARN). *Health Education*, 14(6):58–60.

- Bowling, A. (2005). Mode of questionnaire administration can have serious effects on data quality. *Journal of public health (Oxford, England)*, 27(3):281–91.
- Breckler, S. J. (1994). A Comparison of Numerical Indexes for Measuring Attitude Ambivalence. *Educational and Psychological Measurement*, 54(2):350–365.
- Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
- Brock, T. C. (1968). Implications of Commodity Theory for Value Change. In Greenwald, A. G., Brock, T. C., and Ostrom, T. M., editors, *Psychological Foundations of Attitudes*, pages 243–275. Academic Press, New York.
- Brug, J. (1999). Computer-tailored nutrition education: differences between two interventions. *Health Education Research*, 14(2):249–256.
- Brug, J., Campbell, M., and van Assema, P. (1999). The application and impact of computer-generated personalized nutrition education: A review of the literature. *Patient Education and Counseling*, 36(2):145–156.
- Brug, J., Glanz, K., Van Assema, P., Kok, G., and van Breukelen, G. J. P. (1998). The Impact of Computer-Tailored Feedback and Iterative Feedback on Fat, Fruit, and Vegetable Intake. *Health Education & Behavior*, 25(4):517–531.
- Brug, J., Hospers, H. J., and Kok, G. (1997). Differences in psychosocial factors and fat consumption between stages of change for fat reduction. *Psychology and Health*, 12(5):719–727.
- Brug, J., Oenema, A., and Campbell, M. (2003). Past, present, and future of computer-tailored nutrition education. *American Journal of Clinical Nutrition*, 77((suppl)):1028–1034.
- Bubb, J. A. (2007). *Weight Loss Persuasive Technology and Eating Disorder*. PhD thesis, University of North Carolina at Chapel Hill.
- Burger, J. M. (1999). The Foot-in-the-Door Compliance Procedure: A Multiple-Process Analysis and Review. *Personality and Social Psychology Review*, 3(4):303–325.
- Burger, J. M., Messian, N., Patel, S., Del Prado, A., and Anderson, C. (2004). What a coincidence! The effects of incidental similarity on compliance. *Personality and Social Psychology Bulletin*, 30(1):35–43.
- Burgoon, J. K., Parrott, R., Le Poire, B. A., Kelley, D. L., Walther, J. B., and Perry, D. (1989). Maintaining and Restoring Privacy through Communication in Different Types of Relationships. *Journal of Social and Personal Relationships*, 6(2):131–158.

- Cacioppo, J. T. and Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin*, 115(3):401–423.
- Cacioppo, J. T. and Petty, R. E. (1982). The need for Cognition. *Journal of Personality and Social Psychology*.
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., and Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, 119(2):197–253.
- Cacioppo, J. T., Petty, R. E., Kao, C. F., and Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology*, 51(5):1032–1043.
- Campbell, M. K., Bernhardt, J. M., Waldmiller, M., Jackson, B., Potenziani, D., Weathers, B., and Demissie, S. (1999). Varying the message source in computer-tailored nutrition education. *Patient education and counseling*, 36(2):157–69.
- Caplan, R. (1990). How fingerprints came into use for personal identification. *Journal of the American Academy of Dermatology*, 23(1):109–114.
- Chartrand, T. L. and Bargh, J. A. (1999). The chameleon effect: the perception-behavior link and social interaction. *Journal of personality and social psychology*, 76(6):893–910.
- Chartrand, T. L., Maddux, W. W., and Lakin, J. L. (2005). Beyond the Perception-Behavior Link: The Ubiquitous Utility and Motivational Moderators of Nonconscious Mimicry. In Hassin, R., Uleman, J. S., and Bargh, J. A., editors, *The new unconscious*, volume In R. Hass, chapter 13, pages 334–361. Oxford University Press.
- Chatterjee, S. and Dev, P. (2009). Persuasive Technology, Fourth International Conference, PERSUASIVE 2009, Claremont, California, USA, April 26-29, 2009. Proceedings. In Chatterjee, S. and Dev, P., editors, *PERSUASIVE*, volume 350 of *ACM International Conference Proceeding Series*. ACM.
- Chittaro, L. and Zangrando, N. (2010). The Persuasive Power of Virtual Reality: Effects of Simulated Human Distress on Attitudes towards Fire Safety. In Ploug, T., Hasle, P., and Oinas-Kukkonen, H., editors, *Persuasive Technology*, volume 6137, pages 58–69–69. Lecture Notes in Computer Science, Springer Berlin / Heidelberg.
- Cialdini, R. (2001). *Influence, Science and Practice*. Allyn & Bacon, Boston.

- Cialdini, R. (2004). The science of persuasion. *Scientific American Mind*, 284:76–84.
- Cialdini, R. B. (2005). Don't Throw in the Towel: Use Social Influence Research. *American Psychological Society Observer*, 18(4):33–34.
- Cialdini, R. B., Trost, M. R., and Newsom, J. T. (1995). Preference for consistency: The development of a valid measure and the discovery of surprising behavioral implications. *Journal of Personality and Social Psychology*, 69:318–328.
- Clee, M. A. and Wicklund, R. A. (1980). Consumer Behavior and Psychological Reactance. *Journal of Consumer Research*, 6(4):389–405.
- Consolvo, S., Markle, K., Patrick, K., and Chanasyk, K. (2009). Designing for persuasion: mobile services for health behavior change. In *Proceedings of the 4th International Conference on Persuasive Technology*, Persuasive '09, pages 11:1—11:1, New York, NY, USA. ACM.
- Conway, L. G. and Schaller, M. (2005). When authorities' commands backfire: attributions about consensus and effects on deviant decision making. *Journal of Personality and Social Psychology*, 89(3):311–326.
- De Anda, D. (2006). Baby Think It Over: evaluation of an infant simulation intervention for adolescent pregnancy prevention. *Health social work*, 31(1):26–35.
- de Heer, J. and Poiesz, T. (1998). Dynamic characteristics of motivation, ability and opportunity to process commercial information. *Advances of Consumer Research ACR*, 25:532–537.
- de Vries, H. and Brug, J. (1999). Computer-tailored interventions motivating people to adopt health promoting behaviours: Introduction to a new approach. *Patient Education and Counseling*, 36(2):99–105.
- Deci, E. L. (1975). *Intrinsic motivation*, volume 2009 of *Perspectives in social psychology*. Plenum Press.
- Deci, E. L., Koestner, R., and Ryan, R. M. (2001). Extrinsic Rewards and Intrinsic Motivation in Education: Reconsidered Once Again. *Review of Educational Research*, 71(1):1–27.
- Dennett, D. C. (1987). *The Intentional Stance*, volume 30. MIT Press.
- Deutsch, M. and Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *Journal of Abnormal and Social Psychology*, 51:629–636.

- Dijkstra, A. (2005). Working mechanisms of computer-tailored health education: evidence from smoking cessation. *Health education research*, 20(5):527–39.
- Dijkstra, A. and De Vries, H. (1999). The development of computer-generated tailored interventions. *Patient Education and Counseling*, 36(2):193–203.
- Efron, B. and Morris, C. (1975). Data Analysis Using {Stein’s} Estimator and its Generalizations. *Journal of the American Statistical Association*, 70(350):311–319.
- Eisend, M. (2008). Explaining The Impact Of Scarcity Appeals In Advertising: The Mediating Role of Perceptions of Susceptibility. *Journal of Advertising*, 37(3):33–40.
- Falbe, C. M. and Yukl, G. (2008). Consequences for Managers of Using Single Influence Tactics and Combinations of Tactics. *The Academy of Management Journal*.
- Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: data reanalysis and commentary on Bargh, Chaiken, Govender, and Pratto (1992). *Journal of Personality and Social Psychology*, 64(5):753–758; discussion 759–765.
- Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R E Petty and Krosnik, J. A., editors, *Attitude strength: Antecedents and consequences*, pages 247–282. Erlbaum.
- Ferster, C. (1957). *Schedules of reinforcement*. Prentice-Hall, Englewood Cliffs N.J.
- Festinger, L. (1954). A Theory of Social Comparison Processes. *Human Relations*, 7(2):117–140.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.
- Flavell, J. H., Flavell, E. R., Green, F. L., and Korfmacher, J. E. (1990). Do Young Children Think of Television Images as Pictures or Real Objects ? *Media*, 34(4):399–419.
- Fogg, B. J. (1999). Persuasive technologies: Introduction. *Communications of the ACM*, 42.
- Fogg, B. J. (2002). *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann.

- Fogg, B. J. (2009a). A behavior model for persuasive design. *Proceedings of the 4th International Conference on Persuasive Technology Persuasive 09*, page 1.
- Fogg, B. J. (2009b). Creating Persuasive Technologies : An Eight-Step Design Process. *Technology*.
- Fogg, B. J. (2010). Fogg Behavior Grid. *Stanford University*, pages 8–10.
- Fogg, B. J. and Eckles, D. (2007). Mobile Persuasion: 20 Perspectives on the Future of Behavior Change. In Fogg, B. J. and Eckles, D., editors, *Mobile Persuasion*, pages 1–166. Stanford Captology Media.
- Fogg, B. J. and Nass, C. (1997). 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*, number March, pages 331–332. ACM.
- Franklin, V. L., Greene, A., Waller, A., Greene, S. A., and Pagliari, C. (2008). Patients's Engagement With "Sweet Talk": A Text Messaging Support System for Young People With Diabetes. *Journal of Medical Internet Research*., 10:e20.
- Froggatt, P. (1969). Questionnaire Design and Attitude Measurement. *British Journal of Industrial Medicine*, 68(3):249.
- Fromkin, H. L. (1970). Effects of experimentally aroused feelings of undistinctiveness upon valuation of scarce and novel experiences. *Journal of Personality and Social Psychology*, 16(3):521–529.
- Fuegen, K. and Brehm, J. W. (2004). *The intensity of affect and resistance to social influence*, pages 39–63. Lawrence Erlbaum.
- Gelman, A. (2005). Analysis of Variance: Why It Is More Important than Ever. *The Annals of Statistics*, 33(1):1–31.
- Gelman, A. and Hill, J. (2007). *Data Analysis Using Regression and {Multi-level/Hierarchical} Models*. Cambridge University Press.
- Gerber, B. S., Stolley, M. R., Thompson, A. L., Sharp, L. K., and Fitzgibbon, M. L. (2009). Mobile phone text messaging to promote healthy behaviors and weight loss maintenance: a feasibility study. *Health Informatics Journal*, 15(1):17–25.
- Goldstein, D. G. and Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1):75–90.

- Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, 35(3):472–482.
- Gormezano, I., Prokasy, W. F., and Thompson, R. F. (1987). The rabbit in pavlovian conditioning. In Gormezano, I., Prokasy, W. F., and Thompson, R. F., editors, *Classical Conditioning*, pages 419–432. Laurence Erlbaum Associates.
- Goy, A., Ardissono, L., Petrone, G., Brusilovsky, P., Kobsa, A., and Nejdl, W. (2007). *The Adaptive Web*, volume 4321 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Grediagin, A., Cody, M., Rupp, J., Benardot, D., and Shern, R. (1995). Exercise intensity does not effect body composition change in untrained, moderately overfat women. Technical Report 6, 62nd Medical Group, Ft Lewis, Wash. 98433, USA.
- Gretzel, U. and Fesenmaier, D. (2006). Persuasion in Recommender Systems. *International Journal of Electronic Commerce*, 11(2):81–100.
- Guadagno, R. E., Asher, T., Demaine, L. J., and Cialdini, R. B. (2001). When Saying Yes Leads to Saying No: Preference for Consistency and the Reverse Foot-in-the-Door Effect. *Personality and Social Psychology Bulletin*, 27(7):859–867.
- Guadagno, R. E. and Cialdini, R. B. (2010). Preference for consistency and social influence: A review of current research findings. *Social Influence*, 5(3):152–163.
- Ham, J., Midden, C., and Beute, F. (2009). Can ambient persuasive technology persuade unconsciously?: using subliminal feedback to influence energy consumption ratings of household appliances. *Current*, pages 1–6.
- Hardin, C. D. and Higgins, E. T. (1996). *Shared reality: How social verification makes the subjective objective.*, volume 3, pages 28–84. Guilford.
- Haugtvedt, C. P., Liu, K., and Min, K. S. (2008). Individual Differences: Tools for Theory Testing and Understanding in Consumer Psychology Research. In *Handbook of consumer psychology*, pages 1161–1176. {CRC} Press.
- Haugtvedt, C. P. and Petty, R. E. (1992). Personality and persuasion: Need for cognition moderates the persistence and resistance of attitude changes. *Journal of Personality and Social Psychology*, 63(2):308–319.
- Haugtvedt, C. P., Petty, R. E., and Cacioppo, J. T. (1992). Need for cognition and advertising: Understanding the role of personality variables in consumer behavior. *Journal of Consumer Psychology*, 1(3):239–260.

- Haworth, L. L. (1991). Dworkin on Autonomy. *Ethics*, 102(1).
- Hays, W. L. (1973). *Statistics for Social Scientists*. Holt, Rinehart & Winston of Canada Ltd.
- Hazen, T. J., Weinstein, E., and Park, A. (2003). Towards robust person recognition on handheld devices using face and speaker identification technologies. *Proceedings of the 5th international conference on Multimodal interfaces ICMI 03*, page 289.
- Hornstein, H. A., Fisch, E., and Holmes, M. (1968). Influence of a model's feeling about his behavior and his relevance as a comparison other on observers' helping behavior. *Journal of Personality and Social Psychology*.
- Hutchinson, J. W., Kamakura, W. A., and Lynch, J. J. G. (2001). Unobserved Heterogeneity as an Alternative Explanation for "Reversal" Effects in Behavioral Research. *Journal of Consumer Research*, 27(3):324–344.
- Iachello, G. (2005). Privacy and proportionality: adapting legal evaluation techniques to inform design in ubiquitous computing. In *Conference on Human factors in computing*, pages 91–100.
- IJsselsteijn, W., de Kort, Y., Midden, C. J. H., Eggen, B., and van den Hoven, E. (2006). Persuasive Technology for Human Well-Being: Setting the Scene. In IJsselsteijn, W., de Kort, Y., Midden, C. J. H., Eggen, B., and van den Hoven, E., editors, *Persuasive Technology, First International Conference on Persuasive Technology for Human Well-Being, PERSUASIVE 2006, Proceedings.*, volume 3962 of *Lecture Notes in Computer Science*, pages 1–5. Springer.
- Inman, J., Peter, A., and Raghubir, P. (1997). Framing the Deal: The Role of Restrictions in Accentuating Deal Value. *Journal of Consumer Research*, 24(1):68–79.
- James, J. M. and Bolstein, R. (1990). The Effect of Monetary Incentives and Follow-up Mailings on the Response Rate and Response Quality in Mail Surveys. *Public Opinion Quarterly*, 54(3):346.
- Jarvis, W. B. G. and Petty, R. E. (1996). The need to evaluate. *Journal of Personality and Social Psychology*, 70(1):172–194.
- Johnson, B. T. and Eagly, A. H. (1989). Effects of Involvement on Persuasion : A Meta- Analysis Effects of Involvement on Persuasion : A Meta-Analysis. *Prevention*.
- Kaptein, M. C. (2011a). Adaptive Persuasive Messages in an E-commerce Setting: The use of Persuasion Profiles. In *Proceedings of ECIS 2011*, Helsinki.

- Kaptein, M. C. (2011b). Building Persuasion Profiles in the Wild: Using Mobile Devices as Identifiers Maurits Kaptein. (CHI 2011 PINC workshop).
- Kaptein, M. C., de Ruyter, B., Markopoulos, P., and Aarts, E. (2009a). Persuading you: Individual differences in susceptibility to persuasion. In *12th IFIP TC 13 International Conference - INTERACT*, pages 24–28, Uppsala, Sweden. ACM Press.
- Kaptein, M. C., de Ruyter, B., Markopoulos, P., and Aarts, E. (2011a). Tailored Persuasive Text Messages to Reduce Snacking. *Transactions on Interactive Intelligent Systems*, IN PRESS.
- Kaptein, M. C., Duplinsky, S., and Go, E. M. (2012). Simultaneous Usage of Multiple Influence Strategies in Online Marketing. *Submitted to: International Journal of Internet Marketing and Advertising*.
- Kaptein, M. C., Duplinsky, S., and Markopoulos, P. (2011b). Means based adaptive persuasive systems. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 335–344, New York, NY, USA. ACM.
- Kaptein, M. C. and Eckles, D. (2010). Selecting Effective Means to Any End: Futures and Ethics of Persuasion Profiling. In Ploug, T., Hasle, P., and Oinas-Kukkonen, H., editors, *Persuasive Technology*, pages 82–93. Springer Berlin / Heidelberg.
- Kaptein, M. C. and Eckles, D. (2012). Heterogeneity in the Effects of Online Persuasion. *Journal of Interactive Marketing*, IN PRESS.
- Kaptein, M. C., Eckles, D., and Davis, J. (2011c). Envisioning Persuasion Profiles: Challenges for Public Policy and Ethical Practice. *ACM Interactions*, 18(5):66–69.
- Kaptein, M. C., J, L., and Saini, P. (2010a). Individual Differences in Persuadability in the Health Promotion Domain. In Ploug, T., Hasle, P., and Oinas-Kukkonen, H., editors, *Persuasive Technology*, pages 94–105. Springer Berlin / Heidelberg.
- Kaptein, M. C., Markopoulos, P., Ruyter, B., and Aarts, E. (2009b). Persuasion in ambient intelligence. *Journal of Ambient Intelligence and Humanized Computing*, 1(1):43–56.
- Kaptein, M. C., Markopoulos, P., Ruyter, B., and Aarts, E. (2010b). Two acts of social intelligence: the effects of mimicry and social praise on the evaluation of an artificial agent. *AI & SOCIETY*, 26(3):261–273.
- Kaptein, M. C. and van Halteren, A. (2012). Adaptive Persuasive Messaging to Increase Service Retention. *Journal of Personal and Ubiquitous Computing*, IN PRESS.

- Kass, A. (2007). Transforming the Mobile Phone into a Personal Performance Coach. In Fogg, B. J. and Eckles, D., editors, *Mobile Persuasion: 20 Perspectives on the Future of Behavior Change*. Stanford Captology Media.
- Kellermann, K. and Cole, T. (1994). Classifying Compliance Gaining Messages: Taxonomic Disorder and Strategic Confusion. *Communication Theory*, 4(1):3–60.
- Kelman, C. and Hamilton, V. L. (1989). *Crimes of obedience: Toward a social psychology of authority and responsibility*. Yale University Press, New Haven.
- Kerr, J., Eves, F., and Carroll, D. (2001). Six-month observational study of prompted stair climbing. *Preventive Medicine*, 33(5):422–427.
- Khaled, R., Fischer, R., Noble, J., and Biddle, R. (2008). A Qualitative Study of Culture and Persuasion in a Smoking Cessation Game. *Persuasive Technology*, 5033:224–236.
- Killingsworth, M. A. and Gilbert, D. T. (2010). A Wandering Mind Is an Unhappy Mind. *Science*, 330(6006):932.
- King, P. and Tester, J. (1997). The Landscape of Persuasive Technologies. *Communications of the ACM*, 42(5):31–38.
- Kirkpatrick, D. R., Pavlik, W. B., and Reynolds, W. F. (1964). Partial-reinforcement extinction effect as a function of size of goal box. *Journal of Experimental Psychology*, 68:515–516.
- Koch, M. and Schlichter, J. (2001). Ubiquitous Computing. In Schwabe, G., Streitz, N., and Unland, R., editors, *CSCWKompendium*, volume 2 of *Lecture Notes in Computer Science*, pages 507–517. Springer.
- Komorita, S. S., Hilty, J. A., and Parks, C. D. (1991). Reciprocity and Cooperation in Social Dilemmas. *Journal of Conflict Resolution*, 35(3):494–518.
- Kostakos, V. (2008). Using Bluetooth to capture passenger trips on public transport buses. *CoRR*, abs/0806.0.
- Kostakos, V. and Neill, E. O. (2008). Cityware : Urban computing to bridge online and real-world social networks. *Analysis*, pages 195–204.
- Kreuter, M. (2000). *Tailoring health messages: customizing communication with computer technology*. Routledge.
- Kroeze, W., Werkman, A., and Brug, J. (2006). A systematic review of randomized trials on the effectiveness of computer-tailored education on physical activity and dietary behaviors. *Annals of behavioral medicine : a publication of the Society of Behavioral Medicine*, 31(3):205–23.

- Lacroix, J., Saini, P., and Goris, A. (2009). Understanding user cognitions to guide the tailoring of persuasive technology-based physical activity interventions. In Chatterjee, S. and Dev, P., editors, *Persuasive Technology, Fourth International Conference, PERSUASIVE 2009*, volume 350 of *ACM International Conference Proceeding Series*, page 9, Claremont, California, USA. ACM.
- Lam, X. N., Vu, T., Le, T. D., and Duong, A. D. (2008). Addressing cold-start problem in recommendation systems. In *Proceedings of the 2nd international conference on Ubiquitous information management and communication - ICUIMC '08*, page 208, New York, New York, USA. ACM Press.
- Lambert, C. L. (2001). Baby think it over. *Pediatrics*, 107(4):806.
- Latané, B. (1970). *The unresponsive bystander: why doesn't he help?* Appleton-Century Crofts, New York.
- Latané, B. (1996). Dynamic Social Impact: The Creation of Culture by Communication. *Journal of Communication*, 46(4):13–25.
- Latané, B. and Bourgeois, M. J. (1996). Experimental Evidence for Dynamic Social Impact : The Emergence n Electronic Groups of Subcultures i. *Journal of Communication*, 46(4):35–47.
- Latané, B., Liu, J. H., Nowak, A., Bonevento, M., and Zheng, L. (1995). Distance Matters: Physical Space and Social Impact. *Personality and Social Psychology Bulletin*, 21(8):795–805.
- Latané, B. and Nida, S. (1981). Ten years of research on group size and helping. *Psychological Bulletin*, 89(2):308–324.
- Leino-Kilpi, H., Välimäki, M., Dassen, T., Gasull, M., Lemonidou, C., Scott, A., and Arndt, M. (2001). Privacy: A review of the literature. *International Journal of Nursing Studies*, 38(6):663–671.
- Leippe, M. R. and Eisenstadt, D. (1994). Generalization of dissonance reduction: Decreasing prejudice through induced compliance. *Journal of Personality and Social Psychology*, 67(3):395–413.
- Lockton, D., Harrison, D., and Stanton, N. (2008). Making the user more efficient: design for sustainable behaviour. *International Journal of Sustainable Engineering*, 1(1):3–8.
- Long, J. D. and 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. joann.long@lcu.edu.

- Lynn, M. (1989). Scarcity effects on desirability: Mediated by assumed expensiveness? *Journal of Economic Psychology*, 10(2):257–274.
- Lynn, M. (1991). Scarcity effects on value: A quantitative review of the commodity theory literature. *Psychology and Marketing*, 8(1):43–57.
- Mackie, D. and Cooper, J. (1984). Attitude Polarization : Effects of Group Membership. *Journal of Personality and Social Psychology*, 46(3):575–585.
- Mackie, D. M. (1986). Social identification effects in group polarization. *Journal of Personality and Social Psychology*, 50(4):720–728.
- MacInnis, D. J. and Jaworski, B. J. (1989). Information Processing from Advertisements: Toward an Integrative Framework. *Journal of Marketing*, 53(4):1.
- Macready, W. G. and Wolpert, D. H. (1998). Bandit problems and the exploration/exploitation tradeoff. *IEEE Transactions on Evolutionary Computation*, 2(1):2–22.
- Maio, G. R. and Esses, V. M. (2001). The need for affect: Individual differences in the motivation to approach or avoid emotions. *Journal of Personality*, 69(4):583–615.
- Martin, R. and Hewstone, M. (2003). Social influence processes of control and change: conformity, obedience to authority, and innovation. In Hogg, M. A. and Cooper, J., editors, *Psychological Foundations of Attitudes*, pages 347–366. {Londo} Sage.
- McFarland, R. G., Challagalla, G. N., and Shervani, T. a. (2006). Influence Tactics for Effective Adaptive Selling. *Journal of Marketing*, 70(4):103–117.
- McGuire, W. J. (1981). Theoretical foundations of campaigns. In Rice, R. and Paisley, W., editors, *Public communication campaigns*, volume 2nd, chapter Chapter 2, pages 41–70. Sage Publications, Inc.
- McGuire, W. J. (1995). Transferring research findings on persuasion to improve drug abuse prevention programs. *NIDA research monograph*, 155:225–45.
- McNaughton, N. (1984). Effects of anxiolytic drugs on the partial reinforcement extinction effect in runway and Skinner box. *The Quarterly journal of experimental psychology B Comparative and physiological psychology*, 36(4):319–330.
- Meschtscherjakov, A., Wilfinger, D., Scherndl, T., and Tscheligi, M. (2009). *Acceptance of future persuasive in-car interfaces towards a more economic driving behaviour*. ACM Press, New York, New York, USA.

- Midden, C., Mccalley, T., Ham, J., and Zaalberg, R. (2008). Using persuasive technology to encourage sustainable behavior. In *Workshop paper at 6th International Conference on Pervasive Computing*, number 1, pages 83–86. Eindhoven University of Technology.
- Midden, C. J., Meter, J. E., Weenig, M. H., and Zieverink, H. J. (1983). Using feedback, reinforcement and information to reduce energy consumption in households: A field-experiment. *Journal of Economic Psychology*, 3(1):65–86.
- Milgram, S. (1974). *Obedience to Authority*. London: Tavistock.
- Miller, A. G., Collins, B. E., and Brief, D. E. (1995). Perspectives on Obedience to Authority: The Legacy of the Milgram Experiments. *Journal of Social Issues*, 51(3):1–19.
- Mintz, J. and Aagaard, M. (2010). The Application of Persuasive Technology to educational settings : Some theoretical from the HANDS Project 1 Introducing the HANDS Project. In Hasle, P., editor, *Persuasive 2010 Proceedings of Poster Papers for the Fifth International Conference on Pervasive Technology*, pages 101–104. Oulu University Press.
- Modigliani, A. and Rochat, F. (1995). The role of interaction sequences and the timing of resistance in shaping obedience and defiance to authority. *Journal of Social Issues*, 51(3):107–123.
- Moe, W. W. and Fader, P. S. (2004). Dynamic Conversion Behavior at E-Commerce Sites. *Management Science*, 50(3):326–335.
- Moon, Y. and Nass, C. (1996). Adaptive agents and personality change: complementarity versus similarity as forms of adaptation. In *Conference companion on Human factors in computing systems*, pages 287–288, Vancouver, British Columbia, Canada. {ACM} Press.
- Murray, D. M., Richards, P. S., Luepker, R. V., and Johnson, C. A. (1987). The prevention of cigarette smoking in children: two- and three-year follow-up comparisons of four prevention strategies. Technical Report 6, Division of Epidemiology, School of Public Health, University of Minnesota, Minneapolis 55455.
- Nail, P. R., Correll, J. S., Drake, C. E., Glenn, S. B., Scott, G. M., and Stuckey, C. (2001). A validation study of the preference for consistency scale. *Personality and Individual Differences*, 31:1193–1202.
- Nass, C. (1997). Social and Natural Interfaces : Theory and Design. *Media*, (March):192–193.

- Nass, C., Fogg, B. J., and Moon, Y. (1996). Can computers be teammates? *International Journal of Human-Computer Studies*, 45(6):669–678.
- Nass, C., Steuer, J., and Tauber, E. R. (1994). Computers are social actors. *Conference companion on Human factors in computing systems CHI 94*, 94:204.
- Noar, S. M., Benac, C. N., and Harris, M. S. (2007). Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin*, 133(4):673–693.
- Nowak, A., Szamrej, J., and Latané, B. (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, 97(3):362–376.
- Ochi, P., Rao, S., Takayama, L., and Nass, C. (2010). Predictors of user perceptions of web recommender systems: How the basis for generating experience and search product recommendations affects user responses. *International Journal of Human-Computer Studies*, 68(8):472–482.
- Oinas-Kukkonen, H. and Harjumaa, M. (2008a). A Systematic Framework for Designing and Evaluating Persuasive Systems. In Oinas-Kukkonen, H., Hasle, P. F. V., Harjumaa, M., Segerståhl, K., and Øhrstrøm, P., editors, *Persuasive Technology, Third International Conference, {PERSUASIVE} 2008, Oulu, Finland, June 4-6, 2008. Proceedings*, volume 5033 of *Lecture Notes in Computer Science*, pages 164–176. Springer.
- Oinas-Kukkonen, H. and Harjumaa, M. (2008b). Towards Deeper Understanding of Persuasion in Software and Information Systems. *First International Conference on Advances in ComputerHuman Interaction*, pages 200–205.
- Oinas-Kukkonen, H. and Harjumaa, M. (2009). Persuasive Systems Design : Key Issues , Process Model , and System Features. *Communications of the Association for Information Systems*, 24(1):485–500.
- Oinas-Kukkonen, H., Hasle, P. F. V., Harjumaa, M., Segerstaahl, K., and Øhrstrom, P. (2008). *Persuasive Technology, Third International Conference, PERSUASIVE 2008, Oulu, Finland, June 4-6, 2008. Proceedings*, volume 5033 of *Lecture Notes in Computer Science*. Springer.
- O’Keefe, D. J. (1994). From Strategy-Based to Feature-Based Analyses of Compliance Gaining Message Classification and Production. *Communication Theory*, 4(1):61–69.
- Olson, M. A. and Fazio, R. H. (2002). Implicit Acquisition And Manifestation Of Classically Conditioned Attitudes. *Social Cognition*, 20(2):89–104.

- Owen, A. B. (2007). The pigeonhole bootstrap. *The Annals of Applied Statistics*, 1(2):386–411.
- Pariser, E. (2011). *The Filter Bubble*. Penguin Press.
- Parkes, A., Henderson, M., and Wight, D. (2005). Do sexual health services encourage teenagers to use condoms? A longitudinal study. *J Fam Plann Reprod Health Care.*, 31:271–280.
- Patrick, K., Raab, F., Adams, M. A., Dillon, L., Zabinski, M., Rock, C. L., Griswold, W. G., and Norman, G. J. (2009). A Text Message Based Intervention for Weight Loss: Randomized Controlled Trial. *Journal of Medical Internet Research*, 11:1–9.
- Payan, J. M. and McFarland, R. G. (2005). Decomposing Influence Strategies: Argument Structure and Dependence as Determinants of the Effectiveness of Influence Strategies in Gaining Channel Member Compliance. *Journal of Marketing*, 69(3):66–79.
- Penedo, F. J., Schneiderman, N., Dahn, J. R., and Gonzalez, J. S. (2004). Physical activity interventions in the elderly: cancer and comorbidity. *Cancer Investigation*, 22(1):51–67.
- Perloff, R. M. (2003). *The Dynamics of Persuasion: Communication and Attitudes in the 21st Century*. Routledge.
- Petty, R. E. and Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19(1):123–205.
- Petty, R. E. and Cacioppo, J. T. (2001). Source Factors and the Elaboration Likelihood Model of Persuasion. *Higher Education*, (Lm):668–673.
- Petty, R. E. and Wegener, D. T. (1999). The elaboration likelihood model: Current status and controversies. In Chaiken, S. and Trope, Y., editors, *Dual-process theories in social psychology*, page 41–72. Guilford Press, New York.
- Petty, R. E., Wegener, D. T., and Fabrigar, L. R. (1994). Cognitive Processes in attitude change. In Wyer, R. S. and Srull, T. K., editors, *Handbook of Social Cognition*, pages 69–142. Erlbaum, Hillsdale, NJ, 2 edition.
- Petty, R. E., Wegener, D. T., and Fabrigar, L. R. (1997). Attitudes and attitude change. *Annual Review of Psychology*, 48(1):609–647.
- Ploug, T., Hasle, P., and Oinas-Kukkonen, H., editors (2010). *Persuasive Technology*, volume 6137 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg.

- Plummer, J. T. (1974). The Concept and Application of Life Style Segmentation. *Journal of Marketing*, 38(1):33–37.
- Preacher, K. J., Rucker, D. D., MacCallum, R. C., and Nicewander, W. A. (2005). Use of the extreme groups approach: a critical reexamination and new recommendations. *Psychological Methods*, 10(2):178–192.
- Preece, J. J. (2010). I Persuade , They Persuade , It Persuades ! *Human-Computer Interaction*, pages 2–3.
- Prochaska, J. O. and Velicer, W. F. (1997). The transtheoretical model of health behavior change. *American journal of health promotion*, 12(1):38–48.
- Rawnsley, M. M. (1980). The concept of privacy. *ANS Advances in nursing science*, 2(2):25–31.
- Reeves, B. and Nass, C. (1996). *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press.
- Rhoads, K. (2007). How Many Influence, Persuasion, Compliance Tactics & Strategies Are There? <http://www.workingpsychology.com/numbertactics.html>.
- Ribeiro, C., Mehrotra, G., Vey, G., Althothali, A., and Dimarco, C. (2011). Improving Health Care with a Virtual Human Sleep Coach Technical Report CS-2011-10. *Sleep (Rochester)*.
- Robben, H. S. and Poiesz, T. B. (1993). The Operationalization Of Motivation, Capacity and Opportunity To Process an Advertising Message. In Raaij, W. F. V. and Bamossy, G. J., editors, *European Advances in Consumer Research Volume 1*, pages 160–167. Provo, UT : Association for Consumer Research.
- Rothschild, M. L. (1999). Carrots, Sticks, and Promises: A Conceptual Framework for the Management of Public Health and Social Issue Behaviors. *Journal of Marketing*, 63(4):24–37.
- Roubroeks, M., Ham, J., and Midden, C. (2011). When Artificial Social Agents Try to Persuade People: The Role of Social Agency on the Occurrence of Psychological Reactance. *International Journal of Social Robotics*, 3(2):155–165.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., and Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic bulletin & review*, 16(2):225–37.

- Rozin, P. and Royzman, E. B. (2001). Negativity Bias, Negativity Dominance, and Contagion. *Personality and Social Psychology Review*, 5(4):296–320.
- Ryan, R. and Deci, E. (2000). Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology*, 25(1):54–67.
- Ryan, R. M. and Deci, E. L. (2006). Self-regulation and the problem of human autonomy: does psychology need choice, self-determination, and will? *Journal of personality*, 74(6):1557–85.
- Sakai, R., van Peteghem, S., van de Sande, L., Banach, P., and Kaptein, M. C. (2011). Personalized Persuasion in Ambient Intelligence: the APStairs System. In *Proceedings of Ambient Intelligence (AmI) 2011*, Amsterdam.
- Schultz, P. W. (1999). Changing Behavior With Normative Feedback Interventions: A Field Experiment on Curbside Recycling. *Basic and Applied Social Psychology*, 21(1):25–36.
- Scott, S. L. (2010). A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658.
- Seta, J. J. and Seta, C. E. (1982). Personal equity: An intrapersonal comparator system analysis of reward value. *Journal of Personality and Social Psychology*, 43(2):222–235.
- Siegfried, T. (2010). 4/20/2010 Odds Are, It’s Wrong - Science News. *ScienceNews*, 177(7):1–13.
- Skinner, B. F. (1976). *About Behaviorism*. Vintage Books, New York.
- Slater, M., Antley, A., Davison, A., Swapp, D., Guger, C., Barker, C., Pistrang, N., and Sanchez-Vives, M. V. (2006). A Virtual Reprise of the Stanley Milgram Obedience Experiments. *PLoS ONE*, 1(1):10.
- Snijders, T. A. B. and Bosker, R. J. (1999). *Multilevel analysis: An introduction to basic and advanced multilevel modeling*. Sage Publications Ltd.
- Snyder, C. R. and Fromkin, H. L. (1980). *Uniqueness: the human pursuit of difference*. Plenum Press.
- Stallworth, B. (2010). Future Imperfect: Googling for Principles in Online Behavioral Advertising. *Federal Communications Law Journal*, (62):465–491.
- Stein, C. (1955). Inadmissibility of the Usual Estimator of the Mean of a Multivariate Normal Distribution. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, pages 197–206. Berkeley: Univ. Cal. Press.

- Svane, O. (2007). Helping , Informing or Coaxing the Consumer ? exploring Persuasive Technology as applied to households energy use. *Policy*, 1(11):1–11.
- Tajfel, H. (1982). *Social identity and intergroup relations*, volume 14. Cambridge University Press.
- Teh, K. C. and Aziz, A. R. (2002). Heart rate, oxygen uptake, and energy cost of ascending and descending the stairs. *Medicine & Science in Sports & Exercise*, 34(4):695–699.
- Tesser, A. (1993). The importance of heritability in psychological research: the case of attitudes. *Psychological Review*, 100(1):129–142.
- Tormala, Z. L. and DeSensi, V. L. (2009). The Effects of Minority/Majority Source Status on Attitude Certainty: A Matching Perspective. *Personality and Social Psychology Bulletin*, 35(1):114–125.
- Tormala, Z. L. and Petty, R. E. (2004). Resistance to persuasion and attitude certainty: the moderating role of elaboration. Technical Report 11, Department of Psychology, Indiana University, Bloomington, IN 47405, USA. ztormala@indiana.edu.
- Torning, K., Hall, C., and Oinas-Kukkonen, H. (2009). Persuasive System Design : State of the Art and Future Directions. *Information Systems Journal*, page 1.
- Tynan, A. C. and Drayton, J. (1987). Market segmentation. *Journal of Marketing Management*, 2(3):301–335.
- van Bell, D. and Smulders, K. (2009). Development of the connectedness inventory. *Unpublished manuscript*, pages 1–23.
- Verhallen, T. M. M. and Robben, H. S. J. (1994). Scarcity and preference: An experiment on unavailability and product evaluation. *Journal of Economic Psychology*, 15(2):315–331.
- Walter, H. (2001). *Neurophilosophy of free will: from libertarian illusions to a concept of natural autonomy*. MIT Press.
- Watson, J. M. and Strayer, D. L. (2010). Supertaskers: Profiles in extraordinary multitasking ability. *Psychonomic Bulletin & Review*, 17(4):479–485.
- Webster, D. M. and Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67(6):1049–1062.
- Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 3(3):94–104.

- West, S. G. (1975). Increasing the attractiveness of college cafeteria food: A reactance theory perspective. *Journal of Applied Psychology*, 60:656–658.
- West, S. G., Gunn, S. P., Hay, W., and Lang, L. (1975). Increasing the Attractiveness of College Cafeteria Food : A Reactance Theory Perspective her assistance in conducting the experiment and. *Psychology*, 60(5):656–658.
- Wilcox, R. R. (1981). A Review of the Beta-Binomial Model and Its Extensions. *Journal of Educational Statistics*, 6:3–32.
- Yeo, A. C., Rahim, M., and Ren, Y. Y. (2009). Use of Persuasive Technology to Change End- Users 'IT Security Aware Behaviour : A Pilot Study. *SciencesNew York*, 1(1):48–54.
- Zanker, M., Bricman, M., Jessenitschnig, M., Damiani, E., Jeong, J., Howlett, R., and Jain, L. (2009). *New Directions in Intelligent Interactive Multimedia Systems and Services - 2*, volume 226 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Zhu, F. and Zhang, X. M. (2010). Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics. *Journal of Marketing*, 74(2):133–148.
- Ziliak, S. T. (2008). The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives. *Journal of Economic Literature*, 47(2):499–503.

Articles by Maurits Kaptein

Journal Articles:

1. Kaptein, M. C., & van Halteren, A. (2012). Adaptive Persuasive Messaging to Increase Service Retention. *Journal of Personal and Ubiquitous Computing*, IN PRESS.
2. Kaptein, M. C., Fernandez, N., Castaneda, D., & Nass, C. (2012). Extending the Similarity-Attraction Effect: The Effects of When-Similarity in Mediated Communication. *Journal of Computer-Mediated Communication*, IN PRESS.
3. Kaptein, M. C., & Eckles, D. (2012). Heterogeneity in the Effects of Online Persuasion. *Journal of Interactive Marketing*, IN PRESS.
4. Kaptein, M. C., de Ruyter, B., Markopoulos, P., & Aarts, E. (2012). Tailored Persuasive Text Messages to Reduce Snacking. *Transactions on Interactive Intelligent Systems*, IN PRESS.
5. Kaptein, M. C., Eckles, D., & Davis, J. (2011). Envisioning Persuasion Profiles: Challenges for Public Policy and Ethical Practice. *ACM Interactions*, 18(5), 66-69.
6. Kaptein, M. C., Markopoulos, P., Ruyter, B., & Aarts, E. (2010). Two acts of social intelligence: the effects of mimicry and social praise on the evaluation of an artificial agent. *AI & SOCIETY*, 26(3), 261-273.
7. Kaptein, M. C., Markopoulos, P., Ruyter, B., & Aarts, E. (2009). Persuasion in ambient intelligence. *Journal of Ambient Intelligence and Humanized Computing*, 1(1), 43-56.
8. Kaptein, M. C., Weisscher, A., Terken, J. M. B., & Nelissen, H. (2009). Online contextmapping; using the opportunities of Web 2.0 for the contextmapping procedure. *CoDesign*, 5(4), 213-228.

Peer-reviewed Conference Proceedings:

1. Kaptein, M. C., & Robertson, J. (2012). Rethinking Statistical Methods for HCI. In *Proceedings of the 2011 annual conference on Human factors in computing systems, CHI 2012*. IN PRESS.

2. Sakai, R., Peteghem, S. van, Sande, L. van de, Banach, P., & Kaptein, M. C. (2011). Personalized Persuasion in Ambient Intelligence: the APStairs System. In *Proceedings of Ambient Intelligence (AmI) 2011*. Amsterdam.
3. Lindholm, J., Kaptein, M. C., & Parvinen, P. (2012). The Trade-off between Online Community Activity and Consumption. In *Proceedings of HICSS 2012*. Hawaii.
4. Kaptein, M. C. (2011). Adaptive Persuasive Messages in an E-commerce Setting: The use of Persuasion Profiles. In *Proceedings of ECIS 2011*. Helsinki.
5. Kaptein, M. C., & Eckles, D. (2011). Magnitude and Structure of Heterogeneity in the Effects of Influence Strategies. In *Proceedings of Persuasive Technology 2011*. Ohio.
6. Kaptein, M. C., Duplinsky, S., & Markopoulos, P. (2011). Means based adaptive persuasive systems. In *Proceedings of the 2011 annual conference on Human factors in computing systems (pp. 335-344)*. New York, NY, USA: ACM.
7. Kaptein, M. C., Lacroix, J., & Saini, P. (2010). Individual Differences in Persuadability in the Health Promotion Domain. In T. Ploug, P. Hasle, & H. Oinas-Kukkonen (Eds.), *Persuasive Technology (pp. 94-105)*. Springer Berlin / Heidelberg.
8. Kaptein, M. C., & Eckles, D. (2010). Selecting Effective Means to Any End: Futures and Ethics of Persuasion Profiling. In T. Ploug, P. Hasle, & H. Oinas-Kukkonen (Eds.), *Persuasive Technology (pp. 82-93)*. Springer Berlin / Heidelberg.
9. Scherini, T., Melo, P., van Craenendonck, T., Zou, W., & Kaptein, M. C. (2010). Enhancing the sleeping quality of partners living apart. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems - DIS 2010 (pp. 171-179)*. New York, New York, USA: ACM Press.
10. Kaptein, M. C., Nass, C., & Markopoulos, P. (2010). Powerful and consistent analysis of likert-type ratingscales. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI 2010 (pp. 2391-2401)*. New York, New York, USA: ACM Press.
11. Kaptein, M. C., de Ruyter, B., Markopoulos, P., & Aarts, E. H. L. (2009). Simple Ways to Make Friends. In *Proceedings of the 8th International Workshop on Social Intelligence*.

12. Kaptein, M. C., de Ruyter, B., Markopoulos, P., & Aarts, E. (2009). Measuring the Subjective User Experience. In *Proceedings of the 12th IFIP TC 13 International Conference - INTERACT* (pp. 944-945).
13. Kaptein, M. C., de Ruyter, B., Markopoulos, P., & Aarts, E. (2009). Persuading you: Individual differences in susceptibility to persuasion. In *Proceedings of the 12th IFIP TC 13 International Conference - INTERACT* (pp. 24-28). Uppsala, Sweden: ACM Press.
14. Soute, I., Kaptein, M. C., & Markopoulos, P. (2009). Evaluating outdoor play for children. In *Proceedings of the Proceedings of the 8th International Conference on Interaction Design and Children - IDC 2009* (p. 250). New York, New York, USA: ACM Press.
15. Kaptein, M. C., Slegers, K., Nelissen, H., Weisscher, A., & Terken, J. M. B. (2007). The Digital Workbook: Using Web 2.0 for Generative Research Purposes. In *Proceedings of MobileHCI 2007*.
16. Al Mahmud, A., Kaptein, M. C., Moran, O., Van De Garde-Perik, E., & Markopoulos, P. (2007). Understanding Compliance to Privacy Guidelines Using Text-and Video-Based Scenarios. In *Proceedings of HumanComputer Interaction INTERACT 2007* (pp. 156-168). Springer Berlin / Heidelberg.
17. Romero, N., Szóstek, A. M., Kaptein, M. C., & Markopoulos, P. (2007). Behaviours and Preferences when Coordinating Mediated Interruptions : Social and System influence. In *Proceedings of CSCW 2007* (pp. 123-132). Kluwer Press.

Journal Articles Under Submission:

1. Kaptein, M. C., Duplinsky, S., & Go, E. M. (2012). Simultaneous Usage of Multiple Influence Strategies in Online Marketing. *Submitted to: International Journal of Internet Marketing and Advertising*.
2. Kaptein, M. C., & Nass, C. (2012). Nice to Know you: Familiarity and Influence in Social Networks. *UNDER SUBMISSION*.
3. Kaptein, M. C., & Markopoulos, P. (2012). On the Design of Adaptive Persuasive Systems. *UNDER SUBMISSION*.
4. Kaptein, M. C., & Parvinen, P. (2012). Sequential Learning of Profit Maximizing Sales Prices. *UNDER SUBMISSION*.

5. Lindholm, J., Parvinen, P., & Kaptein, M. C. (2012). The Dangers of Engagement: How Substitution in Online Services Leads to Decreasing Revenues. *UNDER SUBMISSION*.

Informal Conference Proceedings:

1. Kaptein, M. C. (2011). Building Persuasion Profiles in the Wild: Using Mobile Devices as Identifiers Maurits Kaptein. (*CHI 2011 PINC workshop*). Vancouver, British Columbia, Canada.
2. Kaptein, M. C., Markopoulos, P., de Ruyter, B., & Aarts, E. H. L. (2010). Social Connectedness and Compliance (*CHI 2010 Workshop on Social Connectedness*). Atlanta, USA.

Summary

Personalized Persuasion in Ambient Intelligence

This thesis examines the use of personalized persuasion in ambient intelligence. Persuasive technologies—systems and services intentionally designed to influence user behavior—are emergent. Applications range from iPhone apps which help users stop smoking, to distributed networks of smart sensors that persuade users to reduce their energy consumption. Many of these systems apply social science knowledge about *influence strategies* to increase the effectiveness of their persuasion attempts.

This thesis first examines user responses to influence strategies. The work shows that while most influence strategies are effective on average, large individual differences exist. The responses to some strategies are *negative* for a large proportion of users despite replicating the average positive effect of the use of these strategies. This *heterogeneity* in responses to social influence strategies proves stable over both time and contexts.

After showing that a proper selection of a single influence strategy leads to more persuasion than combinations of strategies, the thesis develops the idea of creating and utilizing *persuasion profiles*: collections of the estimates, and their associated certainty, of the effects of influence strategies on individuals. The thesis describes how these profiles can be build both via meta-judgmental measures as well as operative measures. The Susceptibility to Persuasive Strategies Scales (*STPS*) is presented for the purpose of creating persuasion profiles based on meta-judgmental measures.

The thesis next examines the *applied* value of persuasion profiles. Via several designs the thesis shows that persuasive systems in which the influence strategy that is used is adapted to individual users outperform non-personalized systems. These ideas are further advanced by proposing and evaluating a method for building personalized persuasive technologies.

Persuasion Profiles will be a core component of persuasive technologies to come. The ambient intelligence scenario makes it possible to build dynamic profiles based on unobtrusive measurements. This in turn will help researchers and designers to measure, predict, influence, and ultimately *understand human responses to persuasion*.

Acknowledgments

Much of the thoughts and work presented in this thesis would not have been possible without the help of my supervisors, co-workers, colleagues, students, and friends.

I would like to thank my core committee: Prof. Emile Aarts, Dr. Boris de Ruyter, and Prof. Panos Markopoulos. Without your help this thesis would not have been what it is now. I would also like to acknowledge the help of my reading committee: Prof. Cees Midden, Prof. Huib de Ridder, and Prof. Harri Oinas-Kukkonen.

During my PhD I spend over a year at Stanford University. There, Prof. Clifford Nass taught me about science and life in general. Thank you! Also, during my stay at Stanford Mike Nowak, Abhay Sukumaran, Hellen Harris, Prof. Jeremy Bailenson, Prof. David Rogosa, Dr. Martha Russel, Dr. Helen Roque, Adelaide Dawes, Erina DuBois, Abe Downey, Katy Harrison, Will Bishop, and Alvin Lacson made it fun and meaningful. Special thanks to Dean Eckles whom I greatly enjoyed collaborating and BBQ-ing with.

At Philips Research I would like to thank Dr. Aart van Halteren, Dr. Joyca Lacroix, Dr. Gijs Geleijnse, Dr. Martin Pijl, Privender Saini, Tess Speelpenning, Dr. Saskia van Dantzig, Elke Deamen, Jia Du, Tommi Maätta, Roger Holmes, and everyone else I have worked with. A special thanks Dr. Joris Janssen who made Philips Research a lot more fun.

At the Technical University of Eindhoven I would like to thank Prof. Berry Eggen, Prof. Jean Bernard Martens, Prof. Jacques Terken, Dr. Jahved Kahn, Dr. Georgios Metaxas, Dr. Evangelos Karapanos, Nikos Batalas, Iris Stoute, Alex Uyttendaele, Dr. Nathalia Romero, Qonita Shabab, Abdullah al Mahmud, and all others in the UCE group and at the USI program. I greatly enjoyed working with you and discussing science.

Many people have helped me out academically, collaborated with me, or provided me with the opportunity to talk about my work. I would especially like to thank Prof. Per Hasle, Dr. Judy Robertson, Steven Duplinsky, Arjan Haring, Jay de Groot, Dimitra Retsina, Prof. Samir Chattarjee, Thomas Visser, Prof. Petri Parvinen, Jerry Lindholm, Dr. Daan van Bel, Jan Kallenbach, and Essi Pöyry.

I would also like to thank my friends for being supportive of this project (either pushing me to start it, to continue it, or to finish it). In no particular order, thank you: Ruud Niesink & Barto Nijenhuis (thanks for being my paranymphs!), Joyce Kemperman, Julia Nijenhuis, Puk Nijenhuis, Tommy Haccou, Gijs Haccou, Eva Haverdil, Peter van Gent (thanks for helping out on the cover design!), Mariëlle Keteelaar, Peter Kolks, Bram van Gemert, Geert van der Velden, Leon van Kervel, Joost Lorenz, Dennis Jesserun, Maurice de Kunder, Laurens Koot, Vincent Gijssen, Daniel Wienke, Anton Thiel (& everyone else at Waalhalla!), Damiaan Winkelman, Niels Kingma, Hugo Uijlenbroek, Janneke van Ekert, Coen de Bruyn, Sanna Rijpma, Noortje van Laanen, Jaap Gielink, Sietske Oenema, and Samir Hansen.

Also, I would especially like to thank my family: Adrie Kaptein-Vrolijk, Bartjan Kaptein, Reinier Kaptein, Claire Millet-Kaptein, Thomas Kaptein, and Lise Kaptein. Thanks for supporting my “career” choices throughout and providing me with a decent upbringing. Also thanks to Jozef Geurtzen, Ans Geurtzen-van Ree, Naline Geurtzen, Gaard Kets, and Ayla Geurtzen for your support.

Most probably I forgot to include a number of you in the list above. My apologies.

Finally, I would like to thank Rosa Geurtzen. She is cool.

Curriculum

Maurits Clemens Kaptein was born on the 9th of March 1983 in Doetinchem, the Netherlands. After my graduation as Master of Science in Economic Psychology at the University of Tilburg in the Netherlands I completed the Post-master User-System Interaction program at the Technical University of Eindhoven. After obtaining the PDEng. degree I worked as a research development manager at De Vos and Jansen Market research in Nijmegen, the Netherlands. In October 2008 I started my PhD. project which has been carried out at the technical University of Eindhoven, the Netherlands, Philips Research Laboratories, Eindhoven, the Netherlands, and Stanford University, Stanford, USA.

With more than 25 publications (conference proceedings and journal articles) and a patent in the area of persuasion profiling, I have been an active researcher in Design, Psychology, Economics, and Statistics. During these research projects I have collaborated with research teams in Philips Research (NL), Philips Directlife (NL), Facebook (USA), PersuasionAPI (NL), Stanford University (USA), Delft University (NL), the Aalto School of Economics (FIN), and Herriot-Watt University (UK). My core research interests are persuasive technologies and research methodology. I have presented my research during scientific conferences (e.g. CHI, Interact, Persuasive Technology, HICSS) and at several invited presentations (e.g. MediaX Stanford (USA), Aalto School of Economics (FIN), TEDx (NL)). Besides my role as a researcher I am the chairman of the board of the “Stichting Skateboarding Nijmegen”, a non-profit responsible for the exploitation of the largest concrete indoor skateboard park in Europe.

Index

- 8-step design process, 13
- ability, 16
- adaptation of persuasive attempts, 54
- adaptive persuasive systems, 6, 8, 98, 140, 166
- affiliate online store, 157
- Ambient Intelligence (AmI), 3, 12, 112, 172
- ambient persuasive systems, 98, 172
- ANOVA, 31
- Application Programming Interface (API), 158
- APStairs, 140
- attitude, 2, 14
 - ambivalence, 16
 - change, 2, 14
 - perceived attitude basis, 2
 - preferred attitude basis, 2
 - strength, 16
- Authority, 3, 18, 21
- autonomy, 179
- average treatment effect (ATE), 7, 29, 98, 181
- Baby think it over (BTIO), 10
- BARN, 9
- Bayesian Information Criteria (BIC), 35
- behavior, 16
 - change, 2, 16
 - sustaining, 17
- Beta Binomial model, 107
- Beta distribution, 107
- Binomial distribution, 107
- Bluetooth scanning, 103, 142
- borrowing strength, 110
- card-sorting, 142
- central processing (see ELM), 19
- choice effects, 58
- click-through, 75
- Cochran-Mantel-Haenszel test, 120
- Commitment, 22
- compliance, 20, 61
 - index, 61
 - score, 61
 - tactics, 20
- computer-tailored
 - health education, 4, 12, 26, 175
 - interventions, 26
 - persuasion, 102, 175
- conditioning, 17
 - classical, 2, 17
 - operant, 2, 17
- Consensus, 19, 21
- Consistency, 22
- correlations, 38
 - between strategies, 38
- covariance, 31
 - Σ , 31
- decision rule, 107

- demographics, 34, 40
- dissonance theory, 22
- dual processing model, 19
- dynamic adaptation, 102, 104, 140
- e-commerce, 157
- e-selling, 157
- effect size, 182
- Elaboration Likelihood Model (ELM), 19, 24, 80, 179
 - central processing, 19, 80, 179
 - peripheral processing, 19, 80, 179
- emotional state, 41
- empirical p -value, 36
- end-independence, 111, 176
- ethics, 176
- expert systems, 4
- extreme group approach (EGA), 117, 122
- foot-in-the-door, 25
- Gaussian distribution, 30
- heterogeneity, 5, 7, 29, 31, 98, 181
- hierarchical models, 8, 30
- Hygiene Guard, 11
- identification, 8, 103
- immediacy, 19
- impact, 6
- individual differences, 10, 23
 - measurement, 25
- individual persuasion, 181
- individual-level effect, 3, 26, 29, 181
 - estimates, 145
- individual-level theory, 7, 181
- influence strategy, 5, 14, 20, 79
 - choice, 58
 - designing implementations, 124, 150
 - implementation, 33
 - multiple, 7, 58, 100
 - random presentation, 128
 - single, 58
 - susceptibility measurement, 79
 - taxonomy, 10, 20, 173
- interactive systems, 2
- interactive technologies, 2
- intra-individual differences, 45
- item loadings, 83, 88
- item-ranking task, 59
- Javascript tracker, 158
- JSON, 174
- key performance indicator, 161
- Kinder-Kleertjes, 158
- Kustomsignals, 11
- Lark, 3, 10
- Liking, 18, 22
- log-likelihood, 35
 - χ^2 model comparisons, 35
- MCMC sampling, 36
- means to an end, 4
 - means adaptation, 111
- measurement, 8, 104
 - behavior, 118
 - click-through, 75, 161
 - email reminders, 149
 - intention, 118
 - interest, 118
 - meta-judgmental, 25, 79, 170

- operative, 25, 30, 170
- snacking behavior, 123, 129
- stair usage, 142
- meta-analysis, 166
- meta-judgmental measures, 25, 79, 136
 - constraints, 136
- Metabolic Equivalent of Task (MET), 140
- Milgram experiment, 3
- mimicry, 18
- mobile devices, 123
- mood, 41
- motivation, 16
 - external, 17
 - internal, 17
- multi-armed bandit problem, 108
- multi-level models, 8, 30, 110, 181
- MyZeo, 10
- Need for Affect, 24
- Need for Closure, 24
- Need for Cognition (*NfC*), 24, 34, 91, 175
- Need for Uniqueness, 22
- Need to Evaluate, 24
- negativity bias, 100
- Nike +, 10
- non-disclosure, 112
- null-model, 30, 130
- Oomph, 182
- operative measures, 25, 30, 54
- opportunity, 16
- Peitho, 1
- perceived usefulness, 62
- peripheral processing (see ELM), 19
- personal equity theory, 22
- personalization, 123
 - counter-, 123
- personalized persuasion, 95, 97, 137, 169
 - computer-tailored, 102
- personalized persuasive systems, 8
- persuadability, 30
- persuasion, 1, 4, 12
 - adapting, 4
 - in ambient intelligence, 10, 12, 172
 - over time, 182
 - principle, 3, 5
- Persuasion API, 157
- persuasion profile, 100, 139, 170
 - combinations with other profiles, 174
 - creation, 100, 172
 - decision rule, 107
 - disclosure, 175
 - dynamic, 102, 104, 139
 - estimation, 106
 - ethical considerations, 176
 - identification, 103, 140, 149, 158, 172
 - measurement, 104, 142, 152, 159, 172
 - representation, 103, 141, 149, 158, 172
 - sharing, 173
- persuasive intent, 77
- Persuasive Messaging System (PMS), 149
 - PMS server, 152
- Persuasive Systems Design, 13
 - model (PSD), 13
- persuasive technology, 2, 3, 9, 12, 98, 169

- personalized, 169
- Philips DirectLife, 3, 10, 13
- physically active, 140
- posterior distribution, 107
- Preference for consistency, 25, 80
- principal component analysis, 83, 88
- principle of proportionality, 178
- prior distribution, 30, 107
- privacy, 177
- promoting stair usage, 140
- random effects, 8, 30
- randomized probability matching, 108, 153, 160
- reach, 6
- reactance, 22, 137
 - theory, 22
- Reciprocity, 23
- reinforcement, 17
 - scheme, 17
- repetition, 18
- representation, 8, 103
- resistance to persuasion, 112
- revolving door, 105
- sales tactic, 58
- Scarcity, 22
- shrinkage, 106
- Smart meter, 11
- SMS intervention, 123
- social dilemma, 23
- social influence strategy (see influence strategy), 14
- social proof, 19, 21
- sources, 18
 - multiple, 18
 - single, 18
- stability over time, 47
- state, 31
- Stein estimation, 106
- strategy $\times$... interactions, 99
- survival task, 59
- susceptibility, 5, 79, 116, 170
 - high, 117
 - low, 117
- Susceptibility to Persuasive Strategies Scale (*STPS*), 34, 79, 116, 123, 171
 - composite scores, 89
 - correlation with *NfC*, 91, 122
 - full version, 85
 - internal reliability, 89, 93
 - short, 80
 - sub-scales, 85
- tailoring, 11, 26, 175
- technology-mediated communication, 6
- text messages, 123
- The Media Equation, 2
- trait, 31, 41, 172
- ubiquitous computing, 3, 12, 178
 - privacy, 178
- upper-bound, 31

List of Figures

1.1	Graphical representation of the structure of this thesis. The <i>insight generation</i> chapters explore several questions regarding influence strategies using an experimental approach. In the <i>case studies</i> chapters these findings are validated in several field studies.	6
2.1	Two axis model as presented in Kaptein et al. (2009b). .	15
3.1	Comparison of heterogeneity in the effects of influence strategies with the average effects of those strategies. The solid black vertical lines are the estimated average effects of each strategy, as compared with the control message. The black curves are the estimated normal distribution of strategy effects for the population, while the gray curves are the density of the estimates of the strategy effects for this sample. Estimates are from Model <i>C</i>	37
3.2	Estimated influence strategy effects for each participant in Model <i>C</i> , as compared to the control messages. Note that for some individuals, the estimated effect of one strategy is negative, while the estimated effect of another strategy is positive.	39
3.3	Appearance of the bookstore in Study 2. The authority strategy is implemented in the orange box just below the title of the book. The evaluation questions and the tasks for participants appeared in the bottom bar.	42
3.4	Summed raw book evaluations for two participants selected to illustrate variation and consistency in responses to influence strategies over the three sessions.	44
3.5	Appearance of the Musicstore Study. The vertical bar on the left is used to present assignments and questions to participants while they can freely browse the online store.	48

4.1	The effects of choice and disclosure on compliance. Shown are the estimated marginal means and standard errors of the compliance scores. On the x-axis, are the three choice levels, and the two separate lines represent the two disclosure conditions.	63
4.2	The effects of choice and disclosure on the perceived usefulness of the advice. Shown are the estimated marginal means and standard errors of the usefulness scores. . . .	64
4.3	The effects of the use of single or multiple strategies, either authority or consensus and congruent or incongruent, on compliance.	67
4.4	The effects of the use of single or multiple strategies, either authority or consensus and congruent or incongruent advice on compliance in Study 3.	72
4.5	Model based click-through behavior based on the different advertisement versions. Solid lines represent average of the click-through for single strategy advertisements (black) and multiple strategy advertisements (gray). Dotted lines represent individual advertisements.	76
5.1	Scree plot showing the eigenvalues (y-axis) of each of the extracted components (x-axis). A clear increase is visible from component 6 upwards indicating a proper fit of a six factor solution.	89
6.1	Example of a persuasion profile. Dots represent the estimated effect of the respective persuasive strategies, while the bars represent the certainty around this estimate. . .	101
6.2	Graphical representation of a user entering the building (from below). The user is identified, and a message is displayed on the screen. Next, the user chooses between the sliding door (left) and the revolving door (right). . .	105
6.3	Progression of point estimates and certainty of 2 messages for four users, one from each group.	110
7.1	Overview of the average number of snacks eaten each day by our participants separated for the three experimental conditions.	129

7.2	Overview of the number of snacks eaten each day by our participants with superimposed the regression lines of the effects over time of the different conditions.	131
7.3	Overview of the perceived healthiness of daily nutrition score for participants with superimposed the regression lines of the effects over time of the different conditions. .	133
7.4	Overview of the usefulness ratings over time. Added are the estimated regression lines of the effect of time. . . .	134
8.1	Schematic setup of our adaptive persuasive system. Bluetooth scanners are present at the building entrance and at each of the floors.	144
8.2	Overview of the estimated effectiveness of different social influence strategies for two users of the APStairs system.	146
8.3	Comparison of the success-rates in the two experimental conditions during the 5 work-weeks of deployment of the system.	147
8.4	Overview of kinder-kleertjes.com: Left the home page, in the middle a product presentation page. The implementation of the persuasive strategies are shown on the right with their position on the product presentation page indicated by the arrow.	159
8.5	Progression of between-user (overall average) effectiveness estimates of the three strategies and indications of the certainty of the estimates. Initially the estimates are very variable but over time they stabilize. The dates on the x-axis are jittered to increase readability. Error bars are conservative and based on between user variance not on the total number of observations.	162
8.6	Estimated effectiveness of the different persuasive strategies for two randomly selected users (out of those who viewed more than 8 products)—x-axis denotes the number of products viewed.	163
8.7	Progression of between-user (overall average) effectiveness estimates of the three strategies and indications of the certainty of the estimates. As in the first evaluation the products displayed accompanied by a social influence strategy <i>on average</i> have higher conversion rates.	164

- 8.8 Comparison of estimated conversion rates for consumers in the holdout condition—those visiting the no strategy version of the website—and consumers in the test condition—those looking at the adaptive persuasive system. 165
- 8.9 Forest plot of the meta-analysis using a fixed effect model of the two studies comparing adaptive and random strategy selection. 167

List of Tables

3.1	Influence strategies and their respective implementations as used in Study 1.	33
3.2	Table comparing the preferred model (<i>C</i>), in which each of the strategy effects vary by person, with the model without varying effects of strategies by person (<i>A</i>), and the model which includes heterogeneity only for the control vs. strategy contrast (<i>B</i>). Presented are the degrees of freedom (<i>df</i>), Bayesian Information Criteria, <i>BIC</i> , the log-likelihood, <i>logLik</i> of each of the models, and the χ^2 and <i>p</i> -value for the model comparisons.	36
3.3	Estimates of fixed effects in the preferred model. Using the control messages as the reference, each of the point estimates of the average effect of the influence strategies on book evaluations is positive. Empirical <i>p</i> -values computed with draws from the posterior using MCMC (see Baayen et al., 2008).	36
3.4	Summary of random effects in Model <i>C</i> . The square root of the diagonal elements of $\hat{\Sigma}_\beta$ are the first four rows of the standard deviation column, while the remaining entries are presented as correlations to the right.	38
3.5	Table examining the use of demographic and meta-judgmental measures. While both aid significantly in explaining the variance in responses to social influence strategies, explicitly modeling heterogeneity is still preferable according to the log-likelihood comparisons.	40
3.6	Influence strategies and their respective implementations as used in Study 2.	43
3.7	Table comparing Model <i>E</i> with a model without varying effects of strategies by person (Model <i>D</i>). The comparison shows stable heterogeneity in responses to influence strategies over the three sessions.	45

3.8	Summary of random effects of Model <i>E</i> . The square root of the diagonal elements of $\hat{\Sigma}_{\beta}^{(s)}$ and $\hat{\Sigma}_{\beta}^{(t)}$ are rows one to three and four to six of the standard deviation column respectively, while the remaining entries of each are presented as correlations to the right.	46
3.9	Table comparing the Model <i>D</i> with models using demographic or meta judgmental measures of personality. . .	47
3.10	Influence strategies and their respective implementations as used in the Bookstore Study.	49
3.11	Influence strategies and their respective implementations as used in the Musicstore Study.	50
3.12	This table shows a comparison of a model without any effects of strategy with models that include an overall strategy effect or a strategy specific effect. It is clear that the model with varying specific strategy effects by individual is preferred.	53
3.13	Summary of the random effects of Model <i>H</i> . The square root of the diagonal elements of Σ_{β} are the first three rows of the <i>SD</i> column. The remaining entries are presented as correlations (<i>Corr</i>).	53
3.14	Table comparing the Model <i>F</i> with models using demographic or meta judgmental measures of personality. . .	54
4.1	Post-hoc comparisons of the experimental conditions in Study 2. SS=Single strategy, MS=Multiple strategies, Auth=Authority, Cons=consensus, Congr=congruent, Incongr=incongruent.	68
4.2	Post-hoc comparisons of the experimental conditions in Study 3. SS=Single strategy, MS=Multiple strategies, Auth=Authority, Cons=consensus, Congr=congruent, Incongr=incongruent.	72
4.3	The number of views and the click-through rate of the advertisements used in Study 4.	75
5.1	The 12 susceptibility items—presented per influence strategy—that compose together the short- <i>STPS</i>	82
5.2	Component loadings based on a principal components analysis with oblimin rotation for the 12 items of the short- <i>STPS</i> . Loadings smaller than .3 are suppressed. .	84

5.3	Items used in the development of the <i>STPS</i> . Items marked with * are selected after the principal component analysis.	86
5.3	Items used in the <i>STPS</i> (...continued)	87
5.3	Items used in the <i>STPS</i> (...continued)	88
5.4	Component loadings based on a principal components analysis with oblimin rotation for 32 items of the <i>STPS</i> . Loadings smaller than .3 are suppressed.	90
5.5	Overview of the composite scores of the <i>STPS</i> . Presented are the mean, standard deviation, and Cronbach's α of each of its sub-scales.	91
5.6	Model comparisons of the preferred model of Study 1, Chapter 3, and a model with interactions of <i>STPS</i> scores on the appropriate strategies and the strategy in use. . .	92
6.1	Overview of the simulated data for the 4 different user groups.	109
7.1	Results study 1a: Percentage of respondents responding favorably.	119
7.2	Results study 1b: Percentage of respondents responding favorably.	120
7.3	Results study 2: Percentage of respondents responding favorably.	122
7.4	The messages used in the intervention. For each of the four social influence strategies used in this trial three implementations are used.	125
7.5	Comparing the null model with models including a time effect and different time effects for each condition. . . .	131
7.6	Overview of the fixed effects of the model including an interaction between time and condition to predict snacking behavior (Model <i>D</i>). Empirical p-values are based on MCMC simulations.	132
7.7	Overview of the fixed effects of the model including an interaction between time and condition predicting perceived healthiness of nutrition.	133
7.8	Overview of the fixed effects of the model including an interaction between time and condition predicting perceived usefulness of the received message.	134

8.1	Influence strategies and their implementations as used in the APstairs system.	143
8.2	Influence strategies and their implementations in persuasive messaging system.	151
8.3	Overview of the data from the persuasive messaging system evaluation.	154
8.4	Model Comparisons used for the analysis of the persuasive messaging system evaluation.	155
8.5	Coefficients of the fixed effects of Model <i>D</i> . The table shows that over time (<i>Frequency</i>) the success of the email reminders decreases. This decrease over time is significantly smaller in the adaptive condition than in the baseline condition.	156
8.6	Fixed effects of the logistic regression model predicting the drop-out probability of users in the PMS trial. . . .	156