UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 155

Feedback learning and multiple goal pursuit in an electricity consumption task

MONA GUATH

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2018

ISSN 1652-9030
ISBN 978-91-513-0341-3
urn:nbn:se:uu:diva-348821

Dissertation presented at Uppsala University to be publicly examined in Room 13:026, von Kraemers Allé 1A/1E, Uppsala, Thursday, 7 June 2018 at 10:15 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Professor Ulrike Hahn.

Abstract

Guath, M. 2018. Feedback learning and multiple goal pursuit in an electricity consumption task. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences* 155. 68 pp. Uppsala: Acta Universitatis Upsaliensis. ISBN 978-91-513-0341-3.

The overall aim with the thesis was to investigate how learning to pursue two conflicting goals (cost and utility) in an electricity consumption task is affected by different forms of feedback, goal phrasing, and task environment. Applied research investigating the efficiency of outcome feedback on electricity consumption via in-home displays points at modest reductions (2-4%). Further, a wealth of cognitive psychological research shows that learning with outcome feedback is not unproblematic. A new experimental paradigm, the simulated household, that captures the cognitive task that confronts people when trying to regulate their electricity consumption, was developed. In three studies, different aspects of the problem of regulating one's consumption was investigated. **Study I,** investigated how different feedback in terms of frequency, detail, and presence of random noise or not affect performance. It also investigated if participants pursued the goals sequentially or simultaneously and if they were able to derive a model of the task. Results showed that frequent feedback was beneficial only in a deterministic system and, surprisingly, random noise improved performance by highlighting the most costly appliances. Modelling results indicated that participants pursued goals sequentially and did not have a mental model of the task. **Study II,** investigated if a short feedforward training could replace or complement outcome feedback. Results indicated that the performance with one of the feedforward training schemes lead to comparable performance to outcome feedback only. The best performance was obtained when this feedforward scheme was combined with outcome feedback. **Study III,** investigated if the sequential goal pursuit observed in Study I was related to interpretation of the task or cognitive limitations by specifying goals for cost and/ or utility. Further, it investigated the reason for the cost prioritisation. Results indicated that the sequential goal pursuit derives from cognitive constraints. Together, the results from the studies suggest that people pursue the goals sequentially and that instant outcome feedback may harm performance by distracting people from the most important and costly appliances to the appliances that allow large variability in use.

Keywords: feedback, multiple goal pursuit, function learning, electricity consumption, optimisation

Mona Guath, Department of Psychology, Box 1225, Uppsala University, SE-75142 Uppsala, Sweden.

ISSN 1652-9030
ISBN 978-91-513-0341-3
urn:nbn:se:uu:diva-348821 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-348821)

Pour Adèle et Lucie.

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I Juslin, P., Elwin, E., Guath, M., Millroth, P., & Nilsson, H. (2016). Sequential and myopic: On the use of feedback to balance cost and utility in a simulated electricity efficiency task, *Journal of Cognitive Psychology, 28(1),* 106-128.

II Guath, M., Millroth, P., Juslin, P., & Elwin, E. (2015). Optimising Electricity Consumption: A case of function learning. *Journal of Experimental Psychology: Applied, 21(4),* 326-341.

III Guath, M., Juslin, P., & Rackwitz, R. *Why do people pursue goals sequentially when they try to balance the cost and the utility in an electricity consumption task?* Manuscript in preparation.

The contribution of Mona Guath to the studies included in this thesis was as follows:
Study I: Involved in planning, designing, analyzing data, and writing together with supervisor and co-authors.
Study II: Planned and designed the study, analysed the data, and wrote the manuscript with contribution of supervisor and co-authors.
Study III: Planned and designed the study, analysed the data, and wrote the manuscript with contributions from supervisor and co-author.

Contents

Abbreviations

IHD	In-home display
MCL	Muliple-cue learning
MCPL	Multiple-cue probability learning
FTS	Function training scheme
RL	Reinforcement learning
EBM	Exemplar-based model
CAM	Cue-abstraction model
MGPM	Multiple goal pursuit model
MAUT	Multi-attribute utility theory
DDM	Dynamic decision making
NDM	Naturalistic decision making
SA	Survival analysis

Introduction

From a very early age people make use outcome feedback when regulating behaviour, be it overt as social and motor behaviour, or inner processes as manifested in cognitive behaviour. It is therefore common to assume that merely giving knowledge of results, that is, outcome feedback, (Klüger & DeNisi, 1996) is the most influential way of changing and modifying behaviour. Indeed, the earliest studies in feedback learning adopted an associationist view (Hume, 1738; 1975), emphasising the importance of feedback as the sole learning mechanism. Consequently, outcome feedback is often given continuously, across many domains, with the purpose of shaping behaviour in the desired direction. In combination with an exploding technological development, outcome feedback is nowadays given on behaviours ranging from heart rate to profile views on social networks. In short, it is ever present.

In the same vein, the design of the apparatus delivering electricity feedback to the consumers (in-home displays: IHDs) often departs from a naïve assumption that outcome feedback automatically translates into more efficient use of electricity. This view, however, is problematic since there is evidence that instant outcome feedback on electricity consumption is not always an efficient way of reducing the consumption (e.g., Klopfert & Wallenborn, 2011; Krishnamurti, Davis, Wong-Parodi, Wang, & Canfield, 2013; Schleich, Klobasa, Brunner, Gölz., Götz. & Sunderer, 2011, KEMA, 2009; ESMIG, 2011; CER, 2011; EDRP, 2011). The result is not very surprising from a cognitive psychological view, with a wealth of research (e.g., Brehmer, 1980; Hammand, Summers, & Deane, 1973; Remus, O'Connor, & Riggs, 1996) showing that learning from outcome feedback in complex tasks is not unproblematic.

The purpose of this thesis was to investigate people's cognitive ability to use electricity more efficiently with outcome feedback. However, previous cognitive psychological research paradigms are structurally different from the task that confronts people when trying to regulate their electricity consumption. Whereas previous research has focused on how people relate several cues to *one* criterion with outcome feedback, regulating one's electricity consumption involves balancing *two* conflicting goals: the cost and the comfort given by the consumption. A new experimental paradigm was developed, in order to investigate the cognitive task that confronts people when using outcome feedback to control their electricity consumption. Further, the paradigm brings together two important domains, energy saving and cognitive psychology, with the purpose of better understanding if and how instant outcome feedback is efficient in learning to regulate one's electricity consumption. The focus on laboratory experiments and cognitive abilities entail that the investigation is primarily concentrated on the intellectual ability to learn to regulate one's electricity consumption. The studies do not address aspects that are present in a real world environment, for instance, motivation and ability to maintain a behaviour over time. The latter aspects are, indeed, important, but difficult to investigate in the experimental task that is used in the studies presented in the thesis.

The ability to use an IHD is intimately connected to the ability to use outcome feedback to learn the relation between different electricity consuming activities and their cost. In addition, the consumer must balance two goals, the comfort of the consumption against its cost. The capacity to use an IHD is hence related to different areas of research, namely, feedback learning, function learning, multiple cue learning (MCL), and multiple goal pursuit. Therefore, I will begin with a discussion of these areas of research and then present three empirical studies on people's cognitive ability to learn from outcome feedback in the new experimental paradigm. Finally, many of the challenges that confronts people in this task, such as, feedback learning and balancing two conflicting goals, are also related to other everyday situations where people must balance conflicting goals. Accordingly, the results may also be relevant in a wider, general psychological perspective.

Applied research

The demand for more efficient ways of communicating feedback on energy consumption has several motivations. The most pressing issue, though, is the goal of 20% energy efficiency improvement by 2020 (Climate and Energy Directives, 2009/28/CE to 2009/31/CE) as well as the fact that private consumption makes up 40% of the carbon dioxide emissions (Stern, 2011). Further, many companies try to position themselves at the market by offering energy providers and consumers devices and technologies that can measure electricity consumption more efficiently. In other words, there exists large corporate economic incentives for large-scale deployment of technological devices that communicate consumption behaviour that may need to be balanced by scientific scrutiny. The most common way to communicate feedback on electricity consumption to consumers is by means of in-home displays (IHDs), providing the consumer with instant (aggregated) outcome feedback on hourly, weekly or monthly energy usage (Smart Meter, 2017). Early studies on the efficiency of consumer electricity feedback pointed at reductions of 20% (Darby, 2006). More recent studies (e.g., Klopfert & Wallenborn, 2011), however, point at more modest reductions, ranging from 2-4%. Previous studies have investigated consumers' attitudes to IHDs (e.g., Anderson & White, 2009; Karjalainen, 2011), employed qualitative investigation of IHDs (e.g., Hargreaves, Nye, & Burgess) and reviewed existing literature (e.g., Buchanan, Russo, & Anderson, 2014; Darby, 2010; Fischer, 2008; Faruqui, Sergici, & Sharif, 2010). To my knowledge, there exists only one experimental study investigating feedback learning in the context of IHDs (Krishnamurti et al., 2013). However, the study does not investigate people's ability to use feedback to achieve goals, focus is on how feedback affects knowledge of appliance specific consumption from pre- to post-test. One way of expanding and understanding the previous research is to turn to cognitive psychological research on feedback learning and multiple goal pursuit.

Feedback learning as the fundament of all knowledge

> ”To this I answer, in one word, from experience…” (John Locke, 1690/1997)

The dominant view during the first part of the 20th century was that feedback was the fundament of all learning. The view originates from associationism that is based on the principles of an organism's causal history, an idea that dates back to Hume (1738; 1975). The concept is central to early 19th century theories of learning, particularly, classical conditioning (Pavlov, 1927), the law of effect (Thorndike, 1927) and radical behaviourism (Skinner, 1945). Both behaviourism and classical conditioning excluded the consciousness from the analysis, since it was a private experience, and, hence, unobservable (Kihlstrom & Park, 2002). In the late 1950s, however, evidence accumulated for the importance of cognition in the learning process. Pavlov's theory of associations between stimuli was expanded to include expectations in Rescorla Wagner's model (Rescorla & Wagner, 1972). The introduction of high-speed computers drew the researchers' attention to the information processing within the individual, notably Atkinson and Shiffrin's (1968) multi-store memory model and hierarchical goal systems in control theory (Wiener, 1948a; Wiener, 1948b; Powers, 1973). Both theories used the computer as a model for human information processing to account for the cognitive processes in learning and memory. Further, research in attention and memory showed that attention plays an important role in whether an event will be consciously remembered (Broadbent, 1958; Deutsch & Deutsch, 1963). Relatedly, research that emphasised the importance of knowledge structures rather than mere correlation between stimulus and response questioned the associationist account of learning (Bruner, Goodnow, & Austin, 1956). Finally, the importance of goals and self-efficacy in the learning process, was highlighted by Locke (1968) and Bandura (1977). These ideas cleared the way for modern research in cognitive psychology, emphasising the role of mental representations.

In their analysis of different aspects that might affect feedback learning, Kluger and DeNisi (1986, p. 255) proposed a broader definition of feedback interventions: *actions taken by (an) external agent(s) to provide information regarding some aspect(s) of one's task performance.* This includes both what is elsewhere known as cognitive feedback, that is, presenting information about relations rather than outcomes (Balzer, Doherty, & O'Connor, 1989) and, pure outcome feedback (also referred to knowledge of results, KR). This definition permits the authors to include feedback interventions from many different psychological fields, encompassing early behaviourist theories to goal setting theory (Locke & Latham, 1990), control theory (e.g., Carroll & Kay, 1988; Carver & Scheier, 1981; Powers, 1973), multiple cue probability learning (MCPL: Balzer, et al., 1989), and social cognition (Bandura, 1991). Below, the theories and psychological concepts that are relevant are reviewed and connected to the efficiency of feedback interventions.

The Cognitive Revolution I: The Role of Mental Representation

Single cue judgments: function learning

Function learning describes how people learn the mapping of a continuous input variable x by a continuous function F into a single criterion y (McDaniel & Busemeyer, 2005). In a typical function learning experiment, the participants engage in quite extensive training, where they are given an x value and asked to give an estimate of the criterion value y upon which outcome feedback is given. Following the training, the participants are given a transfer test, which is either of an interpolation type, with values in the same range as during training, or extrapolation type, with values outside the training range (DeLosh, Busemeyer, McDaniel, 1997). With extensive training, people are indeed able to learn different function forms (e.g., linear, exponential, quadratic: DeLosh, et al., 1997), but they demonstrate systematic biases. Several cognitive models have been proposed to account for the mechanisms behind function learning, with two distinct theoretical camps: rule-based theories (e.g., Brehmer, 1974; Koh & Meyer, 1991) and associative accounts (DeLosh, et al., 1997; Busemeyer, McDaniel, & Byun, 1997). Rule-based models rely on the idea that people learn an explicit function, from which they derive the criterion value y. Associative accounts assume that people learn to associate a given x value with a y value, and new x values are assessed against their similarity with previous y values (Lucas, Griffiths, Williams, & Kalish, 2015). More recently, a number of hybrid models have been proposed to better account for people's mental models of the task (e.g., EXAM: extrapolation-association model, DeLosh et al., 1997; POLE: Population of Linear Experts Model: Kalish, Lewandowsky, & Krushke, 2004; Lewandowsky, Roberts, & Yang, 2006). In essence, the models combine rules and associations in such a way that they capture how people combine a continuous cue to a continuous criterion in different settings. In sum, function learning is typically concerned with how people learn functions with one input value and to what extent they are able to extrapolate their knowledge to a new range as compared with a known range of values. A related area of research is multiple-cue learning, which it is concerned with how people learn to map several cues, as opposed to one, to one continuous criterion variable. Multiple-cue-learning is the type of learning that best describes the learning task that confronts the participants in the research paradigm used in the studies in the thesis.

Multiple cue judgments: multiple-cue-learning (MCL)

Statistical modelling of MCL

> ”In one word: Not from experience” (Brehmer, 1980, p. 223).

The first studies on multiple cue learning derives from Brunswik’s (1952) lens model, describing how people’s judgments are based on incomplete sensory cues. To investigate the relationships he applied statistical modelling, specifically, regression analysis. His multi-attribute perceptual model was later adapted to a social judgement theory (SJT: Hammond, 1955) describing how people arrive at social judgments by combining multiple (probabilistic) cues to one continuous criterion. This is investigated in an experimental paradigm called multiple-cue probability learning (MCPL), involving the use of one or more cues to infer a criterion variable that is imperfectly correlated (i.e., probabilistically related) with the cues (Brehmer, 1980). There are three components that need to be learned to arrive at an accurate judgment in MCPL (Brehmer, 1979, 1994): i) The functional relationship between the cues and the criterion; ii) The optimal weighting to ascribe to the different cues; iii) The relations between the cues and the best way to integrate them (i.e. additively or multiplicatively). In contrast to early research on feedback learning, studies in MCPL (e.g., Hammond & Summers, 1972; Hammond, Summers, & Dean, 1973; Balzer et al., 1989) show that outcome feedback sometimes impedes learning of complex relations. An even more negative view on the role of feedback, is presented by Brehmer (1980) who concludes that people are unable to learn complex rules from feedback in probabilistic environments. The argument is that for complex relations, such as combining different cues to a single criterion, people need information about how the cues are related (cognitive feedback) in order to make accurate estimates (Todd & Hammond, 1965). Kluger and DeNisi (1996) reason along the same line: feedback is efficient if it leads to task learning. Only offering a positive or negative result (outcome feedback) may lead people to set a goal to achieve positive feedback, without learning the task relations.

The cognitive processes in MCL

More recent research on multiple cue judgment has investigated how the functional relationship (linear or nonlinear) and task environment (deterministic or probabilistic) affect the cognitive representations that underlie multiple-cue judgment. One representation, the cue-abstraction-model, (CAM: Einhorn, Kleinmuntz, & Kleinmuntz, 1979) involves abstracting a rule by which each cue is combined linearly according to its contribution to the criterion,. The other representation, the exemplar based model (EBM: Medin & Shaffer, 1978) involves memorising exemplars of specific combinations of cues and criterion. A relatively large body of literature by now demonstrates that people shift systematically between these representations as i) a function of task characteristics (e.g., Hoffmann, von Helversen, & Rieskamp, 2014; Juslin, Karlsson, & Olsson, 2008; Juslin, Olsson, & Olsson, 2003; Karlsson, Juslin, & Olsson, 2007; Karlsson, Juslin, & Olsson, 2008; Pachur & Olsson, 2012; Platzer & Bröder, 2013; von Helversen & Rieskamp, 2009) and, ii) the decision maker's inclination to use exemplar memory or abstract cue-criterion relations (e.g., Hoffmann, von Helversen, & Rieskamp, 2014; Little & McDaniel, 2015; von Helversen, Mata, & Olsson, 2013). Instead of using statistical analysis, the authors apply computational models attempting to capture the cognitive representations in order to investigate people's judgments. As pointed out by Brehmer (1980), people are inclined to search for linear additive rules when asked to combine cues to a continuous criterion, hence abstracting cue-criterion relations (e.g., Juslin et al., 2008). However, when the relationship between cues and criterion is nonlinear or when the task environment is probabilistic, rules are not easily applied, instead people tend to use exemplar memory (Juslin et al., 2008). In addition, when combining the cues to a continuous criterion, people tend to integrate them additively as a weighted average (Brehmer, 1994; Hammond & Stewart, 2001) regardless of whether the relationship is linear or nonlinear (Karelia & Hogarth, 2008). Finally, there is much evidence that updating is sequential as conceptualised by the Sigma model (Juslin et al., 2008), implying that people update their criterion estimate by adding the effect of one weighted cue at a time.

Feedback frequency and probabilism

Brehmer (1980) points out that most people adopt inductive reasoning to form a mental model of the world. However, this is not sufficient, because one must also engage in deductive reasoning, that is, hypothesis testing. According to Brehmer, however, most people do not consider a probabilistic hypothesis. Rather, they confine themselves to deterministic hypotheses, leading them astray when trying to learn the relationships in probabilistic tasks.

This result is echoed in more recent research on feedback frequency in probabilistic environments. The advent of information technologies has opened up the opportunity to give people instant feedback on their performance. It is generally assumed that more feedback is better, a view that has been contested in a few studies (Lam, DeRue, Karam, & Hollenbeck, 2011; Lurie & Swaminathan, 2009). Lurie and Swaminathan (2009) investigated how feedback frequency affected performance in a study on feedback frequency, low-and high variance environment (corresponding to low and high random noise), and decision frequency. Across all four experiments, participants with frequent feedback performed worse on a newsvendor task, which was attributed to how participants accessed information. Frequent feedback made participants pay greater attention to random information and failing to compare information across trials. Even with less frequent decisions the participants were affected by the noise. Only less frequent feedback resulted in better performance.

How learning affects model choice

What representation people choose to adopt in an MCPL task is not only affected by the environment and the function form, it is also affected by the training offered prior to the test phase. Pachur and Olsson (2012) investigated how paired comparison training and direct criterion learning during training affected the performance and what computational model best accounted for the participants' behaviour in the test phase. In direct criterion learning, the participants were asked to combine the cues into continuous estimate of the criterion variable. Direct criterion learning is based on the idea that in order to learn to estimate a continuous criterion variable the feedback must include metric properties (e.g., Brown & Siegler, 1993). In the paired comparison training, the participants were asked which of two stimuli had the highest criterion estimate. Paired comparison training relates to the idea that people are able to extract metric criterion information based on pairwise comparison (Decision by Sampling: Stewart, Chater, & Brown, 2006). Several studies (e.g., Juslin et al., 2008; Klayman, 1988) have shown that highlighting how the differences between cues are related to the criterion by pairwise comparison, promotes abstraction of the cue-criterion relationships. In Pachur and Olsson

(2012), the participants in the paired-comparison training produced more accurate estimates both on old and new items than those trained with direct-criterion participants, despite not having produced a continuous estimate during training. Modelling results indicated that direct criterion training resulted in a substantially higher use of exemplar memory, whereas paired comparison training induced a higher proportion of rule-based representations. Only when presented with a non-linear statistical environment, did all participants shift to using exemplar memory.

Yet another way of conveying criterion relationships to people, is experimentation training, relying on the idea of causal relationships as described in Holyoak and Cheng's (2011) review. They argue that learning causal relationships are driven by prior assumptions about relationships that induce people to choose sparse and strong causal models and that the learning process is best described by Bayesian inference. There is also evidence (Lucas & Griffiths, 2010) that people can infer functional form from causal relationship based on covariation and that people's inferences are sensitive to noise. Hence, it is possible that learning might be enhanced with multiple cue learning emphasising the causal relations between the cues and the criterion.

To sum up, MCL studies show that people can learn the relationship between several continuous cues and a continuous criterion, but they have difficulty with learning in probabilistic tasks. Moreover, when relying on rule based relations, people tend to integrate cues linearly and additively, regardless of the true relationship. When people cannot rely on rule based relations, for instance, in non-linear environments they use exemplar memory. All MCL studies, presented here, investigate how people maximise one goal (accuracy). In real life, however, people are often confronted with multiple goals that are to be fulfilled simultaneously. I now turn to research in cognitive and social psychology that has given more attention to this subject.

The Cognitive Revolution II: The Role of Goals and Attention

Attention is intimately connected with information processing, as conceptualised in, for instance, dual-processing accounts of reasoning and judgment. Generally, dual-process theories differentiate between automatic, effortless and intuitive processes in System I, and conscious, controlled and analytic processes in System II (Evans, 2008). Evans (2008) further posits that depending on the task, the experience and the activated goal, people will process information to different extent in the two systems. Goals are the reference points for almost all behaviour (Fischbach & Ferguson, 2013). Thus, when a goal is activated, people's attention will be directed to related stimuli, both consciously and unconsciously (Dijksterhuis & Arts, 2010). As a consequence, not all stimuli that are attended to are done so consciously, because the central executive functions, which coordinate different input, have limited resources of attentional capacity (Engle, 2002). In the following I describe how single and multiple goal pursuit have been addressed in cognitive psychology and social psychology, and how goal performance is related to cognitive capacity.

Goals in cognitive psychology

Cognitive psychology has mainly addressed the pursuit of a single goal: i) In research on categorisation learning (Ashby & Valentin, 2017) the categorisation accuracy after immediate outcome feedback is emphasised; ii) In multiple cue probability learning (Brehmer, 1994; Juslin et al., 2008; Karelia & Hogarth, 2008) the accuracy of the predicted criterion variable after cognitive or outcome feedback is stressed as the main goal of the task; iii) Finally, in tasks relating to expected utility, as in multi-attribute utility theory (MAUT, e.g. Keeney & Raiffa, 1976) the outcome variable is the maximal subjective expected utility of one choice after integrating the weighted relative importance of multiple goals. Although MAUT addresses multiple goals, it evades the problem of multiple goal pursuit by reframing the task as maximization of a single goal function.

There are, however, studies in cognitive psychology that have addressed how people pursue multiple (conflicting) goals. Vancouver, Weinhardt, and Schmidt (2010) developed a multiple goal pursuit model (MGPM) that departs from control theory by representing goal pursuit in terms of feedback control systems. According to control theory (e.g., Carver & Scheier, 1998; Powers, 1973; Vancouver, 2008) the decision to follow a goal is determined by the discrepancy between the current state and reference state (the goal). Faced with two conflicting goals, a person choses to follow the goal with the highest subjective utility. A key feature in the model is expectancy, which is the difference between the subjective experience of the resources acquired and the resources available to reach a goal. The model predicts that when two goals are perceived as equal, the goal with the greatest discrepancy will be prioritised at the beginning, but as the deadline approaches, people tend to switch to the goal with the least discrepancy. A recent study (Ballard, Yeo, Loft, Vancouver, & Neal, 2016) extended the MGPM by integrating Decision Field Theory (DFT: Busemeyer & Townsend, 1993) and adding variables that enabled the model to account for more complex decisions, avoidance goals, uncertainty, and individual differences. For instance, depending on their experienced time sensitivity people will either shift between goals during the pursuit (concurrent strategy), regardless of discrepancy, or choose the goal with the highest likelihood of achievement at any time point (sequential strategy).

Translated to the paradigm in the studies in the present thesis, it is reasonable to expect that goal discrepancy and time sensitivity will affect how the participants pursue the goals (cost and comfort) in the experimental paradigm.

Goals in social psychology

Social psychology has taken a broader perspective on goal pursuit when incorporating motivation and attention as determinants of goal pursuit. Of specific interest here, many theories have addressed multiple goal pursuit. Goal-system theory (Kruglanski, Shah, Fishbach, Friedman, Chun, & Sleeth-Keppler, 2002; Koptez, Kruglanski, Arens, Etkin, & Johnson, 2012) brings together the concept of cognition and motivation when conceptualising goals and means as cognitive constructs. The architecture of a goal system affects the goal pursuit. For instance, multiple goals may be associated with one means, leading to a simultaneous goal pursuit: Bicycling to work satisfies both an exercise goal and a financial goal. However, multi-final means and goals are associated with goal-dilution (Zhang, Fishbach, & Kruglanski, 2007), leading to attenuation of the strength of the relationship between the means and the goals.

Another way of pursuing multiple goals is sequential goal pursuit, that is, focusing all attention on one goal at a time. One approach to pursuing goals sequentially is goal-shielding (Shah, Friedman, & Kruglanski, 2002; Lindenberg & Steg, 2007), where the focal goal inhibits the activation of the other goal(s). For instance, a person may choose to focus on academic success while inhibiting the fitness goal (Kopetz, et al., 2012). Sometimes alternative goals are activated simultaneously, leading to attenuated attention to the focal goal, and hence decreased commitment and performance (Orehek & Vazeou-Nieuwenhuis, 2013). Relatedly, when monitoring their progress on a goal when engaging in goal-shielding, people are affected by both their affective signals and goal attainment. Upon perceiving sufficient progress, people may either increase attention to the alternative goal (e.g., Fishbach, Dahr, & Zhang, 2006), or increase attention to the focal goal (e.g., Fishbach & Labroo, 2007). Predicting the impact of feedback during goal pursuit is far from trivial. It depends on many factors: goal framing, decision rules, and the number of goals in mind (Orehek & Vazeou-Nieuwenhuis, 2013).

Sometimes two (or more) goals may be in conflict as in two-task environments that demand the fulfilment of two goals that are incompatible. For instance, a professor is given additional teaching hours at the same time as being requested to publish another paper. Incompatible goals often lead to reduced task performance due to attentional limits or time pressure (Locke, Smith, Erez, Chah, & Schaffer, 1994). The solution of goal conflicts is affected by goal importance, goal difficulty, self-efficacy, planning, and affect (Sun &

Frese, 2012). In addition, research has found that when dual-task expectancy is high, the goal with the larger discrepancy is prioritised, and, conversely, when dual-task expectancy is low, the goal with the least discrepancy is prioritised (Latham & Locke, 2006).

To sum up, when two goals are in conflict, as in the task presented for the participants in the studies presented below, people seem to engage in sequential goal pursuit, because of failure to find a common means. Concretely, the participants will either begin pursuing the cost goal or the comfort goal, and then attend to the other goal. An alternative way of handling multiple goal pursuit is by finding multi-final means, that is, means that fulfil two or more goals. However, multi-final means are associated with goal-dilution leading to attenuation of the strength of the relationship between the means and the goals. Translated to the current task, a simultaneous goal pursuit may lead to an inferior goal attainment than a sequential goal pursuit.

Goals and representation in complex contexts: dynamic decision making

The experimental paradigm in the studies presented in this thesis has some similarities to dynamic decision making (DDM), and therefore I briefly review the paradigm and some results from studies within in that field. DDM tasks have three common characteristics: i) a series of actions are taken to achieve an overarching goal; ii) the actions are interdependent in that earlier decisions affect later actions; iii) in addition to the agent's actions, the system is also affected by random changes (Edwards, 1962; Brehmer & Dörner, 1993). A typical DDM task is presented as a cover story, for instance, You (1989) asked the participants to imagine themselves being a psychiatrist treating patients with a proactive drug. The task was to balance the fictive patients' health state, and the participants could rely on previous drug treatment and previous health state as well as current output of their acts. Results showed that even after extensive training, the participants were unable to control the system. Brehmer (1992) proposed that the explanation for people's poor performance on DDM tasks is a mismatch between the participants' mental model and the true model of the system. More specifically, when feedback is delayed, the participants are prevented from understanding the effects of the nonlinear terms. Another key to understanding people's performance on DDM tasks is provided by an individual difference approach (Funke, 1991). Participants who performed well set integrative goals and collected and evaluated information systematically, whereas those who performed poorly focused on one specific goal at a time. In sum, research on DDM suggests that, due to the complexity of the task, people have great difficulty pursuing goals in dynamic environments. In order to perform well, people must focus on general aspects and not get lost in the details.

Naturalistic decision making

Unlike DDM, naturalistic decision making (NDM) investigates how people make decisions in a real world environment (Klein, 2008). There are several NDM theories, but they all depart from the view that people rely on prior experience and simulation of outcomes. In short, the basis for a decision is a blend of intuition and analysis, for instance, fire-ground commanders simulated if a given action would put out the fire (Klein, Claderwood, & Clinton-Cirocco, 1986). If the simulation was successful, they choose that action, else they continued searching for alternatives. Another central component in NDM is satisficing, experienced decision makers consider the first option they find satisfactory. In other words, they do not optimise (Klein, Wolf, Militello, & Zsambok, 1995). In sum, NDM is an alternative to normative decision making theories, that are better suited for explaining how people make decisions in a real world environment.

Learning strategies

Reinforcement Learning

Reinforcement learning (RL) is a computational approach to learning from interaction with an environment that is goal-directed. More specifically, RL involves learning to map situations to actions so as to maximise the numerical reward signal. The agent is not told what to do, but must discover which actions yield the highest reward by interaction with the environment (Sutton & Barto, 1998). Further, actions may affect both the immediate as well as future rewards and in order to learn about the environment, thus optimising the reward, the agent must make a trade-off between exploration and exploitation. If the agent gains information about the task or environment, the behaviour is regarded as an exploring activity, whereas, exploitation is associated with receiving a reward (Cohen et al., 2007).

Other than the agent and the environment there are four sub-elements in a RL system: a policy, a reward signal, a value function and, sometimes, a model of the environment. The reward signal defines the goal, which is to maximise the total reward during the interaction time with the environment (Sutton & Barto, 1998). Importantly, the agent can alter the environmental state by means of its action and hence also the reward, however the value function remains constant. In other words, the reward defines the short-term consequences of the actions, whereas the value-function defines the long-term consequences (Sutton & Barto). More recent RL models have investigated the conflict between two goals, as conceptualised by the exploration-exploitation trade-off. Gureckis and Love (2009a) studied participants' ability to maximise long-term awards in an dynamic decision making task, concluding that in order to prioritise long-term rewards people needed to identify the correct model of the task environment. The model of the environment allows the agent to draw inferences about how the environment will respond to certain actions. Model-based methods for solving RL problem are sometimes called reflective (model-based), which are opposed to simpler and myopic reflexive (model-free) methods (Knox, Otto, Stone, & Love, 2012). Applied to the present studies, RL can be used to explain the optimisation problem that confronts the participants when balancing the cost and comfort generated from their electricity consumption. Their actions may either be characterised as model-based (reflective) or model-free (reflexive), depending on their actions in the system.

Gradient descent

Gradient descent is a formulation and solution to a mathematical optimisation problem (Snyman, 2005, p.1). Concretely, an iterative algorithm is used to find the minimum or maximum of a function. Mathematically, this is done by computing the first and second order derivative of the function (Snyman, p.6). In the present context, gradient descent may be an alternative strategy to function learning, as a way of optimising the cost and utility in the simulated household task. Concretely, the participant would search the space of electricity consumption for the minimum total cost at the maximum total utility, presumably by adjusting the appliances sequentially. This should be contrasted with function learning, where the strategy of the participant is to learn how the cost and utility are related to the appliances that are high in cost and utility. Note, though, that it would be computationally and cognitively intractable for a human to include all 18 appliances in the experimental task (used in the studies) when adopting either strategy.

Feedforward Learning

So far, I have described how people learn from feedback, but there is also evidence that people can learn from feedforward. Feedforward refers to task information transmitted to the subject by instructions, whereas feedback refers to the trial-by-trial information provided by task outcomes (Björkman, 1972). More specifically, feedforward gives the participant the opportunity to create a mental model of the task, and hence it should be easier to use the information provided by feedback (Newell, Lagnado, & Shanks, 2007). This is also emphasised by Klayman (1988), arguing that an important aspect of judgment is cue discovery. Further evidence for the importance of participants' ability to make inferences about the cue-criterion relationship comes from Castellan (1974) who showed that binary cognitive feedback was not sufficient to improve performance. Newell and colleagues (2007) go even further, when asserting that participants perform best when top-down information, from feedforward or instructions, is combined with bottom-up information from feedback.

Summary of introduction

In the introduction I have reviewed previous research that is relevant for the task of regulating one's electricity consumption. Research in multiple-cue judgement showed that outcome feedback is often not sufficient for people to learn the relations between cues and criterion in non-linear and probabilistic tasks. Further, people rely on different representations (rule-based or exemplar memory) depending on the task environment and cue relationships. In order to abstract cues, they must create a model of the task. However, under certain conditions, as when the cues are related non-linearly to the criterion, people cannot abstract rules and resort to exemplar memory. If they use exemplar memory, people are sometimes not able to generalise their knowledge outside the training range (extrapolating), which has implications for the task of regulating one's electricity consumption. In addition, the learning strategy (e.g., pairwise comparison or direct criterion learning) also affects the representation (rule-based or exemplar memory). One way of facilitating for people to create a model of the task is to give them more informative feedback, cognitive feedback. Another way is to provide them with feedforward training, that is, information transmitted prior to the test about how the cues are related. Although the feedback is informative, too frequent feedback in a probabilistic environment may impede learning, because people are prone to interpret random changes as causal changes. Hence, it is more advantageous to give less frequent feedback in a probabilistic environment.

Most cognitive psychological research in feedback learning and multiple-cue learning have emphasised one goal. The task presented in the studies here, however, contains two conflicting goals. There is, however, some cognitive psychological research on multiple goal pursuit. Results pointed at the importance of goal discrepancy and that goal prioritisation depends on the time remaining of the task. Further, research in dynamic decision making (DDM) underscored the importance of integrative goals when trying to adjust goals in a dynamic environment rather than adjusting one goal at a time. In addition, research in naturalistic decision making (NDM) showed that people rely on prior experience when making decisions in a naturalistic milieu, and that the decisions are a blend of analysis and intuition. Applied to the current context, one can expect that time pressure, goal discrepancy, and ability to use integrative goals (adjust the most costly appliances) lead to better performance, whereas focus on isolated goals leads to poorer performance. Further, the participants' prior knowledge about electricity consumption will affect how they approach the task.

In contrast to cognitive psychology, there is a wealth of research in social psychology on multiple goal pursuit. In essence, it shows that depending on the context and the individual, people may either adopt a sequential or a simultaneous goal pursuit. The choice of goal strategy depends on several factors: cognitive limits, motivation, and task context, just to mention a few. When adopting a simultaneous goal pursuit strategy people often fall prey to goal dilution, resulting in a worse performance on both variables. A sequential strategy is also associated with factors that may deteriorate the performance, for instance, when engaging in goal-shielding the non-focal goal may draw attention from the focal goal.

Finally, given that the task involves learning two functions, I presented two alternatives for how people can approach the task: function learning and gradient descent: function learning describes how people learn to relate a continuous cue to a continuous criterion whereas gradient descent describes how people, or algorithms, find the minimum or maximum of a function by iterative adjustments. Both alternatives are plausible, but the results from the studies indicate that, at least in some conditions, the participants seem to have some notion about the functional relationship between consumption and cost/utility. Yet another plausible strategy for approaching the task is by maximising the reward by minimising the error as described in reinforcement learning (RL). In RL, the agent may either take a model-based approach, similar to function learning, or a model-free approach, as proposed in gradient descent. As in multiple-cue learning and dynamic decision making, the agent needs to have a model of the task in order to take model-based approach.

In the following, I will give a presentation of the research paradigm that was used in all studies, as well as the statistical analyses and dependent measures.

Method: The Simulated Household - A New Paradigm

In all studies we used the same experimental paradigm, consisting of a simulated home that was presented to the participants on a computer screen. The task for the participant, to optimally balance the cost of the electricity consumption against its "utility", approximates the conditions relevant to an IHD (in-home display). In other words, they need to use electricity efficiently, in the sense of adhering to a limited budget for the cost while still obtaining as much utility as possible. The problem is illustrated in Panel B in Figure 1, where the utility of the "electricity consumption" is plotted on the y-axis against the cost on the x-axis. In Panel B, Figure 1, the maximum utility obtainable at a given cost is assumed to disclose a diminishing marginal utility for further consumption. The actual utility obtained at a cost by a participant is illustrated with a dot. Two idealised directions of improved electricity efficiency are also illustrated: either to increase the utility obtained at a given cost ("optimisation") or to decrease the cost of the utility obtained ("saving"). Hence, if the participants have met the cost budget, optimisation involves moving vertically upwards, increasing the utility while maintain the cost, and saving involves decreasing the cost while maintaining the same utility. If the participant has not yet met the cost budget, optimisation is conceptualised as moving horizontally to the left or upwards to the left, that is, maintaining the same utility or increasing the utility while decreasing the cost.

The task is an expansion of an MCL task in at least three respects: i) There are more cues, in total 18 appliances (temperature, lightning, etc.), that are to be mapped to the criteria (one criterion for cost and one for utility); ii) The participant must balance two goals (cost and utility) against each other. Further, the goals compete for the participant's time and attention, and they partially conflict because increased utility often leads to an increased cost. In addition, cost and utility are expressed on different scales (cost in SEK and utility in points) and with different function form (linear cost and nonlinear utility); iii) Finally, the participant must both learn to predict and control the system with her behaviour.

Depending on the experimental design, the participants spent 28-120 simulated days in the simulated household, which lasted approximately one hour.

The participant determines the level of electricity consumption for each day, while trying to adhere to the budget for the maximum cost, minimum utility, or both the cost and the utility. When a person consumes electricity in real life, the signal for the utility of consumption is "internal" and idiosyncratic. In order to objectively observe the balancing of the cost and the utility goal, the utility signal is "externalised" and provided to the participants in the program. Thus, the utility functions (one for each appliance) are the same for all participants and they refer to the electricity consumption of a fictive inhabitant of the house. The utility $u(t_i)$obtained by consumption t_i of appliance i (i = 1…18) was:

$$u(t_i) = w_i \cdot t_i^{\alpha_i}/{r_i}^{\alpha_i}, \tag{1}$$

Where w_i is the linear weight in the overall summed utility ($\Sigma_{i=1}^{18} w_i = 1$), r_i is a ceiling on the allowable consumptions and α_i is a parameter for the curvature of the utility function relevant for appliance i. Equation 1 defines utility functions with diminishing marginal return, where the appliances differ both in the rate of the diminishing marginal return (α_i) and in their weight (w_i) in the total utility. The parameters in Appendix A: Table A1 were selected to approximate realistic functions. For example, the washing machine was given a large weight in the total utility but with a fast decreasing marginal utility, reflecting the fact that using the washing machine once or twice a week generates quite a lot of utility, but using it more does not. Other activities were associated with a more linear utility function, for instance using the computer, but with a smaller total weight. Every additional hour of computer use presumably generates more utility, but compared with other activities, like heating and water use, it does not have a very large weight in the total utility. The total utility U was the sum of the utility of each of the 18 appliances. The utility was provided as a fictive unit of "utility points" in all experiments except in Experiment 2 in Study III, where the utility points were transformed to market value expressed in SEK.

Feedback on the cost was based on a fixed price of 1.40 SEK per kWh. If the cost budget was exceeded, the total cost was red-lighted in the bill that was presented after each day (or after 10 days in the less frequent feedback conditions). With detailed feedback the participants were given the cost for each of the 18 appliances, in addition to the total cost, for aggregated feedback only the total cost was presented. The total cost C was the sum of the consumption cost ($c(t_i)$ of the individual appliances defined in Equation 2.

$$C = \sum_{i=1}^{18} c(t_i). \tag{2}$$

In the deterministic conditions, the cost of a specific consumption was always exactly the same as specified by $c(t_i)$. In the probabilistic conditions, each day, a normally and independently distributed random error, with a standard deviation equal to 5% of the cost, was added to the cost of each appliance. The noise represents the multitude of factors that may affect the feedback at a certain moment and that are potentially unknown to the participant. Aspects that may affect the feedback include imprecision of the feedback instrument (IHD), interactions between appliances, and exogenous factors (e.g., outdoor temperature).

All participants, thus, received information about the utility obtained for each appliance $u(t_i)$ as well as the total utility U on each simulated day. Feedback about the cost was either provided as the total cost C (aggregated feedback) every day or every 10th day (less frequent feedback), or as the total cost C in addition to the separate cost $(c(t_i))$ for each appliance (detailed feedback) every day or every 10th day.

The task is to learn to use the electricity efficiently by balancing the cost and the utility of the electricity consumption according to the instructions in each condition. For every new day, the participant adjusts the indoor temperature, hours of lightning in the different rooms, the hot water, etcetera. The utility from each appliance and the total utility that day are presented on the screen, pictured in Panel A Figure 1, indicating how much comfort each activity gives.

A.

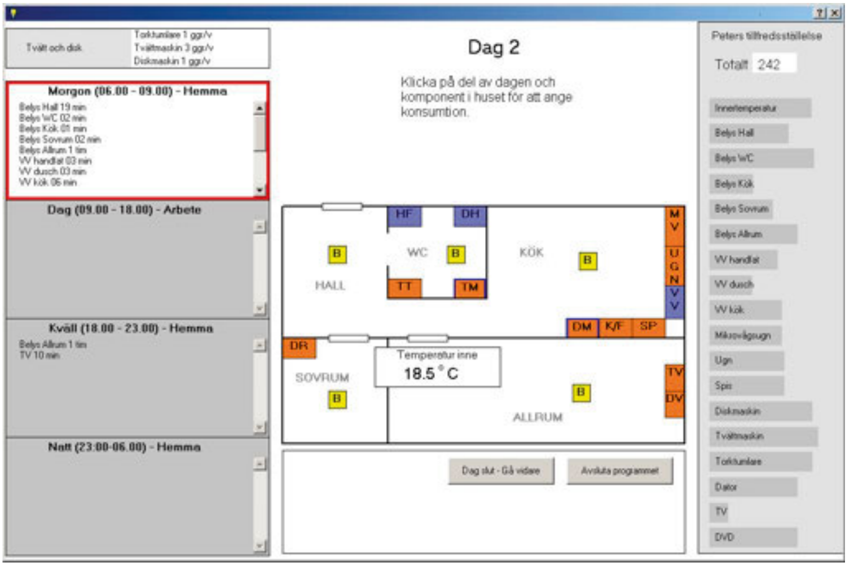

B.

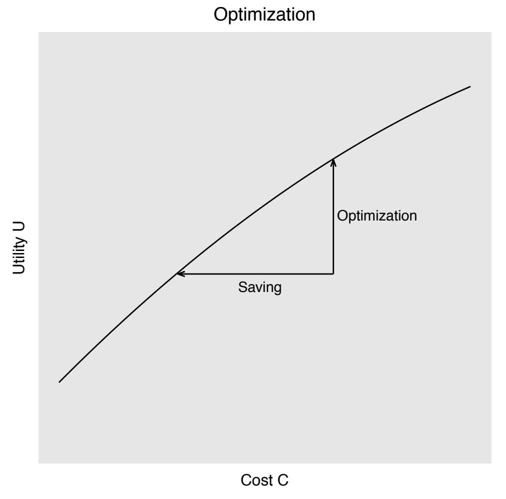

Figure 1. Panel A: The computer display that confronts the participants in the simulated household task. On each simulated day the participant indicates the daily use of the electricity consuming appliances and the fictive household inhabitant's utility from the consumption is conveyed by horizontal bars on the right side of the display. Feedback about the cost of electricity consumption is presented after a simulated day in a separate display (not shown in the figure in Panel A). Panel B: Schematic illustration of the decision problem that confronts the experimental participants, which is to maximize the utility obtained by the fictive household inhabitant given the cost expended on electricity consumption. The intersection between the lines illustrates a possible state when the cost budget is met and the two principal directions for improved electricity efficiency, saving and optimization.

Dependent variables

The dependent variables that were investigated in all studies were total cost, total utility and electricity efficiency compared with the maximum utility obtainable at a given cost. Maximum obtainable utility at each total cost was given by searching for the distribution of costs across the 18 electricity appliances that maximised the overall utility. The function relating to the maximum utility to the total maximum cost (defined by the cost budget) was approximated by a polynomial: $y = .6634 + .0046 \cdot x - 1.784 \cdot 10^{-5} \cdot x^2$. The function is plotted in Figure 1 in Panel B.

The measurement of the dependent variable varied and evolved between the studies, because of the difficulty of measuring performance in a 2D-space composed of arbitrary and incommensurable units (e.g., how does one weight the gain of one arbitrary unit of utility against a loss in terms of the equally arbitrary unit of money). Keeping one variable constant by a budget for the cost and then measuring performance in the other variable (utility) contingent on a satisfied cost budget, as in Study 1, is one solution, but it does not capture the dependence between the variables. Another approach, that we applied in Study II, in Experiment 2, was to measure the distance to the optimal curve and counting the number of budget followers. A third way is to use Survival Analysis, a statistical method that measures at what time point when and if the participants reach one or both goals. As with all research, the methodology evolves as more knowledge is accumulated, and we now believe that survival analysis is the soundest approach to this measurement problem. In Study III, two additional dependent variables were introduced: i) working memory score (Aospan), which was intended to investigate if cognitive ability affected performance, and ii) mean direction score, which was introduced to measure the direction of the movement in the cost-utility space.

Statistical analyses

In Study I, the participants behaviour was modelled with an error-driven computational model with two parts: an action part and an attention part. The action part captures how the outcomes (cost and utility) are attended to, if they are addressed. The attention part captures when and how the outcomes are addressed. In that sense, it is the key part for testing whether the participants act sequentially or simultaneously and whether their actions are model-free (reflexive) or model-driven (reflective). This pattern is represented by two parameters: $\beta_{sequential}$ that captures the degree of sequentiality and $\beta_{reflection}$ that captures the capacity to attend to cost also on trials when no feedback is given. If the participants act sequentially, they first attend to the cost and only when the cost budget is met (cost error is zero) do they begin to adjust the utility. If they, on the other hand, act simultaneously they will attend to both the cost and the utility error, regardless whether they have met the cost budget or not. Finally, those participants who are only given feedback every tenth day, act reflexively if they only attend to the utility feedback (given every day) on the days without cost feedback. If they, other hand, attend to the cost on the days without no cost feedback, they are classified as reflective.

In Study III, the participants' behaviour was analysed with survival analysis, which is a method for analysing binary outcomes. It is concerned with studying the time between the entry in the training and the first occurrence of a subsequent event, in the present context, first goal achievement of either cost, utility, or both goals. In Experiment 3, we used a Kaplan-Meier analysis that estimates the cumulative proportion of surviving individuals for each measurement time and a log-rank test is used to compare the groups. If there are several explanatory variables, as in Experiment 1, survival is analysed with Cox Proportional Hazard regression. The proportional hazard, that is, the probability for an individual to experience an event within a time interval, is regressed on the explanatory variables. The sign of the output of the regression will tell whether the variable increases (positive) or decreases (negative) the risk of death, or, in the present context, goal achievement. If the proportional hazard assumption is violated (the survival curves cross), as in Experiment 2, it is recommended to use accelerated failure models (AFT).

In Study III, the participants' behaviour was also analysed with direction scores. In order to investigate how they moved in the cost-utility space in Figure 1B, their movement direction was scored as the delta (difference) between each trial in cost and utility to account for four directions:

i. Score 1 was given for a positive delta for both variables (movement upwards to the right) that is, more utility, but higher cost;
ii. Score 2 was given for a positive delta for utility and a negative delta for cost (movements upwards to the left), that is, more utility at a lower cost;
iii. Score 1 was given for a negative delta for cost and a negative delta for utility (movements downwards to the left), that is, less utility at higher cost;
iv. Score 0 was given for a positive delta for cost and a negative delta for utility (movements downwards to the right).

The direction scores is an alternative to the computational model that was used in Study I. Since feedback was not given every tenth day neither in Study II nor Study II, there is only one parameter left in the model. Further, the parameter capturing whether the participants act sequentially or simultaneously does not take the movement direction into account. For this reason, we decided not to continue using the model in Study III. As to Study II, the prime interest was to investigate whether the participants' performance was enhanced with feedforward training, and therefore the model was not used in that study either.

Aims of the thesis

The overarching goal of the thesis is to investigate how learning to pursue multiple goals in the simulated household is affected by goal phrasing, different feedback, and task environment. The task attempts to mimic an electricity consumption task in a two-goal environment where the participant must balance the cost and the utility generated by their consumption. Even with feedback, cutting consumption is far from trivial, since it demands people to optimize across two goals: cutting the cost, while obtaining sufficient comfort (utility). Albeit information processing has been extensively studied in cognitive psychology, little attention has been given to how the cognitive limits affect the ability to use feedback to control multiple goal pursuit. Further, the task is a multiple-cue task, and results from multiple-cue learning literature show that depending on the task environment and the cue-criterion relationship, different feedback resolution and frequency will be advantageous. For instance, while frequent feedback is beneficial in a deterministic environment, it is not necessarily so in a probabilistic environment; because people tend to interpret the noise as causal relationships. People use rule-based representations as default and only shift to exemplar memory when it is not possible to derive linear additive rules, for instance, when the cue relationship is non-linear. In addition, depending on what representation people adopt, different feedback is more or less advantageous, and their ability to extrapolate is affected by the type of representation. However, whereas previously studied multiple-cue tasks involve a handful of cues, the present task includes no less than 18 cues, and therefore it is unclear whether previous results are applicable in this context. In this thesis, I bring together these domains and investigate multiple-goal pursuit and people's cognitive ability to simultaneously minimize the cost of electricity use while maximizing the utility of its consumption in a complex multiple-cue task.

Study I, investigated how feedback frequency and detail affected performance in a deterministic and probabilistic version of the electricity consumption task. It also investigated the degree of simultaneous/sequential goal pursuit and whether participants were able to derive a mental model of the task. Results showed that best performance was obtained with detailed feedback, but that participants performed better with less frequent feedback in a probabilistic environment as opposed to frequent feedback in a deterministic environment.

Finally, the participants attended goals sequentially and adopted a model-free approach.

Study II investigated the possibility to enhance the learning in the electricity consumption task. More specifically, it investigated three function-training-schemes that emphasise different learning strategies by giving the participants feedforward training prior to the simulated household task. The manipulation was intended to give the participants a model of the task, and hence enhance the performance. Results showed that a combination of a function-training schemes emphasising metric relations and outcome feedback produced the best performance.

In Study I, the participants were only given one explicit goal, a cost goal. Study III, investigated how giving the participants a defined cost and/or utility goal affected the performance on the respective variables. Further, it investigated whether the participants were able to pursue multiple goals simultaneously. Results from Study I indicated that the majority of the participants pursued the goals sequentially, beginning with the cost goal, but the design of the study did not permit an investigation of the cause of this behaviour. Study III investigated whether the sequential goal pursuit was related to the task itself or whether it related to a cognitive inability in the participants. Results showed that participants with instructions to satisfy both goals outperformed those with single goal instructions, however, they performed worse on each separate goal. Our interpretation is that attending to two goals simultaneously imposes a cognitive load that leads to a cognitive constraint, resulting in a poorer performance on each goal separately. Further, participants with two explicit goals did not act more simultaneously than those who only received one explicit goal. The results did, however, not give a satisfying answer to the question relating to the prioritisation of cost.

Summary of studies

Study I: Sequential and myopic: On the use of feedback to balance cost and utility in a simulated electricity efficiency task

Aims

The present study investigated the effect of different feedback conditions on the ability to regulate electricity consumption in a deterministic and probabilistic version of the simulated household. In addition, it examined whether the participants pursued the cost and utility goal simultaneously or sequentially. Previous research shows that people tend to integrate nonlinear relationships linearly (e.g., Hammond & Stewart, 2001) and that they have difficulties interpreting probabilistic feedback (e.g., Brehmer, 1980). Further, research on multiple-goal pursuit shows that people often focus one goal at a time, whereas the others fall in the background (Lindenberg & Steg, 2007). It is, however, possible to focus on two goals simultaneously (Orehek & Vazeou-Neuwenhuis, 2013), depending on the possibility to address the goals with the same means and if the goals are activated simultaneously.

We predicted that performance would be worse in the probabilistic task as compared with the deterministic task. Second, we predicted that detailed feedback would be superior to aggregated feedback. The former gives the participant more information of how to regulate the different electricity appliances, for example, about what appliances are most costly and high in utility. This advantage is present in both the deterministic and probabilistic task. Third, we predicted that frequent feedback should be advantageous when the system is deterministic, because it allows the participants to test causal models (Holyoak & Cheng, 2011). However, in a probabilistic system we predicted the opposite: it is advantageous to receive less frequent feedback (every 10th day) since this reduces the effect of the noise. Finally, we use computational modelling to address i) whether participants pursue the goals sequentially, attending to the cost and the utility in a sequence, or simultaneously, attending to both the cost and utility from the start; ii) Whether the participants' learning processes were model-free (reflexive) or model-based (reflexive).

Method

The study was a 2x2x2x10 mixed factorial design, with feedback frequency (every day/every 10th day), feedback type (detailed/aggregated) and task type (deterministic/probabilistic) as between-subjects variables, and repeated measurement across 12 blocks à 10 simulated days each. One-hundred participants (71 females) were randomly distributed to one of the between-subjects cells of the design. All instructions stated that the cost should not exceed 2000 SEK per month, and that the utility should be maximised, but different feedback was given for each condition.

Results

Results showed that different feedback was advantageous depending on whether the task was deterministic or probabilistic. For the deterministic task, the participants performed best when given specific feedback frequently. For the probabilistic task, the participants also benefitted from getting specific feedback, but they performed better when the feedback was delivered every tenth day. This is illustrated in Figure 2, the second column in the second row in Panel A illustrates the best performing deterministic condition. The second row and in the first column in Panel B illustrates the probabilistic equivalent. Participants in the probabilistic condition produced significantly lower mean cost, both across all trials and separately in the last 12 days.

A. B.

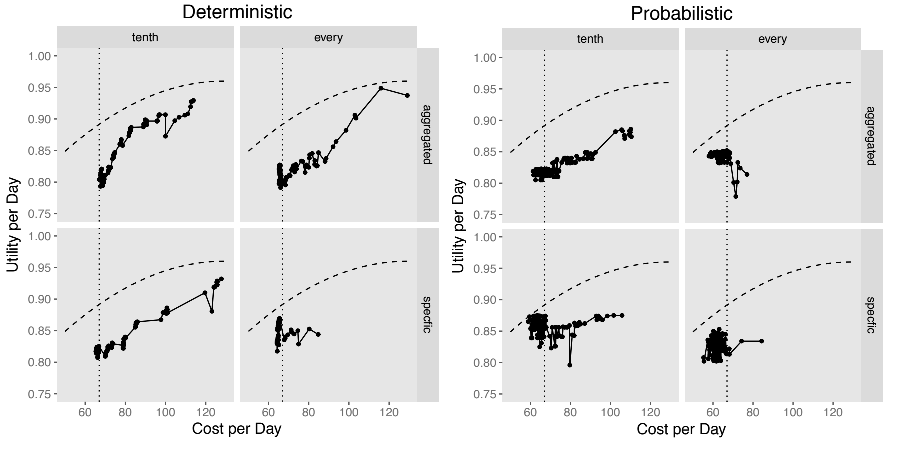

Figure 2. Data showing mean utility obtained per training trial as a function of mean cost per trial, separately for each of the feedback conditions in the deterministic condition ($N = 15$) and probabilistic conditions. The upper two rows illustrate performance in conditions where participants only received feedback about the total cost ("aggregated"). The lower two rows illustrate performance for participants that also received feedback about the cost for each individual appliance ("specific"). In the conditions presented in the first columns feedback was delivered every tenth trial (i.e., simulated day in the household), in the second columns feedback was delivered on every trial. The vertical dotted line is the budget constraint (67 SEK/day) and the curved dotted line is optimal utility function.

Further analyses showed that the deterministic conditions displayed a behavioural pattern of excessive fine-tuning of the appliances. More specifically, they displayed a higher number of manipulations for appliances with a high exponent. That is, appliances allowing a large freedom in regard to how much utility they generated. Participants in the probabilistic conditions, on the other hand, displayed fewer manipulations; focusing on appliances with a high utility and cost weight, that is, the most important appliances. The noise, thus, highlighted what appliances were most important for the cost and the utility.

The computational analysis indicated that participants acted sequentially, beginning with adjusting cost, and reflexively, only adjusting the cost on trials where feedback is given about the cost. The reflexive behaviour is depicted in Figure 2, in the first rows in each Panel, where the participants' trajectories are characterized by bursts. Each time they get feedback on the cost, they make large movements lowering the cost, but without feedback on the cost they tend to move in a small circle. Finally, although the results for $\beta_{reflection}$ are homogeneous, there are some participants (N = 4) whose results from $\beta_{reflection}$ are not. In fact, a subset of the participants have the ability to be reflective and attend to the cost goal also on trials where no feedback about the cost is available.

Discussion

We investigated how feedback type and presence of noise affected the ability to optimise the use of electricity in a simulated household. There were three main results: First, as expected, detailed appliance specific feedback produced better performance in both the deterministic and probabilistic task; Second, the optimal feedback frequency varied depending on the task: participants in the probabilistic task benefitted from feedback every 10^{th} day, whereas feedback every day was more beneficial for participants in the deterministic task; Third, participants in the probabilistic system produced significantly lower mean cost in order to create a safety margin for the fluctuating price. Interestingly, they also produced significantly higher mean utility than participants in the deterministic system. This suggests that, contrary to our expectations, noise did not have adverse effect on the performance, it seems to have drawn the participants' attention to the most important appliances. Concretely, instead of fine-tuning the use of lightning and TV, they focused their attention on factors with higher weight, for instance hot water. More generally, the results support the conclusion that people are indeed able to learn regulating their electricity consumption from feedback, even when the feedback is noisy.

As to the modelling results, they confirmed the impression that the participants addressed the goals sequentially, and acted only on explicit feedback, reflexively. Interestingly, there were a number of participants who were able to act reflectively, future research should investigate what factors contribute to the ability to attend to a variable without getting feedback.

Study 2: Optimizing Electricity Consumption – A Case of Function Learning

Aims

The present study investigated if feedforward training presented as function training schemes (FTSs) is an alternative or a complement to outcome feedback for learning to control one's electricity consumption. Previous research in feedback learning emphasises the importance of feedback providing the recipient with information about the relations of the cues rather than the outcome, cognitive feedback (Balzer et al., 1989). An alternative to cognitive feedback is feedforward, that is, task relevant information presented prior to performing the task (e.g., Newell, et al., 2007). Further, research in function learning (Pachur & Olsson, 2012) shows that depending on the feedforward training people will use different representations: i) forming explicit rules of how variables are related (e.g. "heating the house contributes to 50% of the total cost") or ii) memorizing previous situations "exemplars" (e.g. "using the washing machine twice as much as usual lead to an increase of 100 SEK on the electricity bill"). Finally, results from a previous study with the current experimental paradigm (Juslin, Elwin, Guath, Millroth, & Nilsson, 2016) showed that frequent feedback with noise was difficult to interpret. In the present study, the FTSs are given as a feedforward training, a way of providing the participant with a model of the task that facilitates the interpretation of the outcome feedback.

We examined three FTS emphasising different aspects of learning a functional relation:

i. *Paired comparison training* emphasises the ordinal relations between different electricity consuming activities, based on the theory of decision by sampling (Stewart et al., 2006). On this account teaching people to rank-order the activities according to their contribution to the total cost should be an efficient way of learning the function relations.
ii. *Direct criterion training* concentrates on the metric properties for the various electricity consuming activities. This assumption derives from Brown and Siegler's (1993) idea that in order to represent a variable, it is necessary to have knowledge of both its ordinal and distributional properties.
iii. *Experimentation training* invites the participant to experiment with the electricity consuming activities. Research pertaining to causal nets (Holyoak & Cheng, 2011) suggests that people can learn causal relations in this way.

Experiment 1

In Experiment 1 we investigated how the FTSs alone and together with outcome feedback affected participants' ability to control their electricity consumption in the simulated household. In all FTSs, appliances relating to lightning and hot water were lumped together to two categories (hot water and lightning) resulting in totally 12 appliances instead of 18. To the extent that different aspects of the FTSs are important for learning to control the electricity consumption, participants trained in that scheme should excel in the test phase.

Method

The experiment was divided into four parts:

i. Pretest: involving 30 days in the house with explicit instructions about the cost goal.

ii. Training consisting of either FTS, presented above, or probabilistic outcome feedback training (control group) during 120 days with 2000 SEK budget goal. The control group is given feedback on the total monthly cost and appliance specific feedback, whereas the FTS groups only are given feedback on the contribution of each specific appliance.

iii. Posttest 1 was identical to the Pretest with a budget set to 2000 SEK per month.

iv. Posttest 2 was also identical to the Pretest, but this time the budget was set to 1500 SEK per month.

The design was 4x2 mixed factorial with FTS as between-subjects variable and test-phase as within-subjects variable, where the comparison between the Pretest and Posttest 1 informs about the learning in each condition and the performance in Posttest 2 about the ability to generalise to a new budget in each condition. The dependent variable was the electricity consumption (cost and utility) with a particular eye to the optimization.

Results

Results showed that there were no significant differences in the Pretest between the conditions. Figure 3 illustrates performance on Posttest 1, where direct criterion training and outcome feedback training were the best performing conditions (50% budget followers each), whereas participants in the experimentation training and paired comparison training performed poorer (27% and 12% budget followers). In Posttest 2, the best performing conditions were outcome feedback training with (41% budget followers) and experimentation training (45% budget followers). In the paired comparison training and direct criterion training there were only 9% and 22% budget followers respectively. The participants with outcome feedback training performed equally well (app. 50% budget followers) on both Posttests, however, they did not outperform the best performing FTSs. In sum, Experiment 1 showed that the participants with direct criterion training and experimentation training performed at the same level as the participants with outcome feedback training in Posttest 1 and Posttest 2 respectively.

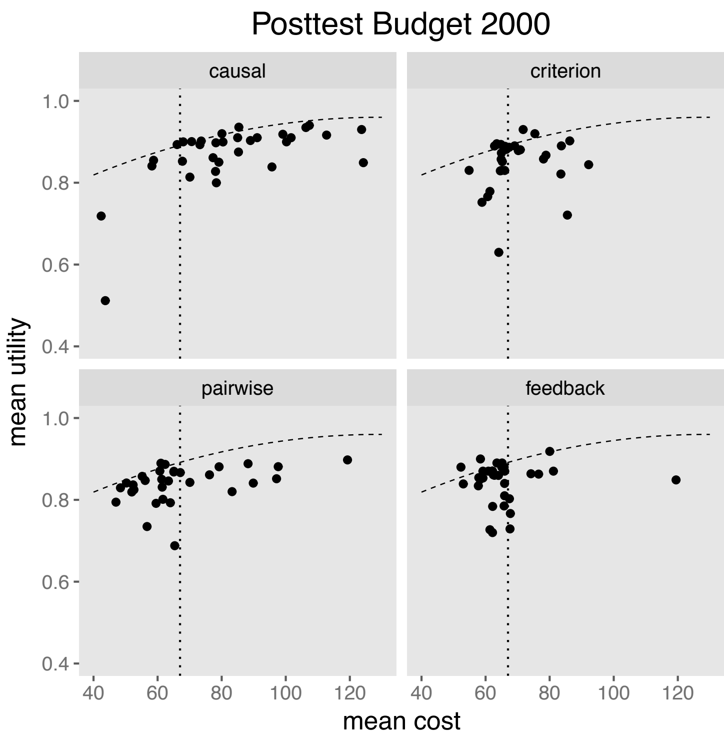

Figure 3. Mean utility for each participant plotted as a function of the corresponding mean cost in each of the four training conditions of Experiment 1 for budget 2000. The vertical line represents the budget, whereas the curve represents an approximation of the maximum utility obtainable as a function of the cost across the use of the 18 appliances. The point at which the curve and the line intersect represents the maximized utility while satisfying the budget, and the area to the right of the budget line represents exceeding the budget.

Experiment 2

In Experiment 2, we further investigated direct criterion training, alone and together with outcome feedback. Although the experimentation training was promising, we decided not to continue investigating it, given that the training was difficult to implement. We suspected that the participants in the direct criterion training used exemplar memory during training, which did not permit them to generalise to a new budget. If this is true, direct criterion training may be useful in combination with outcome feedback, by providing an informative model of the task from which the feedback information may be interpreted (Newell et al., 2007). Further, we applied a stricter stopping criterion than in Experiment 1 for direct criterion learning with 4/5 correct answers in each Block (12 in total).

Method

The experimental design was a 2x2 mixed factorial design with direct criterion training/not as between-subjects variable and budget (2000 SEK/1500 SEK) as within-subjects variables. Participants with direct criterion training first completed a training session (approximately 15 minutes), then both conditions made a Pre-feedback test either with a budget of 2000 SEK or 1500 SEK during 10 days, and finally all participants spent two sessions in the simulated household (80 days) beginning with either a budget of either 2000 SEK or 1500 SEK. In order to reduce the initial large variance, we entered a default value for all appliances corresponding to a total cost of 2730 SEK (the mean for the initial settings for all groups in Experiment 1).

Results

In Experiment 2, we analysed performance in terms of distance to the optimal performance, that is, the vertical distance to the curve of optimal performance in Figure 4. The reason for introducing a new dependent measure was that it better captured the large difference in variance between the conditions. In addition, the measure better reflects how well the participants can balance the cost and the utility.

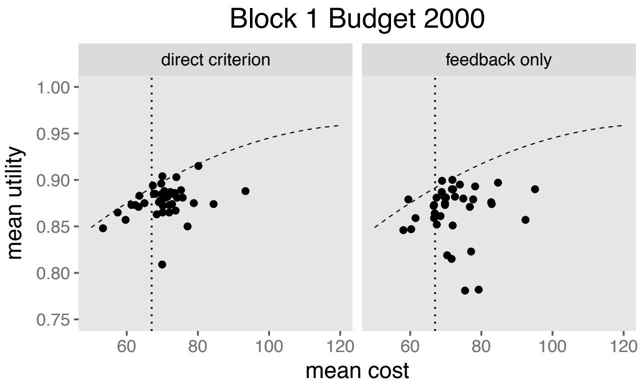

Figure 4. Mean utility for each participant plotted as a function of the corresponding mean cost in each of the two training conditions of Experiment 2 with a budget of 2000 SEK in Block 8. The vertical line represents the budget, whereas the curve represents an approximation of the maximum utility obtainable as a function of the cost, across use of the 18 appliances. The point at which the curve and the line intersect represents where the utility is maximized while satisfying the budget, and the area to the right of the budget line represents exceeding the budget.

In the Pretest, the mean distance was significantly smaller for the participants who received direct criterion training and they also had significantly smaller variance in performance. In the Posttests (1500 SEK and 2000 SEK) we aggregated the days into 8 Blocks (10 days each), and we report the performance in Blocks 1 and 8. For the budget of 1500 SEK the mean distance was not significantly different in either of the blocks. However, the variances for the direct criterion condition were significantly smaller in both Blocks 1 and 8. For a budget of 2000 SEK the mean distances in Blocks 1 and 8 were significantly smaller for the participants with direct criterion training. Further, the same group also had significantly smaller variances in both Blocks 1 and 8. The effect of variance is quite large: the variance with only outcome feedback is seven to eight times larger. This is translated to a greater number of participants with direct criterion training being near the optimal performance curve (cf. Figure 4).

To investigate the cause of the differences in the variance in performance, we performed a mixed factorial ANOVA with mean number of manipulations as dependent variable, condition as between-subjects variable, and cost category (low, medium, high cost appliances) and Block as within-subjects variables. Results showed at an interaction between cost category and condition: participants with direct criterion training made significantly more manipulations in the high cost category.

Discussion

We explored three FTSs, in two experiments, as a potential alternative or complement to outcome feedback from IHDs, with the goal to teach people how to control their electricity consumption. Results from Experiment 1 showed that participants with FTSs were able to learn to control their electricity consumption in the simulated household, despite never getting feedback on the total cost. When asked to generalise to a different budget, however, the performance dropped for all FTSs except experimentation training. In previous research, the inability to generalise has been related to exemplar memory as opposed to rule-based representations (Karlsson et al., 2007; DeLosh et al., 1997). This result is also echoed in Pachur and Olsson (2012), in which participants with direct criterion training tended to use exemplar memory whereas learning by paired comparison training invited rule-based representation. Participants with experimentation training also performed well when generalising to a new budget, but in practical terms it proved more difficult to administrate and implement this FTS than direct criterion training. It should be noted, though, that experimentation training may be a valid option for FTS, and therefore it should be further investigated.

In Experiment 2, the best performance was obtained by combining direct criterion learning and outcome feedback. Our interpretation is that the direct criterion training was useful for learning the relative importance of different electricity appliances, whereas outcome feedback was useful for implementing the insights into a useful model of the system. An analysis of the behavioural differences between the conditions revealed that the participants with direct criterion training focused on the high cost appliances, in the same way as the participants in the probabilistic conditions in Study I focused their attention on appliances with higher weight. In other words, a probabilistic system and feedforward training directed the participants' attention to the most important appliances (high in utility and cost), resulting in a better overall performance.

In both experiments, the direct criterion learning performed worse with a budget of 1500 SEK. Supposedly, the participants with direct criterion training used exemplar memory, which is more difficult to generalise. Another, complementary, explanation is that a budget of 2000 SEK was closer to the participants' every-day consumption pattern, which is supported by the fact that the mean initial settings in Experiment 1 was 2730 SEK.

In conclusion, we propose that it may be useful to take into account the extensive research in function learning and multiple-cue probability learning when designing IHDs and other types of feedback on electricity consumption. Outcome feedback is not always easy to interpret, even more so when it is noisy and frequent. Therefore FTSs may be prove to be useful to draw people's attention to the most salient and costly appliances, before starting to use an IHD.

Study 3: Why do People Pursue Goals Sequentially when Optimising Cost and Benefit in an Electricity Consumption Task?

Aims

The present study investigated the causes of the sequential goal pursuit, observed in Study I. Results showed that participants tend to address the goals sequentially, first reducing the cost to the stated budget, and thereafter increasing the utility obtained from this cost. This raises two questions:

i) The sequentiality question pertaining to whether the sequential behaviour observed in the participants derives from the instructions and the perceptions of the specific task ("The Task Adaption Hypothesis") or if it is due to cognitive limits, for instance working memory capacity ("The Cognitive Constraint Hypothesis").

ii) The prioritization question: given sequential goal pursuit, why do the participants generally start with the cost goal and then move to the utility goal (rather than the other way around)?

Study III addresses these questions, by adopting a different goal phrasing and statistical analysis. As in the previous studies, Study III used the simulated household task as experimental paradigm in all three experiments. Dependent measures in all experiments were the ability to satisfy the cost goal, utility goal, and both goals, and the trajectory in the cost-utility space. The Cognitive Constraint Hypothesis (CCH) predicts that pursuing two goals simultaneously has a "performance-price", measured by Survival Analysis, and that those participants who are given two goals are not more disposed to implement actions that lead towards both goals simultaneously, measured by a direction score in the cost-utility space (see Figure 2).

Experiment 1

In Experiment 1 we also define an explicit budget for the utility goal, hence equating the specificity and the salience of cost and utility. We also added a working memory task (Aospan: Unsworth, Heintz, Schrock, & Engle, 2005; Unsworth & Engle, 2008), to investigate if cognitive limitations are responsible for the sequential approach. All participants were instructed to minimize the electricity cost while maximising the utility from its consumption. Further, by explicitly defining a budget for either the cost and/or the utility, we investigated if the goal specificity affected i) the sequential goal pursuit, and ii) the cost prioritisation. If the sequential behaviour derives from the task, we expect it to disappear when both goals are explicit. If it derives from a cognitive inability, the effect will remain, even with both goals defined. If the previously observed initial prioritisation of cost derives from the fact that only the cost budget was explicitly defined, we expect the effect to disappear in the conditions with an explicit utility budget or both goals defined.

Method

The participants were randomly assigned to one of six conditions, created by a 2x3x28 mixed factorial design with Explicit Time Limit on the training trials or not and Goal Focus (Focus Cost, Focus Utility and Focus Both) as between-subjects variables, and Trial (simulated day) as within-subjects variable. All participants completed 28 days in the simulated house. However, half of them were told that they were given 28 simulated days to complete the task (explicit time limit) whereas the other half was given no explicit time limit. The time limit was included to evaluate if performance was affected by imposing a time pressure (cf., Vancouver et al., 2010).

Results

In all experiments, the data were analysed as follows: First, we performed a survival analysis (SA, see Method for more details) for each goal (cost goal, utility goal, and both goals) for each condition. For each simulated day, SA measures the proportion of participants that have not reached the goal(s). Second, we analysed the direction of movement with direction scores (see Method for details). Finally, we performed a mixed ANOVA with mean score as dependent variable, and condition as between-subjects variable and Block (simulated days aggregated in 4 or 6 Blocks) as within-subjects variable.

The explicit time limit did not have any effect on the performance, hence the groups were aggregated. The SA showed that participants performed best on the goal for which they were given specific instructions, for instance, participants given cost goal instructions had the highest cost goal satisfaction (see Figure 5). Although participants with both goals instructions performed best on both goals satisfaction, they performed worse on each separate goal (cost satisfaction and utility satisfaction). In other words, attending to two goals had a price. The scoring analysis revealed that participants who were given both goals instructions did not have a higher ability to move in both directions simultaneously. We correlated working memory and satisfaction of each goal (cost goal, utility goal, both goals) on day 28. Results showed that there was only a weak negative correlation between cost and working memory, implying that working memory correlated with better performance on the cost goal. In sum, giving the participants instructions to follow both goals did not lead to less sequentiality, and the cost was still the goal that was prioritized first.

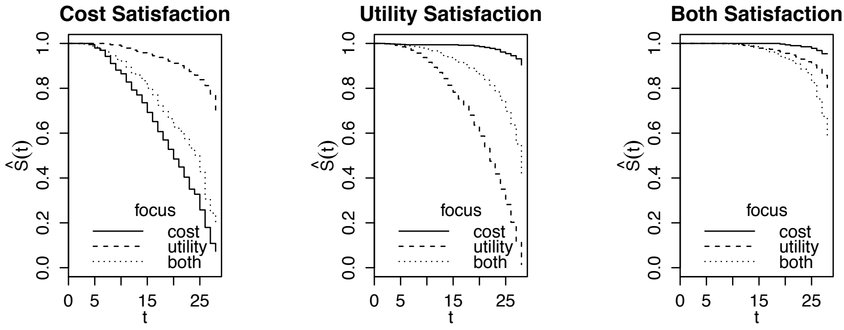

Figure 5. The y-axis depicts the proportion of the participants that have not yet satisfied the cost goal (left panel) the utility goal (middle panel), and both goals (right panel) in each of the focus conditions, Focus Both, Focus Cost and Focus Utility, after each of the 28 training trials *t* (simulated days in the household) (x-axis).

Experiment 2

Experiment 2 further investigated the reasons for the sequential goal pursuit and goal prioritization for participants with both goals specified. Since the ability to pursue both goals increased during training, we investigated if the inability to attend to both goals simultaneously could be cured by extending the training to 60 trials. In Experiment 1, the unit of utility (points) was initially unfamiliar to the participant, whereas the cost unit (SEK) was well-known. Evaluability (Hsee & Zhang, 2010) refers to the extent to which a person has relevant information to evaluate a value. In order to increase the evaluability of the utility variable, we rephrased the utility goal for half of the participants into market value by conceptualizing the metrics of utility as the market price, the rent (in SEK), and the cost (in SEK) as the costs that were associated with renting out the house. If the initial prioritisation of cost, observed in the previous studies, was due to the fact that the cost variable was expressed in a more evaluable unit, this effect should decrease in the condition that was given cost and utility expressed in the same unit (SEK).

Method

The participants in the market conditions were told that they were landlords who were to let out the house for the highest market price possible (minimum 2850 SEK) with as small cost (maximum 2000 SEK) as possible. The experiment involved a 2x3x60 mixed design with Utility Wording (Utility Points or Utility Market) and Goal Focus (Focus Cost, Focus Utility or Focus Both) as between-subjects variables, and Trial (simulated day) as within-subjects variable.

Results

SA showed that, contrary to our hypothesis, utility expressed as market value did not improve utility satisfaction, however, it did improve cost satisfaction in the conditions that were focused on the utility goal (36%) or both goals (61%). As to the satisfaction of both goals, the participants with utility as market value produced 70% poorer both goals performance than those given utility as points. In sum, utility expressed as market value did not improve utility goal satisfaction, but it did improve performance on the non-focal goal for participants given one goal and the cost goal for participants given both goals instructions. Finally, as in the previous experiment, the cost goal was prioritized initially, regardless of utility wording (points/market value).

Analysis of mean scores for movement in the cost-utility space showed at a main effect for utility wording, with a significantly higher mean score for Utility Points. In other words, when utility was expressed as money, the participants displayed less simultaneous movements. In sum, expressing utility as market value did not cure the participants' tendency to begin adjusting the cost goal, and an extended training did not produce a more simultaneous behaviour in participants given both goals instructions.

Experiment 3

A possible explanation for the initial prioritisation of cost is that the utility function is non-linear, in contrast to the linear cost function, and hence more difficult to learn (e.g., Brehmer, 1980). In Experiment 3, we further investigated the initial prioritization of reductions in cost by transforming the non-linear utility function for half of the participants, while the other half was given a non-linear utility function. We only included one Goal Focus, Focus Both, since our prime interest was to see whether a linear utility function would affect the cost prioritization for participants who were given two explicit goals. The design was a mixed factorial design, 2x60, with Utility Condition (linear/non-linear) as between-subjects variable and Trial (simulated day) as within-subjects variable. The number of trials was 60, as in Experiment 2. To the extent that the previously observed cost prioritisation depends on the nonlinear character of the utility variable, this effect should be reduced in the condition in which both the cost and the utility are linear functions.

Results

SA showed that participants with linear utility satisfied the cost goal somewhat better (cf. the left Panel in Figure 5) and the participants with non-linear utility satisfied the utility goal to a higher extent. However, regardless of utility function the participants satisfied both goals equally well (cf. the right Panel in Figure 5), in other words, transforming the utility to a linear function did not improve the participants' both goals satisfaction. An analysis of the mean scores showed that participants with non-linear utility were slightly better at moving in both directions simultaneously (M_{linear} = 1.31 vs. $M_{non\text{-}linear}$ = 1.37).

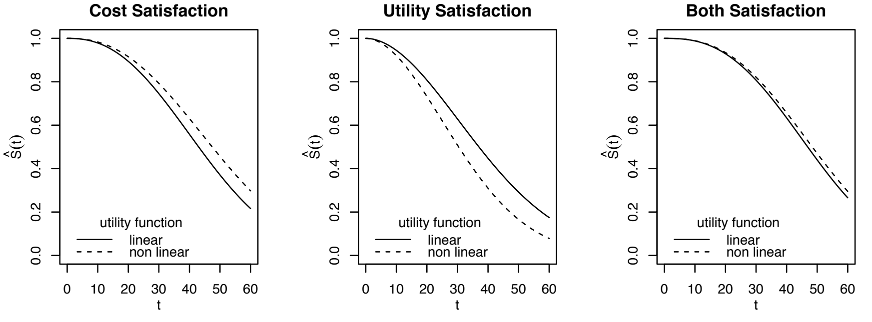

Figure 5. The y-axis depicts the proportion of the participants that have not yet satisfied the cost goal (left panel) the utility goal (middle panel), and both goals (right panel) in each of the utility function groups, Linear Utility and Non-Linear Utility, after each of the 60 training trials *t* (simulated days in the household) (x-axis).

Discussion

In Study 3 we investigated whether the sequential rather than the simultaneous goal pursuit derived from a cognitive constraint or was related to the task. Second, we investigated the reason why the participants begun adjusting the cost value rather than the utility value.

In Experiment 1, results showed that the participants with instructions to satisfy both goals outperformed those with single goal instructions, however, they performed worse on each separate goal. Our interpretation is that attending to two goals simultaneously imposes a cognitive load that leads to a cognitive constraint, resulting in a poorer performance on each goal separately. Further, there was only a weak correlation between working memory capacity and both goals satisfaction. One interpretation is that regardless of working memory capacity, pursuing two goals simultaneously may have been beyond the participants' ability in any event.

In Experiment 2, the utility unit was transformed to market value (in SEK) for half of the participants and the other half was given utility expressed as points. Results showed that the transformation was advantageous for the participants attending to one goal, but not for two goals where it contributed to worse performance. The performance of the participants with two explicit goals was impeded by getting utility expressed as money. Supposedly, the two goals were diluted (Zhang, et al., 2007), which may happen when the goals are not distinct enough and share common means. In this case, both goals were represented by sums of money and the means to achieve the goals is adjustment of electricity consumption.

In Experiment 3, we made a last attempt to answer the prioritization question, by transforming the nonlinear utility to a linear function. Results from Experiment 3 did not support this hypothesis: participants performed equally well on both goals satisfaction and there were only small differences in the satisfaction on the utility and cost. We could not find any evidence that the function form affected the order in which the participants begun adjusting the variables, and therefore we could not draw any conclusions about the cause of the order in which they were adjusted.

To sum up, results from three experiments support the view that the ability to pursue two goals simultaneously is inferior to that of satisfying each goal separately. Further, this inability seems to derive from cognitive limitations rather than task specific features. Regardless of instructions and training period, the participants begun with adjusting the cost, a pattern which could not be explained by giving a utility budget (Experiment 1), by making the utility more evaluable (Experiment 2), or by transforming the utility to a linear function (Experiment 3). The implications of the results is that people seem to be inclined to process goals sequentially, and hence interventions that target two goals should be designed in such way that facilitates attention of one goal at a time. In the case with electricity consumption, the best results is obtained if people first are encouraged to attend to the cost, and then adjust the utility that is generated by that cost.

General discussion

Summary of results

Three studies investigated how learning to pursue multiple goals in a simulated household is affected by goal phrasing, feedback type, and task environment. Results indicated that people are indeed able to learn to regulate the electricity consumption in the simulated household, despite the complexity of the task, involving 18 cues that are to be mapped to conflicting criteria (cost and utility). Further, in contrast to previous findings in multiple-judgement research, results from Study I indicated that the participants' performance was not negatively affected by probabilistic outcome feedback. Less frequent feedback with random error drew the participants' attention to the appliances high in cost and utility, leading them to create a mental model of the task. In Study II, the best performance was achieved with outcome feedback in combination with feedforward training, causing the participants to identify the relative importance of different energy consuming activities. This is in line with previous research pointing at the importance of a correct model of the task in order to create rules about how the cues are related to the criterion (e.g., Balzer et al., 1989). Finally, results from Study III indicated that the performance was improved by getting both goals explicitly phrased. However, although the participants were able to pursue two conflicting goals, they were not able to do so simultaneously. Rather, they began by attending to the cost goal, and then the utility goal, demonstrating an inclination for sequential goal pursuit.

Discussion of the results

The results and the future direction of studies will be discussed at two levels: First, I discuss aspects that are relevant to the applied setting implied by the task content and cover story used in the experiments, that is, how the results relate to the use of feedback from an IHD (in-home display). Second, I consider the results from a more general psychological perspective, interpreting the task as a generic cognitive task to study processes and abilities in regard to multiple goal pursuit and feedback learning. From this perspective, the specific content of the task is more arbitrary, and utility and cost can in principle be replaced by any two goal variables, such as revenue and cost in an organisation with a number of sub-sections. Finally, I discuss the limitations of the studies, pertaining to both the applied setting and a general psychological perspective, ending with implications and future research.

Applied setting

Learning from an IHD

In all studies the participants displayed an ability to regulate their consumption in the simulated household. The experimental design did not permit us to investigate in detail how the participants approached the task of regulating all appliances, but the results from Study I and II indicate that at least some of the participants developed a mental model of the task enabling them to focus on the appliances with a high utility and cost weight. This gives some support to the view that IHDs may be efficient way of teaching people how to regulate their electricity consumption. However, these results do not imply that mass installment of IHDs is the most cost effective way to achieve this goal. On the contrary, the results from Study II, indicate that the learning from an IHD can be replaced by a feedforward training, particularly if the training is opted at the consumers' individual consumption pattern and goals. This view is supported with results from Krishnamurti et al. (2013) who conclude that the ability to interpret the feedback is central for learning from an IHD. Accordingly, a set of heuristics or a feedforward training may show to be more effective for cutting the consumption. Further, results from Study III emphasise the importance of explicit goals, which is confirmed in Buchanan et al.'s (2014) review, who conclude that feedback itself is not enough for changing a behaviour. In sum, an IHD may make the invisible electricity visible, but changing a behaviour is more complex. It involves interpreting the feedback, setting goals, implementing new consumption behaviours, and, importantly, keeping up the motivation to continue with the new behaviour. Motivation has not been addressed here, but it is, of course, paramount in an applied setting.

Factors contributing to a good performance

Overall, the participants performed well, most of them were able to satisfy both the cost and the utility goal by the end of the training. Contrary to what was indicated by previous research, the performance was not affected negatively by probabilistic feedback. This is good news for the applied use of IHD, given the inaccurate measurements of IHDs (Hargreaves et al., 2010). Indeed, Study I established that adding noise to the feedback resulted in a more optimal performance, leading the participants to focus on the most costly appliances. A possible explanation for this behaviour is given in a study using a DDM (Dynamic Decision Making) task (Gureckis & Love's, 2009b), where low noise on the reward function increased the participants' explorative behaviour, which, in turn, resulted in a better performance as compared with the no-noise condition. It seems as if noise, both in a DDM task and the present MCPL task, induces the participant to create a better model of the system.

Another aspect that helped the participants identifying the appliances that contributed most to the cost and utility was less frequent feedback. In Study I, the results showed that deterministic and frequent feedback caused the participants to engage in excessive fine-tuning of the appliance with a linear utility function (as opposed to appliances with utility functions with a diminishing marginal return). In contrast, the participants in the conditions with less frequent feedback were able to discover what activities had the largest weights in cost and utility. The results are echoed in Krishnamurti et al.'s study (2013), where aggregated feedback lead to better performance. They conclude that: "[..] people would fare better from just receiving projected monthly costs for their appliances, as they have difficulty extrapolating from current cost and energy use to monthly cost." (p. 455). At a more general level, frequent feedback may cause people to fall prey to the "covariation illusion" (Denrell & Le Mens, 2011). In real life, this may cause people to focus on appliances that give immediate co-variation in consumption (e.g., vacuum cleaner), not paying attention to other correlations. In the current task, it caused the participants to focus on appliances that have no ceiling their utility function (e.g., computer). In both cases, people draw the false conclusion that those appliances contribute most to the cost and utility.

General psychological view

Sequential goal pursuit and cost prioritisation

Results from the studies point at a propensity to address the goals sequentially, first the cost goal and then the utility goal. In Study I, modelling results indicate that the participants pursued goals sequentially. Although, the question was not addressed explicitly in Study II, post-hoc analyses indicated at a similar movement pattern as that observed in Study I. Finally, the results in Study III, where the sequential goal pursuit was explicitly addressed, were unambiguous: the participants first begun with satisfying the cost goal, and then tried to increase the utility. In addition, the direction scores did not indicate that participants with two explicit goals acted more toward both goals than those with one explicit goal. In Study III, the interpretation of the sequential goal pursuit is that it pertains to a cognitive constraint. The argument is that neither transforming the utility to the more familiar unit of money, nor transforming the utility function to a linear function, increased the simultaneous goal pursuit. However, it is possible that there exists individuals who are able to pursue goals simultaneously: post-hoc analyses of data in Study III show that some individuals move more simultaneously than the results at group level indicate.

The number of participants was too small, though, to make any further analyses.

The results in Study III did not give a satisfying answer to the prioritisation question. One possible explanation is given by time sensitivity as conceptualised in Ballard et al.'s (2016) model for multiple goal pursuit (MGPM*). In contrast to the explicit time pressure that was used in Study III, Ballard et al.'s time sensitivity parameter measures how the *experienced* time pressure affects the participant: A person who perceives time-pressure will always prioritise the goal with the highest likelihood of being fulfilled, engaging in a sequential goal pursuit, and, contrarily, a person who does not perceive time pressure will shift between both goals at any time during the pursuit. Accordingly, if the cost was perceived as a goal with a high likelihood of being fulfilled, the participants with high time-sensitivity would prioritise that goal, hence begin by adjusting the cost, and continue doing so during the training until the budget is satisfied. A complimentary explanation is offered by looking at the exponents associated with each of the eighteen utility functions. Most exponents have an $\alpha \leq .1$, which entails an almost horizontal curve, which in turn makes it very difficult for the participant to detect any changes in utility. Hence, the participants may have concluded that it was difficult or even impossible to change the utility variable. In addition, it is plausible that people focus more on cost than utility due to loss aversion (Tversky & Kahneman, 1991).

Non-linearity

A nonlinear utility function was chosen for two reasons, first it captures the quality of diminishing marginal utility of the variable, and second, it permitted us to investigate the participants' ability to control and predict a nonlinear variable. It is, however, possible that the participants used a linear strategy when estimating the utility: i) A wealth of research point at people's inclination to use a linear model when estimating nonlinear functions (e.g. Soyer & Hogarth, 2015); ii) The nonlinearity of the utility functions may have been too small to affect performance, and therefore although using a linear model the participants may achieve quite accurate judgments. Relatedly, transforming the utility to a linear function did not enhance the participants' performance (cf., Experiment 3 in Study III). The present paradigm does not permit modelling the data with CAM (cue-abstraction model) and EBM (exemplar-based model), since it involves too many cues. Results from Study II were interpreted to mean that the participants with direct criterion learning used exemplar memory, and, hence, they were unable to generalise to a new budget. Participants given outcome feedback, were, however, capable of generalising, presumably using a rule-based representation. Future studies should investigate if people given outcome feedback and direct criterion learning rely on exemplar memory and rule-based representation, and, in addition, what representation people are induced to use with causal learning (the third FTS).

Creating a model of the task

The participants' behaviour was investigated with an RL (reinforcement learning) inspired model in Study I. The model resembles RL models in that it emphasises actions as a measurement of performance and it uses the prediction error to account for future actions. Specifically, it focused on whether participants were able to derive a mental model of the task or not (reflective/reflexive), and whether they could take both goals into account in one action or not (simultaneous/sequential). A future model should take a more detailed approach, for instance, by taking into account whether the participants adjust high/low weight utility appliances, since this is an indicator of whether they have a mental model of the task. Further, a future model should investigate the exploration-exploitation trade-off (e.g., Mehlhorn et al., 2015) by looking at actions that maximise the optimal balance between cost and utility, instead of looking at the reward-prediction error on each trial for each separate variable. Another promising expansion of the paradigm is to give the participants cues, on each trial, that may enable them to create a mental model of the task. Results from a study (Gureckis & Love, 2009a) using a DDM task showed that by giving the participants cues about the system state they were able to build a mental model of the system, which in turn enhanced performance.

Another aspect that pertains to whether the participants derived a mental model of the task or not is what strategy they used to regulate the appliances in the simulated house. One possibility, which was investigated in Study II, was that they used a strategy akin to function learning, in essence, focusing on learning what are the most important appliances. Another possibility, is that they used a strategy akin to gradient descent (Snyman, 2005), that is, fine-tuning the appliances sequentially in order to find the minimum cost and maximum utility. It is, indeed, possible that the participants in the deterministic environment in Study I approached the task in such a fashion. Post-hoc analyses of the data show that they made significantly more manipulations than participants in the probabilistic environment. However, this strategy caused the participants to act myopically, engaging in excessive fine tuning of the less important appliances that allowed a lot of covariation between their use and the utility. As a consequence, they performed worse than the participants in the probabilistic environment who focused on the most important appliances, and presumably derived a better mental model of the task.

Finally, the causal relations training provided in Study II, deserves further investigation. A promising expansion is to provide the participants with information about their past feedback. In Hogarth, Mukherjee, and Soyer (2013) the best performing simulation condition involved providing participants with feedback of past trials (outcome and success) giving them a larger sample to draw from when producing their estimates, resulting in more accurate judgments.

Limitations

There are, obviously, limitations with investigating electricity consumption in an experimental environment, one being the short time scale: the participants learn within an hour to control their electricity consumption, which in real life takes weeks or months. Further, participants are given the same externalised utility signal whereas in real life they are guided by their own internal utility functions. However, it could be argued that the task is by no means easier in real life. On the contrary, it is even more complex to learn to control one's electricity consumption in a busy error-perturbed real-world situation. As discussed previously, with instant feedback and without a model of the system, people may fall victim to the "co-variation illusion" (Denrell & Le Mens, 2011).

A limitation that is relevant from both the applied an the more general psychological perspective is that the analyses did not include analyses at an individual level. This is indeed a weakness, both from a methodological point of view, but also when considering the results pointing at relevant individual differences in performance. In Study I, a subset of the participants were classified as reflective by the model, thus taking the cost into account even when no feedback was given. Further, in Study III, post-hoc analyses pointed at individual differences in the directions scores, indicating that some participants were indeed able to move more simultaneously in the cost-utility space. Taken together, the results indicate that there are pertinent differences in how the participants approach the task, which in turn points at the possibility that people may profit from different types of facilitating cues in this task, and supposedly also in a real world setting.

Relatedly, the participants' prior knowledge about electricity consumption was not measured. In hindsight, this is lapse, since research in causal learning shows that the prior has implications for how people learn a task (Yeung & Griffiths, 2015) and people use knowledge about causal relationships from experience as a guide in causal learning (Lucas & Griffiths, 2010). Presumably, participants were guided by previous knowledge of electricity consumption when learning to regulate the electricity consumption. Data on their individual priors would have contributed to the understanding of how they approached the task.

Finally, we changed the dependent measurements and the number of days spent in the simulated household between the studies (and within the studies). This makes it difficult to compare results between the studies, and ideally, we should have continued with the same dependent measures and training in all three studies. However, as knowledge was accumulated, more insight was gained in what settings and measures were more adequate for investigating peoples' behaviour. For instance, spending 120 days (as in Study I) in the simulated house proved to be redundant and hence the training was shortened. Also, the insight about how to best measure the dependent variables evolved as data was gathered, and hence the dependent measures were changed.

Implications and future directions

Applied setting

It would be of interest to examine if function training schemes, which aim at teaching people the weight of each appliance, are suitable for teaching people to control electricity consumption in a real world environment. They could either be given in combination with an IHD or alone. In the former case, people would be offered a short training when installing the IHD software. The feedforward training schemes could either offer simple heuristics of what appliances are high in cost and utility, or a simulation tool permitting people to discover how cost and utility interact. An important question is whether the training can replace detailed outcome feedback, because, presently, it is unfeasible and costly to deliver detailed, appliance specific feedback. Ideally, people would answer a few questions that would generate their own utility function, which would be applied during the training. It would also be of interest to give people less frequent feedback from their IHD, despite the opportunity of getting instant feedback, since frequent feedback is detrimental for learning in probabilistic environments.

Another promising expansion of the paradigm in an applied setting is to investigate how motivation affect the performance in the simulated household. Previous research has shown that economic incentives may counteract positive attitudes towards environmental behaviour (Thøgersen, 2003) or demotivate people to act in an environmentally friendly way (e.g., Stern & Kirkpatrick, 1977; Thøgersen, 1994). One reason is that economic incentives appeal to extrinsic values that are counteractive to the intrinsic values that are associated with beliefs that promote sustainable behaviour (Evans, Maio, Corner, Hodgetts, Ahmed, & Hahn, 2013; Bolderdijk, Steg, Geller, Lehman, &

Postmes, 2012). Adding a social dimension to the feedback seems to be beneficial. However, a disadvantage with social norms feedback is the counteractive boomerang effect (Fischer, 2008), leading good performers to regress to the mean. The effect has been curbed by, for instance, the presentation of an emoticon together with the feedback (Miller & Prentice, 2016), that is, a positive reinforcement. The mechanisms behind the boomerang effect are, however, unclear, and so are the effects of competition within and between individuals and its relation to social norms. In order to evade the boomerang effect, and benefitting from motivational power of norms, future studies should isolate competition from social norms and investigate the effects of competing individually vs. toward others.

General psychological view

An important and potentially fruitful expansion of the paradigm from a general psychological aspect would be to further investigate the reason for the prioritisation of cost. Is it due to loss aversion, or is it due to the small exponents, leading the adjustments in the utility to become undiscernible? If it is loss aversion, it has implications for the design of the an IHD and possibly also for other appliances presenting information involving cost and utility. For instance, appliances giving information about how many calories different training generates. Given the results in the presented studies, it is possible that people would prioritize the cost, that is, the effort it takes to exercise, rather than the utility, that is the output – better fit and calorie loss. This in turn, would potentially put people off from exercising. Another hypothesis was that the cost prioritisation was tied to the difficulty of pursuing each goal. This may have affected the order in which goals are pursued, as conceptualised in Ballard et al.'s (2016) model for multiple goal pursuit. Uncovering the reason for the prioritisation of cost may potentially advance the knowledge of multiple goal pursuit from both an applied and general psychological point of view. This would, for instance, enable the design of more efficacious learning environments for similar tasks as well as moderating people's work flow (e.g., which goal they prioritise) in organisations.

Another future direction of research is to compare people's performance on a conflicting goal task with multiple cues with the optimal performance. Research on Bayesian updating of learning (Steyvers, Lee, & Wagenmakers, 2009) in a choice task shows that it is indeed a useful way of analysing people's behaviour. For the present task, however, it would be quite demanding to calculate the optimal behaviour, but given the potential benefits it may be worthwhile. Comparison between the participants' performance with the optimal behaviour, would possibly give insight in the individual processes of pursuing conflicting goals. Further, with a predefined optimal behaviour, one could define, for each decision, whether it is more beneficial to explore or exploit each variable. This would, in turn, clear the way for a model taking into account whether it is best to explore or exploit, and in that way, gain insight in the participants' behavioural limitations. For instance, do people exploit, for instance by gradient descent, to a greater extent in a deterministic than in probabilistic environment? And how does that relate to the optimal performance in each environment?

Concluding remarks

The findings in the thesis add to the large body of feedback learning research, showing that outcome feedback alone is not sufficient for modifying behaviour. Particularly, instant outcome feedback causes people to fall prey to the "covariation illusion", leading them to chase "gradient descent" rather than trying to get an overall view. Further, people pursue two conflicting goals sequentially, adjusting one goal at a time, first the easiest or the most relevant goal and then the more difficult. Optimisation is thus achieved sequentially, rather than simultaneously. In addition, the studies showed that the results from cognitive psychological research are also relevant for a more applied setting, electricity consumption. This is not very surprising, but given that much research is conducted within each sub-discipline, this thesis contributes to the growing number studies aiming at connecting different areas within psychology.

Acknowledgments

I would like to thank Peter Juslin, my supervisor, whose fantastic support and profound knowledge was crucial for completing this thesis; Anders Winman, my co-supervisor, for knowledgeable comments on my research; Ebba Elwin for wisely pointing me in the right direction in the research jungle; Håkan Nilsson for his sharp comments and knowledge of R; Marcus Lindskog for wise comments on my research; Leo Poom for sharing his insightful perception psychologist views; Ronald van den Berg for introducing me to Bayesian statistics; Philip Millroth for fruitful discussions about prospect theory; Britt Stikvoort for creative discussions about applied cognitive research; Elina Stengård and Joakim Sundh for their mathematical insights, and August Olsson for sharing interesting research. I also want thank Mathilda Frick for valuable discussion early mornings; Johanna Motilla-Hoppe for being a supportive room companion; Johan Lundin-Kleberg for intelligent discussions about life and science; Andrea Handl for inspring talks about science and language. Many thanks also, my friends, Cecilia Eriksson, Carolina Orozco-Laakso, Lina Aquadro, Milena Haykowska for believing in me when I did not; My parents, Annelie and Saad, for always having encouraged my interests; My brother Vic for seeing things as they are; My lovely daughters Adèle and Lucie for reminding me about the importance of detail; Finally, Stefan, for showing me a new way of thinking, both in life and science .

References

Anderson W, White V. *Exploring consumer preferences for home energy display functionality.* Technical report; 2009.

Ashby, F.G., & Valentin, V.V. (2017). Multiple Systems of Perceptual Category Learning: Theory and Cognitive Tests. In H. Cohen, & C. Lefebvre, (Eds.) *Handbook of Categorization in Cognitive Science* (2nd ed.), (pp. 157-188). https://doi.org/10.1016/B978-0-08-101107-2.00049-X

Atkinson, R.C., & Shiffrin, R.M. (1968). Human memory: A proposed system and its control processes. In K.W. Spence & J.T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (vol. 2, pp. 89-195). New York: Academic Press.

Ballard, T., Yeo, G., Loft, S., Vancouver, J.B., & Neal, A. (2016). An Integrative Formal Model of Motivation and Decision Making: The MGPM*. *Journal of Applied Psychology, 101(9),* 1240-1265.

Balzer, W. K., Doherty, M. E., & O'Connor, R., Jr. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin, 106,* 410-433.

Bandura, A. (1977). Social Learning Theory, Alexandria, VA: Prentice Hall.

Bandura, A. (1991). Social cognitive theory of self regulation. *Organizational Behavior and Human Decision Processes, 50,* 248-287.

Björkman, M. (1972). Feedforward and feedback as determiners of knowledge and policy: Notes on a neglected issue. *Scandinavian Journal of Psychology, 13,* 152–158.

Bolderdijk, J.W., Steg, L., Geller, E.S., Lehman, P.K., & Postmes, T. (2012). Comparing the effectiveness of monetary versus moral motives in environmental campaigning. *Nature Climate Change, 3,* 413-416.

Brehmer, B. (1980). In One Word: Not From Experience. *Acta Pscyhologica, 45,* 223-241.

Brehmer, B. (1974). Policy conflict, policing consistency and interpersonal understanding. *Scandinavian Journal of Psychology, 15,* 273-276.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica 81,* 211-41.

Brehmer, B. (1992). Dynamic decision making: human control of complex systems. *Acta Psychologica, 81(3),* 211-241.

Brehmer, B. (1994). The psychology of linear judgment models. *Acta Pscychologica, 87,* 137-154.

Brehmer, B. Preliminaries to a psychology of inference. (1979). *Scandinavian Journal of Psychology, 20,* 193-210.

Brehmer, B., & Dörner, D. (1993). Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior, 9,* 171–184.

Bremer, B. (1980). In one word: Not from experience. *Acta Psychologica, 45(1-3)*, 223-241.

Broadbent, D (1958). *Perception and Communication.* London: Pergamon Press.

Brown, N.R., & Siegler, R.S. (1993). Metrics and Mapping: A Framework for Understanding Real-World Quantitative Estimation. *Psychological Review, 100(3),* 511-534.
Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A Study of Thinking.* Wiley, New York.
Brunswik, E. (1952). *The conceptual framework of psychology.* Chicago: University of Chicago Press.
Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432– 459.
Busemeyer, J. R., McDaniel, M. A., & Byun, E. (1997). The abstraction of intervening concepts from experience with multiple input-multiple output causal environments. *Cognitive Psychology, 32,* 1-48.
Carroll, J. M., & Kay, D. S. (1988). Prompting, feedback and error correction in the design of a scenario machine. *International Journal of Man-Machine Studies, 28,* 11-27.
Carver, C. S., & Scheier, M, F. (1981). *Attention and self regulation: A control theory to human behavior.* New York: Springer-Verlag.
Castellan, N.J. (1974). The effect of different types of feedback in multiple-cue probability learning. *Organizational Behavior and Human Performance, 11,* 44-64.
CER (2011), Smart Metering Project Phase 1 Updates, Irish Commission for the Energy Regulation, retrieved from: http://www.cer.ie/en/information-centre-reports-and-publications.aspx.
Darby, S. (2010). Smart metering: What potential for house-holder engagement? *Building Research and Information, 38(5),* 442–457.
DeLosh, E.L., Busemeyer, J.R., & McDaniel, M.A. (1997). Extrapolation: The Sine Qua Non for Abstraction in Function Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23(4),* 968-986.
Denrell, J., & Le Mens, G. (2011). Seeking positive experiences can produce illusory correlations. *Cognition, 119,* 313-324.
Deutsch, J.A.; Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review. 70*, 80-90.
Dijksterhuis, A., & Aarts, H. (2010). Goals, Attention, and (Un)Consciousness. *Annual Review of Psychology, 61,* 467-490.
EDRP (2011), Energy Demand Response Project : A final analysis, AECOM, Report commissioned by Ofgem on behalf of DECC, retrived from: http://www.ofgem.gov.uk/Sustainability/EDRP/Pages/EDRP.aspx
Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors, 4,* 59-73.
Einhorn, H.J., Kleinmuntz, D.N., & Kleinmuntz, B. (1979). Linear Regression and Process-Tracing Models of Judgment, *Psychological Review, 86(5),* 465-485.
Engle, R. (2002). Working Memory Capacity as Executive Attention. *Current Directions in Psychological Science, 11,* 19-23.
ESMIG (2011), Empower Demand. The potential of smart meter enabled programs to increase energy and systems efficiency: a mass pilot comparison. Authors : J. Stromback, C. Dromacque & M. H. Yassin (VaasaETT).
Evans, J.B.T. (2008). Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *Annual Review of Psychology, 59,* 255-278.
Evans, L., Maio, G.R., Corner, A., Hodgetts, C.J., Ahmed, S., & Hahn, U. (2013). Self-interest and pro-environmental behaviour. *Nature Climate Change, 3,* 122-125.

Faruqui A, Sergici S, Sharif A. (2010). The impact of informational feedback on energy consumption – a survey of the experimental evidence. *Energy, 35(4)*, 1598-1608.
Fischer, C. (2008). Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency, 1,* 429–434.
Fishbach, A, Dahr, R., & Zhang, Y. (2006). Subgoals as substitutes or complements: The role of goal accessibility. *Journal of Personality and Social Psychology, 91(2),* 232-242.
Fishbach, A., & Labroo, A.A. (2007). Be better or be merry: how mood affects self-control. *Journal of Personality and Psychology, 93(2),* 158-173.
Funke, J. (1991). Solving complex problems: Exploration and control of complex systems. In: Sternberg, R. J., Frensch, P. A. (eds.) *Complex Problem Solving: Principles and Mechanisms*, (pp.185-222). Hillside, NJ: Lawrence Erlbaum Associates.
Guath, M., Millroth, P., Juslin, P., & Elwin, E. (2015). Optimising Electricity Consumption: A case of function learning. *Journal of Experimental Psychology: Applied, 21(4),* 326-341.
Gureckis, T.M., & Love, B.C. (2009a). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition, 113,* 293-313.
Gureckis, T.M., & Love, B.C. (2009b). Learning in noise: Dynamic decision-making in a variable environment. *Journal of Mathematical Psychology, 53,* 180-193.
Hammond, K. R. (1955). Probabilistic functioning and the clinical method. *Psychological Review, 62(4),* 255-262.
Hammond, K. R., & Stewart, T. R. (2001). *The Essential Brunswik: Beginnings, Explications, Applications.* Oxford University Press.
Hammond, K. R., & Summers, D. A. (1972). Cognitive control. *Psychological Review, 79(1),* 58-67.
Hammond, K. R., Summers, D. A., & Deane, D. H. (1973). Negative effects of outcome-feedback in multiple-cue probability learning. *Organizational Behavior and Human Performance, 9,* 30–34.
Hargreaves, T., Nye, M., Burgess, J., 2010. Making energy visible: a qualitative field study of how householders interact with feedback from smart energy monitors. *Energy Policy 38(10),* 6111–6119.
Hoffmann, J.A, von Helversson, B., & Rieskamp, J. (2016). Similar task features shape judgment and categorization processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 42(8),* 1193–1217.
Hoffmann, J. A., von Helversen, B., & Rieskamp, J. (2014). Pillars of judgment: How memory abilities affect performance in rule-based and exemplar-based judgments. *Journal of Experimental Psychology: General, 143(6),* 2242-2261.
Hogarth, R.M., Mukherjee, K., & Soyer, E.(2013). Assessing the Chances of Success: Naïve Statistics Versus Kind Experience. Journal of Experimental Psychology: Learning, Memory, and *Cognition, 39(1),* 14-32.
Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inference as a rational process: The new synthesis. *Annual Review of Psychology, 62,* 135–163.
Hsee, C.K., & Zhang, J. (2010). General Evaluability Theory. *Perspectives on Psychological Science, 5(4),* 343-355.
Hume, D. (1738/1975). *A Treatise of Human Nature*, L.A. Selby-Bigge (ed.), 2nd ed., revised by P.H. Nidditch, Oxford: Clarendon Press.
Juslin, P., Elwin, E., Guath, M., Millroth, P., & Nilsson, H. (2016). Sequential and myopic: On the use of feedback to balance cost and utility in a simulated electricity efficiency task, J*ournal of Cognitive Psychology, 28(1),* 106-128.

Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition, 106,* 259–298.
Juslin, P., Olsson, H., & Olsson, A. C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General, 132(1),* 133-156.
Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: Knowledge partitioning and function learning. *Psychological Review, 111,* 1072–1099.
Karjalainen, S. (2011). Consumer preferences for feedback on household electricity consumption. *Energy and Buildings, 43(2-3),* 458-467.
Karelia, N., & Hogarth, R. M. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin, 134(3),* 404-426.
Karlsson, L., Juslin, P., & Olsson, H. (2007). Adaptive changes between cue abstraction and exemplar memory in a multiple-cue judgment task with continuous cues. *Psychonomic Bulletin and Review, 14,* 1140– 1146.
Karlsson, L., Juslin, P., & Olsson, H. (2008). Exemplar-based inference in multi-attribute decision making: Contingent, not automatic, strategy shifts? *Judgment and Decision Making, 3(3),* 244-260.
KEMA (2009), Cost-benefit Analysis of the introduction of smart meters in the Brussel Capital Region, Version 0.3 (final), KEMA International BV, 6 April 2009, page 34.
Keeney, R. L., & Raiffa, H. (1976). *Decisions with Multiple Objectives.* New York: Wiley.
Kihlstrom, J.F., & Park, L. (2002). *Cognitive Psychology, Overview.* In Encyclopaedia of the Human Brain, Vol. 1. (pp. 839-853).
Klayman, J. (1988). Cue discovery in probabilistic environments: Uncertainty and experimentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 317-330.
Klein, G. (2008). Naturalistic Decision Making. *Human Factors, 50(3)*, 456-460.
Klein, G. A., Calderwood, R., & Clinton-Cirocco, A. (1986). Rapid decision making on the fireground. *Proceedings of the Human Factors and Ergonomics Society 30th Annual Meeting, 1*, 576–580.
Klein, G. A., Wolf, S., Militello, L., & Zsambok, C. (1995). Characteristics of skilled option generation in chess. *Organizational Behavior and Human Decision Processes, 62(1)*, 63–69.
Klopfert, F., & Wallenborn, G. (2011). *Empowering consumers through smart metering.* Report for BEUC, The European Consumer Organization.
Kluger, A.N., & DeNisi, A. (1996). The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory. *Psychological Bulletin, 119(2),* 254-284.
Knox, W.B., Otto, A.R., Stone, P. & Love, B.C. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology, 2.* doi: 10.3389/fpsyg.2011.00398
Koh, K., & Meyer, D. E. (1991). Function learning: Induction of continuous stimulus-response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 811-836
Kopetz, C.E., Kruglanski, A.W., Arens, Z.G., Etkin, J., & Johnson, H.M. (2012). The dynamics of consumer behavior: A goal systemic perspective. *Journal of Consumer Psychology, 22,* 208-223.
Krishnamurti, T., Davis, A.L., Wong-Parodi, G., Wang, J., & Canfield, C. (2013). Creating and in-home display: Experimental evidence and guidelines for design. *Applied Energy, 108,* 448-458.

Kruglanski, A.W., Shah, J.Y., Fishbach, A., Friedman, R., Chun, W.Y., & Sleeth-Keppler, D. (2002). A Theory of Goal Systems. In M. P. Zanna (Ed.) *Advances in Experimental Social Psychology,* (Vol 34, pp. 331-378), San Diego: Academic Press.

Lam, C. F., DeRue, D. S., Karam, E. P., & Hollenbeck, J. R. (2011). The impact of feedback frequency on learning and task performance: Challenging the “more is better” assumption. *Organizational Behavior and Human Decision Processes, 116,* 217–228.

Latham, G. P., & Locke, E. A. (2006). Enhancing the benefits and overcoming the pitfalls of goal setting. *Organizational Dynamics, 35(4),* 332-340.

Lewandowsky, S., Roberts, L., Yang, L-X. (2006). Knowledge partitioning in categorisation: Boundary conditions. *Memory & Cognition, 34(8),* 1676-1688.

Lindenberg, S., & Steg, L. (2007). Normative, Gain and Hedonic Goal Frames Guiding Environmental Behavior. *Journal of Social Issues, 63(1),* 117-137.

Little, J. L., & McDaniel, M. A. (2015). Individual differences in category learning: Memorization versus rule abstraction. *Memory & Cognition, 43(2),* 283–297.

Locke, E. A. (1968). Toward a theory of task motivation and incentives. *Organizational Behavior and Human Performance, 3(2),* 157-189.

Locke, E. A., & Latham, G. P. (1990). *A theory of goal setting and task performance.* Englewood Cliffs, NJ: Prentice Hall.

Locke, E. A., Smith, K. G., Erez, M., Chah, D. O., & Schaffer, A. (1994). The effects of intra-individual goal conflict on performance. *Journal of Management, 20*, 67-91.

Locke, J. (1690/1997). *An Essay Concerning Human Understanding*. R. Woolhouse (ed.). Bungay: Penguin Books.

Lucas, C.G., & Griffiths, T.L. (2010). Learning the Form of Causal Relationships Using Hierarchical Bayesian Models. *Cognitive Science, 34,* 113-147.

Lucas, C.G, Griffiths, T.L., Williams, J.J., & Kalish, M.L. (2015). A rational model of function learning. *Psychonomic Bulletin Review, 22(5)*, 1193-1215.

Lurie, N., & Swaminathan, J. M. (2009). Is timely information always better? The effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes, 108,* 315–329.

McDaniel, M.A., & Busemeyer, J.R. (2005). The conceptual basis of function learning and extrapolation: Comparison of rule-based and associative-based models. *Psychonomic Bulletin and Review, 12(1),* 24-42.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207-238.

Mehlhorn, K., Newell, B.R., Todd, P.M., Lee, M.D., Morgan, K., Braithwaite, V.A., Hausman, D., & Fiedler, K. (2015). Unpacking the Exploration-Exploitation Tradeoff: A Synthesis of Human and Animal Literatures. *Decision, 2(3),* 191-215.

Miller, D.T., & Prentice, D.A. (2016). Chaning Norms to Change Behavior. *Annual Review of Psychology, 67,* 339-361.

Newell, B., Lagnado, D., & Shanks, D. (2007). *Straight choices: The psychology of decision making.* New York, NY: Psychology Press.

Orehek, E., Vazeou-Neuwenhuis, A. (2013). Sequential and Concurrent Strategies of Multiple Goal Pursuit. *Review of General Psychology, 17,* 339-349.

Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology, 65(2)*, 207–240.

Pavlov, I.P. (1927). Conditioned Reflexes: *An Investigation of the Physiological Activity of the Cerebral Cortex* (translated by G.V. Anrep). London: Oxford University Press.

Platzer, C., & Bröder, A. (2013). When the rule is ruled out: Exemplars and rules in decisions from memory. *Journal of Behavioral Decision Making, 26(5),* 429–441.

Powers, W.T. (1973). Feedback: Beyond Behaviorism. *Science, 179(4071)*, 351-356.

Remus, W., O'Connor, M., & Griggs, K. (1996). Does Feedback Improve the Accuracy of Recurrent Judgmental Forecasts? *Organizational Behavior and Human Decision Processes, 66(1),* 22-30.

Rescorla, R.A., & Wagner,A.R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A.H. Black & W.F.Proktsy (Eds.), *Classical conditioning 1/: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Schleich, J., Klobasa, M., Brunner, M., Gölz, S., Götz, K., & Sunderer G. (2011). Smart metering in Germany and Austria – results of providing feedback information in a field trial, *Working Paper Sustainability and Innovation No. S 6/2011, Fraunhofer ISI, Karlsruhe.* Retrieved at: http://www.isi.fraunhofer.de/isi-de/service/presseinfos/2011/pri11-13.php

Shah, J. Y., Friedman, R., & Kruglanski, A. W. (June, 2002). Forgetting all else: On the antecedents and consequences of goal shielding. *Journal of Personality and Social Psychology, 83,* 1261−1280.

Skinner, B.F. (1945). The operational analysis of psychological terms. *Psychological Review, 52(5),* 270-277.

Smart meter in home display - how it works. (2017, December 8). Retrieved from: https://www.uswitch.com/gas-electricity/guides/smart-meter-in-home-display/

Snyman, J.A. (2005). *Practical Mathematical Optimization: An Introduction to Basic Optimisation Theory and Classical and New Gradient-Based Algorithms.* New York: Springer.

Soyer, E., & Hogarth, R. (2015). Fooled by experience. *Harvard Business Review, 93(5),* 72-77.

Stern, P. (2011). Contributions of Psychology to Limiting Climate Change. *American Psychologist, 66(4),* 303-314.

Stern, P. C., & Kirkpatrick, E. M. (1977). Energy behavior: Conservation without coercion. *Environment, 10,* 10–15.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53,* 1–26.

Steyvers, M., Lee, M.D., & Wagenmakers, E-J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology, 53,* 168-179.

Sun, H.S., Frese, M. (2012). Multiple Goal Pursuit in Locke, A. & Latham, G.P. (Ed.), *New Developments in Goal Setting and Task Performance* (pp. 177-194). New York, NY: Routledge, Taylor and Francis Group.

Sutton, R., & Barton, A. (1998). *Reinforcement learning: An introduction.* Cambridge: MIT Press.

Thøgersen, J. (1994). Monetary incentives and environmental concern. Effects of a dif- ferentiated garbage fee. *Journal of Consumer Policy, 17,* 407–442.

Thogersen, J. (2003). Monetary Incentives and Recycling: Behavioural and Psychological Reaction to a Performance-Dependent Garbage Fee. *Journal of Consumer Policy, 26,* 197-228.

Thorndike, E. L. (1927). The law of effect. *American Journal of Psychology, 39,* 212-222.

Todd, P.M., & Hammond, K.R. (1965). Differential effects of feedback in two multiple-cue probability learning tasks. B*ehavioral Science, 10,* 429-435.

Tversky, A., & Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. The Quarterly Journal of Economics, 106(4), 1039-1061.

Unsworth, N., & Engle, R.W. (2008). Speed and Accuracy of Accessing Information in Working Memory: An Individual Differences Investigation of Focus Switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34(3),* 616-630.

Unsworth, N., Heitz, R.P., Schrock, J.C., & Engle, R.W. (2005). An automated version of the operation span task. *Behavior Research Methods, 37,* 498-505.

Vancouver, J. B. (2008). Integrating self-regulation theories of work motivation into a dynamic process theory. *Human Resource Management Review, 18,* 1–18.

Vancouver, J.B., Weinhardth, J.M., & Schmidt, A.M. (2010). A Formal, Computational Theory of Multiple-Goal Pursuit: Integrating Goal-Choice and Goal-Striving Processes. *Journal of Applied Psychology, 95(6),* 985-1008.

von Helversen, B., & Rieskamp, J. (2009). Models of quantitative estimations: Rule-based and exemplar-based processes compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35(4),* 867-889.

von Helversen, B., Mata, R., & Olsson, H. (2013). Do children profit from looking beyond looks? From similarity-based to cue abstraction processes in multiple-cue judgment. *Developmental Psychology, 46(1),* 220–229.

Wiener, N. (1948a), Time, Communication, and the Nervous System. *Annals of the New York Academy of Sciences, 50,* 197–220.

Wiener, N. (1948b). *Cybernetics: Control and Communication in the Animal and the Machine.* Wiley, New York.

You, G. (1989). *Disclosing the decision-maker's internal model and control policy in a dynamic decision task using a system control paradigm.* MA thesis, Purdue University.

Zhang, Y., Fishbach, A., & Kruglanski, A. W. (2007). The dilution model: How additional goals undermine the perceived instrumentality of shared path. *Journal of Personality and Social Psychology, 92(3),* 389−401.

Appendix A: Table A1

Table 1. Parameters of the Utility and Cost Functions for each of the 18 Electricity Consuming Appliance in the Simulated Household. The Relative Weight (w) refers to the Importance of the Appliance for the Overall Utility of the Simulated Person, the Exponent (α) to the Nonlinear Rate of Change in of the Utility Function for the Appliance, and Ceiling (r) to the Maximum Utility for the Appliance.

Electricity appliance	Relative weight w	Exponent α	Ceiling r	Crowns/ Unit c
Indoor temperature	0.31	0.1	24	2 / degree
Lighting Hall	0.02	0.1	8	.03 / hour
Lighting WC	0.05	0.01	2	.03 / hour
Lighting Kitchen	0.05	0.15	8	.04 / hour
Lighting Bedroom	0.02	0.1	8	.04 / hour
Lighting Living room	0.05	0.15	8	.05 / hour
Hot water: WC sink	0.04	0.25	0.33	35 / hour
Hot water: shower	0.12	0.5	0.33	35 / hour
Hot water: kitchen sink	0.04	0.25	0.5	35 / hour
Micro Oven	0.02	0.25	0.5	1.3 / hour
Oven	0.02	0.1	2	2.6 / hour
Stove	0.03	0.1	2	1.3 / hour
Dishwasher	0.08	0.1	7	2.6 / hour
Washing machine	0.08	0.1	7	2.6 / hour
Tumble drier	0.01	0.1	7	2.6 / hour
PC	0.02	0.1	8	.13 / hour
DVD	0.01	0.1	8	.01/ hour
Television	0.03	0.5	8	.32 / hour

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 155

Editor: The Dean of the Faculty of Social Sciences

A doctoral dissertation from the Faculty of Social Sciences, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences".)

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2018

Distribution: publications.uu.se
urn:nbn:se:uu:diva-348821