

Choice Architecture for Human-Computer Interaction

Anthony Jameson

German Research Center for Artificial Intelligence, Germany
jameson@dfki.de

Bettina Berendt

KU Leuven, Belgium

Bettina.Berendt@cs.kuleuven.be

Silvia Gabrielli

CREATE-NET, Italy

silvia.gabrielli@create-net.org

Federica Cena

University of Torino, Italy

cena@di.unito.it

Cristina Gena

University of Torino, Italy

cgena@di.unito.it

Fabiana Vernerio

University of Torino, Italy

verneriof@di.unito.it

Katharina Reinecke

University of Michigan, U.S.A.

reinecke@umich.edu

Contents

1	Introduction	3
1.1	What Is Choice Architecture for HCI?	3
1.2	Hasn't It Already Been Done?	5
1.3	Preview of the Rest of This Publication	11
2	Types of Preferential Choice in HCI	13
2.1	Macro- vs. Micro-Level Choices	13
2.2	Generic Choice Problems	14
2.3	Preview of Sections on Content-Specific Types of Choice	19
3	Choice Patterns: The ASPECT Model	21
3.1	The Need for a Comprehensive View of Human Choice	21
3.2	Introduction to the ASPECT Model	23
3.3	Preview of the ASPECT Choice Patterns	25
3.4	Relationship to Two Modes of Processing	31
3.5	Ecological Rationality	32
3.6	What Constitutes a Good Decision for Choosers?	34
4	Choice Support Strategies: The ARCADE Model	39
4.1	<i>Access Information and Experience</i>	40
4.2	<i>Represent the Choice Situation</i>	44
4.3	<i>Combine and Compute</i>	46

4.4	<i>Advise About Processing</i>	47
4.5	<i>Design the Domain</i>	49
4.6	<i>Evaluate on Behalf of the Chooser</i>	50
4.7	Alternative Goals in Applying the ARCADE Strategies	52
5	Attribute-Based Choice	59
5.1	Introduction to the Pattern	59
5.2	Thinking in Advance About Evaluation Criteria	61
5.3	Winnowing	62
5.4	Choosing From a Manageable Set of Options	65
6	Consequence-Based Choice	71
6.1	Introduction to the Pattern	71
6.2	Recognizing That There Is a Choice Opportunity	75
6.3	Situation Assessment	77
6.4	Deciding When to Choose	79
6.5	Identification of Options	81
6.6	Anticipation of Consequences	83
6.7	Evaluation of Anticipated Consequences	88
6.8	Time Discounting	91
6.9	Dealing With Uncertainty	94
7	Experience-Based Choice	97
7.1	Introduction to the Pattern	97
7.2	Recognition-Primed Decision Making	99
7.3	Habit-Based Choice	103
7.4	Choice Based on Instrumental Conditioning	106
7.5	Affect-Based Choice	110
8	Socially Based Choice	113
8.1	Introduction to the Pattern	113
8.2	Overview of Forms of Social Influence	114
8.3	Social Examples	116
8.4	Social Expectations	120
8.5	Explicit Advice	122

9	Policy-Based Choice	127
9.1	Introduction to the Pattern	127
9.2	Research on Time Bracketing	127
9.3	Dimensions of Variation Among Policies	131
9.4	Support for the Generation of Possible Policies	132
9.5	Support for the Evaluation of Possible Policies	134
9.6	Support for the Execution of a Policy	134
10	Trial-and-Error-Based Choice	139
10.1	Introduction to the Pattern	139
10.2	Research on Exploration Strategies	144
10.3	Support for Exploration	147
10.4	Research on Learning From Feedback	149
10.5	Combating Typical Problems With Feedback	151
11	Choice in Online Communities	157
11.1	Introduction	157
11.2	Choices About Whether to Participate	159
11.3	Consequence-Based Choices in On-Line Communities	162
11.4	Socially Based Choices in On-Line Communities	165
11.5	Policy-Based Choices in On-Line Communities	169
11.6	Trial-and-Error-Based Choices in On-Line Communities	171
11.7	Concluding Remarks on Choice in Online Communities	175
12	Choices Concerning Privacy	177
12.1	Introduction	177
12.2	Consequence-Based Choices About Privacy	184
12.3	Attribute-Based Choices About Privacy	193
12.4	Socially Based Choices About Privacy	196
12.5	Policy-Based Choices About Privacy	197
12.6	Trial-and-Error-Based Choices About Privacy	205
12.7	Concluding Remarks on Privacy-Related Choices	210
13	Concluding Remarks	211
13.1	More Focused Analyses	211
13.2	Extension to Decision Making by Groups	212

13.3 Application to Other Types of Choice in HCI 213
13.4 Shouldering Responsibility for the Future of Human Choice 213

Acknowledgments 215

References 217

Abstract

People in human-computer interaction have learned a great deal about how to persuade and influence users of computing technology. They have much less well-founded knowledge about how to help users choose for themselves. It's time to correct this imbalance. A first step is to organize the vast amount of relevant knowledge that has been built up in psychology and related fields in terms of two comprehensive but easy-to-remember models: The ASPECT model answers the question "How do people make choices?" by describing six *choice patterns* that choosers apply alternately or in combination, based on Attributes, Social influence, Policies, Experience, Consequences, and Trial and error. The ARCADE model answers the question "How can we help people make better choices?" by describing six general high-level *strategies for supporting choice*: Access information and experience, Represent the choice situation, Combine and compute, Advise about processing, Design the domain, and Evaluate on behalf of the chooser. These strategies can be implemented with straightforward interaction design, but for each one there are also specifically relevant technologies. Combining these two models, we can understand virtually all existing and possible approaches to choice support as the application of one or more of the ARCADE strategies to one or more of the ASPECT choice patterns.

After introducing the idea of choice architecture for human-computer interaction and the key ideas of the ASPECT and ARCADE models, we discuss each of the ASPECT patterns in detail and show how the high-level ARCADE strategies can be applied to it to yield specific tactics. We then apply the two models in the domains of online communities and privacy. Most of our examples concern choices *about the use of* computing technology, but the models are equally applicable to everyday choices made *with the help of* computing technology.

1

Introduction

1.1 What Is Choice Architecture for HCI?

If you work in human-computer interaction, you are probably a *choice architect*—even if you have been as unaware of that role as Molière’s “bourgeois gentleman” was of having spoken prose all his life.

As Thaler and Sunstein [2008] wrote when introducing the term: “A choice architect has the responsibility for organizing the context in which people make decisions” (p. 3). And users of today’s ever-present computing technology are constantly making small choices and large decisions:

1. Sometimes, the main purpose of an interactive system is to help people make a particular type of choice: Think of e-commerce websites and of apps for helping people choose healthy food.

2. Even if the main purpose is different—as with a navigation system that helps you follow a route from one place to another—the user often has choices to make about details—such as which of the several proposed routes to follow. Helping people make these “microchoices” (2.1) better is one (often not obvious) way of enhancing the user experience.

3. Finally, just about any interactive system, regardless of its purpose, requires its users to make some choices about how to operate the system: Which of these two text entry methods should I use to enter text right now? How might I configure this application so as to make it more convenient to use? And might I be better off using some other application instead of this one?

In all of these cases, the fact that the choice is “up to the user” does not release the designers from their responsibility as choice architects to “organize the context” so that users can easily make choices that they will ultimately find satisfactory. But fulfilling this responsibility is easier said than done, if we want to go beyond reliance on designer intuition and familiar design patterns. Good choice architecture for human-computer interaction (HCI) must ultimately be based on a solid understanding of two complex topics:

- The psychology of choice and decision making: How do people go about making choices in their everyday lives, with or without computing technology?
- Strategies and technologies for supporting everyday choice: What are the general ways in which it’s possible to help people make better choices; and how can these be applied in the context of—and with the help of—today’s interactive computing technology?

This publication aims to equip readers with a coherent understanding of both of these topics, along with an ability to pursue them in more depth by following up on the references. Figure 1.1 gives a preview of the two complementary models that we call the ASPECT and the ARCADE models after their two acronyms: The letters in ASPECT stand for the six *choice patterns* that we introduce to cover the phenomena of everyday choice and decision making. The letters in ARCADE stand for the six high-level *choice support strategies* that we have distilled from previous research and practice.

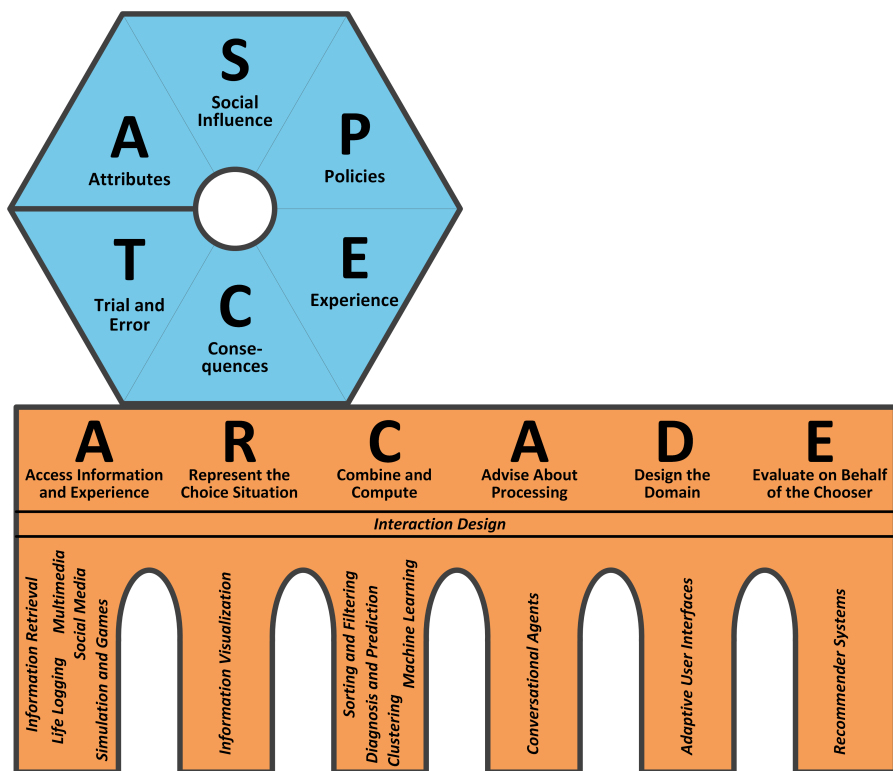


Figure 1.1: High-level overview of the ASPECT and ARCADE models of choice patterns and choice support strategies.

1.2 Hasn't It Already Been Done?

The idea of combining psychology and computing technology to help people make better choices is not new. So why does the HCI field need a new conception of choice architecture? We will explain by first introducing two general conceptual distinctions and then considering in turn several related lines of research and practice.

1.2.1 Preferential vs. Nonpreferential Choice

Many of the “choices” that have received the most attention in the HCI field are *nonpreferential* choices: A user wants to choose the steps (e.g., clicks on particular icons) that are required to achieve a particular goal, such as turning on change-tracking mode in his¹ word processing application. With nonpreferential choices, the question is not what the chooser *prefers* to do but rather what she *has* to do if she wants to achieve a particular goal.

With *preferential choice*—for example, “Shall I turn on change tracking or simply use the commenting functionality to recommend changes to my coauthors?”—a user can prefer one option over another one even though neither one is objectively right or wrong. A preferential choice can be influenced by factors such as the value that the chooser assigns to particular anticipated consequences, the policies the chooser wants to follow, and social expectations that the chooser wants to conform to—a multifaceted set of considerations that will be discussed in connection with the six ASPECT choice patterns.

1.2.2 Persuasion vs. Choice Support

It is also worthwhile to distinguish between two goals that a choice architect can have when attempting to influence a person’s choices: *persuasion* versus *choice support*. It is true that neither of these concepts is easy to define crisply and that there are multiple equally reasonable alternative definitions for each concept. Still, there is an important high-level difference between them:

- We will use the term *persuasion* when the goal of the choice architect is to increase the likelihood that the chooser will choose a particular option (e.g., fruit salad instead of cake); or choose an option from some particular class (e.g., fruits and vegetables); or adopt a particular goal (e.g., eat in a more health-conscious way).

¹To avoid clumsy formulations like “him or her” when using personal pronouns in a generic way, we will alternate between the masculine and feminine forms on an example-by-example basis.

- One possible definition of *choice support* runs as follows: The goal is to help the chooser make the choice in such a way that, from some relevant perspective, the chooser will be satisfied with the choice. One candidate for a “relevant perspective” is: “after learning about the consequences of the choice and taking the time to reflect on all important aspects of it”. But other definitions can be argued for. In fact, a first step toward getting better at supporting choice is to understand better what constitutes a “good choice” from the point of view of the chooser (see the discussion in 3.6 below).

These two goals of persuasion and choice support can be pursued simultaneously in various ways. Sometimes, persuasion is used even when the top-level goal is that of choice support. A doctor who tries to persuade a patient to stop smoking presumably believes that the patient will ultimately approve of this choice from some relevant perspective. And in fact maybe the patient has arrived at this conclusion himself and begged the doctor to “persuade” him to perform the specific actions required to stop smoking.

Conversely, even if your top-level goal is to induce a chooser C to choose a particular option O that is in your own interest—for example, the option of buying your software application—adopting choice support as a subgoal can be a good strategy, for either of two reasons:

- You are convinced that C , given high-quality, unbiased choice support, will conclude for herself that O is her best option.
- There are various specific ways of executing O (e.g., various ways of using your software application); and you think that by helping C to choose the specific ways that are best for her, you will increase the likelihood that she will find it attractive to execute O .

Because of these and other interrelationships, techniques for persuasion and choice support can be compared to the black and white keys on the piano (Jameson, 2013): There are some tunes that you can play on just the black keys or on just the white keys; but if you know

how to use all of the keys together, your range of possibilities is vastly increased.

1.2.3 Thaler and Sunstein's Conception of Choice Architecture

Thaler and Sunstein [2008], who coined the term *choice architecture*, present a synthesis of psychological research (chaps. 1–4) that overlaps at many points with the synthesis in our newer ASPECT model, along with six “principles of good choice architecture” (chap. 5), captured with the acronym NUDGES, which suggest how to help people make better choices in everyday life. The remaining 13 chapters of this stimulating and influential book discuss in detail how their principles can be applied in a variety of areas of life, such as personal finance and health.

The relevance of this work for the HCI field is somewhat limited by the fact that Thaler and Sunstein do not devote particular attention to computing technology, either as a means for supporting everyday choice or as a domain in which choices need to be made. Also, as is understandable for a best-selling book, the synthesis of psychological research and the NUDGES principles do not have the clearly articulated structure and explicit grounding in previous literature that is required in a solid foundation for HCI researchers and practitioners. Work that has built on Thaler and Sunstein's conception (e.g., Johnson et al. [2012]) has begun in both of these respects to make the idea of choice architecture more relevant to HCI, but there are still many gaps for the present work to fill.

It is instructive to relate the concept of a *nudge*, which lies at the center of Thaler and Sunstein's conception of choice architecture, to the two conceptual distinctions just introduced above. On close inspection, we can see that the term *nudge* has several different meanings even in these authors' own book:

1. It often refers to a mild form of persuasion intended to bias a person's choice in the direction of a particular option while still being largely compatible with the goal of choice support in that the suggested option seems to be at least reasonably good for the

chooser and in any case the chooser is not compelled to choose it.² One of the types of nudge that they suggest (see, e.g., chaps. 5, 6, and 11)—the careful design of default options (cf. 6.2 below)—clearly illustrates this interpretation of the concept of a nudge.

2. Other forms of nudge that they propose—such as structuring complex choices, giving informative feedback, and helping people to “map” information onto concepts that are meaningful for them—can be useful approaches to supporting preferential choice that do not necessarily involve bias toward any particular option. We will be discussing these forms of choice support (along with many others) at many points in the present publication, relating them to the ASPECT and ARCADE models.
3. Finally, several of the forms of nudge can be seen as approaches to supporting *nonpreferential* choice. Under the category “Expect error”, the authors present ideas, which will look familiar to readers from the HCI field, about how to help people to avoid doing the objectively wrong thing (e.g., forgetting to attach a document to an email message). Their examples of the nudges in the previous category likewise sometimes concern nonpreferential choice.

The existence of these very different meanings limits the usefulness of the term *nudge* as a way of communicating about tactics for choice support and persuasion. In particular, we may be inclined to agree readily that “people could use a nudge” when we think of the broad meaning that includes any sort of intervention to support or influence choices; but when doing so we can be interpreted as having accepted, in the narrow meaning of the term, a vision of a world in which people’s choice processes are constantly being intentionally biased in subtle ways, often without their awareness.

²They write: A nudge is “any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives” (p. 6).

1.2.4 Persuasive Technology

When HCI people hear the idea of “helping people make better choices”, they often think of persuasive technology: a line of research and practice which was introduced mainly by B. J. Fogg (2003) and which has since become widely represented both in the research literature and in practical systems and interface design methodologies. Like the present publication, Fogg’s seminal book systematically combines research from psychology with a framework for making use of the research results in interactive computing technology. Many others have expanded and fleshed out Fogg’s framework, and persuasive technology constitutes an important part of a choice architecture for HCI.

A limitation is that persuasive technology focuses squarely on persuasion, as opposed to choice support, as a way of influencing people’s choices. It therefore does not provide direct guidance to choice architects who are pursuing the goal of choice support. For this purpose, we need to exploit and organize (in the ASPECT and ARCADE models) a vast amount of literature on choice and choice support that is seldom taken into account in the persuasive technology area.

Paradoxically, our inclusion of concepts and research results that are not oriented toward persuasion may well provide new ideas even to readers who are interested exclusively in persuasion. The reason is that just about every tactic that is designed with the goal of supporting choice can also be (mis)applied in an intentionally biased way (4.7.1). In the present work, we will focus our attention almost entirely on choice support efforts that are not characterized by intentional bias; readers more interested in persuasion will find it easy enough to work out biased versions of any new ideas that they acquire here.

1.2.5 Recommender Systems

A major computing paradigm that can be seen as supporting everyday nonpreferential choice is that of *recommender systems* (see, e.g., Jannach et al., 2011; Ricci et al., 2011). These systems aim to support and influence users’ choices concerning products to buy, documents to read, and a variety of other types of item. As we will see in Section 4,

recommender systems essentially implement one of the six ARCADE strategies for choice support, *Evaluate on Behalf of the Chooser*: They typically apply any of a variety of algorithms to predict how satisfied a given chooser would be with particular options. In some cases, an algorithm of this sort can be seen as realizing a variant of one of the six ASPECT choice patterns. For example, some variants of the popular paradigm of *collaborative filtering* (see, e.g., Ekstrand et al., 2011) can be seen as automating a variant of the socially based choice pattern (3.3.4; Section 8), since they make use of information about choices or evaluations made by people who are similar to the current chooser.

1.2.6 Other Contributing Technologies

There are a number of other areas of computer science which, like persuasive technology and recommender systems, contribute techniques that can be used as part of a choice architecture. A number of these are discussed in Section 4 in connection with the ARCADE strategies, which help to explain how they fit into the picture.

1.3 Preview of the Rest of This Publication

Section 2 introduces the several types of choice problem that will yield most of the examples for the present publication. Section 3 offers a compact but broad overview of how people make everyday choices, introducing the ASPECT model. Section 4 introduces the other major part of our conceptual framework, the six high-level ARCADE strategies, giving initial examples of their application and discussing the most important technologies that can be used to realize these strategies. Each of the subsequent six major sections looks at one of the ASPECT choice patterns in more depth, summarizing key ideas from psychological research and discussing how the ARCADE strategies can be applied to support choosing according to the pattern. In the final two main sections, we illustrate how the ASPECT and ARCADE models can help to enhance understanding of choice processes in two important contexts: online communities and privacy, respectively. The final brief section

lists several directions in which the foundation laid in this work can be extended in future work.

2

Types of Preferential Choice in HCI

As was mentioned above, we will focus in this work on preferential choices that users make about the use of computing technology, leaving the extension to other types of choice mainly as an exercise for the interested reader. In this brief section, we introduce a general conceptual distinction and three generic, domain-independent types of choice that are frequently encountered in connection with computing technology use.

2.1 Macro- vs. Micro-Level Choices

It will be useful to bear in mind a distinction introduced by Fogg [2003] in his seminal book on persuasive technology: He distinguished between *macro-level* persuasion, which is done by systems that are specifically designed to persuade people to do things—for example, an application or website designed to encourage more efficient use of energy; and (b) *micro-level* persuasion, which can occur at various points in any type of interactive system, when the designer intends to persuade the user to perform a particular action or adopt a particular attitude (e.g., the

user may be persuaded to back up his work at regular intervals, where doing so is just a small part of the overall operation of the application).

Analogously, we can distinguish between systems that are designed primarily to support choice and systems that apply choice support tactics at various points where a nontrivial preferential choice needs to be made. Whereas systems of the former type are often based explicitly on particular approaches to choice support, in the latter case the choice support will often be the result of a decision by the interaction designer that is not based on any explicit principles, though it may be informed by previous practice and experience of the designer.

2.2 Generic Choice Problems

Though opportunities to make preferential choices crop up constantly with just about every type of interactive system, there are three generic classes of choice that are worth distinguishing, because of their frequency of occurrence and because they have attracted a fair amount of attention in HCI research. Table 2.1 introduces them to facilitate reference to them at various points later in this publication.

2.2.1 Choice of a System to Use

One common type of choice is that of whether to use a given system at all, or which of a set of possible systems to use. The most extensive line of research that has looked into this question is research on *technology acceptance*. A good entry point to this literature is the influential article by Venkatesh et al. [2003], which presented the Unified Theory of Acceptance and Use of Technology (UTAUT), a model that integrates eight previously developed models, including the especially widely studied *Technology Acceptance Model* (TAM; see, e.g., Venkatesh and Davis, 2000). These models in turn drew their inspiration from more general theories from social psychology and sociology, such as the precursors of the Reasoned Action Approach of Fishbein and Ajzen [2010].

Table 2.2 gives an impression of the basic nature of the models in this area by depicting the four main variables in the UTAUT model

Table 2.1: Three general types of preferential choice that have been studied in HCI.

Generic Choice Problem	Selected Research Issues
Decision about whether to use a given system	What variables influence people's decisions about whether to use a given system if it is made available to them (usually: within an organization)? What are the causal relationships among these variables? How can these variables be measured?
Choice of a method from a set of alternative methods	When more than one method is available for a particular subtask, how do users decide which one to use? Why do even experienced users sometimes persist in using inefficient methods?
Configuration decision	How do people decide whether and when to configure an application? What difficulties do they encounter when making configuration choices?

that are thought to influence a person's intention to use a given system and her actual use of the system, along with examples of questionnaire items typical of those used to measure these variables. The model also includes claims about several variables that moderate the influence of these main variables: *gender*, *age*, *experience*, and *voluntariness of use*.

Though some of these questions are reminiscent of questions from usability scales such as SUS (System Usability Scale, Brooke, 1996), the overall goal of the model and the associated measuring instruments is not to assess usability but rather to predict whether potential users (typically, employees in a given company) will actually use a given system (e.g., a new videoconferencing system) if it is made available to them. Note that most of the questions related to the variables *Social Influence* and *Facilitating Conditions* concern considerations other than usability.

Researchers and practitioners in the HCI field usually want to go beyond *predicting* whether people in a given target group will use a given (type of) system, to attempt to *improve* the system (and/or related resources) to increase the likelihood that the system will be used

Table 2.2: The four main variables in the UTAUT model and typical questionnaire items used to measure them. (Based on parts of Figure 3 and Tables 9–12 of Venkatesh et al., 2003.)

Performance Expectancy

Using the system in my job would enable me to accomplish tasks more quickly.

Using the system would improve my job performance.

Using the system would make it easier to do my job.

...

Effort Expectancy

Learning to operate the system would be easy for me.

My interaction with the system would be clear and understandable.

I would find the system to be flexible to interact with.

...

Social Influence

People who influence my behavior think that I should use the system.

People who are important to me think that I should use the system.

...

Facilitating Conditions

I have control over using the system.

I have the resources necessary to use the system.

I have the knowledge necessary to use the system.

The system is not compatible with other systems I use.

...

and the success of its use. Still, the large amount of information collected in the technology acceptance literature about variables related to choices about system use and about ways of measuring these variables can help to stimulate and structure thinking about this class of choices. Researchers in this area regularly introduce new variables and new perspectives that shed light on different aspects of acceptance decisions (see, e.g., Bagozzi, 2007; Loraas and Diaz, 2009).

2.2.2 Choice of a Method

In all but the simplest interactive systems, there is often more than one method available for achieving a given goal. Whenever the user can choose freely between two or more methods, the choice is preferential. Card, Moran, and Newell introduced in their GOMS model (described most completely in Card et al., 1983; see also Kieras, 2008) a notation for such cases: The two or more available methods are described as part of the model for a given task, and it is assumed that each user has learned a *selection rule* for making the choice (e.g., “Use the mouse instead of the cursor keys if the target is more than a couple of inches away on the screen”); this assumption is plausible given that the GOMS model assumes that users have considerable experience with the system and the tasks in question.

In the intervening years, some researchers have looked at the ways in which users learn selection rules on the basis of experience with the methods in question (see, e.g., Gray and Boehm-Davis, 2000) and at the considerations that users take into account when choosing among methods (see, e.g., Young and MacLean, 1988; Jameson and Klöckner, 2005). Other researchers have investigated situations in which users systematically fail to use suitable methods that are available to them (see, e.g., Carroll and Rosson, 1987; Bhavnani and John, 2000; Bhavnani et al., 2008; Charman and Howes, 2003).

2.2.3 Configuration Choices

A usually more complicated type of choice that users can make concerns whether, when, and how to configure an application to suit their own tastes and needs. Over the years, researchers have repeatedly found this type of problem to be challenging for most users (see, e.g., Mackay, 1991; McGrenere et al., 2007). It has attracted increased attention in recent years because of the practically important problem of configuring privacy settings in social networking sites and ubiquitous computing systems (see Section 12).

One way to see why configuration choices are so challenging is to consider an example of a relatively simple configuration problem. As

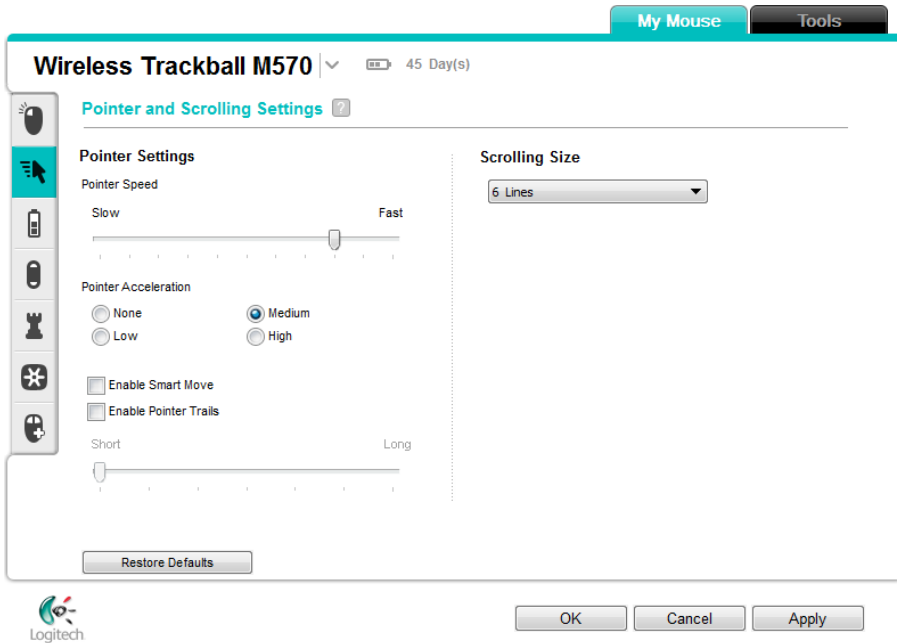


Figure 2.1: Configuration screen for a Logitech trackball.

is illustrated by the screenshot in Figure 2.1, with most computers it is possible to configure the relationship between the movement of the mouse and the movement of the cursor in terms of at least two parameters: speed and acceleration. In the example configuration screen, that of a widely used brand of mouse, there are 4 possible accelerations and 11 possible speeds, for a total of 44 combinations. So in principle the user should try to figure out which of these 44 combinations works best for him. Even if trying out one combination took only a few seconds, presumably few users would have the patience to try out even a large sample: Among other things, the user would have to keep track of which configuration has worked best so far and how well it has worked, so that he can determine whether the next configuration is better. Worse yet, it would actually take much more than a few seconds to determine reliably how well a given parameter combination would

work for a given user: The user would have to try the configuration out on a representative sample of tasks which require different types of mouse movement, and he would have to take into account the fact that his performance with a given configuration is likely to improve over time.

2.3 Preview of Sections on Content-Specific Types of Choice

In addition to these three generic types of choice, we will discuss in separate main sections of this publication (Sections 11 and 12) and in occasional examples in the sections before them two types of choice defined in terms of their content: choices concerning contribution to online communities and choices about privacy, respectively. Readers who would like at this point to get a broader view of the types of choices that arise in connection with computing technology may want to read the introductory subsections (11.1 and 12.1) of these sections before turning in the next two main sections to the general ASPECT and ARCADE models.

3

Choice Patterns: The ASPECT Model

3.1 The Need for a Comprehensive View of Human Choice

Psychological research about how people make preferential choices has received limited attention in HCI so far.¹ One reason may be the fact that there is no single relevant theory in psychology that could be straightforwardly adapted to the needs of the HCI field. Though dozens of books and hundreds of articles from relevant psychological research exist, they come from several research traditions that only partly overlap with and refer to each other. Our discussion will draw from these areas, among others: judgment and decision making (see, e.g., Hastie and Dawes, 2010; Koehler and Harvey, 2004; Lichtenstein and Slovic, 2006; Schneider and Shanteau, 2003; Newell et al., 2007; Weber and Johnson, 2009), naturalistic decision making (Klein, 1998), the Reasoned Action ap-

¹Two thorough book-length syntheses of cognitive psychology research for HCI (Gardiner and Christie, 1987; Johnson, 2010) include hardly any references to the sort of psychology literature cited in the present publication. A recent exception is the book of Payne and Howes [2013], which proposes “a utility-maximization approach to understanding human interaction with technology” and accordingly cites psychological works that are relevant to this perspective.

Table 3.1: Overview of the choice patterns in the ASPECT model. (*C* = the chooser.)**Attribute–Based Choice***Conditions of Applicability*

- The options can be viewed meaningfully as items that can be described in terms of attributes and levels
- The (relative) desirability of an item can be estimated in terms of evaluations of its levels of various attributes

Typical Procedure

- (Optional:) *C* reflects in advance about the situation–specific (relative) importance of attributes and/or values of attribute levels
- *C* reduces the total set of options to a smaller *consideration set* on the basis of attribute information
- *C* chooses from among a manageable set of options

Consequence–Based Choice*Conditions of Applicability*

- The choices are among actions that will have consequences

Typical Procedure

- *C* recognizes that a choice about a possible action can (or must) be made
- *C* assesses the situation
- *C* decides when and where to make the choice
- *C* identifies one or more possible actions (options)
- *C* anticipates (some of) the consequences of executing the options
- *C* evaluates (some of) the anticipated consequences
- *C* chooses an option that rates (relatively) well in terms of its consequences

Experience–Based Choice*Conditions of Applicability*

- *C* has made similar choices in the past

Typical Procedure

- *C* applies recognition–primed decision making
- or *C* acts on the basis of a habit
- or *C* chooses a previously reinforced response
- or *C* applies the affect heuristic

Socially Based Choice*Conditions of Applicability*

- There is some information available about what relevant other people do, expect, or recommend in this or similar situations

Typical Procedure

- *C* considers *examples* of the choices or evaluations of other persons
- or *C* considers the *expectations* of relevant people
- or *C* considers explicit advice concerning the options

Policy–Based Choice*Conditions of Applicability*

- *C* encounters choices like this one on a regular basis

Typical Procedure

- [Earlier:] *C* arrives at a policy for dealing with this type of choice
- [Now:] *C* recognizes which policy is applicable to the current choice situation and applies it to identify the preferred option
- *C* determines whether actually to execute the option implied by the policy

Trial–and–Error–Based Choice*Conditions of Applicability*

- The choice will be made repeatedly; or *C* will have a chance to switch from one option to another even after having started to execute the first option

Typical Procedure

- *C* selects an option *O* to try out, either using one of the other choice patterns or (maybe implicitly) by applying an *exploration strategy*
- *C* executes the selected option *O*
- *C* notices some of the consequences of executing *O*
- *C* learns something from these consequences
- (If *C* is not yet satisfied:) *C* returns to the selection step, taking into account what has been learned

proach (Fishbein and Ajzen, 2010); and research on fast and frugal heuristics (Gigerenzer and Todd, 1999), habitual behavior (Wood and Neal, 2007), behavioral economics (Ariely, 2008; Iyengar, 2010; Thaler and Sunstein, 2008), self-control (Rachlin, 2000), compliance tactics (Cialdini, 2007), and the neuropsychological bases of choice (Glimcher and Fehr, 2014).

3.2 Introduction to the ASPECT Model

In an effort provide a compact synthesis that will be useful for researchers and practitioners in HCI, we introduce in this work a novel way of conceptualizing choice and decision making processes: We define six basic *choice patterns*, of which Table 3.1 gives an overview. Each of these patterns is sometimes found in its pure form, especially in research, which tends understandably to focus on relatively pure variants of the phenomena under investigation, often creating this purity through experimental manipulations. In real life, these patterns sometimes likewise appear in pure form, but often they are mixed together in various ways.

An advantage of distilling out these six basic patterns is that it becomes more straightforward to think about how to support choice when it occurs according to each pattern. As can be seen in Table 3.1, each pattern comprises a set of typical processing steps that are mostly different from the steps found in the other patterns.²

As can be seen in Figure 3.1, the first letters of these six patterns make up the word ASPECT—which is fortuitous, since each pattern can be viewed as an “aspect” of choice, in the (relatively infrequent) sense of the word referring to a “way in which a problem, idea, etc. may be considered: ‘to consider a problem from every aspect’”.³

Each of these basic patterns will be discussed in detail in one of the Sections 5 through 10. We will summarize relevant psychological research on the pattern and discuss ways of supporting choice processes that fit the pattern.

The inclusion of these six patterns makes our summary of choice and decision making research broader (though necessarily less deep) than all

²Readers who are already familiar with other ways of conceptualizing psychological research on choice are advised to view the ASPECT model as they would view a new textbook on the psychology of choice that organizes the material in a novel way, not as a new theory that aims to compete with existing theories. The criterion for judging the value of this type of exposition is the extent to which it helps researchers and practitioners in the HCI field—and maybe in other areas as well—to think realistically and productively about the choices that are made by users of computing technology.

³Collins English Dictionary, Complete Unabridged 11th Edition, 2011.

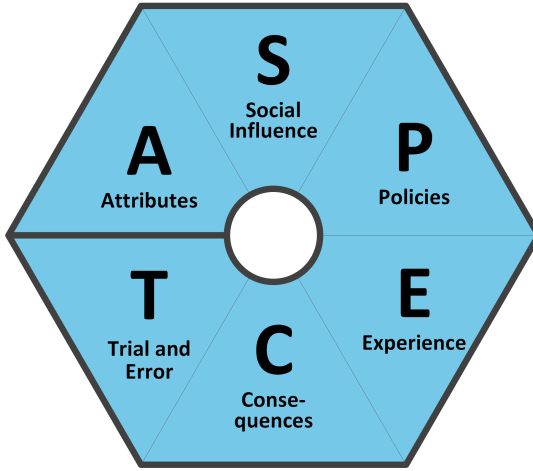


Figure 3.1: Graphic intended to aid recall of the 6 choice patterns distinguished in the ASPECT model.

accounts known to us that are currently available in the literature, even in textbooks such as those of Hastie and Dawes [2010], Newell et al. [2007], and Jungermann et al. [2010]. It is understandable and desirable that when a researcher in one of the relevant areas writes even what is purported to be a fairly comprehensive account of work on choice and decision making, she tends to devote most attention to the topics to which she herself has made important contributions, leaving other topics in the background. For the purpose of serving the HCI field, it is important to be as balanced and comprehensive as possible. For example, if a researcher or practitioner took into account only 2 of the 6 choice patterns shown in Table 3.1, he would have to try to understand and support processing that occurs in accordance with the other 4 patterns in ways that in fact make sense only for the 2 patterns that he was aware of.

If the reader is wondering “How do we know that these six choice patterns exist and no others?”, our answer is this: These six patterns have proven useful and adequate in several years of studying the research literature and applying its concepts and results to examples of

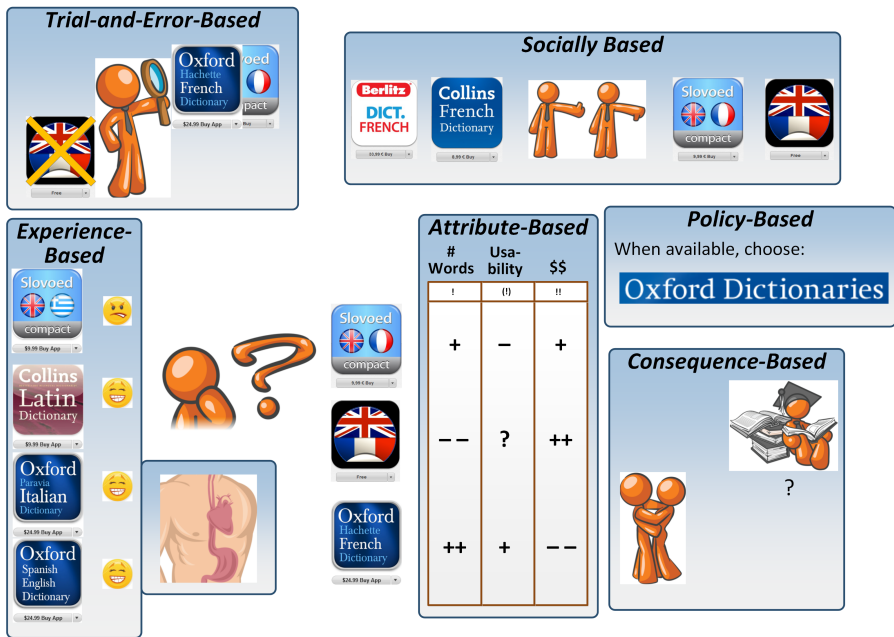


Figure 3.2: Illustration of the six ASPECT choice patterns. (Discussion in text.)

choice processes both within and outside of the HCI field. It would surely be possible to formulate an alternative model with 5 or 7 choice patterns by defining the patterns differently; but it seems unlikely that this alternative model would do a much better job of covering the relevant phenomena.

3.3 Preview of the ASPECT Choice Patterns

To give a more concrete idea of the patterns and the relationships among them, we will discuss how each pattern could be applied in the following situation: An English-speaking tourist is about to visit France, so he would like to buy a French-English dictionary for his iPhone. The Apple App Store offers a number of such dictionaries, of which several are visible in Figure 3.2.

3.3.1 Attribute-Based Choice

If our tourist applies the attribute-based pattern, he will view each dictionary as an object that can be described in terms of various evaluation-relevant *attributes*. Each object has a *level* with respect to each attribute, such as a particular number of words, which the user may or not be aware of in advance. The chooser assigns an (absolute or relative) *value* to (at least some of) the levels of the attributes. Roughly speaking, the chooser will select the dictionary that seems most attractive in terms of the values of its attributes. But there are many ways of doing so, ranging from thoroughly considering each object's values on all of its attributes to considering only a small sample of the attribute information and taking the object that looks best in terms of the sample.

Useful entry points to the literature on attribute-based choice include Payne et al. [1993], Hsee [2000], Pfeiffer [2012, chap. 2], and Bhatia [2013].

3.3.2 Consequence-Based Choice

A different way of thinking about an option is to consider the concrete consequences of choosing it. So instead of contemplating the large number of words that are offered by a particular dictionary, our tourist might consider how well he could use it to meet French people during his vacation—or, in the longer term, to read the great works of French literature. Consequence-based choice raises a partly different set of issues than attribute-based choice: The chooser needs to deal with uncertainty about what consequences would occur and the fact that they may occur in the distant future. And there can be a considerable variety of possible consequences, including feelings such as the regret that he may feel if he later finds out that he could have chosen a much better option.

Useful entry points to the literature on consequence-based choice include works on the most prominent descriptive model, *prospect theory* (Kahneman and Tversky, 1979; Wakker, 2010) and works that focus on support for consequence-based choice (e.g., French et al., 2009).

3.3.3 Experience-Based Choice

The two preceding choice patterns can involve some quite elaborate reasoning about the merits of the available options. The remaining four patterns describe how people use quite different approaches to arrive at choices more quickly and with little effort (though it is possible to apply them in elaborate ways as well).

Experience-based choice occurs when the chooser's past experience with the choice situation and/or with particular options directly suggests some particular option (without any need for the chooser to think about its attributes or consequences). For example, if our dictionary chooser has had positive experience with the products of a particular publisher, he is likely to have a good feeling when he thinks about purchasing another product from the same publisher, even if he does not remember the previous experiences. Or he may have fallen into the habit of purchasing products from a particular publisher, even without any particularly rewarding experiences. Experience-based choice is often quick and effortless, and it accounts for a large proportion of our everyday choices. A limitation is that the past experiences do not necessarily offer good guidance in the current situation. For example, the positive affect associated with a particular publisher may be due to skillful advertising rather than to any relevant positive experience of the chooser. And habits can be picked up that do not yield particularly good outcomes.

Entry points to the literature on experience-based choice include: Klein [1998], Betsch and Haberstroh [2005a], Plessner et al. [2008], Wood and Neal [2007], and Gigerenzer [2007].

3.3.4 Socially Based Choice

Especially when the previous three patterns do not straightforwardly suggest a particular choice, the chooser may allow himself to be guided by the examples or expectations of others.

If many other people have tried a given dictionary and rated it positively, their ratings can be seen as a summary of a great deal of

relevant experience that it would be impractical for the current chooser to acquire himself.

In addition to providing such *social examples*, other people can provide *social expectations*: For example, a given brand may be viewed as “cool”, while another is seen as politically incorrect. It is understandable that people should be inclined to conform to social expectations of relevant persons, for various reasons.

Entry points to the literature on socially based choice include: Ariely [2008, chap. 4], Cialdini [2007, chaps. 4 and 6], Fishbein and Ajzen [2010, chap. 4], and Thaler and Sunstein [2008, chap. 3].

3.3.5 Policy-Based Choice

Often, the choice process can be seen as comprising two phases, which may be separated considerably in time: In the first phase, the chooser arrives at a *policy* for making a particular type of choice (e.g., “When buying a dictionary for your smartphone, always choose the Oxford dictionary if there is one available”). Later, when faced with a specific choice to make, the user simply applies the policy.

Acting according to policies can have several benefits: It requires less effort at the time of making the individual choices; and it yields a degree of consistency that may be valuable in some situations (e.g., when the choices need to be anticipated by other relevant persons). Perhaps least obviously and most importantly, a policy can (implicitly) take into account important considerations that are often otherwise not visible. For example, always choosing the same publisher’s dictionary should increase the average usability of these dictionaries for the individual user, since the chooser will become skilled at using this publisher’s dictionary interface.

The ways in which choosers arrive at policies in the first place are diverse; they can in principle involve any of the other choice patterns. For example, a farsighted user may reason about the consequences of always choosing the same publisher’s dictionary, taking into account emergent consequences such as the usability benefit just mentioned.

Policy-based choice has been mainly discussed in the literature on organizational decision making, where policies play a more obvious role than they do with individual choice. March [1994, chap. 2] discusses the relationship between what we are calling policy-based and consequence-based choice. Relevant research on individual choices has been conducted in connection with the concept of *choice bracketing* (see, e.g., Read et al., 1999b) and in connection with self-control (see, e.g., Rachlin, 2000).

3.3.6 Trial-and-Error-Based Choice

Especially if none of the other patterns leads readily to a choice, a chooser may simply (perhaps randomly) choose an option and see how well it works out. For example, our dictionary chooser might download the free dictionary and quickly look up a few words, judging whether it seems necessary to spend money on one of the other dictionaries.

One of the tricky questions here, with which a chooser can use some support, is that of the order in which the options should be tried. For example, it may be a good idea to try one of the nonfree dictionaries first, if it seems much more likely that this dictionary will be satisfactory. The other key aspect of this pattern where users can be supported is the interpretation of feedback that the chooser gets when trying out an option. The chooser needs to infer from the results of the trial whether to stick with the current option; and if not, how his strategy for choosing the next option to try should be adapted. Interpreting feedback in this way is not in general straightforward.

It is useful to view the trial-and-error-based pattern as being applied even in some cases where the chooser does not go all the way in executing the chosen option. For example, our dictionary chooser might “try out” a dictionary in the weaker sense of closely examining its description in the app store and carefully reading the reviews. The choice process and the appropriate forms of support are largely similar to those that arise when more thorough trials are involved.

This pattern is often applied in combination with others: The chooser may initially choose on the basis of one of the other patterns, believing that his selected option is the best one—only to notice on the

basis of unexpected feedback that the originally chosen option wasn't as good as expected. So the process becomes one of trial and error even though it was not originally intended as such.

Trial-and-error-based choice has been studied from various perspectives in the psychological literature, mostly not associated with the term "trial and error". Useful entry points include: Rakow and Newell [2010], Cohen et al. [2007], Pirolli [2007], Lindblom [1979], and Zwick et al. [2003].

3.3.7 Combinations of Choice Patterns

The six choice patterns are often used in combination. Explicit discussions of forms of combination are rather rare in the psychological literature (which is not surprising, given that these six patterns have not previously been distinguished as such); but see, for example, Johnson and Busemeyer [2005] for a simulation model of the transition from the consequence-based to the experience-based pattern; March [1994, chap. 2] for a discussion of the relationship between the consequence-based and the policy-based patterns; and Mukherjee [2010] for a model of how the consequence-based and experience-based patterns can be applied in parallel. Many studies do indirectly yield ideas about forms of combination, as does everyday experience:

- Simultaneous application of patterns

Sometimes, the chooser applies two choice patterns and compares the results. If the results differ, there is a conflict that needs to be resolved somehow (e.g., with reference to the amount of confidence that each pattern gives rise to). We can all, for example, remember choice situations in which our experience-based "gut feeling" conflicted with the result of a careful consequence-based analysis, indicating that the two patterns had been applied in parallel and perhaps largely independently of each other.

- Invocation of one pattern by another one

Within one pattern, another pattern can be invoked as a sort of subroutine; for example March [1994, chap 2] points out that

the decision policies of a basically policy-based organization can include particular forms of consequence-based evaluation of options. Similarly, an attribute-based chooser who doesn't know how desirable a particular feature of an option is can rely on the judgments of others to answer this question (borrowing from the socially based pattern). In another common combination, one choice pattern is used to reduce a large number of options to a manageable set and then a different pattern (e.g., trial and error) is used to choose among these options (cf. 5.3).

- Application of a choice pattern on the meta-level to determine which choice pattern to apply on the domain level.

Given that people apply different choice patterns on different occasions, they are in at least some sense “choosing” among choice patterns. For example, a chooser might in effect reason “If I take the time to think carefully about which restaurant to go to, I won't have enough time left to eat lunch. So I'll just go right away to the same restaurant that I chose last time.” Here, consequence-based choice is being applied on the meta-level to select a particular subpattern of the experience-based pattern for choosing among the options in the choice domain. It seems more typical for these meta-level choices to be experience-based (e.g., habitual).

It will be interesting to see further descriptive research on the ways in which choice patterns are combined in naturally occurring everyday choice. As choice architects, though, we also have the opportunity to consider what combinations of patterns *could* usefully be applied if we provided adequate support for them.

3.4 Relationship to Two Modes of Processing

A distinction that is largely orthogonal to the distinction of the six choice patterns in the ASPECT model concerns two modes of processing, each of which can be found in conjunction with each of the patterns. Throughout the psychological literature and even in everyday parlance, we find different variants of a general distinction between two modes of

Table 3.2: Some features typically associated with the intuitive and analytic modes of choosing.

Intuitive	Analytic
Automatic, unconscious	Deliberative, self-aware
Fast, effortless	Slow, effortful
Concrete, emotional	Abstract, largely affect-free
Parallel, high-capacity	Serial, low-capacity
Exploits automatic learning from experience	Exploits prior acquisition of beliefs
Changes only with repetitive or intense experience	Can change immediately

thinking, which we here call *intuitive* and *analytic*, respectively, since these terms seem to work naturally in HCI contexts. In the scientific literature, many other pairs of terms are found; it appears as if no two writers use the same terms.⁴ On close inspection, the conceptions used by these different authors differ in some respects, but for our purposes it is more important to understand the common ideas underlying this distinction. Table 3.2 lists some of the characteristics of each mode that are repeatedly mentioned.

The intuitive mode of processing occurs most prominently in the experience-based pattern, where often (almost) all of the processing occurs in this mode. But in all of the other patterns as well, a good deal of processing can occur in this mode, as when the chooser roughly assesses the likelihood that a particular consequence will occur or gets a sense of the social expectations that prevail in a particular situation.

3.5 Ecological Rationality

As preparation for the next main section, in which we will survey possible forms of choice support, we should now consider from a psycho-

⁴See, for example, Epstein [2010]; Hogarth [2001]; Thaler and Sunstein [2008, chap. 1]; Kahneman [2011].

logical perspective what constitutes a “good” choice process in the first place.

Before considering what everyday choosers themselves think about this issue, we should mention a shift in the thinking of scientists who have studied choice and decision making. Traditional notions of what constitutes a good choice are that a decider should (a) apply a decision procedure that is normatively justifiable (e.g., consistent with the laws and principles of logic, probability, and expected utility) and (b) choose the action that will maximize desirable (and minimize undesirable) outcomes under idealized conditions (see, e.g., Gigerenzer and Todd, 1999, chap. 1; Gigerenzer, 2007, chap. 5). But inspired by thinkers like Simon [1956], researchers have become more impressed by the extent to which animals and humans can function quite effectively by using choice procedures that are *ecologically rational*: justifiable in the sense that they work well in the environment in which they are applied and make good use of the chooser’s limited time and cognitive resources. For example, Simon [1956] discusses the choice processes required by two hypothetical species of animals that forage for food in two different types of environment: (a) an environment in which food is distributed randomly, making systematic search strategies useless; and (b) an environment in which there are cues available to the animal which signal the regions in which it makes sense to look for food. In each of these environments, there is a relatively simple strategy for the animal that works just as well as any more sophisticated strategy would and that is much simpler than general normative models of decision making, which involve notions such as estimating and maximizing overall utility.

Looking at our six choice patterns (Table 3.1), we can see that each one seems likely to work better under some conditions than under others. For example, being guided by the choices that other people make (the socially based pattern) could hardly be advocated as a universally applicable normative model for decision making; but it can be the best procedure available if (a) other people have successfully learned how to make the type of choice in question, (b) they are sufficiently similar to the chooser so that it is reasonable for her to try to benefit from their experience; and (c) reasonably reliable information about their

choices is available. The fact that each pattern is ecologically rational under some circumstances helps to explain why these different patterns exist. It also implies that, to support everyday choice, we should aim to support each of these patterns in situations where it works well. In particular, we should not try to impose any particular procedure as a substitute for combinations of these patterns; such a procedure might look great on paper, but it would probably not be well adapted to the many environments in which people make everyday choices.

3.6 What Constitutes a Good Decision for Choosers?

3.6.1 Factors That Influence Satisfaction With a Choice

In addition to aiming to achieve ecological rationality, we should consider what ordinary people view as constituting a good choice process. A number of researchers have investigated this question (see, e.g., Bettman et al., 1998; Hastie, 2001; Yates et al., 2003). Although specific answers to this question vary, the following statements are widely accepted:

1. Choosers want their decision to yield a good outcome.

This point isn't as straightforward as it may seem, because what counts as a good outcome depends in turn on a variety of factors, as we will see. In particular, the criteria by which a chooser evaluates an outcome can be quite different after the outcome has occurred—and especially much later in time (cf. March, 1994, chap. 6). To take a familiar case, an immediately gratifying outcome (e.g., watching an amusing video on YOUTUBE) may be evaluated highly as it is being experienced but evoke only regret at a later point in time. A much more subtle example is given by Hsee [2000]: A person who has carefully compared two sets of stereo speakers in a store chooses the set with slightly higher sound quality, even though its appearance is displeasing, on the grounds that sound quality is the most important attribute of speakers. When the buyer uses the speakers at home, he can no longer perceive the slight superiority in sound quality of his speakers, but he is constantly disturbed by their appearance and hence he regrets his choice. Note that the result could be different if the buyer regularly

had a chance to compare his speakers with others (e.g., if his spouse had bought the more physically attractive set of speakers). An implication is that choosers and those who want to support choosing should try to take into account the context in which the chosen option will be experienced, including the subtle features such as whether the competing options will still be available for comparison. One approach is to aim to achieve similarity between (a) the way in which the options are presented to the user at choice time and (b) the way in which the chosen option will be experienced after the choice has been made (cf. also Read et al., 1999a).

In the present publication, we will generally view a “good outcome” as one that the chooser is (or would be) satisfied with in retrospect, after having acquired the most relevant knowledge and experience. Admittedly, this assumption is subject to debate, and other assumptions could be substituted for it.

2. Choosers don’t want to invest time and effort in the decision making process itself that is out of proportion to the resulting benefits.

For example, when installing a new application, a user who is asked which specific components should be installed will often choose the option “Everything” simply to save the time of deciding about the individual components, since the possible benefits of choosing any other option (e.g., saving a few megabytes of hard disk space) do not seem to justify the investment of even a few seconds of time to make a more deliberate choice.

This tendency to want to save time and effort is more important than you might expect, because of the special salience of the effort that is invested in a choice process, relative to the benefits that can be obtained with a better choice. For example, even if 2 minutes invested in configuring a new application could end up making the user’s experience more pleasant during many hours of use of the application, these future benefits may be insufficient to motivate the user to think carefully about the configuration problem: Whereas the tedium and frustration of the configuration work is immediately experienced, the benefits are uncertain and largely distant in time (cf. Bettman et al., 1998); and as we will see in 6.8 and 6.9, consequences like these tend

to be discounted heavily. An implication of these facts for supporting everyday choice is that any measure which reduces time and effort is likely not only to be appreciated by users but also to encourage users to improve their choice process in ways that they would otherwise not consider worthwhile.

The following two desiderata of a choice process are much more specific than the first two, but they can likewise help to guide the selection of forms of support for choice:

3. Choosers prefer to avoid unpleasant thoughts.

Some ways of thinking about a decision can involve distressing thoughts, as when a driver faces a choice between (a) ignoring an incoming text message from her boss and (b) driving less safely for a while in order to respond to the message. The chooser may be motivated to think about the decision in a way that avoids such thoughts (e.g., by convincing herself that she can respond to the boss's message without taking the slightest risk, or by ignoring safety considerations entirely).

4. Choosers often want to be able to justify the decision that they have made to other persons—or to themselves.

Justifiability is often simply a necessary condition for being able to implement a decision (cf. Lerner and Tetlock, 2003). For example, even if a business person would really like to buy the latest iPad for professional use, he is likely to request a less flashy tablet instead if he thinks that this choice is more likely to be approved by his company's purchasing department. But even just the prospect of being able to convince another person or oneself that a decision was sound can cause people to look for justifiable decisions (see, e.g., Shafir et al., 1993).

Consequently, one way of supporting preferential choice is to make it easy for the user to come up with a satisfying justification of whatever option is best for her, for example, by supplying a justification explicitly (as is done by many recommender systems; see Tintarev and Masthoff, 2010).

3.6.2 Choice Overload as an Illustration

The importance of these different determinants of satisfaction with a choice is illustrated by the oft-discussed phenomenon of *choice overload*: the fact that, at least in some cases, choosers end up less satisfied when they have been offered a large number of options than when they have been faced with a smaller number.⁵ The question of how prevalent this phenomenon is and under what conditions it occurs is a complex one that is still being researched; but there is considerable agreement on some of the possible reasons why choice overload can occur, and these reasons illustrate the factors listed above:

- All other things being equal, choosing from a larger number of options requires more time and effort, which may not be compensated for by an improvement in the quality of the result.
- With a larger number of options, it is more likely that there will be competing options of similar attractiveness, which can vary in their relative attractiveness with every shift in the chooser's attention to different priorities. Oscillating between competing alternatives can be a frustrating as well as a time-consuming process.
- For the same reason, it tends to be harder for the chooser to justify—to himself or others—whatever choice he finally does make. This difficulty is exacerbated when there are options that the chooser did not manage to consider at all.

In situations in which choice overload seems likely to occur, an obvious remedial tactic is to limit the number of options made available to a chooser. From the point of view of choice architecture, limiting a priori the number of options is just one of many possible tactics

⁵The report by Iyengar and Lepper [2000] on the “jam study” is the most widely cited (and overinterpreted) example of apparent choice overload. Both Iyengar [2010] and Schwartz [2004] offer popular book-length discussions of issues related to choice overload. Scheibehenne et al. [2010] present a compact but comprehensive meta-analysis and synthesis of ideas. Oulasvirta et al. [2009] studied choice overload in a typical HCI context, finding evidence of the phenomenon in connection with the use of web search engines.

for supporting choice—one which is derived from the general ARCADE strategy *Design the Domain* (4.5).

The next main section will show how a comprehensive conception of approaches to choice support can help to generate a variety of tactics for helping choosers to overcome challenges—of which choice overload is but one of many—that can arise during the application of the ASPECT choice patterns.

4

Choice Support Strategies: The ARCADE Model

Against the background of the overview of human choice that has been presented, how can designers of interactive systems help people make better choices?

Just as it is important to have a comprehensive view of how people make choices (cf. 3.1), choice architects also need to bear in mind the full range of possible approaches to supporting choice. Otherwise, we might overuse one or two approaches even in cases where another approach would be much more effective. Interaction designers in particular are in a position to deploy a wider range of strategies and tactics than most others who try to help people make better choices, such as authors of self-help books and even professional consultants.

The ARCADE model (summarized in Figure 4.1) distinguishes six high-level strategies that have been used in various combinations and implemented in various ways, with and without computing technology. The model takes into account many ideas contributed by earlier analyses of approaches to choice support (see, e.g., Fischhoff, 1982; Payne et al., 1993, chap. 7; Larrick, 2004; Yates et al., 2003; Johnson et al., 2012). Each of the ARCADE strategies has been discussed in at least some previous literature (though not with the same

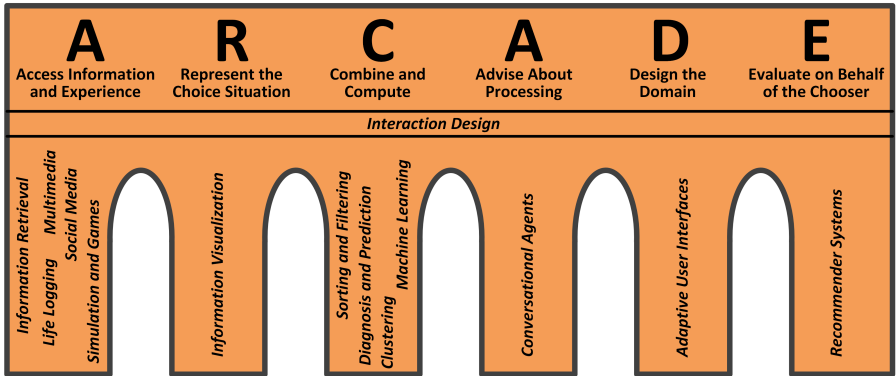


Figure 4.1: The six high-level strategies for choice support in the ARCADE model and some of the technologies for realizing them.

designation), but to our knowledge they have never all been presented together in a coherent way.

Each of these strategies may be too abstract in itself to be very suggestive to interaction designers; but when we consider how the strategies can be applied to particular choice patterns, we will be able to derive a large number of fairly concrete *support tactics*.

4.1 Access Information and Experience

4.1.1 Basic Idea

The most obvious way of supporting choice is to supply the chooser with relevant information. A bit less obviously, we can provide access to relevant *experience*, which can help in roughly similar ways.

In fact, this strategy is so obvious that there can be a temptation to respond to any perceived need for choice support by displaying potentially relevant information and assuming that the chooser will somehow benefit from it. The limitations of this approach can be seen especially clearly in domains, such as automotive user interfaces, where the chooser has very limited capacity to process such information at the time of choosing. One advantage of thinking in terms of the six choice



Figure 4.2: Examples of how information relevant to consequence-based choice can be provided straightforwardly in a user interface with text or graphics.

patterns is that they indicate to us what particular types of information and experience can be used in what ways; they therefore help us focus on the most useful ways of making information and experience available.

Here are some examples of different types of information and experience that can be useful in connection with different choice patterns; many more examples will be given in the sections on the individual patterns:

- *Attribute-Based Choice:* Information about each option in terms of its levels on the various attributes.
- *Consequence-Based Choice:* Information about what's going on in the current situation and what consequences particular actions are likely to have (see, e.g., Figure 4.2); a preview of the experience of executing a particular action.
- *Experience-Based Choice:* Reminders of past experience with particular options.
- *Socially Based Choice:* Information about choices made by relevant other people, or about their expectations concerning the chooser's choices.
- *Trial-and-Error-Based Choice:* Information about (not immediately visible) consequences of an option that the chooser has executed.

4.1.2 Straightforward Interaction Design

As with all of the other five ARCADE strategies, there are ways of providing access to information and experience in an interactive system through what we will call “straightforward interaction design”: design patterns and elements that are typical of the interaction style being applied, which do not require any particularly sophisticated or intelligent technology.

A lot of choice-relevant information is conveyed in user interfaces via natural language (e.g., English) text, which can be shown in various places in a graphical user interface or spoken in a speech-based interface. In the example shown at the bottom of Figure 4.2, the information being provided concerns a general fact that is known to the interaction designer; hence no particular technology for acquiring the information is needed.

Often, nonverbal ways of conveying information are more effective, as with the simple graphical representations at the top of Figure 4.2, which inform the user unobtrusively about potentially choice-relevant facts (e.g., her likelihood of getting a good connection to the mobile internet if she tries).

Where the relevant information or experience concerns what it’s like to use a part of a system, it is often possible simply to offer the user the opportunity to use that part of the system (e.g., to try out an unusual text entry method briefly before deciding whether to use it to enter some text).

4.1.3 Relevant Technology

A large proportion of the computing technology that has become widely used in recent years can be seen as helping to enable the provision of information and experience, such as access to databases via the worldwide web and mobile apps; and semantic technology for linking data from diverse sources. The following particular technologies can be singled out for mention:

- Web search engines

Even a chooser who is armed simply with a (mobile) web browser and a search engine may be able to access most or all of the types of information and experience that are relevant to a particular choice, ranging from objective facts about products offered for sale to vivid experience reports about the consequences of undergoing a particular medical treatment. In that sense, a web browser with a search engine can be seen as a general-purpose choice support system—one which does, however, leave it up to the chooser to decide what particular information and experience to access and what use to make of it.

- Multimedia technology

Access to experience is often most effectively provided via media such as photos, audio tracks, and video. Hence the increasing capabilities of multimedia technology to identify and deliver media relevant to a particular information need, combined with the increased ability of users to consume media even on mobile devices, offer great potential for choice support—provided that the role of such media in the choice process is well understood.

- Simulation and gaming

Knowing what it's like to perform a given action or to experience a particular consequence sometimes requires more than just the passive consumption of media: Interacting with a simulation of reality can be more effective; such simulations are often embedded in games. An example is the simulation game used by Chittaro [2012] to teach air travelers what it is like to be involved in an airplane crash without having learned proper evacuation procedures.

- Social media

Technologies that enable computer users to publish their opinions, knowledge, and experience make it possible for other users to access not only information relevant to the socially based choice

pattern but also a wide range of other information and experience.

- Digital life logging

Much of the experience that is relevant to a person's choices lies in his own past; but there is no guarantee that a chooser will think of relevant experiences or recall them accurately (cf. Kahneman, 2011, chap. 35). The many life logging systems that now exist for helping people to record their experiences were mostly developed with other goals in mind; but helping people to access past experiences that are relevant to current choices is a promising way of putting this technology to good use (see, e.g., Plate et al., 2006).

4.2 Represent the Choice Situation

4.2.1 Basic Idea

One of the key points emphasized in the Thaler and Sunstein [2008] conception of choice architecture, as well as in some of the psychological research that it is based on, is that the particular way in which a choice situation is represented to the chooser can strongly influence how she goes about making the choice and which option she will choose, even when the information being presented remains essentially the same.

Thaler and Sunstein [2008, Introduction] discuss the case where the choice architect is the manager of a cafeteria: The way in which particular foods are arranged on the shelves will inevitably have some influence on which foods the visitors to the cafeteria tend to choose. As the authors emphasize, there is no completely “neutral” way of arranging the options. For example, it is inevitable that some dishes will tend to be encountered before others and to be more perceptually salient. We will return to this point in 4.7.1.

We can see different possible effects of representation of the choice situation in different choice patterns:

- *Attribute-Based Choice*: Presenting information about options in a neat tabular form makes it easier to engage in attribute-based

processing; the order in which options and attributes are shown will influence the amount of attention that is paid to them.

- *Consequence-Based Choice*: A suggestive representation of the choice situation can help the chooser predict quickly what will happen next (see, e.g., Figure 4.2). The various different types of consequences that the chooser may be interested in can be represented (if at all) in more or less salient ways. Even the values that the chooser assigns to particular consequences are often influenced by the exact ways in which they are represented (see 6.7.2 for examples).
- *Experience-Based Choice*: The way in which a choice situation is represented to the chooser can help determine which (relevant or irrelevant) experiences the situation evokes in the chooser's mind.
- *Trial-and-Error-Based Choice*: A convenient representation of the results of the chooser's experience in trying out options can make it easier for the chooser to judge which option is emerging as the winner.

4.2.2 Straightforward Interaction Design

Designing appropriate representations is a key part of good interaction design. Designers typically think carefully about how to represent items visually (or via other sensory channels) and about how to arrange them (e.g., on a visual display) so as to suggest appropriate interpretations, processing, and actions. These skills and experience can be a valuable resource to interaction designers as they fulfill their role as choice architects—provided that they understand the ways in which different representations can influence processing within the various choice patterns. A number of examples will be given in the next six main sections of this publication.

4.2.3 Relevant Technology

More complex techniques are explored in the area of information visualization (see, e.g., Ware, 2004; Card, 2012), where sophisticated com-

putation and computer graphics are used to create representations of information that support particular types of visual and cognitive processing especially well. To give just one example, Pu and Faltings [2000] provided a sophisticated visualization of various options for flying from one city to another which supported attribute-based processing more effectively than a traditional tabular representation could have.

4.3 *Combine and Compute*

4.3.1 Basic Idea

This third general choice support strategy involves performing on behalf of the user mechanical tasks such as calculations and other manipulations of available information.

- *Attribute-Based Choice*: Decision aids in e-commerce often enable the chooser to sort and filter options according to their levels on particular attributes (see, e.g., Pfeiffer [2012]). Note that the system is not providing any new information; it is helping with the manipulation of available information.
- *Consequence-Based Choice*: A mobile app for helping to choose food will normally go beyond providing information about the various options: It is likely to compute quantities such as the percentage of the chooser's target calorie budget or the amount of weight that the chooser would gain if he consumed the food.
- *Socially Based Choice*: Ratings and opinions supplied by other persons can be clustered and summarized so that they can be more easily processed by the chooser.
- *Policy-Based Choice*: If the system knows the chooser's relevant policy for a given choice, it may be able to determine which of the available options are consistent with that policy.

As these examples show, application of the strategy *Combine and Compute* is often, but not always, combined with the accessing of information that is not already available to the chooser.

4.3.2 Straightforward Interaction Design

Useful computations can often be performed with straightforward algorithms such as those for sorting items or for standard arithmetical operations.

4.3.3 Relevant Technology

In individual domains, the strategy *Combine and Compute* can involve the application of sophisticated algorithms. One characteristic feature of many decision support systems (see, e.g., French et al., 2009, chap. 14; Clemen and Reilly, 2013) is the presence of sophisticated algorithms for diagnosing the current situation (e.g., determining the likelihood that a patient has a particular disease) and predicting future outcomes (e.g., the likelihood that particular medical treatments will be successful). This sort of heavyweight computation is more typical of systems that are designed to support extensive, deliberate decision making in organizations, as opposed to everyday choice; but there is no reason in principle why scaled-down versions of these techniques should not be applied, when applicable, to everyday choices, especially in view of the increasing computational capacity of today's personal computing devices.

4.4 Advise About Processing

4.4.1 Basic Idea

Traditionally, an important way of helping people make better choices has been to give them advice as to how to proceed. A famous example is Benjamin Franklin's 1779 letter to Jonathan Williams (discussed, e.g., by Gigerenzer, 2007, p. 4), in which he outlined a systematic way of writing down arguments for and against particular options and integrating the information to arrive at a decision. Decision analysts who help their clients make decisions (see, e.g., Clemen and Reilly, 2013; Keeney, 1992) often recommend particular procedures, such as first reflecting on their fundamental values before they start to evaluate partic-

ular options; or systematically considering arguments *against* whatever option the decision maker tends to favor at a given moment.

Popular books about how to make decisions in everyday life tend to be full of less formally articulated advice. For example, the book by Schwartz [2004], which helped to popularize the concept of *choice overload* (3.6.2), advises choosers not to worry too much about finding the best possible option but to be satisfied with a solution that is adequate. (Further advice from this book is mentioned in 6.6.3.)

In terms of our six choice patterns, this type of advice can be seen as urging the chooser to apply a particular choice pattern (or combination of patterns) and to execute individual steps in particular ways.

4.4.2 Straightforward Interaction Design

The most obvious way to provide procedural advice is with texts like “Try out the free version before buying the app!”. But since an interaction designer typically has a lot of control over the choice environment, it is often easy and more natural to influence a chooser’s procedures implicitly by making some procedures more salient or easier to execute—for example, by providing an easily accessible link to the free version of an app, so as to facilitate the trial-and-error-based pattern; or by prominently displaying user ratings and reviews, encouraging the application of the socially based pattern. That is, tactics derived from other ARCADE strategies can often be interpreted as implicit procedural advice. We will normally refer to the strategy *Advise About Processing* only when the procedural advice is reasonably explicit, but the reader should remain aware of the various implicit forms that procedural advice can take. In fact, in order to *avoid* influencing the procedures employed by the chooser, an interaction designer would often need to make a special effort to enable the chooser to employ several different procedures with approximately equal ease.

4.4.3 Relevant Technology

A more sophisticated way of offering verbal procedural advice is to deploy conversational agents. These systems engage in a dialogue with the user that has some features of human-human dialogue, and they

sometimes involve human-like features such as speech and animated faces. Virtual sales agents were popular in e-commerce and other websites around the turn of the century (see, e.g., André et al., 1999), and they seem to be regaining some popularity at the time of this writing. Although an agent of this sort can't compete very well with other interaction styles when it comes to applying the other five ARCADE strategies, maybe a careful consideration of how they can be used for providing procedural advice can help them to gain a foothold as components of choice support systems.

4.5 *Design the Domain*

4.5.1 Basic Idea

In connection with the strategy *Represent the Choice Situation*, we mentioned how an interaction designer has many opportunities to represent a choice situation in ways that support choice. Going a step further, especially where choices are being made about the use of computing technology, the interaction designer can even change the basic reality in the domain in question in a way that makes it easier to make choices.

4.5.2 Straightforward Interaction Design

A typical example is sometimes found in the area of privacy choices: If users find the set of privacy configuration options to be confusing and have difficulty choosing among them, the most effective approach may be to reconceptualize the set of privacy options so that users are offered a different (probably smaller) set of options that they can more easily understand and choose from (cf. 12.2.3).

A somewhat more specific variant of this strategy is that of minimizing (a) the number of choices that users need to make in the first place and (b) the number of options that are available when they do need to choose. Steve Jobs (see, e.g., Isaacson, 2011) was a powerful advocate of this strategy, even in the early days of the Macintosh, when he opposed the inclusion of a second mouse button and cursor keys; and

this principle is still a subject of heated debate between supporters of the Apple and Android approaches to designing smartphones.

Here are two examples concerning different choice patterns:

- *Trial-and-Error-Based Choice*: Applications (and methods used within them) can be designed in such a way that they can be tried out easily—for example, with little effort or commitment.
- *Consequence-Based Choice*: The possible actions that users can perform can be designed in such a way that their consequences are easy to predict. This principle is a general HCI design principle whose benefits are not restricted to cases involving preferential choice; but it does also make it easier for a chooser to evaluate alternative options by anticipating their consequences.

4.5.3 Relevant Technology

Since there are many ways in which a choice domain can be redesigned so as to facilitate choosing, this strategy is not strongly associated with any particular technology. To some extent, however, the technology of adaptive user interfaces (see, e.g., Jameson and Gajos, 2012) can be enlisted in the service of this strategy. For example, a user interface that somehow recognizes that the current user is a novice and switches to a simple “novice mode” can be seen as automatically changing the choice domain—which is in this case the user interface itself—in the way described above.

4.6 Evaluate on Behalf of the Chooser

4.6.1 Basic Idea

This support strategy involves doing some choosing or evaluating on behalf of the chooser and (implicitly or explicitly) advising the chooser to adopt the result. This strategy has been applied for millennia by human advisers; in the realm of technological support, it is found most prominently in *recommender systems* (1.2.5; Jannach et al., 2011; Ricci et al., 2011).

4.6.2 Straightforward Interaction Design

Interaction designers often provide recommendations straightforwardly via natural language texts such as “It is recommended that you close all other applications before continuing” and “[This configuration setting is] Recommended for new users”. A less obtrusive way of conveying the second recommendation is to preselect the option in question as a default.

4.6.3 Relevant Technology

There are various points in the different choice patterns at which a recommender system can make an (absolute or relative) evaluation on behalf of the chooser.

One possibility is for the recommender to take over the entire process of making a choice: from recognizing that a choice should be made through selecting a particular option and then even executing it. An example is a music recommender that fully automatically decides what songs to play for the user. In this limiting case, instead of “choice support” we could speak of replacing the human chooser by a recommender system.

A more common and somewhat less aggressive role for recommenders is to reduce an initially very large set of options to a manageable *consideration set* (see 5.3).

Often, though, recommenders take over particular parts of the choice process and provide some support for the other parts (often applying one or more of the other ARCADE strategies).

- *Trial-and-Error-Based Choice*: A particular type of recommender system, the *critique-based* recommender (see, e.g., McGinty and Reilly, 2011; Ricci and Nguyen, 2007; Reilly et al., 2007) provides support for the trial-and-error-based pattern: The recommender suggests which option to choose next, applying an exploration strategy that it does not necessarily expose to the user.
- *Policy-Based Choice*: A less frequent but potentially important function of recommenders is to recommend a *policy* that applies

to a large number of similar choices; the policy can then be applied either by a system or by the individual chooser. An example would be a system that recommends a particular daily calorie intake for a user who wants to achieve a weight loss goal.

4.7 Alternative Goals in Applying the ARCADE Strategies

The overall goal pursued by a choice architect who applies the ARCADE strategies is to help people make better choices. The architect can also pursue two more specific goals that we haven't mentioned so far and that require the strategies to be applied in somewhat different ways:

1. *Avoid unwanted bias* in the choice architecture.
2. Help the chooser deal more successfully with *similar choices in the future*.

For each of these goals, we consider why it is often worth pursuing and how it requires somewhat different ways of applying the ARCADE strategies.

4.7.1 Goal 1: Reduce and Avoid Bias in the Choice Architecture

To see what sort of bias we are talking about, recall a point that was made in 1.2.2 and 1.2.4:¹ Although the ARCADE strategies have been presented as ways of achieving choice support, each of them can also be applied for the purpose of persuasion: an attempt to induce a person to choose a particular (type of) option. Although the present authors don't want to advocate this type of use of the ARCADE strategies, it is important to be aware of them, since similar effects can arise despite the best intentions of the choice architect.

¹The sort of bias discussed in this subsection exists in the choice environment. By contrast, the term *bias* is most commonly used in connection with the psychology of judgment and decision making to refer to some sort of systematic error that people make, typically while judging things like frequencies and probabilities; see, for example the collections by Kahneman et al. [1982], and by Gilovich et al. [2002]. Discussions of the relationships between biases "in the head" and those in the choice environment can be found in the collection edited by Fiedler and Juslin [2006], in particular in the introductory chapter.

Suppose you are the choice architect and that you would like a chooser C to choose a particular option O . Here are some examples of how you could apply ARCADE strategies in biased ways for this purpose:

- *Access Information and Experience:* Give the chooser access to an unrepresentative sample of information or experience which will incline her to prefer O .

For example, make sure that the positive aspects of O are over-represented and its negative aspects underrepresented.

- *Represent the Choice Situation:* Arrange the choice environment in such a way that O is more likely to be identified as an option or is likely to be seen as the default option.

Example: Make the icon for choosing O bigger and easier to click on than the others.

- *Combine and Compute:* Perform computations which, although not necessarily strictly incorrect, are more likely to lead to a choice of O than other equally valid computations would.

The seemingly limitless possibilities of introducing bias into computations have been popularized by books like *How to Lie With Statistics* (Huff, 1993).

- *Advise About Processing:* Encourage C to apply a procedure that is especially likely to lead to the choice of O .

Example: Encourage affect-based choosing when O evokes positive affect.

- *Design the Domain:* Make options that might compete with O unavailable or costly to reach.

Example: If O is a commercial product, put competing products on separate, hard-to-access pages or screens.

- *Evaluate on Behalf of the Chooser:* Recommend O to C regardless of whether there is reason to believe that C will be satisfied with it.

Example: Use a recommendation algorithm that is biased in favor of recommending things like O .

Biases like these can slip in despite the best intentions of the choice architect: They may already be present in the choice situation (as when information is available to the chooser from sources that lie outside of the control of the choice architect), or the choice architect may introduce them inadvertently. So we should aim to reduce existing biases and avoid introducing new ones.

A general limitation here was pointed out in the discussion (already mentioned in 4.2) by Thaler and Sunstein [2008, Introduction] of the example of arranging food in a cafeteria: While some choice architectures may be clearly more biased than others, there is in general no such thing as complete lack of bias.

This point is important, but it is not obvious what its practical implications are. One possible conclusion is that, given that some amount of bias is inevitable anyway, we might as well introduce a bit of bias that serves some apparently desirable goal that may not be shared by the chooser, such as a bias toward organ donation (cf. the first sense of the term *nudge* discussed in 1.2.3).

In the present work, instead of advocating a particular approach to this question, we will view the minimization of predictable bias as one goal that can be pursued in the application of the six ARCADE choice support strategies. Our readers as choice architects are free to think for themselves about how to deal with any inevitable residual bias, taking into account the relevant details of the choice situation.

4.7.2 Goal 2: Support Future Choices

Up to now, we have been implicitly assuming that the goal of a choice architect is to help a chooser deal with a choice situation that she is currently facing. But it often makes sense to take a longer-term perspective, taking steps now that will help the chooser deal more effectively with choice situations that arise in the future.

Education and Training

The most obvious situation of this sort is education and training in decision making, which can occur anywhere from the school classroom (see, e.g., Hogarth, 2001; Gigerenzer, 2014) to professional training courses (see, e.g., Phillips et al., 2004; Ross et al., 2005). In these contexts, choosers are typically given choice support as they handle a carefully chosen range of example choice situations. The ARCADE strategies are still relevant, but the goal of applying them is different: It is less important that the chooser handles the example choices well than that the chooser learns how to do a better job of handling similar choices in the future. While acknowledging the general importance of education and training in choice and decision making, we do not discuss it in this publication, partly because education and training, at least in their familiar forms, do not in general constitute a natural and feasible type of choice support for everyday choices made by users of computing technology: People seldom take courses to learn how to make the types of choice that we have been discussing.

Support for “Learning While Choosing”

There is a related approach, however, which is feasible even for small everyday choices: Help the chooser to learn while making real everyday choices, so that he can do a better job with similar choices in the future. We will see a number of examples of this approach when discussing support for trial-and-error-based choice (Section 10), which by definition involves learning from the experience of making choices. But with the other five ASPECT choice patterns as well, any application of an ARCADE strategy can lead to some amount of learning; so it can make sense for a choice architect to aim intentionally to support that type of learning, especially if it is known that similar choice situations will arise repeatedly in the future.

We will use as an example here the type of choice that concerns health and well-being (rather than the use of computing technology): a choice about what sort of exercise (specifically, running) to do on a particular day. The screenshot in Figure 4.3 shows feedback that was

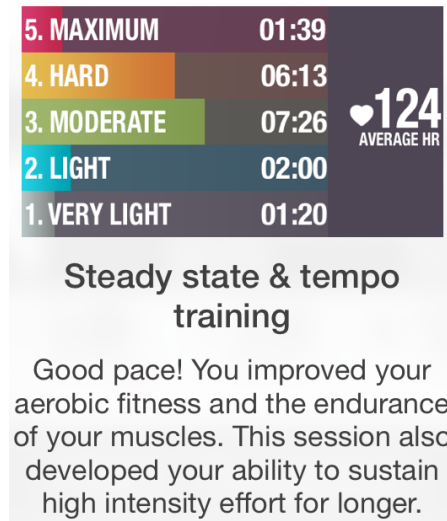


Figure 4.3: Screen shot from the iPhone application POLAR BEAT showing feedback on a running session performed with a heart rate (“HR”) monitor.

given to a runner after the completion of a run. It therefore constitutes support for trial-and-error-based choice; but note that a similar screen could be presented to a runner who was simply *considering* a particular trajectory and speed, previewing for her the likely outcomes of choosing this option.

This example illustrates several tactics that are well suited to supporting similar choices in the future:

- *Access Information and Experience:* Instead of giving the chooser access to just the information and experience that she needs to deal with the current choice, the screen includes textually presented information that helps her to acquire a more complete mental model of the effects of running.
- *Represent the Choice Situation:* The histogram encourages a way of conceptualizing the properties of a running session (in terms of time spent at different levels of heart rate) that can influence

the chooser's thinking about running even when she is not using this decision aid.

- *Evaluate on Behalf of the Chooser*: The screen includes an explicit (positive) evaluation of the depicted run, which may encourage the runner to choose similar runs in the future.

This goal of supporting “learning while choosing” has received relatively little attention in the scientific literature so far, but it deserves close attention of choice architects in HCI, especially because of the fact that most choices in this area occur repeatedly.

5

Attribute-Based Choice

In this and each of the following five main sections, we will look more closely at one of the six ASPECT choice patterns, noting some of the key concepts and results that have emerged in research about the pattern and considering how the ARCADE strategies can be applied.

5.1 Introduction to the Pattern

The attribute-based pattern (summarized in Table 5.1) can be applied whenever the set of available options can usefully be described in terms of attributes (e.g., “price”, “average user rating”) and levels of those attributes.¹ The most obvious cases where this style of choice is relevant is when users are shopping in e-commerce sites (e.g., app stores) that explicitly present available products and services in terms of their attributes.

¹Note that the attribute “average user rating” summarizes information about the evaluations of other persons. So a chooser who takes this attribute into account is also invoking part of the socially based pattern of choice, in the way described in 3.3.7.

Table 5.1: Attribute-Based Choice. (In this and subsequent similar tables, each tactic in the right-hand column can be seen as an instantiation of the general ARCADE strategy listed just before it. These tactics are just a sample from the set of tactics that can be formulated on this level of abstraction. In the left-hand column, a step that is indented to the right should be seen as one possible way of executing the more general step shown above it.)

Steps	Support Tactics
<ul style="list-style-type: none"> – (Optional:) <i>C</i> reflects in advance about the situation-specific (relative) importance of attributes and/or values of attribute levels 	<p><i>Advise About Processing:</i> Encourage <i>C</i> to reflect in this way <i>Represent the Choice Situation:</i> Represent the choice situation in a way that avoids unnecessary bias in favor of particular attributes or levels</p>
<ul style="list-style-type: none"> – <i>C</i> reduces the total set of options to a smaller <i>consideration set</i> on the basis of attribute information 	<p><i>Design the Domain:</i> Make only a limited set of options available in the first place <i>Combine and Compute:</i> Offer computational support for winning strategies <i>Evaluate on Behalf of the Chooser:</i> Recommend a consideration set</p>
<ul style="list-style-type: none"> <i>C</i> considers only the first few items that are presented or <i>C</i> filters (or sorts) the options according to one attribute <i>A</i> at a time, until there are only a manageable number left 	<p><i>Combine and Compute:</i> Filter or sort according to criteria specified by <i>C</i> <i>Represent the Choice Situation:</i> Display winning results so as to enable <i>C</i> to judge how successful the winning has been</p>
<ul style="list-style-type: none"> – <i>C</i> chooses from among a manageable set of options <i>C</i> applies a comparative integration strategy 	<p><i>Access Information and Experience:</i> Provide objective information about the attribute levels for options <i>Represent the Choice Situation:</i> Present information about attribute levels in a way that facilitates mental application of comparative strategies</p>
<ul style="list-style-type: none"> or <i>C</i> applies an additive evaluation strategy <i>C</i> evaluates each of (a sample of) attribute-level combinations 	<p><i>Evaluate on Behalf of the Chooser:</i> Suggest evaluations for particular attribute levels <i>Represent the Choice Situation:</i> Represent attribute levels relative to evaluation-relevant reference points</p>

5.2 Thinking in Advance About Evaluation Criteria

As we will see, there is a wide range of strategies that the chooser can (implicitly) apply with attribute-based choice. Within each strategy, the chooser's evaluation criteria—for example, which attributes he considers most important when looking for a new smartphone and how he evaluates particular attribute levels—will influence the outcome. So a relatively deliberate chooser may want to take some time from the start to settle on a particular strategy and particular evaluation criteria, maybe before looking at any particular options. This sort of reflection is a standard part of decision analysis, for which elaborate procedures and representations have been developed (see, e.g., Keeney, 1992; Clemen and Reilly, 2013, sect. 3). In terms of the ASPECT model, this step is most naturally seen as the construction of a *policy* for making the choice, which will subsequently be executed. The main section on policy-based choice includes a discussion of the various ways in which choosers can arrive at policies (see in particular 9.4).

We note here the following points about the evaluation criteria applied in attribute-based choice:

1. It is not necessary for the chooser to be able to assign some sort of value to every observed level of every known attribute. As we will see shortly, choosers can often arrive at a choice by considering a small subset of the available information about attributes and levels and by applying a small and possibly vague set of evaluation criteria (e.g., “the most important attribute is A , and the most desirable level of A is L ”).
2. It's also not realistic to expect choosers to bring to the choice situation comprehensive, stable evaluation criteria, unless they have a good deal of knowledge and experience with the domain and/or they have gone to the trouble of working out an evaluation policy.
3. A consequence of these first two points is that the chooser's evaluation criteria can evolve during the process of choosing, as she encounters more information about the options; and that her

specific evaluations can be influenced by contextual factors that would normally be considered irrelevant.

A striking example of the third point is provided by a study by Mandel and Johnson [2002]: They found that two different colored backgrounds of an experimental e-commerce website for car buyers, which suggested the values of “safety” or “economy”, respectively, led to measurably different choices of prospective buyers. Results like these underscore the importance of the support strategy *Represent the Choice Situation* and in particular of the need to represent choice situations in ways that do not introduce unintended bias (cf. 4.7.1) .

5.3 Winnowing

5.3.1 Key Concepts

In some real-life situations, the number of options to be considered is limited from the start to a small number that can be easily contemplated at the same time (e.g., 5 or 6). And many experimental studies of attribute-based choice create situations of this sort. But in many real-life situations, a chooser will start off with a set of options that is too large to permit consideration of each one. In these cases, people often use some procedure to *winnow* (Edwards and Fasolo, 2001, p. 593 ff.) the set of options to a manageable number of options, sometimes called a *consideration set*.

Consider, for example, a user who is trying to find a suitable navigation application for use in his car. The initial set of options may be returned by a query in an app store with keywords like “car, navigation”. If this search returns, say, 50 apps, the chooser is likely to start by trying to eliminate most of them.

A typical procedure is to focus on one important attribute and filter out all options whose levels are inadequately desirable with respect to that attribute. As long as the number of options that remain is still too large, the filtering is repeated with another attribute. This procedure is similar to the *elimination by aspects* strategy that Tversky [1972] introduced as a general model of choice.

A winnowing procedure is typically quicker (at least for the chooser) and less reliable than the procedure that the chooser will subsequently apply when focusing on the consideration set (see below). A common phenomenon is therefore the “winnowed-out winner”: An option that the chooser would have chosen if she had considered it carefully is eliminated the winnowing phase. For example, if elimination by aspects is used, an option that rates poorly on the most important attribute will be eliminated even if it is outstanding on all other attributes.

5.3.2 Support Tactics

A commonly applied support tactic is to have the computer take over the mechanical work of winnowing, in accordance with the strategy *Combine and Compute*:

- The application of *filters* is often supported, as in commonly found decision aids for e-commerce websites (see, e.g., Pfeiffer, 2012, especially chap. 5). A large number of variants of the filtering approach have been explored within the the widely applied paradigm of *faceted search* or *faceted browsing* (see, e.g., Tunkelang, 2009; Hearst, 2009, chap. 8). At any given point in time, the user is considering a (possibly very large) set of options that he has reached by searching or filtering. He is typically offered a number of possible ways of filtering the set down to a smaller set (e.g., by specifying a range of acceptability for a numerical attribute or one or more desired levels for an attribute that has discrete levels).
- In a roughly similar way, the user can be offered the option of *sorting* the options according to their levels on particular attributes. If the chooser then attends only to the first options in the sorted list, the effect is similar to that of filtering (cf. Pfeiffer, 2012, chap. 5).

In addition to the strategy *Combine and Compute*, tools like these apply the strategy *Access Information and Experience*, in that they access information about the options (much of which would normally not be known to the chooser).

In practice, their usefulness is sometimes reduced by limitations in the criteria that the tools allow the chooser to specify. Here are some typical problems:

- In some simple tools, the chooser can filter or sort according to only one attribute, which means that the winnowing task can be only partly accomplished if the chooser would have liked to apply other filters as well.
- The chooser may be unable to specify a filtering criterion that she would like to apply. For example, using the “color” attribute, the chooser might want to eliminate several colors while keeping several others, whereas the system allows the specification of only one desired color.

A straightforward way to improve the usefulness of filtering and sorting tools is to determine the winnowing criteria that users are likely to want to apply and ensure that these can be expressed in the user interface.

One benefit of automating winnowing in this way is that it becomes feasible for the chooser to redo the winnowing if the result does not seem very promising. For this purpose, the new consideration set should be displayed in a way that supports a quick assessment of how likely it is to yield a good choice. Faceted browsing systems normally make it easy for the user to try an alternative winnowing principle by removing one of the currently used filters and applying a new filter. A relatively new approach is to enable the user to see the results of two or more winnowing attempts at the same time, so as to judge which one(s) should be pursued further (see, e.g., Buschbeck et al., 2013).

Automatic support for winnowing does not need to be restricted to methods (like repeated filtering) that people could in principle apply themselves. In a system that supports keyword search, the chooser can use search terms to arrive at an initial consideration set (which he may then winnow further). In the extreme case, a recommender system can take over the entire winnowing step, generating a consideration set on the basis of some sort of model of the chooser’s interests. This tactic is shown in Table 5.1 as being derived from the strategy *Evaluate on*

Behalf of the Chooser, because the recommender has to perform some evaluations or comparisons on the chooser's behalf.

5.4 Choosing From a Manageable Set of Options

Once the chooser has done any winnowing, how does she choose among the available options in the consideration set, which may or may not still be fairly numerous?

Over the past few decades, several conceptions of how people handle this task have been studied. There is still no consensus about the relative prevalence of these approaches in realistic situations, but they are all worth taking into account. After all, even if a given approach does not occur frequently in unaided choice processes, it might work well if effectively supported with tactics derived from the ARCADE strategies.

A wide variety of strategies for examining and integrating information about attribute levels have been compared in the research literature (see, e.g., the influential book-length treatment by Payne et al., 1993, or the more compact summaries by Payne and Bettman, 2004, and Hastie and Dawes, 2010, chap. 10). Instead of discussing the many dimensions along which strategies can be compared, we will discuss here one broad distinction between what may be called *comparative* and *additive* strategies.²

The key distinction is whether the chooser, when contemplating the fact that a given option has a given level of a given attribute, needs to assign some sort of quantitative evaluation to that attribute-level combination (in the spirit of "I'll give this dictionary 20 points for having such a large number of words"). If so, the chooser can apply an additive strategy, adding up these quantitative evaluations to arrive at an overall quantitative evaluation of the option in question (even if not all of the attribute-related information about the option is taken into account). Even though we normally don't intuitively feel that we are adding up numbers when we form an impression of the desirability of an item on the basis of its attributes, both behavioral studies (see,

²The most closely related distinction that already occurs in literature such as the works just cited is that between *noncompensatory* and *compensatory* strategies.

e.g., Glöckner, 2008; Dijksterhuis et al., 2009) and neuropsychological studies (see, e.g., Glimcher and Fehr, 2014, chap. 13) have shown that partial evaluations are often represented and added up in the brain without conscious control.

By contrast, a *comparative* strategy presupposes only that the chooser can compare any two levels of a given attribute to determine which level is more desirable (e.g., 50,000 words in a dictionary is a better number than 40,000 words, regardless of how good these numbers are in an absolute sense); and (in some such strategies) that the chooser can judge which of two attributes is more important.

5.4.1 Comparative Strategies

A frequently studied comparative strategy is the *lexicographic strategy*.³ The chooser focuses on the one attribute that he considers most important and chooses the option with the most desirable level on that attribute; if two or more options are equally desirable, the second most important attribute is used in a similar way to break the tie. This strategy is “frugal” in the sense that it normally requires processing of information concerning only one attribute and the processing concerns only the question of which option rates best on that attribute.

The *elimination by aspects* strategy that was discussed in connection with winnowing likewise falls into the category of comparative strategies.

Many strategies of this general sort are *noncompensatory*: An option can be eliminated on the basis of its level on one attribute even though it rates very well on other attributes.

Research on these strategies (often in comparison with additive strategies) has yielded several general conclusions, which mostly make comparative strategies seem more attractive to choosers than they might appear at first glance:

1. The fact that comparative strategies typically involve less precise evaluation criteria fits well with the fact that choosers often have only incomplete, unstable evaluation criteria anyway.

³In recent literature, variants of this strategy have often gone under the name *take the best*, as in the work of Gigerenzer and Todd [1999]

2. How well a given strategy works in terms of yielding a good outcome depends strongly on the specific choice task that is involved, as well as on aspects of the context such as time pressure. In particular, even strategies that make use of only a small proportion of the available information can be surprisingly effective in some choice environments. Intuitively, it can be understood that a lexicographic strategy will work well if (a) there are large differences among the options in terms of the most important attribute and (b) the differences with respect to other attributes are small, or the other attributes are much less important. In an environment like this, a person who tries to apply a more elaborate strategy, such as an additive strategy, will at best obtain no more desirable result and may at worst run out of time before she has processed the most important information. There is a good deal of evidence (see Payne et al., 1993, and Gigerenzer and Todd, 1999) that people are *adaptive* in the sense that they have some ability to choose a strategy that is relatively well suited to the current task and context.

3. Comparative strategies often offer advantages in terms of the third desirable property of good choice procedures that was mentioned in 3.6: the avoidance of unpleasant thoughts. One type of unpleasant thinking that people often like to avoid is thinking about *trade-offs*: If Dictionary 1 has more words but poorer usability than Dictionary 2, C needs to think about which of these differences is more important. This task can be especially difficult if C has not previously thought about the exact relative importance of these attributes and the values of particular levels of them. And after having decided, C will be aware of having chosen a dictionary that has a significant drawback relative to one of the other dictionaries. With most comparative strategies, such as the lexicographic strategy, C can remain blissfully ignorant about any such trade-offs.

Another type of unpleasant thought is related to specific attribute levels that bring to mind unpleasant associations. When choosing among several foods, at least some of which have high caloric content, C may be reminded of the problems that he has keeping his weight under control; if he applies a selective comparative strategy that ignores

the attribute “number of calories”, he will be spared any such negative associations.

4. A choice made with a comparative strategy also tends to be easy to understand and justify: Compare “I chose this dictionary because it has the largest number of words” with an explanation that mentions a number of attribute-level combinations of the options and claims that the sum of these values for one particular option is greater than the corresponding sum for any other option.

5.4.2 Additive Strategies

The classic additive strategy, which is often seen as an ideal that can at best be approached by other strategies, is called *weighted adding*: The chooser assigns a value for each level of each attribute of each option and then adds up all of these values to see which option is best. When applied deliberately, this strategy would in general require a good deal of cognitive effort, and it might even be impossible to execute. But research in recent years has suggested that people can do something like weighted adding quite effectively as long as they don’t pay too much attention (see, e.g., Glöckner, 2008; Dijksterhuis et al., 2009). Still, it would often be superfluous and impractical to take into account every single attribute of every option: Some options can quickly be seen to be inferior, and some attributes are so unimportant that they can safely be ignored. Many computational models of attribute-based choice postulate the use of selective additive strategies in which the attribute/level combinations to be evaluated are *sampled* from those of all options.⁴ The sampling is typically assumed to stop once one option has emerged as being sufficiently superior to the others.

One advantage of an additive strategy is that it can be used to evaluate each option separately, whereas comparative strategies presuppose that options can be compared with each other attribute by attribute, which is normally easy only if they are perceptually present at the same time. The downside is that an additive strategy requires

⁴An influential model of this type has been Decision Field Theory (Busemeyer and Townsend, 1993), which has been extended over the past two decades (see, e.g., Roe et al., 2001; Bhatia, 2013).

the chooser to be able to assign some sort of meaningful value to individual attribute-level combinations, which may be impossible if the chooser has little relevant knowledge or experience. In a classic study, Hsee [2000] showed that nonspecialists had difficulty assigning a value to the fact that a given music dictionary had 10,000 entries (or 20,000 entries) when judging such a dictionary separately. Hence this attribute had relatively little impact on their choices when they contemplated the dictionaries separately—but it had a large impact when the dictionaries were compared with each other directly. Hsee [2000] argued that, if an attribute like this is going to figure in a noncomparative evaluation of an object, it can be important to take measures to increase the *evaluability* of the attribute—for example by providing information about the typical range of values or by using terms like “above average”.

5.4.3 Support Tactics Related to Attribute-Based Strategies

A general approach to supporting this phase of attribute-based choice can be based on the idea of helping the chooser C with the selection and application of an appropriate strategy:

1. Consider what sort of strategy is desirable for the specific choice situation:

- Is C likely to have relatively precise and comprehensive evaluation criteria or only a few vague notions about the relative importances of some attributes and the relative desirability of some of their possible levels?
- Is the domain one for which a simple comparative strategy is likely to yield satisfactory results?
- How much processing is C likely to want to engage in?
- Is it a good idea to enable C to avoid thinking about trade-offs, or is it in the best interests of C to be confronted with important trade-offs?
- How important is it for C to be able to understand and justify (to herself and/or others) the choice she has made?

2. Take measures to facilitate and encourage the application of the desirable strategy. These steps can take various forms:

- *Represent the Choice Situation:* Display the options in a way that makes it relatively easy to apply the strategy.

Comparative strategies tend to be easier to apply if the options are displayed simultaneously in tabular form, with the attributes as rows and the options as columns (or vice versa). By contrast, if each option is described on a separate screen in a paragraph of text, applying a comparative strategy could require great mental effort.

- *Combine and Compute:* Have the system perform mechanical computations that are required by the strategy. Computational support for filtering and sorting has already been mentioned. Some systems (see, e.g., Edwards and Fasolo, 2001; Pfeiffer, 2012) support a weighted additive strategy by computing the necessary weighted sums after the chooser has supplied the relevant parameters (attribute weights and values of the various attribute levels).
- *Evaluate on Behalf of the Chooser:* The system can go beyond the previous tactic by making assumptions (maybe derived using recommender system technology) about the evaluation criteria that C must have in order to apply the strategy (e.g., concerning attribute importances and the values of particular levels). In this case, the system is not just combining available information but also performing some evaluation on behalf of the chooser.

6

Consequence-Based Choice

6.1 Introduction to the Pattern

When applying the consequence-based pattern, the chooser contemplates possible actions and evaluates them in terms of their anticipated consequences.

6.1.1 Status in Scientific Research

This choice pattern is the one which, in both the popular mind and the thought of many scientists, is most closely associated with the idea of rational decision making (cf. March, 1994, chap. 1). In their influential textbook on artificial intelligence, Russell and Norvig [2010] define a *rational agent* as “one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome” (sect. 1.1.4). Payne and Howes [2013] propose for the HCI field a theoretical framework based on the idea that “people choose strategies so as to maximise utility given constraints” (p. xi). From these perspectives, all of the other five ASPECT choice patterns can be seen as indirect ways of finding an option that will yield a good outcome, which may

be helpful when the chooser does not have the information or resources that would be required to evaluate anticipated outcomes directly.

This choice pattern has understandably attracted more attention in research and practice than any of the other choice patterns. Nonetheless, there remain important gaps to be filled if we want to provide a coherent picture of consequence-based choice that will be useful in HCI contexts.

6.1.2 Complications Relative to the Gambling Paradigm

To see why, let's consider the most prominent single descriptive theory of consequence-based choice: *prospect theory* (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Wakker, 2010). This theory is based on numerous studies that were conducted within the *gambling paradigm*. A typical choice presented to participants in such a study is the following one (adapted from Problem 12 of Kahneman and Tversky, 1979, p. 273):

“In addition to whatever you own, you have been given \$2,000. You are now asked to choose between:

- a 50% chance of having to give back \$1000; and
- the certainty of having to give back \$500.”

Everyday choices made by users of computing technology differ from this type of choice in important ways:

1. In the gambling paradigm, the chooser is presented with a small set of explicitly described options and is asked to choose among them.

By contrast, users of computing technology are often not even aware that there is a choice that can be made; and if they do, they may not know what options are available to them.

2. In the gambling paradigm, the consequences are normally a limited number of hypothetical financial gains or losses that can be quantified in a natural way.

By contrast, the consequences of users' everyday choices range from the subjective experience of performing an action to a variety of types of consequence in a dynamically changing physical and social world. And there are often so many different consequences of a given action that it is impractical for a chooser to consider all of them; instead, he must in some way consider a subset of the consequences, much as the chooser in the attribute-based pattern typically considers only a subset of the potentially relevant attribute/level pairs (5.4).

3. Although many of the outcomes in the gambling paradigm are not expected to occur with certainty, they are at least explicitly associated with numerical probabilities.

By contrast, when a user of computing technology is not sure what will happen, she in general has no straightforward way of quantifying her uncertainty.¹ Moreover, some of these consequences may be expected to occur at some time in the more or less distant future, a fact which has implications for the weight that they receive in the choice process.

These gaps can be filled in part with the help of research from paradigms other than the gambling paradigm. But the account in this section will necessarily remain sketchy at some points, leaving opportunities for others to flesh it out in the future.

The first four steps shown in Table 6.1 belong to what is sometimes called the *preselectional phase* of choice or decision making (see, e.g., Betsch and Haberstroh, 2005b). They occur before the chooser has finished identifying a set of options whose consequences are to be anticipated and evaluated.²

¹Discussions of the distinction between precisely quantified *risk* and subjectively experienced *uncertainty* are given, for example, by Fox and Tversky [1998], Gigerenzer [2014], and Wakker [2010].

²The preselectional phase can also occur with the other choice patterns, but adding it explicitly to the descriptions of other patterns would produce redundancy in the set of ASPECT patterns. We discuss it as part of the consequence-based pattern because the steps in this phase are most naturally illustrated with cases in which a person is choosing among possible actions that will have consequences.

Table 6.1: Consequence-Based Choice.

Steps	Support Tactics
– <i>C</i> recognizes that a choice about a possible action can (or must) be made	<p><i>Advise About Processing:</i> Alert <i>C</i> to the opportunity or need to make a choice</p> <p><i>Design the Domain:</i> Ensure that, if the choice opportunity is not recognized, what happens by default will be acceptable</p>
– <i>C</i> assesses the situation	<p><i>Access Information and Experience, Represent the Choice Situation:</i> Provide information about the current situation in a form that makes it easy to see what is going on and what is likely to happen</p> <p><i>Combine and Compute:</i> Compute diagnoses or predictions about the situation</p>
– <i>C</i> decides when and where to make the choice	<p><i>Design the Domain:</i> Avoid forcing a choice to be made in an unfavorable situation</p>
– <i>C</i> identifies one or more possible actions (options)	<p><i>Access Information and Experience:</i> Provide explicit information about possible actions</p> <p><i>Represent the Choice Situation:</i> Represent the situation in such a way that the possible actions are easy to recognize or infer</p>
– <i>C</i> anticipates (some of) the consequences of executing the options	
<i>C</i> anticipates objective consequences	<p><i>Access Information and Experience, Represent the Choice Situation:</i> Conveniently represent information about likely consequences</p> <p><i>Design the Domain:</i> Make consequences inherently easy to anticipate</p>
<i>C</i> anticipates the experience of executing particular actions	<p><i>Access Information and Experience:</i> Enable <i>C</i> to sample the relevant experience</p>
<i>C</i> anticipates possible affective responses concerning the decision process	<p><i>Advise About Processing:</i> Alert to possible negative affective responses (e.g., disappointment, regret)</p>
– <i>C</i> evaluates (some of) the anticipated consequences	
... temporally distant consequences	<p><i>Design the Domain:</i> Avoid cases where the consequences of different options lie at different temporal distances</p> <p><i>Represent the Choice Situation:</i> Make the distant consequences salient enough to be given due attention</p>
... uncertain consequences	<p><i>Represent the Choice Situation:</i> Depict the consequences in a way that suggests an appropriate weighting</p>
– <i>C</i> chooses an option that rates (relatively) well in terms of its consequences	<p><i>Represent the Choice Situation:</i> Represent possible consequences and/or their evaluations in a way that facilitates application of an appropriate integration strategy (cf. integration strategies for attribute-based choice)</p>

6.2 Recognizing That There Is a Choice Opportunity

Here are two examples of how a computer user can recognize a choice opportunity:

- *Method selection*: The user employs a particular method to achieve a goal (e.g., choosing a command from a menu) without being aware that there exists at least one other method for achieving the same goal (e.g., using a keyboard shortcut).

One common support approach is to apply the strategies *Access Information and Experience* and *Represent the Choice Situation* by “advertising” the available methods (e.g., by displaying a visual representation of the keyboard shortcut next to the relevant menu option). This tactic is unfortunately more difficult to apply with some types of interactive system, such as those that are embedded (largely) invisibly in the environment (cf. Bellotti et al., 2002).

- *Configuration choice*: When a user is operating a system that can be configured in some way, he in principle could choose at almost any time whether he wants to change the configuration in an attempt to improve the system’s behavior. But it is in general difficult for him to know when it is worthwhile even to consider such a configuration change.

A designer can apply the strategies *Evaluate on Behalf of the Chooser* and *Advise About Processing* by considering at what point in the operation of the system a user might be interested in changing a configuration setting and alerting her to the choice opportunity. Figure 6.1 shows an example of this approach, in which the user is also offered information about the consequences of the two available options (turning on Wi-Fi vs. leaving it off). Note that the user is being offered, in the context of her navigation app, essentially the same information that is available on the “Settings” screen that was shown in Figure 4.2.

A related approach is to offer the user an opportunity to change a specific setting without having to visit a different part of the




Figure 6.1: Example in which a system alerts a user to an opportunity to make a configuration choice. (The user has just launched the TOMTOM navigation app, which uses the iPhone’s built-in GPS.)

system, as by offering the option “Don’t ask me again” next to a regularly appearing query that the user might find superfluous.

In addition to (or instead of) trying to alert users to a choice opportunity, interaction designers often find it more appropriate to (a) accept the fact that many users will not recognize the opportunity and (b) ensure that whatever option gets chosen *by default* is likely to be acceptable to the user, if not optimal. It is in fact a widely accepted HCI design guideline that default options should be chosen carefully by designers. This approach instantiates both the strategy *Evaluate on Behalf of the Chooser*—because the designer is thinking about the likely desirability of particular options for the users—and the strategy

Design the Domain—because the way in which the system behaves in a particular circumstance is being designed.

It is worth bearing in mind that there are reasons why users are likely to (implicitly) choose the default option even if they are aware of alternative options:

- Accepting the default option often requires less effort than choosing any other one.
- The user is often justified in assuming that the default option has been selected by the designers because it is appropriate in most cases.

Outside of the domain of computing technology, the judicious design of default options constitutes one of the primary and most successful tactics of interventions based on behavioral economics (see, e.g., Thaler and Sunstein, 2008; Johnson et al., 2012). For example, laws that state that every person can be viewed as an organ donor unless he has specified otherwise yield dramatically higher donation rates than laws according to which not donating is the default option (Johnson and Goldstein, 2003).³

6.3 Situation Assessment

6.3.1 The General Concept

In cases where it is important to anticipate the consequences of actions, it is often important to have some understanding of what the current situation is and—if the situation is changing—how it is likely to become in the future. After all, these facts will help to determine what the consequences of actions will be. Outside of the domain of computing technology, clear examples where this type of *situation assessment* (or *situation awareness*) is important include (a) choices about medical

³Appendix D of the *2010 Report to Congress on the Benefits and Costs of Federal Regulations and Unfunded Mandates on State, Local, and Tribal Entities*, by the Office of Information and Regulatory Affairs of the U.S. Office of Management and Budget, includes a detailed set of prescriptions about how to use defaults effectively for the purpose of choice support.

treatments, since the consequences of any treatment will depend on the patient's current state and prognosis; and (b) choices about emergency situations such as those faced by firefighters.

Situation assessment has been extensively studied in the research paradigm of *naturalistic decision making* (see, e.g., Klein, 1998; Maule, 2010), and in the field of human factors engineering (see, e.g., Guastello, 2014). Decision support systems (see, e.g., French et al., 2009) typically include powerful methods for retrieving relevant information about the decision situation (applying the strategy *Access Information and Experience*); performing computations on such information to yield summaries, diagnoses, and predictions (*Combine and Compute*); and presenting the results to decision makers in a way that facilitates the necessary thinking (*Represent the Choice Situation*).

6.3.2 Examples Involving Computing Technology

In connection with the everyday use of computing technology, situations tend on the whole to be less complex and rapidly changing than in the domains just mentioned. Also, users are less able to devote attention to heavyweight situation assessment tools. Still, there are many cases where it is worthwhile to consider how situation assessment can be supported, such as the following:

- A mobile user is wondering whether to dictate an email message using the iPhone's SIRI agent or to type it in with a virtual keyboard. She may want to consider situational factors such as how favorable the current conditions for successful dictation are (e.g., in terms of ambient noise, movement by the user, and the quality of the available microphone).
- A driver is wondering whether and how to respond to a voice message that he has received on his smartphone. He should take into account factors such as the current traffic situation, the state of the road ahead of him, and the urgency of the message (i.e., what will happen if he does not respond to it right now).
- A user of a social networking site is considering communicating some sensitive personal information. She ought to take into ac-

count features of the situation that influence the likelihood of various types of privacy violation (see 12.2.1).

6.3.3 Support Opportunities

The discussion of the last example just mentioned in the section on privacy-related choices (12.2.1) illustrates some of the typical approaches and issues in this area. Although modern computing technology offers many opportunities for supporting situation assessment, the following general points should be borne in mind:

- The strategy *Access Information and Experience* needs to be combined with the strategies *Combine and Compute* and *Represent the Choice Situation*.

The amount of more or less relevant information that can be presented about a given situation typically exceeds the capacity and motivation of users to attend to and make sense of it. Think, for example, of the information that is presented in a website's privacy notice or in the "blue screen of death" that appears when your computer has just crashed. It is therefore in general necessary to determine carefully what sort of representation will really support a relevant situation assessment and to apply whatever techniques of information retrieval, computation, and representation are needed to create it.

- Even a clear understanding of the current situation will not necessarily enable the chooser to make a good choice. Support for the other steps in the consequence-based pattern, such as generating options and anticipating consequences, is likely also to be required.

6.4 Deciding When to Choose

One of the less-studied approaches to supporting choice is to enable the chooser to make the choice at a favorable time and in a favorable situation. This approach is especially worthy of attention in the HCI context, because designers of interactive systems can often strongly

influence when choice opportunities arise and when they have to be dealt with.

Here are some reasons why one situation (“Situation 2”), may be more favorable for making a given choice than another situation (“Situation 1”):

1. In Situation 2, the user will have had a chance to acquire more relevant knowledge of the choice situation.

Suppose, to take an extreme case, that a user of a new application is forced to decide about all of the configuration settings before starting to use the application. He is likely to have neither the required knowledge nor enough interest to make good choices. For this reason, Lederer et al. [2004] argued that users ought to be enabled to set privacy-relevant parameters on the fly, in the context of their use of the system, instead of in a separate configuration session (cf. the discussion in 12.5.2).

2. In Situation 2, contextual factors such as the amount of available time and the extent of distraction (e.g., by noise or by competing tasks) are more favorable than in Situation 1.

For this reason, a driver may be well-advised to decide, before getting into her car, what sort of electronic communication he will engage in while driving. Outside of the car, she is likely to be in a better position to consider all of the available options and to take into account possible consequences such as having an accident. (Note that the previous consideration would by itself suggest the opposite conclusion.)

To be sure, the question of what contextual factors are most favorable for the choice process is in general more complex than it may seem at first glance, partly because of the possibility of applying choice patterns other than the consequence-based pattern, such as the experience-based pattern. For example, Klein [1998, chap. 16] argues that the presence of factors that cause stress does not have as much impact as we might expect on the choices of decision makers like experienced firefighters. Dijksterhuis and colleagues have reported on counterintuitive results of studies in

which people appeared to make better choices when they were distracted (see, e.g., Dijksterhuis et al. [2006]).

Interaction designers often have an opportunity to apply the strategy *Design the Domain* either (a) to ensure that the situation in which the chooser has to make the choice is favorable or (b) to give the chooser freedom to decide when and where to make the choice. Examples of methods enabling on-the-fly configuration were given above in 6.2. The discussion of *commitment mechanisms* below in 9.6 will suggest how it is possible to enable a user realistically to make a choice long before it is time to execute the chosen option.

6.5 Identification of Options

In order to apply the consequence-based pattern, the chooser needs to be aware of at least one possible action whose consequences can be anticipated. Identifying possible options can take various forms, depending on the nature of the available options. We can distinguish four types of option space:

6.5.1 Small Number of Predefined Options

This situation often arises in HCI contexts, since artificial interactive systems often offer the user a small, predetermined set of possible actions. For example, a user of a flight reservation website will probably have no more than three methods available for entering the desired departure date. The issues involved in making users aware of the existence of (some of) these options were discussed above (6.2) in connection with the problem of recognizing the existence of a choice opportunity.

6.5.2 Large Number of Predefined Options

In connection with the attribute-based pattern, we discussed situations in which a user is initially confronted with a very large number of options, such as products in a store or documents in a library. We noticed that a typical first step is to winnow the options down to a manageable subset.

A similar situation can arise within the consequence-based pattern even where there is only one type of action that the user can perform, if this action can be applied to any of a large number of objects. For example, a user who feels like proofreading a WIKIPEDIA article can in general choose from a huge number of articles to work on (cf. 11.3). And a member of a social networking site who wants to share some content with friends and acquaintances can choose any of the numerous *subsets* of his contacts as targets of his sharing action (cf. 12.3). As the discussion of these two cases in the later chapters will show, it can be natural to represent the possible objects of the user's actions in terms of their attributes, to enable the sort of winnowing that we saw in connection with the attribute-based pattern. Once the number of possible objects of the action has become manageable, the chooser can switch back to the consequence-based pattern, anticipating and evaluating the consequences of applying the action to each of the possible objects.

6.5.3 Open-Ended Set of Options

In this situation, the set of possible options is open and in principle unlimited, because the chooser can generate any number of actions. Examples involving the use of computing technology include (a) composing a lengthy contribution to an online community forum and (b) designing your web homepage. Generating options in this situation can involve problem solving, creativity, and domain-specific methods (e.g., methods for generating possible web page designs)—topics which fall largely outside of the scope of choice architecture for HCI.

One general theme from research that is relevant in all of these cases is the concept of *value-focused thinking*. When introducing this concept, Keeney [1992] argued that all too often decision making is *alternative-focused*: The chooser focuses on a limited set of options that have come to her attention somehow (e.g., job offers that a college student has received shortly before graduating), even though many other options are in principle available (e.g., job opportunities that the student could find if she searched for them actively and creatively). The problem is essentially that the chooser's evaluation criteria are being

used only in the phase of evaluating anticipated consequences (which will be discussed below). If these criteria were used also to generate options in the first place, the generated options would tend to have greater overall evaluations for the chooser.

Though the techniques for supporting value-focused thinking which Keeney presents in his book are mostly too elaborate for the support of everyday decisions about computing technology, a few simple strategies are worth considering:

- *Represent the Choice Situation*: Represent the choice situation in such a way that it is clear that the set of options is open-ended and that it is not necessary to restrict attention to any preselected or obvious options.
- *Advise About Processing*: Encourage the chooser to think about his currently relevant goals and values while generating options.

The “Guidelines” page⁴ of the community weblog METAFILTER (also discussed in Section 11) can be seen as encouraging value-focused thinking in these ways: To a large extent, it *prescribes* to users particular values that should guide the generation of posts to the weblog, under headings like “What makes a good thread post to METAFILTER?”. But since a lot of different possible goals of posting and commenting are mentioned, the effect can be to encourage users to consider what their own goals are before they start generating options for action.

6.6 Anticipation of Consequences

We now turn to the question of how to help users of computing technology anticipate the consequences of a possible action. One complication here is the fact that there exist qualitatively different types of consequence that can be taken into account, which can be illustrated with the example of a driver who is considering listening to and dictating email messages while driving, using a smartphone app that she has not tried before.

⁴<http://www.metafilter.com/guidelines.mefi>

- The significant objective consequences of executing the action: How long will it take? Will the result be satisfactory? Might I cause an accident because of being distracted by this task?
- The experience of executing the action and the experiences brought about by the consequences: What will it feel like to dictate the message? Will it be fun, stressful, . . . ? How will I feel if I manage to respond to all of the messages in my inbox?
- Feelings about the choice process: How will I feel about my decision to dictate messages if it leads to an accident?

We will consider each of these types of consequence in turn.

6.6.1 Objective Consequences

Users' actions typically have some consequences that can be evaluated largely independently of the associated experiences: the extent to which the goal of the action was achieved, the costs of executing the action (e.g., the time taken), and side effects.

Straightforward Provision of Information About Consequences

As is the case with situation assessment (6.3), decision support systems for supporting complex decisions often apply sophisticated combinations of the first three ARCADE strategies to help decision makers predict the consequences of actions (see again French et al., 2009). Analogous tactics can be applied on a smaller scale to everyday choices about computing technology.

The strategy *Access Information and Experience* can often be applied straightforwardly with text messages (e.g., “Estimated download time: 3 minutes”) or graphical means such as a familiar icon indicating that clicking on a given link will take the user to an external website. The FATBATT application for managing a laptop's battery⁵ features an indication of the impact that dimming the screen will have on the amount of time that the user can work before having to recharge the

⁵<https://fatbatt.com>

battery; it thereby helps the user to decide whether it's worthwhile to dim his screen at any given moment.

It remains to be seen to what extent more sophisticated combinations of the first three ARCADE strategies such as those that are found in decision support systems will be developed to support consequence anticipation in connection with everyday choices. Section 12, on privacy-related choices, provides several examples of a trend in this direction.

Designing to Make Consequences Inherently Easy to Anticipate

The general HCI principle that an interactive system should be designed so that its behavior is simple and predictable can be seen as an application of the strategy *Design the Domain*: In contrast to the real physical world, where the consequences of an action are largely outside of any designer's control, the interaction designer can often create an artificial world in which consequences are inherently easy to anticipate and evaluate. Typical ways of achieving this goal are adherence to familiar design conventions, use of easily understandable metaphors, and use of other means to support the formation of accurate mental models.

6.6.2 Experience of Executing an Option

In some cases, as with games and entertainment, the experience of executing an option is the most important consequence and the reason for choosing the option in the first place. With more utilitarian tasks, such as booking a hotel room via an e-commerce website, the experience can still be an important consequence that influences users' willingness to choose an option—a fact that underlies the growing emphasis in recent years on the notion of *user experience* in HCI. Also, after an option has been executed, its consequences can be experienced.

There is a line of research in psychology about the problems that people have in anticipating what it will be like to experience something if they haven't experienced it in the past. One complication is caused by people's tendency to adapt their tastes and expectations on the basis of new experience (see, e.g., Wilson, 2002, chap. 7; Gilbert et al., 2002;

Schwartz, 2004, chap. 8); this adaptation is in turn hard for them to predict. For example, you may initially be amused by the cute cartoon character jumping around in the corner of the screen; but how are you going to feel about it half an hour from now?

Even if the chooser has already had the experience in question, it is not always straightforward for her to anticipate what the experience will be like the next time. Kahneman and colleagues (see, e.g., Kahneman, 2011, chap. 35) famously studied people who were subjected to various sorts of unpleasant experience, such as holding their hand in painfully cold water. They found that there was no one-to-one relationship between the amount of pain a person actually experienced and his later recollection of that pain: In some cases participants, when given a choice about which of two experiences to repeat, chose the one which had certainly caused more pain.

Several approaches to supporting the anticipation of experience are already in more or less wide use:

1. In some cases, a straightforward form of support is to apply the strategy *Access Information and Experience* by providing a sample of the experience in question, allowing the user to try out the action before she decides whether to go through with it. This form of support amounts to enabling the small-scale, local application of the trial-and-error-based pattern. But often the experience in question is not available to be offered on a trial basis (e.g., the experience of being admired by the members of your social network after you have made a valuable contribution; or the experience of being a skillful touch typist, if you are currently a novice typist). And even if it is, there may not be time for effects such as adaptation to the new experience to set in.
2. Another variant of the strategy *Access Information and Experience*, which has been popular in the persuasive technology area, involves using simulations to give a foretaste of experiences. A well-known example that does not involve choices about computing technology is the BABY THINK IT OVER infant simulator, which helps teen-aged girls anticipate realistically what it is like to take care of a baby (see, e.g., Fogg, 2003, chap. 4). Another

example is the airplane emergency simulator of Chittaro [2012] that was mentioned above in 4.1.3.

3. A less technically complex method of helping a chooser to anticipate experiences is the use of *narratives*: stories about the experience of other persons (Hibbard and Peters, 2003; Klein, 1998, chap. 11). Such narratives can evoke affective responses and stimulate processing on the intuitive level (3.4). A familiar means of providing such narratives is the provision of users' reviews and testimonials. A single harrowing description of another user's attempt to install a newly purchased software application—and of its disappointing performance once installation was complete—can enable a chooser to imagine vividly what the consequences of purchasing that application might be, including the feeling of regret about having chosen to make the purchase.
4. If the chooser has had similar experiences in the past, the application of the strategy *Access Information and Experience* can involve reminding him, maybe by means like those just listed, of what the experience was like, so as to overcome the tendency for recollections of experience to be distorted in memory.

6.6.3 Affective Responses to the Choice Process

In addition to the affective states that are caused by experience with the consequences of an action, there are often affective states that constitute responses to the choice process on a meta-level (see, e.g., Schwartz, 2004, chap. 7; Mellers, 2000; Isen and Labroo, 2003).

- *Disappointment*: The selected option is found to be not as satisfying as the chooser had expected (maybe because of the sort of unexpected adaptation to the new benefit that was mentioned above).

A chooser who is sophisticated enough to anticipate this possibility might think, for example: “If I pay for this expensive app and it turns out not to be so great, I’ll be disappointed”.

- *Regret*: The chooser ultimately concludes that a different option would have been better than the one that she chose.

A chooser can anticipate this result by thinking something like: “If I buy this app without checking out the alternatives, I’ll kick myself if I find out later that one of the alternatives is much better.”

There appears to have been little research about how to help choosers anticipate affective responses of the sort, but one natural approach is to apply the strategy *Advise About Processing* by somehow encouraging the chooser to think about how he will feel if particular outcomes occur. And simply applying the strategy *Access Information and Experience* to enable the chooser to imagine particular consequences vividly can have a similar effect. Schwartz [2004, p. 231] takes a different approach to choice support with regard to regret: He aims to reduce the extent to which regret occurs in the first place by applying the strategy *Advise About Processing*—recommending, for example, “Adopting the standards of a satisficer rather than a maximizer” and “Reducing the number of options we consider before making a decision”. Though this advice is aimed directly at readers of Schwartz’s popular book, imaginative choice architects who aim to minimize choosers’ regret may be able to think of ways of achieving the same effects via tactics other than explicit advice-giving.

6.7 Evaluation of Anticipated Consequences

We have considered so far how users can be helped to anticipate the consequences of their actions, but what about the ways in which they evaluate these consequences?

6.7.1 Reference Levels

Even within prospect theory (6.1.2), where consequences are assumed to be quantified in terms of money, the process of evaluating consequences is far from trivial. One key insight, which has been found to apply even in many situations in which the consequences are not so naturally quantified, is that outcomes are evaluated relative to a *reference*

level (see, e.g., Kahneman and Tversky, 1979, p. 274). For example, consider a user who is accustomed to filling in a travel expense report in about 5 minutes. If his company has installed a new expense reporting system that will require (at least at first) 10 minutes of work, then the user is likely to view the use of the new system as causing a loss of 5 minutes of her time (in addition to the less quantifiable disadvantage of having to pay more attention while learning to operate the new application). By contrast, another user who is accustomed to spending 15 minutes on this type of task will view the new system as saving him 5 minutes.

One reason why this use of reference levels is important is that “losses loom larger than corresponding gains” (see, e.g., Tversky and Kahneman, 1991). All other things being equal, the “execution time” consequence will have a bigger impact on the choice of the user who loses time than on that of the user who saves time. This tendency, called *loss aversion*, has been invoked to explain a variety of phenomena in people’s money-related behavior (see, e.g., Tversky and Kahneman, 1991). In the HCI area, one implication is the prediction that users will tend to stick with applications and methods that they already use: If switching to another option would lead to some sort of loss and some objectively equally large gain, the loss will tend to have greater weight in the choice process.

It is not always obvious what the relevant reference level is. In the previous example, the user might conceivably compare her expected time for the first use of the new system with the time she spent long ago when using the previous system for the first time. In principle, each chooser can decide for himself which reference level to employ, but there is often a reference level that is naturally suggested by the way in which the problem presents itself to the chooser—for example, by the verbal formulation of the choice problem.

6.7.2 Framing

Research Results

This sort of influence has been studied in laboratory research under the heading of *framing* (see, e.g., Kahneman and Tversky, 1984; Tversky and Kahneman, 1986). To cite the best-known example, people tend to be influenced strongly by whether the consequences of emergency-intervention options are described in terms of people being “saved” vs. people “dying”, even when the situations described in these terms are objectively identical (Kahneman and Tversky, 1984). To take another example, research results suggest that a person is more likely to join a dating website if she is told that she has a 40% chance of finding a partner than if she is told that she has a 60% chance of *failing* to find a partner.

Although the best-known framing effects concern the distinction between gains and losses, a more general point is that the way in which a consequence is evaluated can be strongly influenced by the way in which it is represented—a fact which underscores the importance of the ARCADE strategy *Represent the Choice Situation*.

Implications of Framing for Choice Support

When the choice architect’s goal is to persuade the chooser to select a particular option, a natural tactic is to describe the options in ways that influence the evaluations in the desired direction. When the goal is not persuasion but choice support (1.2.2), the idea is to avoid introducing unnecessary distortion via framing effects—even though it is in general difficult or impossible to find a completely neutral framing of any consequence (cf. 4.7.1). For example, though it’s debatable whether it’s more appropriate to characterize a dating site in terms of its success rate or its failure rate, it is clear that two competing dating sites should be described in the same terms if bias is to be avoided. When options are being evaluated singly rather than jointly (cf. 5.4), attempts to achieve neutrality normally have to be more cumbersome. For example, Kahneman [2011] reports on a U.S. Department of Agriculture guideline which . . .

... permits the inclusion of messages such as “90% fat-free” on the label of meat products, provided that the statement “10% fat” is also displayed “contiguous to, in lettering of the same color, size, and type as, and on the same color background as, the statement of lean percentage.” (p. 414)

6.8 Time Discounting

6.8.1 Research Results

Humans and animals alike tend to prefer a benefit that will come soon to an equal benefit that will occur later in time. That is, they *discount* future benefits.⁶ For example, a member of an online community may be more willing to make a contribution if it appears on a web page immediately, so that its positive consequences (which can take various forms) occur without delay. As is the case with monetary investments, there are various good reasons to discount temporally distant benefits. When the choice architect’s goal is persuasion, there is a straightforward design implication: To encourage a user to choose a particular option, arrange for its benefits to come sooner rather than later (an application of the strategy *Design the Domain*). This tactic was applied by McDowell et al. [2003] to encourage nontechnical users to annotate HTML data for semantic web services.

Time discounting is often due at least in part to uncertainty about whether and when the long-term benefit will really occur (see, e.g., McGuire and Kable, 2013). Accordingly, another strategy for avoiding undue time discounting is to find a way of eliminating such uncertainty, for example simply by reassuring the user about the time at which the future benefit will occur.

Some of the most striking cases where time discounting plays a role concern the frequent situation in which a person can choose between (a) an option that will bring a small benefit soon and (b) another option that will yield a larger benefit at a later time. If people’s discount

⁶Useful collections of articles on phenomena that arise when choices and/or their consequences are distributed over time have been edited by Loewenstein and Elster [1992] and by Loewenstein et al. [2003].




Figure 6.2: Exponential time discounting functions for a “smaller-sooner” and a “larger-later” benefit. (Each of the vertical line segments on the right represents the value of a benefit at the point in time at which it occurs. Each curve represents the discounted value of the anticipated benefit at an earlier point in time.)




Figure 6.3: Hyperbolic time discount functions that cross. (Compare with Figure 6.2. The larger-later benefit is preferred until shortly before the time at which the smaller-sooner benefit will occur.)

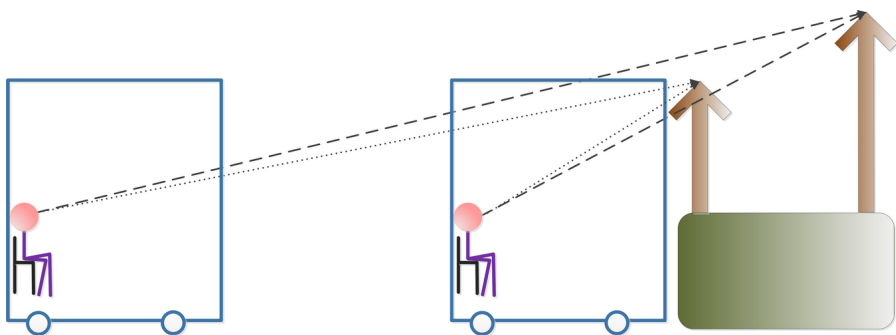


Figure 6.4: Visualization of a perceptual phenomenon that is analogous to the hyperbolic discounting of benefits that are different distances away in time. (When the perceiver is far away from the two trees, the taller tree looks taller; when the perceiver is close, the shorter tree does. The heights of the points at which the dotted lines intersect the “windshield” of the vehicle correspond to the values of a hyperbolic time discounting function. This visualization was suggested by Figure 2.3 of Rachlin, 2000.)

curves were exponential—as is the case with typical discount rates for financial investments and in early normative models of time discounting (see, e.g., Read, 2004)—, people would always show *time consistency* in their preference between the smaller-sooner and the larger-later option: If, when asked on Monday, you prefer a larger benefit on Saturday afternoon to a smaller benefit on Friday afternoon, then you will express the same preference on Friday morning. As is illustrated in Figure 6.2, the discounting curves in question never cross.

A large number of studies with humans and animals have shown, however, that discounting curves are better described by a hyperbolic function (Figure 6.3) than by an exponential one. One implication of the mathematical form of a hyperbolic function (see, e.g., Read, 2004) is that the curves in a problem like the one we are considering can cross. Concretely, in our example, when Friday morning arrives and the small benefit could be obtained almost immediately, the chooser may change his mind and opt for the smaller benefit after all. This particular type of *preference reversal* has been documented countless times in studies with animals (e.g., pigeons) and humans (see, e.g.,

Rachlin, 2000, chap. 2), and it corresponds to our everyday experience that benefits that are tangibly near can loom disproportionately large (see Figure 6.4 for a visualization that aims to make this type of reversal of relative prominence plausible and memorable).

In some cases, such preference reversals can be explained in more concrete ways than in terms of hyperbolic discounting: If a researcher working on a journal manuscript receives an email containing some amusing jokes, the action of continuing to read the jokes will have the concrete advantage that its rewards are already being experienced right now, whereas the benefits of sticking to the article writing can only be imagined.

6.8.2 Support Tactics

People are often aware of the danger of a last-minute preference reversal and are willing to avoid it by *committing* themselves at an early point in time to the option with the larger-later benefit (Rachlin, 2000, chap. 3). Hence commitment strategies are an important support mechanism in connection with choices where time discounting is a major issue. Since the issue of commitment arises even more prominently in connection with the policy-based pattern, its discussion will be postponed to Section 9.

6.9 Dealing With Uncertainty

6.9.1 Relevant Research

The question of how choosers deal with uncertainty about possible outcomes is one of the most thoroughly treated topics in research on the consequence-based pattern (see, e.g., Hastie and Dawes, 2010; Kahneman et al., 1982). For example, prospect theory, which assumes that the probabilities are provided explicitly as numbers, improved on normative models by taking into account some observed regularities in the ways in which people process information about probabilities (e.g., the fact that the difference between a 0.00% probability of having an accident and a 0.01% probability is more important than the difference between a 1.00% probability and a 1.01% probability).

A research paradigm of more recent origin (see Hertwig et al., 2006, and 10.5.4 below) investigates how choosers process probabilities that they assess by observing outcomes of multiple choices, which is a more typical situation in the HCI context. The results contrast in part with those from the gambling paradigm. For example, very low probabilities are often in effect rounded down to .00 simply because the chooser has not observed any positive example.

Overall, it is difficult to derive clear implications of research results about dealing with uncertainty for everyday situations in which a user of computing technology is dealing with a wide variety of consequences, having only vague impressions of how likely they are—impressions which she may have arrived at in various ways. The following statements can serve as rough summaries:

1. All other things being equal, the more likely a consequence is seen as being, the greater weight it will have in the chooser's evaluation of an action.
2. Outcomes with low probability and/or low importance are often simply ignored.
3. Instead of systematically taking into account all consequences of all options, choosers often apply integration strategies like those discussed in connection with the attribute-based pattern (e.g., comparing several alternative methods only in terms of their expected execution time, ignoring all other consequences unless two methods are expected to have equal execution times).

6.9.2 Implications for Choice Support

Researchers and practitioners in the areas of decision analysis and decision support systems (see, e.g., Clemen and Reilly, 2013; French et al., 2009, chap. 8) have devised a vast and impressive arsenal of ways of representing complex decision problems, eliciting realistic probability judgments, and representing the uncertain outcomes of actions in ways that help decision makers to evaluate competing options. Currently, these techniques tend to require more effort and attention than users

of computing technology are likely to want to invest in everyday situations. Maybe more lightweight versions of these methods can be developed that will fit better into a choice architecture for HCI. In the meantime, simple tactics like the following may be worth considering:

1. Sometimes, a system or its designer will have access to reliable information about the probabilities of particular consequences (e.g., the probabilities of hard disk failure for different hard disks). In these cases, it is natural to apply the strategy *Access Information and Experience* by stating the probabilities; but even if this is done, it is probably worthwhile to provide a representation of the problem situation that will make it easy for the chooser to assign appropriate mental weights to the consequences in question; providing the correct numbers cannot be assumed to suffice in itself.

2. Where no precise probability information is available, means such as graphical representations or verbal formulations can be applied within the strategy *Represent the Choice Situation* in an effort to encourage an appropriate weighting of the various consequences.

7

Experience-Based Choice

7.1 Introduction to the Pattern

The two choice patterns that have been considered so far can in principle be applied even if the chooser has never made a similar choice in the past—though in this case the chooser will have more difficulty in making the required assessments and evaluations, being more reliant on external sources of information. In cases where a chooser has repeatedly made similar choices in the past, his past experience can serve as a valuable resource that makes it less necessary to consider things like attribute levels and anticipated consequences. The idea is simply that what you have done in similar situations in the past (with some success) is likely to be a decent option now.

This idea appears in many different forms, and it has been studied in various lines of research. In this section, we consider four subpatterns of the experience-based pattern (see Table 7.1) which together cover most of the relevant phenomena.

- With *recognition-primed decision making*, the chooser retrieves from memory an option that she has used in a previous similar situation.

Table 7.1: Experience-Based Choice.

Steps	Support Tactics
– C applies recognition-primed decision making	
C recognizes the situation as one that has been encountered in the past	<i>Represent the Choice Situation:</i> Ensure that the current situation is recognizably similar to relevant previous situations <i>Access Information and Experience:</i> Remind C of similar past situations
C's experience suggests an option that C has applied successfully in the past	<i>Access Information and Experience:</i> Remind C of what C has done in similar situations in the past (and whether the choices led to success)
C chooses this option	
– or C acts on the basis of a habit	
C chooses the action that is triggered by this type of situation	
(If the associations between triggers and actions are well adapted:)	<i>Design the Domain:</i> Increase the occurrence of triggering situations <i>Represent the Choice Situation:</i> Make situations similar to the usual triggers <i>Evaluate on Behalf of the Chooser:</i> Propose or automatically execute actions that have been triggered by this situation
(If the associations between triggers and actions are not well adapted:)	<i>Design the Domain:</i> Reduce the occurrence of triggering situations <i>Represent the Choice Situation:</i> Make situations dissimilar to the usual triggers <i>Advise About Processing, Represent the Choice Situation, Design the Domain:</i> Encourage C to apply a different choice pattern
– or C chooses a previously reinforced response	
C chooses the action that has been most effectively rewarded in the past in such situations	<i>Design the Domain:</i> Prevent consequences that unduly reward basically undesirable options; and vice versa <i>Design the Domain, Represent the Choice Situation:</i> Try to ensure that situations in which actions have similar consequences resemble each other
– or C applies the affect heuristic	
C generates one or more options	
C chooses an option that evokes positive affect	<i>Advise About Processing:</i> Encourage C to try to distinguish between relevant and irrelevant affective responses to options <i>Represent the Choice Situation:</i> Present options in such a way that they evoke (only) choice-relevant feelings <i>Design the Domain, Represent the Choice Situation:</i> Avoid allowing options to become associated with feelings that do not correspond to actual value for C

- With *habit-based choice*, a frequently performed action is triggered by particular features of a situation.
- With *reinforcement-based choice*, an action is chosen because in the past it has had consequences that were in some way rewarding to the chooser, even if he is unaware of the influence of these consequences on his choices.
- With *affect-based choice*, an action is chosen which evokes positive affective associations, which may or may not reflect relevant positive experiences from the chooser's past.

Just as the six top-level ASPECT choice patterns often occur in combination, the same is true of these subpatterns. For example, a teenager's choice to write text messages while driving can be triggered by particular situations, reinforced by particular consequences, and associated with positive affect. Still, it is helpful to consider each subpattern separately, since each one is associated with characteristic phenomena and support opportunities.¹

For each subpattern, we will first summarize some key concepts and results from research and then consider possible support tactics in terms of the ARCADE model.

7.2 Recognition-Primed Decision Making

7.2.1 Key Points From Research

Origin in Research on Naturalistic Decision Making

One well-known line of research was conducted within the research paradigm of *naturalistic decision making* (see, e.g., Klein, 1998; Maule, 2010). This type of decision making is typified by the situation of a fire brigade arriving at the scene of a burning building: The problem situation is changing rapidly over time, even as the decision makers think about how to deal with the fire; there is considerable stress because of the high stakes and because of environmental factors such as noise and

¹A variety of additional perspectives, organized in terms of the concept of *routines of decision making*, can be found in the collection by Betsch and Haberstroh [2005a].

heat; and on the positive side, the decision makers typically have long experience in dealing with such situations, which makes it unnecessary for them to analyze the problem from first principles.

Klein and collaborators developed the concept of *recognition-primed decision making* in contexts like these. On the basis of previous research on decision making, most of it in the laboratory, Klein had expected that decision makers like fireground commanders would typically consider two possible courses of action before deciding which one to execute—in our terms, that they would apply the consequence-based pattern. They were surprised to find that usually the “decision makers” seemed not to be making decisions at all: Most often, they would look at the situation confronting them, remember a course of action that they had previously applied in one or more similar situations, and proceed to execute that action. As Table 7.2 indicates, a somewhat more complex variant of this basic procedure was observed in cases where a contemplated action was not obviously appropriate: The decision maker would anticipate the consequences of the action by a process called *mental simulation* and if necessary modify it until the mental simulation produced a satisfactory result—in effect introducing and evaluating alternative options that were based on the originally considered option. In a small fraction of cases, the decision makers really did find it necessary to start with (and maybe modify) two or more different options before arriving at a satisfactory course of action. In terms of the ASPECT model, the first procedure sketched in Table 7.2 instantiates the experience-based pattern, while the other two involve an increasing number of elements of the consequence-based pattern.

Some of the characteristics that make recognition-primed decision making ecologically rational (3.5) are: (a) the decision maker has a great deal of experience with previous similar situations; and (b) there is no time available for exhaustive generation and comparison of the alternative options.

These two conditions often apply as well to users of computing technology, though the time pressure is usually due not to a dynamically changing emergency situation but rather to a need to proceed briskly with the activities that really interest the user.

Table 7.2: Three forms of recognition-primed decision making. (Summarized on the basis of Klein, 1998, chap. 3.)

Straightforwardly retrieve an action

Experience the situation.

Recognize it and identify a typical action for that situation.

Execute that action.

Retrieve, evaluate, and modify an action

Experience the situation.

Recognize it and identify a typical action for that situation.

Evaluate that action via mental simulation.

Until it seems likely to work in its familiar form, modify it and evaluate it again.

When satisfied, execute it.

Make sense of the situation and consider more than one action

Experience the situation.

Try to make sense of the situation until you have identified it as matching a familiar pattern.

Generate one or more plausible actions for this type of situation.

Evaluate each action via mental simulation, modifying it if necessary, until you have found one that seems likely to work.

Execute the selected action.

Example From HCI Research

The importance of reusing previous choices was discussed in an influential article by Carroll and Rosson [1987] on the problem of method selection (cf. Table 2.1). The authors began with the observation that computer users often persist in employing a relatively inefficient method to perform a given task even when they have more efficient methods available. One of the two explanations that the authors offered was *assimilation bias*: The authors noted that, if users can immediately think of an adequate method for performing a given task, they are inclined to use that method instead of taking the trouble to search for a better method.

7.2.2 Support Opportunities

One step at which support can be provided is the one at which the chooser recognizes the situation as being one that they have some past experience with. One way to facilitate this recognition is to try to ensure that situations that are similar in terms of their choice implications actually look similar to the user—an application of the strategy *Represent the Choice Situation*. This idea is related to the concept of *encoding specificity* in connection with the psychology of episodic memory (see, e.g., Tulving, 1983, for a psychological treatment and Gardiner, 1987 for a discussion of implications for user interface design): It is easier to recall the answer to a question if the way the question is presented is similar to the ways in which it was presented while you were learning the material in question in the first place.²

For example, if a user has acquired particular privacy-related habits while using photo-sharing sites, it will be helpful if situations that are equivalent in terms of their privacy implications also look the same. This point may be considered the preferential choice version of a point that is more familiar in connection with skill acquisition (which largely involves nonpreferential choice): When a new version of a software application changes the layout of a user interface in such a way that situations now look different from their earlier functional equivalents, users will not be able to apply their acquired skills as fluently, partly because the familiar triggers for their habitual behavior are no longer available. This fact is one source of irritation that follows software version upgrades that make major changes to the user interface (e.g., the introduction in 2013 of Apple's iOS 7). It is simply a bit less obvious that preferential choices can likewise be triggered by superficial environmental features.

A related application of the strategy *Represent the Choice Situation* is to ensure that the key features of the current situation that should cause it to be classified as a trigger for a choice are easily recognizable. If a user is accustomed to taking a particular action when their hard

²For this reason, students tend to perform better on an exam if it is held in the classroom in which they originally learned the material.

disk is almost full, it should help if there is some easily perceivable indication when this state has arisen.

7.3 Habit-Based Choice

7.3.1 Key Points From Research

A familiar way in which people's choices are based directly on their past experience is when they act out of habit. The topic of habits is one of the oldest in psychology, but it continues to be an active area of research, bringing forth new theoretical perspectives, many of which now make use of neuropsychological concepts and research methods (see, e.g., Fu and Anderson, 2006; Bayley et al., 2005). For HCI researchers, a useful synthesis is found in an article by Wood and Neal [2007],³ who characterize habits as follows:

Habits are learned dispositions to repeat past responses. They are triggered by features of the context that have covaried frequently with past performance, including performance locations, preceding actions in a sequence, and particular people. Contexts activate habitual responses directly, without the mediation of goal states. (p. 813)

Though the ability to be triggered independently of any particular goal is a characteristic feature, habits can also interact with goals in various ways (Table 7.3), which are of particular interest to interaction designers who wish to take into account—or influence—habit-based behavior. The ways in which goals control habits are relevant to attempts to help users form appropriate habits or to leverage habits that they already have. The ways in which habits can conflict with goals are relevant, for example, to attempts to induce users not to act in accordance with an existing habit.

The topic of habits has been attracting increasing attention in the HCI field (see, e.g., Oulasvirta et al., 2012). From the point of view of the ASPECT model, focusing a great deal of attention on this particular

³Another useful compact discussion is given in the first few pages of Verplanken et al. [2005].

Table 7.3: Forms of interaction between habits and goals. (Formulated on the basis of Figure 1 of Wood and Neal, 2007.)

Relationship Between Goals and Habits	Example
<i>1. Goals control habits</i>	
A person may intentionally form a habit.	"I'll back up my computer every evening just before leaving the office, so as to get into the habit of backing it up once a day"
A person's goal-directed behavior may lead to the formation of a habit, without the person having any such intention.	"I decided on several days in a row to start my day by checking Facebook messages; and it became a bad habit."
<i>2. Habits give rise to (inferences about) goals</i>	
A person can observe his own habitual behavior and make inferences about his own goals.	"I guess I assign high priority to good spelling and grammar: I always check the language of every email message carefully before sending it off."
These inferences can in turn give rise to new goals.	"... So I guess I should spend more time proofreading my scientific articles before submitting them."
<i>3. Habits can conflict with goals</i>	
A person is sometimes aware that some habitual behavior of hers conflicts with a goal that she has.	"I really have more important things to do at the beginning of each day than checking my Facebook messages."
But this awareness is not in itself enough to overcome the habitual behavior; two strategies are often successful:	
A. Actively and effortfully resist performing the undesired habitual response.	"I'm going to ignore the Facebook notification that just arrived!"
B. Change the situation so that you are no longer exposed to the cues that trigger the behavior.	Disable automatic notification about incoming Facebook messages; disable your entire Facebook account.

subpattern of the experience-based pattern—or indeed on any other single pattern or subpattern—brings with it the danger that alternative explanations and support opportunities will be overlooked (cf. 3.1).

7.3.2 Support Opportunities

Arranging for Appropriate Triggers

The tactic mentioned in connection with recognition-primed decision making of making fundamentally similar situations look recognizably similar is even more applicable in connection with habit-based choice, because in the latter case the chooser typically spends even less time classifying the situation and retrieving a relevant action. Within email clients, one benefit of clearly representing a number of attributes of each message using icons and colors is that this type of display makes it easier for a user to form and execute habits concerning how to deal with particular types of message: A particular combination of visual features associated with a message on the computer screen can come to serve as a trigger for a particular way of dealing with it.

System Automation of Routine Behaviors

The fact that habitual behaviors are relatively predictable makes it possible to automate them in some cases—an application of the strategy *Combine and Compute*. For example, if a user habitually sorts email from immediate family members into a folder called “Family”, she can define a sorting rule in her email client that either sorts these messages automatically or offers to do it on approval by the user. For users who don’t want to take the time to define rules of this type, or who lack the knowledge of how to do so, a reasonably intelligent email agent (see, e.g., McCreath et al., 2005) can automatically learn the rules.⁴

This type of support is different from most we have seen so far in that it doesn’t necessarily help the chooser to make a better choice; the result is after all supposed to be the same as if the chooser had done the sorting manually. Still, there is some choice support in the sense

⁴See Jameson and Gajos [2012, p. 434–435] for a brief overview, with further references, of systems that help users to automate routine tasks.

that the effort required from the user to make the choice is reduced or eliminated. In some cases, the automated procedure can make the choices more consistently than the person whose behavior the choices are based on, but the opposite may also be true, depending on how easily automatable and learnable the chooser's habitual behavior is. In either case, if we recall (from 3.6.1) that one of the characteristics of a good choice process is that it makes effective use of the chooser's resources, then even simply reducing choice effort can be considered worthwhile.

Helping the User to Avoid Undesirable Triggers

Whereas so far we have discussed ways of supporting processes of routinization, we should also note that the most helpful approach can sometimes be to help a chooser overcome routines. The most familiar examples involve bad habits that a person would like to break but has difficulty breaking. A good deal of the research done in the persuasive technology area aims at helping people to abandon unsatisfactory behavior patterns like these and replace them with new ones. This work will be discussed in a bit more detail in Section 9.

One general approach is to encourage and facilitate the use of other ASPECT choice patterns, which are likely to lead to different choices.⁵ An example is the sort of calorie counter mentioned later in 9.3.1, which essentially aims to replace experience-based eating with policy-based and consequence-based eating by making it easier to apply the latter two patterns.

7.4 Choice Based on Instrumental Conditioning

A form of experience-based choice that has most famously been studied in animals is *instrumental conditioning* (sometimes called *operant conditioning*).⁶ This pattern was brought to wide attention in the first

⁵This approach can be seen as a subtle application of the strategy *Advise About Processing*.

⁶Accessible treatments of instrumental conditioning, which provide more detail about all of the points made in this subsection, can be found in psychology textbooks such as those of Lieberman [2012, chaps. 4–6] and Gluck et al. [2014, chap. 5].

half of the 20th century by behaviorist psychologists like Edward Lee Thorndike and B. F. Skinner. The basic idea is that, when a chooser acts in a particular way in a particular situation (e.g., responding to a message that he has received while driving) and the behavior has consequences that are in some sense positive, the probability increases that he will act the same way the next time this situation arises (and vice versa if the consequences are negative, though the relationship is not entirely symmetric).

7.4.1 Key Points From Research

An important point here is that it is often not obvious what consequences of an action serve to increase its probability. Intuitively, we might think that these would be pleasant or otherwise rewarding consequences; but concepts like these are too vague to be useful. Instead, the concept of a *reinforcer* is used: A consequence of behavior is a reinforcer if it increases the probability that the same behavior will be shown in similar situations in the future. (Similarly, a *punisher* is a consequence that tends to decrease the probability.) This conception does not make the idea of instrumental conditioning circular: Once it has been determined what consequences function as reinforcers for a given person (or animal) in a given situation, the falsifiable prediction can be made that these particular consequences will increase the likelihood of that behavior in similar situations in the future.

Though this conception conforms with common sense, its implications and applications are less obvious than you might think, for several reasons:

1. Consequences that do and do not function as reinforcers and punishers are not always the ones that we might intuitively expect. To give a frequently cited example: A teacher who thinks that she is punishing a particular undesirable behavior of a child by scolding him is often actually reinforcing the behavior by devoting more attention to the child than he would receive otherwise. Hence one general lesson from instrumental conditioning is that we often need to be careful to find out what reinforcers and punishers are really operating in a given context.

2. A chooser is often completely unaware of some of the *contingencies of reinforcement* (i.e., relationships between behaviors and reinforcements and punishments) that shape his choices; and when these influences are brought to his attention, the chooser often says that he would not want to be influenced by any such consequences of his behavior. This fact implies that particular reinforcers and punishers are in some cases best viewed as sources of bias or noise that should be eliminated if possible (cf. 4.7.1).

3. The effects that reinforcers have on behavior can depend strongly on their exact placement in time. Every dog owner knows that an immediately presented reward or punishment will generally be more effective than one that is delayed. But less intuitively understandable relationships have also been uncovered. For example, reinforcement that is intermittent and unpredictable (such as that which occurs in gambling) has a surprisingly powerful effect relative to constant or predictable reinforcement, often causing the chooser to keep exhibiting the behavior even when the rate of reinforcement has become very low. Research on *schedules of reinforcement*, most of which was conducted with animals like pigeons, has yielded a wealth of knowledge that is sometimes useful in connection with human choices.

7.4.2 Support Opportunities

When the goal of the choice architect is persuasion rather than choice support (1.2.2), it's pretty obvious how instrumental conditioning can be put to use: Arrange for the computing technology to provide reinforcement and punishment in such a way as to encourage or discourage particular behaviors. This type of intervention falls under the strategy *Design the Domain*, since it involves changing the consequences that particular actions have. And indeed, instrumental conditioning is a frequently applied tactic in persuasive technology (see, e.g., Fogg, 2003, chap. 3) and educational computing (where students are rewarded in various ways for achieving learning goals). More generally, one reason why designers of interactive applications try in so many ways to create rewarding user experiences is that they (whether explicitly or not) expect these experiences to serve as reinforcers that will increase

the likelihood of people choosing to use their applications. Efforts of these sorts sometimes make skillful use of principles of instrumental conditioning. Interaction designers who are not aware that they are applying instrumental conditioning should be able to benefit from acquiring more knowledge about the relevant principles and application methods.

But since in this publication we are focusing on the goal of helping users choose for themselves, it is more relevant here to consider how bias and distortion can be *reduced* (4.7.1) through careful attention to processes of instrumental conditioning.

One approach is to look out for cases where a user action has a reinforcing or punishing consequence that plays more of a role in influencing the user's choices than it deserves to play. Suppose, for example, that a mobile application offers two virtual keyboards for text entry, each of which is well suited for a particular situation (e.g., mobile versus stationary use). Suppose also that the use of Keyboard *A* has an immediate and consistent reinforcing consequence—due to its pleasing visual appearance—that causes people systematically to prefer it to Keyboard *B* even in cases where *B* is superior in terms of important consequences such as text entry speed and error avoidance. We might determine that the pleasing appearance of *A* is having more influence on users' choices than they would probably want it to have on reflection, because its effectiveness as a reinforcer is out of proportion to its actual importance. In this case, we could consider applying the strategy *Design the Domain* to make *A* *less* visually appealing, since doing so would tend to level the playing field for the choice between *A* and *B*. (Less controversially, we could make an extra effort to improve the visual appeal of *B* for the same purpose, even if it didn't seem worthwhile to do so otherwise.)

A different way in which an option can be more reinforcing than it deserves to be is because of schedules of reinforcement. Activities like gambling, playing computer games, and participating in social networks are addictive partly because of the schedules on which they offer reinforcement. If our goal is to help people make satisfying choices between such activities and other available activities, it may be justifiable to ap-

ply the strategy *Design the Domain* by avoiding or prohibiting schedules of reinforcement that incline people strongly to repeat a given behavior. Note that we need not have the goal of biasing people *away from* choosing a particular option; we simply want to prevent an undesired bias from occurring. If the users in question choose to continue to play a particular game or to participate in a social network even after its addictive character has been eliminated, we will have no reason to object to that outcome.

7.5 Affect-Based Choice

Past choices can influence subsequent choices somewhat less directly than by suggesting specific options that can be executed: They can lead to the formation of affective associations with options, which can in turn suggest which option should be chosen. For example, consider a user who has purchased software products from a vendor *X* and has had unpleasant experience with them: When she sees the logo of vendor *X* in a display of possible software products, she is likely to experience some sort of negative feeling. Especially if the bad experience lies far in the past, the user may have no recollection of the specific experience; the negative affective response constitutes a compact summary of the previous experience, which guides the user away from choosing the same software vendor.

7.5.1 Key Points From Research

Making a choice on the basis of affective responses like these is known in common parlance as deciding on the basis of “gut feelings”. A more scientific account (Slovic et al., 2002) characterizes this subpattern as the *affect heuristic*, which the authors explain as follows:

... representations of objects and events in people’s minds are tagged to varying degrees with affect. In the process of making a judgment or decision, people consult or refer to an “affect pool” containing all the positive and negative tags consciously or unconsciously associated with the representations.

Another well-known line of research on this topic was initiated by Antonio Damasio (Damasio, 1994), who introduced the concept of a *somatic marker* to denote affective responses of this type.

Although the affect heuristic immediately suggests processing in an intuitive mode (3.4), note that it is also possible for a chooser to reason analytically about the signals provided by his gut feelings (“I don’t know why, but my gut tells me . . .”).

7.5.2 Support Opportunities

The affect heuristic works best when the feelings evoked by the available options reflect relevant past experience with those options. Unfortunately, affective responses can be due to factors other than relevant previous experience. The most obvious example is the way in which advertisers repeatedly associate attractive stimuli with their products, so that a presentation of the product will evoke positive affect even if the prospective customer has had no direct experience—or negative experience—with the product (see, e.g., De Houwer, 2009).

Here again, the interaction designer who aims to support choice can try to do the opposite thing: Avoid creating irrelevant affective associations with options before the time of choice; and at the time of choice, present options in such a way that they tend to evoke choice-relevant feelings and tend not to evoke choice-irrelevant feelings. In other words, the goal is to maximize the signal-to-noise ratio as far as the evoked affect is concerned.

In some ways, avoiding the evocation of irrelevant feelings is the easier task: Although it is sometimes tempting to present items at choice time together with attractive stimuli such as beautiful smiling young people, unless such stimuli are well suited to evoke relevant responses, they are likely just to add affective noise—especially if the stimuli are distributed unequally among the various available options. Note that stimuli like these can create noise even for choices that the user will face only later, at a time at which the irrelevant stimuli are not present. In fact, in this situation the user may have greater difficulty in appropriately ignoring her positive affective response, because the recognizably irrelevant stimulus that evokes it is not currently present. To avoid this

type of affective noise is to apply the strategy *Access Information and Experience* with the goal of reducing bias (4.7.1).

With regard to the complementary strategy of trying to *increase* the strength of *relevant* affective associations, the same basic idea can be applied as for reminding the chooser of relevant previous experiences: Choose a presentation of the option that is likely to be strongly associated with the user's previous experiences. For example, the feeling of frustration or pleasure that the user experienced while using a particular application is more likely to be invoked by a typical screenshot of the application than by a mere textual listing of its name (strategies *Access Information and Experience* and *Represent the Choice Situation*).

8

Socially Based Choice

8.1 Introduction to the Pattern

In another important choice pattern, the chooser is guided by the social context—specifically the examples and expectations established by other people and the advice that they explicitly give. Like the other choice patterns, this one is more applicable in some contexts than in others. For example, a person who has acquired a new computer for home use might decide using one of the three patterns already discussed what applications to install, what privacy and security settings to choose, and how to communicate with friends. When the same person works at the office, many of these decisions are likely to be influenced by written or unwritten rules, conventions, or social examples.

Socially based choice is the one of the six choice patterns that seems to have benefited most dramatically from recent advances in computing technology. People are connected to other people via technology to an extent that would have been hard to imagine even 10 years ago. This connectivity has the potential to amplify all of the forms of social influence that are discussed in this section. So it might be claimed that improved support for this choice pattern is not needed—or indeed that this choice pattern is currently being facilitated excessively at

Table 8.1: Socially Based Choice.

Steps	Support Tactics
– <i>C</i> considers <i>examples</i> of the choices or evaluations of other persons	<i>Access Information and Experience</i> : Provide examples from relevant people—and/or information that enables <i>C</i> to judge their relevance <i>Combine and Compute, Represent the Choice Situation</i> : Summarize the examples in ways that facilitate their interpretation <i>Evaluate on Behalf of the Chooser</i> : Provide a recommender system that makes use of social information
– or <i>C</i> considers the <i>expectations</i> of relevant people	<i>Access Information and Experience</i> : Provide explicit information about expectations <i>Evaluate on Behalf of the Chooser</i> : Provide tools that take relevant expectations into account (e.g., by eliminating from consideration options that do not conform to them)
– or <i>C</i> considers explicit advice concerning the options	<i>Represent the Choice Situation</i> : Avoid the (intentional or unintentional) introduction of factors which would tend to bias <i>C</i> 's response to the offered advice <i>Access Information and Experience, Represent the Choice Situation, Combine and Compute</i> : Derive and represent information about the credibility of the advice giver(s) <i>Evaluate on Behalf of the Chooser</i> : Bring <i>C</i> into contact with persons who can provide reliable advice

the expense of other patterns, thereby diminishing people's inclination and ability to make choices for themselves. Still, if we carefully consider how social influence works, we can see some opportunities for improved support that have not yet been considered as much as they deserve to be.

8.2 Overview of Forms of Social Influence

The fact that the social environment often exerts a powerful influence on people's choices and decisions is known from everyday experience, and the mechanisms of social influence have been analyzed thoroughly

Table 8.2: Reasons why people can be influenced by social examples, expectations, and norms.

Reason to Choose in Accordance With Social Influence	Example: Using the Company's Social Network
<i>If others set an example (without necessarily expecting you to follow it):</i>	
Their experience is a useful source of information.	"If these coworkers have acquired experience with this social network and are still using it, their experience must have been positive."
You want to enjoy practical benefits of conformity.	"There will be direct practical benefits to being in the same network as my coworkers, such as being able to exchange information with them conveniently."
You want to feel that you belong to their group.	"If I use the social network, I will feel more like a typical employee of this company."
<i>If others expect you to make a particular choice:</i>	
They can reward or punish you.	"If I don't use it, I may be subject to disapproval or even concrete disadvantages."
They have a legitimate reason for their expectation.	"The managers in my company have a right to expect me to do things like this."
The expectation seems likely to be grounded in some sort of insight or experience.	"They wouldn't expect me to use this social network if it were not in my own interest."

in theories from social psychology and sociology. The diverse perspectives are associated with different concepts and terminology.¹

¹Useful summaries, with references to primary literature, are provided by Cialdini and Goldstein [2004], Fishbein and Ajzen [2010, chap. 4], and Schultz et al. [2008]. More popularly written discussions, which include ideas about how to use social examples and expectations to influence or support people's choices, are provided by Ariely [2008, chap. 3]; Cialdini [2007, chap. 4]; and Thaler and Sunstein [2008, chap. 3].

The summary in Table 8.2 summarizes in everyday terms some commonly accepted ideas, which will be discussed in the rest of this section. The key message of this table is that there are several distinct reasons why choosers might want to take into account what other people do and think. To fail to recognize these different reasons for social influence is to miss out on opportunities to enable users to make effective use of social information.

8.3 Social Examples

8.3.1 Key Points From Research

We can speak of *social examples* in cases where the chooser takes into account what other people do or think without being concerned about what these other people think that the chooser himself ought to do.

1. The most obvious reason why it makes sense for a chooser to take social examples into account is that these other people are likely to know something that she doesn't know. For example, the ratings that other users have provided of a given software application are supposed to be based on their experience in using it. This information can therefore serve as a surrogate for experience that the chooser could in principle acquire for herself by applying the trial-and-error-based pattern—though it differs in important ways from such experience. Other types of social example, such as the number of downloads or purchases of the application, reflect experience with the options less directly, since these actions tend to have occurred before the persons in question had acquired much experience with the options.

2. A second reason why social examples can be relevant is that it can be practically useful to do the same thing that other people do. For example, by using a video editing tool that many other people use, you will in general be more likely to find people with whom you can exchange software extensions and ideas about how to use the tool. Where the chooser is interested in this type of consequence, he is actually applying the consequence-based pattern (Section 6) and using social information as a cue for making a prediction about a particular type of consequence.

3. On a more subjective level, following the social examples provided within a particular group can increase a person's sense of identification with that group—a consequence that is sometimes more important than the objective consequences (as when the choice concerns a ringtone for your smartphone).

As can be seen in Table 8.2, these three reasons for following social examples can all be relevant at the same time, a fact which helps to explain why the influence of social examples is often strikingly strong. But it is still worthwhile to distinguish these cases, just as we found it worthwhile to distinguish four subpatterns of the experience-based pattern, because the reasons for following social examples are not always applicable at the same time and they are associated with different support opportunities.

8.3.2 Support Opportunities

Contemporary websites and applications often offer a great deal of information about what other users have done and/or how they have evaluated their experiences. Commercial websites show distributions of product ratings and reviews. Media exchange websites show how often an item has been downloaded or viewed and how many others have “liked” it. Providing this information involves applying the strategies *Access Information and Experience*, *Represent the Choice Situation*, and *Combine and Compute* to help the chooser answer the question “What choices or evaluations have the relevant people made in this situation?”

Although this form of support is now ubiquitous, there are still types of choice to which it has yet to be applied. For example, there appears to have been little follow-up on the suggestion made more than two decades ago by Mackay [1991] on the basis of her study of customization that “users want information about their own use and that of other people with similar job responsibilities and attitudes [on] which they can base their customization decisions” (p. 159). Given that choosing among different possible configurations for a software application is inherently more complex than choosing among competing products that are clearly described on a website, it is regrettable that

socially based choice support for configuration decisions has not been explored more thoroughly.

In contexts where the goal of the choice architect is persuasion—for example, encouraging people to save energy—the provision of social examples has become a widely used technique. For example, house owners may be informed about how much energy other comparable house owners consume. As long as the information presented is accurate and unbiased, it can be seen as a form of choice support rather than persuasion. And in fact, it is sometimes observed that the recipients of this type of information adapt their choices in the “wrong” direction: People who have been consuming less energy than their peers sometimes increase their consumption.

From the point of view of choice support, the goal is not to shift choices in a particular direction but rather to enable choosers to make good use of social information. The discussion above suggests several ways in which this type of information can be made more useful than it often is.

A step in this choice pattern which is less frequently addressed but which is important in connection with the one just mentioned is the step where the chooser in effect wonders “What people should I consider to be relevant to this choice?”. The types of information just mentioned often come from some sample of people that has limited and/or unknowable relevance to the current chooser’s choice. A positive exception is found in e-commerce sites, like **booking.com**, which enable a user to view reviews and ratings that have been supplied by members of a particular subpopulation, such as “business travelers”. Here, in addition to applying the strategy *Access Information and Experience*, the system is applying the strategies *Combine and Compute* and *Represent the Choice Situation* to make the social information available to users in an easily interpretable form.

In principle, this idea could be generalized to other types of information about social examples: A user of a video sharing site could be allowed to specify that information about downloads and ratings should be shown only for persons who are especially relevant to the current user, for example because they share demographic characteris-

tics or have exhibited similar choice patterns in the past. To offer this sort of option would involve borrowing some relevant techniques from the field of recommender systems, mentioned below in this subsection.

Where this sort of person sampling is impractical, a simpler alternative would be to supply some sort of indication of the characteristics of the sample of persons whose responses are being summarized. A report such as “233 likes” could be accompanied by a hyperlink to a graphic showing the distribution of age and nationality of the 233 respondents.

Note that the relevant reference group does not necessarily comprise people who are similar to the chooser. If a chooser desires to identify with or become part of some new group (as in the case of a new user of a complex application who wants quickly to become an expert), then that group may be the most relevant one.

Similarly, if one reason why the user is interested in what other people have chosen is that the user wants to collaborate with these other people (cf. the second row in Table 8.2) then the relevant reference group is the group of potential collaborators, regardless of how similar they may or may not be to the current chooser.

Instead of presenting this type of information to choosers and allowing them to interpret it, a system can use this information itself as a basis for recommending particular options, using algorithms from the recommender systems field that make use of social information in sophisticated ways (see, e.g., Ekstrand et al., 2011). Instead of taking the task of interpreting social information entirely out of the chooser’s hands, the recommender system can take over only the winnowing phase (5.3), proposing a number of options and accompanying them with summaries of the relevant social information, maybe together with other information. In this way, the chooser is enabled to evaluate the social information herself regarding the most promising candidates—and also to apply other choice patterns (e.g., the attribute-based pattern) when making the final choice.

8.4 Social Expectations

8.4.1 Key Points From Research

Note that most of the social information discussed so far does not directly reflect social *expectations*. The fact that 2 million people have downloaded a particular smartphone application does not necessarily imply that anyone would think ill of you if you chose a different application instead.

By contrast, we speak of social *expectations* when it is of interest to the chooser what other people think that he ought to do (even if they will not necessarily find out what he chooses to do). One important class of social expectations comprises norms that specify what forms of behavior in a given culture are considered desirable or acceptable. But there also exist innumerable (often unstated) norms within smaller groups, such as online communities (11.4.2) or working groups.

There are several reasons why it is worthwhile to distinguish social expectations from social examples:

1. It is possible for a chooser to be aware of social expectations even in the absence of social examples. A company's management can make it clear that they expect employees to refrain from accessing social networks during working hours without providing any information about the current prevalence of this behavior (which might in fact deviate strongly from the expectation).

2. The reasons for conforming to social expectations are in part different from the reasons for following examples. As is noted in the second part of Table 8.2, the chooser can fear negative consequences of failure to conform to expectations, for example in the form of disapproval by others; or she may simply consider the expectations to be legitimate and therefore consider it natural to conform to them even in the absence of any particular anticipated benefit. There are sometimes more specific reasons than these two that apply in particular contexts. When, for example, the expectations in question concern privacy-related behavior (cf. 12.4), a chooser may decide to conform to a disclosure-limiting norm in order to avoid upsetting others, even if he

is not concerned about any negative consequences of violation of the norm for himself.

8.4.2 Support Opportunities

Here are some ways in which social expectations are sometimes reflected in current systems, which thereby address the question “How would the relevant people want me to choose?”:

- Explicit statements of norms that are intended to govern behavior in a given domain.

For example, discussion forums often describe explicit norms that the forum owner expects to be fulfilled (see 11.4.1).

- More specific explicit statements in reviews or blogs reflecting social expectations (e.g., “Anybody who uses this app must be weird”).
- Social feedback about a particular choice or behavior.

If a user of a social network receives negative feedback about a contribution that she has made—even only in the form of ratings such as “likes” and “dislikes”—then this feedback can be interpreted as reflecting an expectation that the type of contribution that she has made is socially desired or disapproved of.

- Often, only social examples are presented, and it is left to choosers to judge for themselves whether corresponding expectations exist. (“Most people seem to be using App *A*; do they expect me to use it as well?”)

Given the importance of social expectations in some contexts, the practice of providing explicit information about expectations may be worth considering even in contexts in which it is not commonly applied. For example, instead of providing information on how other people have rated an item themselves, a system might provide information about how people evaluate the choice of an item when it is made by someone else (e.g., by summarizing answers to questions like “What do you think of people who buy this app?”).

A relatively easy-to-convey type of social expectation is the kind that corresponds to an explicit policy of a group to which a user belongs. For example, a company may explicitly encourage its employees to use particular software products or to apply particular security practices. Straightforward methods for making people aware of such expectations include explicitly reminding the users of the expectation (e.g., “Please don’t print this email unless it is really necessary”; “Use the train for your business trip unless there is some special reason to use your car”). But if the users also make use of other forms of choice support (e.g., a travel website for planning a business trip) such statements of expectations may be hard to integrate with other considerations. A more ambitious possibility is to include the policies and expectations in question as part of the overall choice support procedure. For example, a travel website can be adapted to take a company’s travel policies into account, either eliminating incompatible options from consideration or explicitly marking them as such.

8.5 Explicit Advice

A type of “social influence” that is studied in other areas of social psychology concerns recommendations by individual persons (or in some cases, by institutions such as professional associations). In contrast to more diffuse social expectations, such recommendations are explicit messages concerning the advisability of adopting some option that are provided by a more or less identifiable individual and that may (or may not) be directed specifically at a particular chooser. In connection with the use of computing technology, typical recommendations would be (a) reviews of software available on the web or in app stores or (b) advice given by a member of a user group on a web forum about how best to use a given application.

8.5.1 Key Points From Research

The most relevant psychological research here is research on persuasion (see, e.g., Cialdini, 2007; Johnson et al., 2005; and Bohner et al., 2008

for useful syntheses).² It would be impractical to summarize even briefly this vast research literature, but we can mention a couple of especially relevant topics.³

Part of this literature covers empirical phenomena and associated tactics that are mainly relevant in the case where the goal of the advice giver is to induce a person to comply with the advice, with little regard for whether doing so is in that person's own interest. Examples of such *compliance tactics*, as discussed in the popularly written and widely used compendium of Cialdini [2007], include:

- *Liking*: Make the advice giver appear likable.
- *Scarcity*: Create the impression that the recommended option will be available for only a limited time.
- *Reciprocation*: Create a sense of obligation in the advice taker by doing something nice for him.
- *Commitment and consistency*: Create the impression that the advice taker would be acting inconsistently with her own previous behavior or commitments if she didn't follow the advice.

A related line of research that has special relevance to choice architecture for HCI concerns the *credibility* of an advice giver. To what extent is the advice giver perceived as being knowledgeable and unbiased? Studies of the credibility of entire websites (see, e.g., Fogg, 2003, chap. 7) have revealed a number of specific cues that people use to judge a site's credibility, including some cues whose relevance is not intrinsically obvious, such as the presence of typographical errors and the recency of the last update.

²A different line of research, which concerns situations in which there is interaction between a professional advisor and a client, is exemplified by Jungermann and Fischer [2005].

³The interested reader is referred for deeper discussion to the literature on persuasive technology, where authority figures and simulated social agents constitute important means of persuasion; see, for example, Fogg [2003, chap. 5].

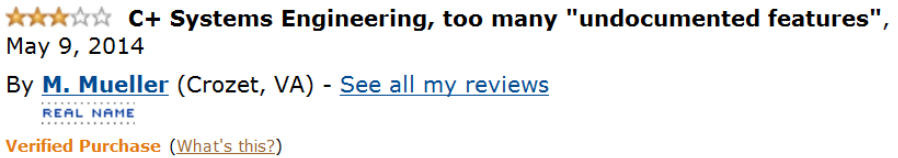


Figure 8.1: Partial screenshot from amazon.com showing information relevant to the credibility of a product reviewer.

8.5.2 Support Opportunities

One way to exploit knowledge about manipulative compliance tactics is to try to ensure that they are not unintentionally introduced in a way that biases the choice toward one option over others (cf. 4.7.1). For example, if two or more persons are presented who offer conflicting advice concerning the available options, we can try to ensure that the advice givers are represented as being about equally likable and that the other compliance-inducing factors, if present at all, are present to an approximately equal extent across the advice givers. This general approach can be seen as an application of the strategy *Represent the Choice Situation* with the goal of minimizing bias.

With regard to the credibility variable, the situation is more complex, since it is in general more difficult to equalize true credibility among advice givers. Here, the goal is to enable a chooser to judge credibility *accurately*: neither underestimating credibility and thereby neglecting a potentially useful source of information nor overestimating it and potentially falling for a biased sales pitch. A number of methods for enabling credibility assessment of individuals have been developed, some of which can be found in conjunction with product reviews, such as: the number of reviews contributed, the possession of any special status such as that of an expert reviewer, and information as to whether the person has actually purchased the product about which advice is being given (see, e.g., Figure 8.1).

An alternative to the provision of credibility information is to help users get in touch with people who are especially likely to be credible. Systems called (among other things) *expert finders* can help with

both preferential and nonpreferential choices (see, e.g., the summary in Jameson and Gajos, 2012, p. 444). Since they take over some of the work of evaluating potential advice givers, they can be seen as realizing the strategy *Evaluate on Behalf of the Chooser* with regard to the subproblem of finding a good source of advice.

With some social networks, such as LINKEDIN, helping people to find credible sources of advice is one of the core functions offered.

9

Policy-Based Choice

9.1 Introduction to the Pattern

The discussion so far has presupposed implicitly that a chooser is making one choice at a time. But it often makes sense for a chooser to think ahead and plan a number of interrelated choices (Table 9.1): A user who has just installed new backup software might decide in advance that he will make an incremental backup every evening just before leaving work. A FACEBOOK user may decide once and for all to ignore all notifications of the form “X has uploaded a new photo”. In cases like these, the choices in question are essentially instantiations of the same choice problem occurring at different times.

We will refer to this type of advance plan concerning multiple future choices as a *policy*, even though this term is more frequently used in connection with organizational decision making, where the policy-based pattern is especially important.

9.2 Research on Time Bracketing

The most directly relevant research has been conducted under the heading of *choice bracketing* (see, e.g., Read et al., 1999b). Choice bracket-

Table 9.1: Policy-Based Choice.

Steps	Support Tactics
– [Earlier:] <i>C</i> arrives at a policy for dealing with this type of choice	<p><i>Design the Domain:</i> (If policy-based choice is clearly desirable:) Make it easy (or even necessary) to specify a policy that the system will execute</p> <p><i>Access Information and Experience:</i> Make <i>C</i> aware of social expectations that suggest possible policies</p> <p><i>Evaluate on Behalf of the Chooser:</i> Recommend an appropriate policy</p> <p><i>Advise About Processing:</i> Encourage <i>C</i> to test possible policies against relevant examples</p>
– [Now:] <i>C</i> recognizes which policy is applicable to the current choice situation and applies it to identify the preferred option	<p><i>Evaluate on Behalf of the Chooser:</i> (If <i>C</i> has communicated the policy to the system:) Apply the policy on <i>C</i>'s behalf to determine the preferred option</p>
– <i>C</i> determines whether actually to execute the option implied by the policy	<p><i>Design the Domain:</i> Enable <i>C</i> to change the environment so as to make deviation from the policy impossible</p> <p><i>Design the Domain:</i> Enable <i>C</i> to arrange to be rewarded for adherence or punished for deviation</p>

ing in its general form involves considering a *set* of choices rather than individual choices; the set does not necessarily have to be distributed over time, but for concreteness we will focus on this case, which is called *time bracketing*. A type of choice for which policy-based choice is especially practically important involves maladaptive behaviors such as unhealthy eating and drug use, where there is typically a contrast between the action that is most attractive in the short run and the sequence of actions that is most beneficial in the long run. Rachlin [2000, chap. 3] analyzes choices of this type.¹

9.2.1 Benefits of Broad Time Bracketing

Taken together, this research has brought to light several typical benefits of what Read et al. [1999b] call *broad time bracketing* (Figure 9.1).

¹Situations where broad bracketing is possible often also involve time discounting (6.8): The benefits of healthy eating appear later than the appealing taste of the unhealthy snack. But the issues just discussed cannot all be reduced to time discounting. Rachlin [2000] uses the terms *complex ambivalence* and *simple ambivalence*, respectively, to distinguish between the two cases.




Figure 9.1: Visualization of the distinction between narrow and broad time bracketing. (Discussion in text.)

These can be illustrated with the example of a user of a new smartphone who can choose between two virtual keyboards each time she enters text: the traditional QWERTY keyboard or an unfamiliar keyboard that has been optimized for one-handed text input. If she employs narrow time bracketing (see the top of Figure 9.1), she will make this choice every time she has a bit of text to enter. With broad time bracketing (bottom of Figure 9.1), she will adopt a policy such as (a) “Always use the QWERTY keyboard”; (b) “Always use the alternative keyboard”; or (c) “Use the alternative keyboard when you have a lot of text to enter”.

One benefit is that a sequence of choices can have important properties that the chooser cannot see when contemplating the individual choices. For example, if the user consistently employs the alternative keyboard, she will initially enter text more slowly and with greater

mental effort than with the QWERTY keyboard; but if she persists long enough, the alternative keyboard will eventually become easier and faster to use than the QWERTY keyboard. Similarly, the user's tastes can change: After a while, she will probably find the appearance of the alternative keyboard less strange and distracting.

Another emergent property of a sequence of choices is the amount of variety associated with it: A user might prefer to alternate between the use of a trackball and the use of a mouse in order to avoid one-sided use of his hand and arm muscles. In general, when bracketing broadly, choosers opt for more variety than they do when bracketing narrowly (see, e.g., Read et al., 1999a).

A more obvious benefit of establishing a policy is that, once it has been established, the process of making each individual choice is in general simpler, since the chooser needs only to see which option is implied by the policy. In HCI contexts, this benefit is often magnified by the possibility of specifying the policy in such a way that it can be (semi)automatically executed by the computer; this form of choice support will be discussed below.

Finally, the application of a policy tends to lead to behavior that is more predictable, both by the chooser and by others. A user who files away email according to a consistent policy will probably have an easier time refinding messages than a user who thinks about each filing action separately. And a user who always uses the same word processor to produce documents makes it easier for collaborators to know what type of material to deliver to her.

9.2.2 Limitations of Policy-Based Choice

Despite these potential benefits of broad time bracketing, policy-based choice also has its limitations, which will be illustrated in passing in the discussions of this pattern below. Most obviously, it is not always possible or practically feasible to specify a policy that regularly yields satisfactory choices (though advances in the use of machine learning in interactive systems are creating new possibilities; cf. 12.5.2). Even if it is possible, the effort required to come up with a policy or perhaps to specify it for the computer may not be justified by the benefits—

especially if the policy needs to be updated regularly because of changes in the environment or in the chooser's goals.

9.3 Dimensions of Variation Among Policies

As used in the ASPECT model, the term *policy* covers a rather wide and partly unfamiliar range of phenomena. To give a clearer idea of its scope, we discuss two dimensions along which policies can differ.

9.3.1 What Is the Content of a Policy?

1. In some cases, a policy takes the form of a rule that specifies what particular option will be chosen each time the choice situation in question arises, possibly as a function of particular parameters of the situation (e.g., “Ignore any photo notification unless it comes from a family member”).

2. In other cases, the policy doesn't refer to any particular option but rather specifies a procedure that is to be applied to make the choice, along with specific criteria (e.g., “When deciding what email messages to reply to first, give priority to work-related messages and especially those written by high-level managers”).

3. Some policies specify neither options nor procedures but rather *goals* that are to be met repeatedly. Some of the clearest examples come from the myriad of applications, whose popularity has been increasing rapidly over the past year or two, that allow users to monitor their achievement of goals such as eating more healthily or getting more exercise.² A policy could specify that the chooser should consume no more than 1800 calories or take no fewer than 10,000 steps each day. A benefit of this type of policy is that it replaces a long-term goal, whose achievement will not be noticeably influenced by any particular action that the chooser takes, with a concrete, short-term goal that can guide individual choices.

²Related phenomena involving financial choices have been studied under the label of *mental accounting* (see, e.g., Thaler, 1999).

9.3.2 Does a Policy Necessarily Concern Separate Choices?

Although a policy normally covers separate choices, as in the examples given above, sometimes a chooser will formulate a policy for dealing with individual parts of a larger choice task. Consider, for example, a user who is considering buying a new type of smartphone. He may decide to look at one smartphone each day for several days, each time deciding whether or not to buy the phone in question. He may decide in advance on a policy for evaluating each smartphone (e.g., in terms of attribute importances or desired levels of attributes; cf. Section 5). Even though the multiday decision process can be seen as concerning a single decision—that of which smartphone to buy—the issues that arise with this type of policy are mostly the same as those for policies that concern distinct choices.³

Finally, it is sometimes helpful to view as examples of policy-based choice even situations in which only one choice is being made, if the option resulting from that choice is going to be executed in the future (as when a person decides today that she is going to participate in a particular webinar in 2 weeks). What this type of situation has in common with other policy-based choice is that the time of forming the intention and the time of executing the option are separated considerably in time, a circumstance that raises the question of whether the chooser will actually execute the option in question when it's time to do so.

9.4 Support for the Generation of Possible Policies

Consider a new TWITTER user who says to himself: “Now that I’ve joined TWITTER, I’d better come up with some rules for deciding what types of tweet to write, when to write them, when to read tweets sent by other people, . . .”. One problem is that the number of possible policies

³Interrelated choices can also concern different questions, as when a user plans a sequence of qualitatively different actions to migrate their old PC to a newly purchased one. In these cases, we would normally speak of a *plan* rather than of a policy. These cases will not be considered in this work, since the topic of planning is too large to fit into the scope of this publication.

with regard to each of these questions can be very large, even if there are just a few options for each individual choice.

Fortunately, choosers don't always have to start from scratch when generating possible policies. Here are some things that they can often start with:

1. *Related policies that the chooser already has:* A user who joins a new online community may already have developed a policy in a previous community that determines what sorts of contribution she makes or what sorts of personal information she discloses. She can then adapt this policy to fit the new community.
2. *Social expectations and advice:* As will be discussed later in connection with online communities (11.4.1) and privacy-related choices (12.5), people like managers of social networking sites often explicitly recommend or require particular policies. In cases where these policies are not sufficiently accessible to users, a form of choice support is to apply the strategy *Access Information and Experience* to help advertise the policies or to give feedback on the extent to which a given user is conforming to them.
3. *Policy specification interfaces:* Where a system offers a user interface for specifying a policy that is to be executed (semi)automatically, this interface will in general make it easy to specify some types of policy, more difficult to specify some others, and impossible to specify anything else. By deciding which types of policy to support and encourage, the interaction designer is applying the strategy *Evaluate on Behalf of the Chooser*, probably making more use of experience with different policies than any single user is likely to have (cf. 12.5.2).
4. *Experience-based patterns of choice:* Section 7 presented four ways in which recurrent patterns of similar choices (sometimes called *routines*) can arise. If the chooser becomes aware of a pattern like this (e.g., that he always proofreads an email message before sending it off), he may decide to adopt this routine as a policy (cf. Table 7.3). Alternatively, if he identifies the routine as an undesir-

able one, he may formulate a policy with the goal of abandoning the routine.

A possible form of support here is to enable the chooser to become more aware of the routines that he has acquired (*Access Information and Experience*). This sort of support is already offered by innumerable mobile and web-based applications for monitoring activities like eating and exercising.

9.5 Support for the Evaluation of Possible Policies

Once the chooser has in some way thought of a possible policy (e.g., “Always proofread an email message before sending it off”), she may want to evaluate it somehow before adopting it.

One possibility is to generate example choice situations in which the policy would apply. This in turn can be done in two ways:

1. Remember relevant choice situations from the past: “If I had proofread that message to the CEO, it would have saved me a lot of embarrassment.”
2. Generate hypothetical (perhaps borderline) cases: “What if I’m writing under great time pressure and the message isn’t so important?”

As an alternative to this sort of mental evaluation, the chooser can apply the trial-and-error-based pattern: Start applying the policy, see how well it works in practice, and where necessary adjust it on the basis of experience (e.g., after it has yielded a bad result). An example of a system that explicitly supports this sort of trial and error is the *LOCYOUTION* system discussed in 12.6.2. A limitation is that it can require a good deal of feedback on individual cases to get a reliable idea of how well a particular policy works.

9.6 Support for the Execution of a Policy

When a policy specifies a goal that the chooser is supposed to achieve on a regular basis, one simple but effective form of support is to remind

the chooser of the goal and help him to monitor his achievement of it. For example, a calorie tracking app will help the user to see how many calories remain for him to consume today, applying the strategies *Access Information and Experience*, *Combine and Compute*, and maybe also *Evaluate on Behalf of the Chooser* (if the system offers specific recommendations about how to achieve the daily goal). This sort of self-monitoring is often viewed as a tactic of persuasion (see, e.g., Fogg, 2003, chap. 3), but it can equally well be viewed as a form of choice support: Even the most highly motivated dieter has to decide somehow what food to eat each day, and adhering to a particular policy is one effective way of doing so.

Often, even after a chooser has explicitly adopted a sensible policy, she can have difficulty in executing the individual choices specified by the policy when the time comes to do so. Even a user who has resolved to ignore a particular type of FACEBOOK notification may find it difficult to resist the temptation to click on a notification of this type if it arrives while he is engaged in a tedious activity. The key problem is that the perspective of the chooser faced with an individual choice is quite different from the perspective of a chooser who is contemplating a policy—which is in fact an important reason why policies are valuable in the first place.

One way around this dilemma is to delegate the execution of the policy to another person or to the computer, a form of support that was mentioned above.

Another important approach is to set up some sort of *commitment mechanism* that will increase the likelihood that the chooser will stick to a policy. The classic example is Odysseus' request to be tied to the mast of his ship so that he would not succumb to the temptations created by the calls of the Sirens. As this example illustrates, commitment mechanisms are relevant not only for the enforcement of policies but also as a way of enforcing one-time decisions to perform a particular action later in time, as was discussed in 6.8.2 in connection with preference reversals involving smaller-sooner and larger-later benefits.

The implementation of commitment mechanisms constitutes a major part of persuasive technology. The fact that we are now consid-

ering commitment mechanisms as a form of choice support illustrates once again how choice support and persuasion are often best applied together. From the point of view of choice support, a commitment mechanism is a way of enabling a chooser to implement policy-based thinking, which might otherwise be infeasible and which offers many potential benefits in terms of the quality of the resulting choices.

Here are two types of commitment mechanism (discussed in Rachlin [2000, chap. 2]):

- *Elimination of the options ruled out by the policy:* Often, the most effective mechanism involves the general strategy *Design the Domain*: Build in an opportunity for the user to change her environment in such a way that it is impossible for her to deviate from her policy. For example, the highly immersive game WORLD OF WARCRAFT provides a flexible set of parental controls, which some adult users apply to constrain their own use of the system.⁴
- *Punishment for noncompliance:* When it isn't feasible to prevent deviations from the policy entirely, an alternative is to arrange for deviations to be punished (or adherence rewarded). Although people and systems sometimes arrange for financial or material punishment, the threat of social disapproval is often easier to arrange. When the chooser's policy was derived from social expectations, no special action may be needed to get the social environment to supply the necessary punishments and rewards: If an active WIKIPEDIA user doesn't adhere consistently to the policies concerning editing behavior that are enforced by the WIKIPEDIA community, he is likely to suffer negative consequences, possibly including eventual exclusion from the community. Where the policy has been devised by the chooser herself and adherence is not expected by anyone else, the chooser can create the necessary social expectations by announcing her policy to relevant people. A drawback of the punishment-based approach is that sometimes the chooser will ultimately find it more attractive to accept the punishment than to stick to the policy. In this case, she fails to

⁴<http://www.wowhead.com/forums&topic=178118/balancing-real-life-and-wow>

derive the benefits of the policy and in addition suffers the disadvantage of the punishment.

10

Trial-and-Error-Based Choice

10.1 Introduction to the Pattern

Despite their ability to apply all of the five choice patterns that have been discussed so far, people often find while making a choice that they don't really know what option to choose. In many cases, an appropriate procedure is simply to start by trying out one of the options, especially if you are reassured by the knowledge that you can switch to another option if the first one doesn't seem satisfactory.

Trial and error is especially applicable as a choice pattern in HCI contexts, partly because computer users are constantly being confronted with new options and partly because it is possible to design interactive systems in such a way that trial and error can be conducted effectively and without some of the costs and drawbacks that it often has in the natural world.

10.1.1 Dimensions of Variation

There are a number of ways in which instantiations of this choice pattern (summarized in Table 10.1) can differ. The first three dimensions concern the very question of what we mean by "trial and error":

Table 10.1: Trial-and-Error-Based Choice.

Steps	Support Tactics
<ul style="list-style-type: none"> - <i>C</i> selects an option <i>O</i> to try out, either using one of the other choice patterns or (maybe implicitly) by applying an <i>exploration strategy</i> 	<p><i>Advise About Processing</i>: Explicitly recommend an exploration strategy</p> <p><i>Access Information and Experience</i>: Make available information that is relevant to exploration strategies (e.g., about the reversibility of actions)</p> <p><i>Design the Domain, Represent the Choice Situation</i>: Design and/or arrange the options in a way that facilitates various exploration strategies and avoids bias in favor of particular strategies</p> <p><i>Evaluate on Behalf of the Chooser</i>: Recommend an option to try next</p>
<ul style="list-style-type: none"> - <i>C</i> executes the selected option <i>O</i> 	
<ul style="list-style-type: none"> - <i>C</i> notices some of the consequences of executing <i>O</i> 	<p><i>Access Information and Experience</i>: Provide information about consequences that are otherwise not immediately perceivable, avoiding bias toward favorable or unfavorable consequences</p> <p><i>Access Information and Experience, Represent the Choice Situation, Combine and Compute</i>: Provide a representation of the resulting situation that helps <i>C</i> to evaluate its desirability</p> <p><i>Evaluate on Behalf of the Chooser</i>: Suggest an evaluation of the resulting situation</p>
<ul style="list-style-type: none"> - <i>C</i> learns something from these consequences 	<p><i>Represent the Choice Situation</i>: Represent information about consequences (perhaps including those of previous trials) so as to facilitate interpretation and evaluation of them</p> <p><i>Combine and Compute</i>: Provide a diagnosis of reasons for observed consequences (if feasible)</p>
<ul style="list-style-type: none"> - (If <i>C</i> is not yet satisfied:) <i>C</i> returns to the selection step, taking into account what has been learned 	

1. To what extent is the chooser intentionally engaging in trial and error?

In the most obvious instantiations of this pattern, the chooser has no idea in advance which of several available options should be chosen, so he tries out one or more options with the explicit goal of learning about the option(s) before making a choice.

At the other end of the spectrum, the chooser arrives at a choice using one or more of the other five ASPECT patterns and has no

reason to doubt in advance that the chosen option is a good one. It is only when negative feedback is obtained that the chooser sees a need to try a different option. And even on the second attempt, the chooser may think that he now knows what will be best. As long as the chooser does not see in advance a need to learn from feedback, he cannot be said to be intentionally engaging in trial and error; but if he does learn from the feedback that he receives, he is in effect applying the second part of the trial-and-error-based pattern. Accordingly, we view this situation as a special limiting case of the trial-and-error-based pattern.

In between these two extremes, the chooser believes to some extent that a particular option will be better than others but is prepared to learn from feedback and maybe try different options later.

2. To what extent do the trials of options accomplish a single task?

The attempts using different options can accomplish completely separate tasks (as when different text entry methods are tried on different occasions); or they may be seen as accomplishing one task (as when a user switches text entry methods midstream while entering a particular block of text).

3. How extensive is the experience that the chooser acquires with the option that she tries out?

Let's refer to as the *full experience* of an option all of the experience with the option that the chooser would consider relevant when judging retrospectively whether it was a good idea to select that option. For example if it's a matter of selecting a website for acquiring audiobooks, the full experience might be that of subscribing to the service for a year, acquiring and listening to a number of audiobooks, and making use of the website's customer service when the user encounters a technical problem. When trying out such a website, each user is likely to acquire only *partial experience*, such as browsing the site for a while and downloading and listening to a free trial audiobook. Even simply visiting the

site for a few seconds to get a first impression of it before moving on to the next competing website can be seen as an example of trying a site out, on the grounds that a bit of relevant partial experience is being acquired (for example, how much fun it is to look at the site).

By contrast, if all that the chooser does is acquire information about an option (e.g., by reading a textual description of it or reviews by other users), we would not speak of “trying the option out”. Instead, we would say that the chooser is acquiring information about one or more options before choosing.¹

10.1.2 Favorable and Unfavorable Conditions

We can also distinguish four dimensions that concern the question of how favorable the choice environment is for application of the trial-and-error-based pattern:

- How much cost is involved in switching between different options as opposed to sticking with one option?

In some situations, it is difficult or impossible to return to an option that the chooser has tried in the past and not immediately accepted.

- How serious are the negative consequences when an option yields an unsatisfactory result?

Trying out an option can sometimes have significant negative consequences like having to pay a high price for something that is not going to be used or disclosing sensitive personal data (see 12.6.1). In more favorable situations, either negative consequences cannot occur at all or they can be “undone” by the chooser.

¹Even in this latter case, some of the points made in this section apply—for example, concerning the question of the order in which particular options should be investigated. Despite this more general applicability of the discussion, we will focus for concreteness on cases where a chooser is acquiring at least partial experience with one or more options.

- In particular, to what extent will any benefits obtained through the use of an option be lost once that option has been abandoned?

When an option is abandoned in favor of another one, it is possible that something useful that was achieved with the abandoned option can be retained even after the switch; or the results of using that option may disappear soon as the option is abandoned.

An example of the latter case would be when a user tries out a new password management application, entering some passwords into it, only to discover that (a) he is not satisfied with the application and (b) there is no way to export the passwords that he has entered so that he can access them in some other password management application.

- To what extent is the choice environment changing over time in such a way that an option that was desirable at one point in time may not be desirable at a later time?

Consider a user who is configuring a newly purchased mouse using the configuration screen shown in Figure 2.1: If she experiments with different configurations while working on a technical drawing that is continually getting more complex, then the answer to the question of which configuration is best may change even as the user is engaging in trial and error.

10.1.3 HCI Examples

Some instantiations of the trial-and-error-based pattern from the wide range just defined can be seen in the following examples:

An application selection choice: A user in an app store sees 10 apps that allegedly have the functionality that he requires. Being unable to judge on the basis of the descriptions and reviews which one is most suitable for him, he downloads one app at a time and tries it out (if possible using a free trial version) until he has found a satisfactory app, has given up, or has run out of money. The time spent using a given app may result in some lasting benefit, or everything that was done with it may have to be discarded.

A method selection choice: A user is not sure which method of text entry will work best in a given situation. She starts using Method *A*, sees how it goes, switches to Method *B*, and then maybe switches back to *A* if it seems to have been better after all. In this case, the user will presumably have been entering some text with each method, so the cost of trial and error need not be very high.

A configuration choice: A user starts with the default configuration of a new application, tweaks the configuration when some unsatisfactory experience occurs (or some possibility for improvement suggests itself), and continues to do so, perhaps over a long period of time. The main difficulties in this case are likely to be that the consequence of a configuration change are often hard to evaluate, since they depend on future actions and situations; and (b) the space of options is in principle huge, since each specific combination of parameter values may have to be treated as a separate “option”, if the effects of the individual configuration parameters are not independent of each other.

As can be seen in Table 10.1, there are two main parts of this pattern, which we will discuss separately:

1. What *exploration strategy* should the chooser employ?
2. How should the chooser make use of feedback about the consequences of options that he has tried out?

10.2 Research on Exploration Strategies

Exploration strategies have been investigated in various lines of research in different disciplines:

10.2.1 Research on the Relationship Between Exploration and Exploitation

A fairly recent trend in laboratory research on judgment and decision making (see, e.g., Rakow and Newell, 2010; Hertwig et al., 2006) concerns *decisions from experience*:² It looks at variants of the gambling

²This paradigm is sometimes labeled *experience-based choice* in the relevant literature, but we have appropriated that term here as the name of an ASPECT pattern.

paradigm (cf. 6.1.2) in which the probabilities of particular outcomes are not explicitly provided; instead, the chooser has to infer them by observing the outcomes of individual choices (made by herself or someone else). For example, instead of being told that Option *A* offers a 10% chance of winning \$12 while Option *B* guarantees a win of \$1, a participant is allowed to click repeatedly on two buttons corresponding to the two options and to observe the resulting rewards. An important issue in this sort of situation is the tension between *exploration* and *exploitation*: In order to learn efficiently which of the two options is preferable, a chooser should in principle systematically “explore” both of them, trying them out until it is clear which one is better—a process that may take some time, as in the example just given. But in practice, once a chooser has the impression, say, that Option *B* is better, there is a temptation to “exploit” this insight by consistently choosing *B*.

Another example of the exploration/exploitation tension is the *production bias* that was described by Carroll and Rosson [1987]: the tendency of computer users to stick with procedures that they already know instead of experimenting with new ones that might turn out to be better. Whether or not you want to call it a “bias”, computer users have to assign some priority to getting things done as opposed to learning new things.

The exploration/exploitation relationship has attracted a good deal of attention from neuropsychologists, who have identified brain regions associated with exploration and exploitation, respectively (see, e.g., Cohen et al., 2007, for a review; and Hills et al., 2010).

In the field of machine learning, the question of how to allocate resources to exploitation of known solutions vs. exploration of new solutions has been analyzed carefully because of its inescapable importance in this area, especially in the area of reinforcement learning (see, e.g., Russell and Norvig, 2010, chap. 21). No completely general solution appears to exist, but interesting results have been achieved in environments characterized by particular assumptions, such as the stationary *multi-armed bandit* environment, in which each of a number of slot machines yields rewards with some particular probability distribution.

10.2.2 Studies of Exploration and Interactive Search in HCI

Some types of exploration behavior in information environments show characteristics of trial and error. For example, a user searching for information on a particular topic may read a web page and start following links to subpages from that web page. At some point, the user may determine that the current line of exploration is unlikely to lead to success and abandon it, returning to a choice point visited earlier (e.g., clicking on a different search result to visit a different website). The best-known model of this general sort of behavior is the *information foraging* model of Pirolli and collaborators (see, e.g., Pirolli, 2007; Pirolli, 2003; Fu and Pirolli, 2007). For other studies along these lines, see Young [1998], and Brumby and Howes [2008].³ One recurrent idea in this area is that the explorer continually reevaluates the promise of the current line of exploration by processing cues in terms of their *information scent*. A frequent recommendation to interaction designers is that they should try to provide easily interpretable and valid information scent. Increasing information scent can involve either applying the strategy *Design the Domain* by making options inherently more assessable in terms of their potential—for example designing web pages so that their content immediately reveals what they are about and what sorts of other pages they lead to—or it can involve the strategy *Access Information and Experience*, as when thumbnail previews of webpages are displayed next to the hyperlinks that lead to them. More generally, the idea of providing valid cues about the potential of particular options is one way of helping to support trial and error, in that it makes it easier for the chooser to see which option is best tried next.

10.2.3 Studies of Decision Making in Organizations

In studies of organizational decision making, it has been noted that decision makers often engage in a process of conservative exploration called *muddling through* (Lindblom, 1959; Lindblom, 1979), in which they incrementally adapt a policy, often in response to environmental

³Broader overviews of exploratory search in information retrieval are given by Wilson et al. [2010] and White and Roth [2009].

changes, whereas they could in principle more systematically and boldly explore the space of possible policies.

10.2.4 Studies of Sequential Search

One frequently studied type of problem is sometimes called the *secretary problem*, because it is typified by the task of testing a sequence of candidates for a secretary position. An example of a specific formulation of a problem of this type is the one given by Zwicky et al. [2003]:

Alternatives are inspected in a random order, one at a time, and only the rank order of the current alternative relative to the ones that have already been observed can be ascertained. At each period, the consumer may either accept the current alternative, continue to search and pay a fixed cost, or recall an alternative that has already been inspected. A recalled alternative is assumed to be available with a known probability. The consumer's goal is to select the overall best alternative from the fixed set. (p. 503)

Making particular assumptions, it is possible to specify how long the chooser ought to search and what an optimal exploration strategy is.

10.3 Support for Exploration

10.3.1 Encouraging Particular Exploration Strategies

Not surprisingly, because of the wide variety of situations that fall under the trial-and-error-based pattern, it is not possible to derive general recommendations concerning particular exploration strategies from the research just summarized. But this research does show that the question of what is a good exploration strategy is in general not an easy one and that consequently the selection of an exploration strategy is a step at which choice support can be especially valuable. An interaction designer who thinks about possible exploration strategies for a particular situation can often think of a strategy which (a) is suitable for that situation but (b) would probably not readily occur to a busy computer

user. The designer can then either build support for the strategy into the system (*Design the Domain*) or recommend it to the user (*Evaluate on Behalf of the Chooser*).

Consider, for example, an app store that offers 10 applications that serve a particular purpose (e.g., the French-English dictionaries from 3.3). Suppose we can assume that the user wants to try applications in turn until he has found one that exceeds a given threshold of satisfactoriness; and that a priori all applications have an equal chance of being satisfactory for him. Under these assumptions, the best exploration strategy is to try out the applications in increasing order of their cost, since in this way the amount of money the user has to pay before finding a satisfactory application can be expected to be minimized. A user might well think of such a strategy without any support, but consider a slightly more complex case, where the more expensive apps have higher probabilities of being satisfactory (as reflected, for example, in their higher user ratings). Then the best exploration strategy should take into account these probabilities, so the determination of the best order of trying out the options is more complex than in the previous situation.

An example in which an exploration strategy is recommended to new users of an online community is given in 11.6.1.

10.3.2 Creating Favorable Conditions for Exploration

An alternative to encouraging a particular exploration strategy is to create generally favorable conditions for exploration. The favorable and unfavorable conditions listed above (10.1.2) for trial and error situations may be partly under the control of the choice architect: For example, applying the strategy *Design the Domain*, the designer may be able to minimize the costs of unsuccessful trials and make it easy for the chooser to switch from one option to another without a large cost or loss of intermediate results.

10.4 Research on Learning From Feedback

10.4.1 Types of Feedback That Can Be Obtained

In connection with the consequence-based choice pattern, we considered several qualitatively different types of consequence that a computer user may want to anticipate and evaluate. Similarly, it is worthwhile to consider what sort of information a user can acquire by trying out an option and what sort of thing she can learn from it. As we did in 6.6, we will consider the example of a driver who is experimenting with listening to and dictating email messages.

1. “To what extent did the consequences that I anticipated actually occur?”

For example, the driver may note that he didn’t enjoy processing email as much as he had expected but at least he didn’t cause an accident. This sort of feedback can (a) help the chooser update the subjective likelihoods that he associates with particular consequences; (b) make him aware of consequences (e.g., unexpected side effects) that he had never thought of; and (c) more generally help him to update his mental model of the causal relationships in the situation.

2. “Do I now evaluate the consequences that occurred in the same way that I would have evaluated them while anticipating them in advance?”

Even if the driver processes just as many email messages as he had expected to, he may end up with a different assessment of how important it was to get these messages processed quickly.

3. “Did anything happen that was not especially important in itself but that gives me clues about whether the option I tried was a good one?”

The driver might notice that he had driven less steadily than normal and had noticed a stop sign only just in time to stop. Though these are not consequences that have much weight in themselves, they do serve as hints that the chooser might well cause an accident in the future if he continues to execute the same action. Especially when the consequences that are of primary interest are improbable, temporally distant, or hard to perceive, this sort of feedback about intermediate

results can be more informative than feedback about the significant consequences.

4. “How did other people respond to my action?”

Trying something out is one way to find out about relevant social expectations (e.g., those of the driver’s family members) that may not otherwise be knowable.

10.4.2 Norman’s Model of Action

The model of action introduced by Norman [1986], which is well known in the HCI field, is worth bearing in mind in this context, even though it was not specifically intended to illuminate processes of preferential choice. In his discussion of the *gulf of evaluation*, Norman distinguishes the phases of perceiving, interpreting, and evaluating the results of an action. Each of these phases can be seen as a way in which a chooser may have difficulty in learning from experience in making a certain type of choice. For example, a person who has acted on a decision to contribute one paragraph to a WIKIPEDIA article will probably never know how many people have read the paragraph or how much they benefited from it. The author may well notice the changes that other WIKIPEDIA contributors make to the paragraph, but she may interpret them unrealistically and thereby arrive at an inappropriate evaluation of her original decision to contribute the paragraph. The following subsections will look in more detail about such difficulties and ways dealing with them.

10.4.3 Models of Learning From Feedback

A good deal of the psychological research that is relevant to trial-and-error-based learning has aimed to model in great detail how particular types of feedback are processed in the human mind and brain (see, e.g., Newell et al., 2007, chap. 11; Glimcher and Fehr, 2014, chaps. 15–18). This research has yielded a remarkable convergence of insights from theoretical work about learning algorithms, psychological observations in the laboratory, and neuropsychological research into the neural substrates of learning (e.g., the role of the neurotransmitter dopamine). So far, though, it is hard to derive from this research implications about

how to help people learn from feedback when making everyday choices. The most important difficulties that people have in this context appear to come not from subtle aspects of the ways in which they process feedback but rather in the properties of the feedback itself, which have been studied in a largely separate type of research.

10.4.4 Research on Properties of Feedback

Hogarth [2001] introduced the influential notion of an unfavorable, or *wicked* learning structure as opposed to a *kind* structure when calling attention to differences among environments in terms of the ease of learning by trial and error.⁴ On a more general level, research on problems with the information samples that people have to deal with has shown that many of the difficulties that people have in dealing with feedback (and other types of information sample) are due to limitations of these samples (see especially the collection edited by Fiedler and Juslin, 2006).

Some specific ways in which a learning environment can be “wicked” are considered in turn in the following subsection, along with ideas about how to help choosers deal with these difficulties by applying support tactics suggested by the ARCADE model.

10.5 Combating Typical Problems With Feedback

The first few typical difficulties in processing feedback involve some kind of difficulty in perceiving the feedback; the remaining ones concern factors that make it difficult to interpret the feedback and derive lessons from it with regard to future choices.

10.5.1 Invisible Consequences

Examples of events that the chooser may simply be unable to perceive are: (a) damage caused by a computer virus as a consequence of an unwise decision to open a malicious email message; and (b) favorable or unfavorable subjective responses of other persons to an action by a

⁴A shorter exposition can be found in Hogarth [2008].

member of an online community that are not expressed in communications to that member.

A natural way of combating this difficulty, which corresponds to general HCI guidelines, is to ensure that the relevant feedback is perceivable. Doing so can involve straightforward provision of information (e.g., about damaged files), via the strategy *Access Information and Experience* or more complex approaches such as giving members of an online community an opportunity to provide feedback about perceived inappropriate behavior of other members (*Design the Domain*).

Another way of applying the strategy *Design the Domain* is to try to ensure that there exist no consequences that users should know about but that are inherently hard to perceive. Consider, for example, a company firewall that is known to have succeeded in completely blocking viruses from the company's computers: Users could then choose which emails to open and which websites to visit without even thinking about possible damage from viruses.

10.5.2 Consequences Not Clearly Related to Actions

Sometimes, the observed events are not clearly related to the action of the chooser that caused them. For example, a user who notices that his computer has run out of available RAM may have no idea what caused the problem. Here, what is needed is some indication that relates the consequence to the action that caused it. Methods of diagnosis that can be applied to support situation assessment (6.3) may be applicable in these cases.

10.5.3 Small Differences Among Consequences of Different Options

Sometimes, the consequences of different options differ in subtle ways that can be noticed only through careful comparison.

For example, Text Entry Method 2 may be 10% faster than Method 1, but the user may not be able to notice this consequence of switching to Method 2; after all, she is not conducting a controlled experiment. An extended example of this point comes from the research of Bhavnani and colleagues: Bhavnani and John [2000] studied

in depth expert users of computer-aided design systems who persisted in using inefficient methods: Among other things, they tended not to take advantage of the opportunity that their systems offered to perform an operation on multiple objects at one time. For example, when they needed to create three identical objects, they would draw them separately instead of drawing one object and making two copies of it. One of the authors' explanations for the persistent use of inefficient methods concerned the fact that the users did not obtain clear feedback that revealed the inefficiency: The quality of the resulting drawings was in general identical, and the difference in execution times was not easy to notice from experience.⁵ In view of this and other obstacles to spontaneous learning of the more efficient procedures, Bhavnani and his collaborators concluded that explicit training was required (see, e.g., Bhavnani et al., 2008).

By contrast, Gray and Boehm-Davis [2000] showed that, under more favorable learning conditions, users can sometimes take into account a difference between alternative *microstrategies* that involves only milliseconds of execution time.

The interaction designer who wants to ensure that differences of this general sort are recognizable might apply the strategy *Combine and Compute* to provide summary feedback that accumulates over time, somewhat like the results of a controlled experiment. For example, the system might summarize the text entry rates with two different text entry methods in a way that made even small differences noticeable. This type of feedback is already provided by many applications in the health and well-being domains, such as the one shown in Figure 4.3.

10.5.4 Low-Probability Consequences

Sometimes, feedback consists in information about whether a particular low-probability event has occurred or not, such as a major hard disk failure, identity theft due to inadequate security measures, or an accident due to texting while driving. As studies of “decisions from

⁵Some users never tried the more efficient method in the first place, which is a problem concerning their exploration strategy (10.2) rather than the perception of consequences.

experience” (mentioned in 10.2.1) have confirmed (e.g., Hertwig et al., 2006), since a large amount of feedback is required to arrive at an accurate estimate of a low probability, it is in general infeasible for any individual chooser to learn very low probabilities from experience. So it makes more sense to support other choice patterns, such as the consequence-based pattern (providing probability information derived from other sources) or the socially based pattern (calling attention to relevant social examples and expectations).

Another approach is to provide feedback not about the low-probability consequence itself but about more frequent (though less significant) consequences that are correlated with it (cf. the third type of feedback listed in 10.4.1 above) For example, a multitasking driver can be given feedback whenever he drives in a way that increases the likelihood of an accident (e.g., occasionally moving out of his lane or moving too close to the driver in front of him). Or an imaginative choice architect, applying the strategies *Access Information and Experience* and *Combine and Compute*, could invent a new type of feedback of this nature (e.g., a dangerous driving index), which need not correspond to any naturally occurring consequences of a driver’s actions. In a similar vein, in 12.6.2 we will discuss how it can be more feasible to give feedback about situations associated with privacy *risks* than about actual privacy *violations*.

10.5.5 Biased or Distorted Feedback

For example, feedback from members of a social network on the appropriateness of contributions is often distributed in an unrepresentative way: Maybe only inappropriate contributions elicit any feedback, in which case a contributor is likely to underestimate the benefits that others are deriving from her contributions. Similarly, a user who reads reviews on an e-commerce site may wonder whether extreme evaluations are being overrepresented. Where this pattern arises, a choice architect may want to consider how to balance the distribution of feedback. For example, in some situations buttons labeled “Like” and “Dislike” might require so little effort to use that the distribution of the

responses received is more representative of the actual distribution of responses.

Of course there are limitations to efforts like these, and some amount of bias is likely to remain. One possible response is to make choosers aware of the bias in the feedback (e.g., by explaining or visualizing the source of the bias). A problem with this approach is that, even if a person is aware of bias in a sample, he may have no way of knowing with any precision what the feedback would be like if it were not biased. (“If there were a ‘Dislike’ button beside the ‘Like’ button, how many ‘Dislikes’ would I be getting?”) If the system or its designer has some reasonable way of correcting for the bias (e.g., by exploiting information collected in the past with other users), it might be reasonable to try to present a debiased sample of feedback in some way—at the risk of introducing bias of another sort.

11

Choice in Online Communities

11.1 Introduction

One of the most thoroughly discussed and studied types of choice and decision made by computer users concerns their decisions about whether and how to contribute to an online community.¹

Online communities (see, e.g., De Souza and Preece, 2004; Kraut and Resnick, 2012b; Preece, 2000) come in many forms, including: those devoted to the sharing of media such as photos (e.g., FLICKR); those whose goal is to maintain an artifact of common value (such as WIKIPEDIA); those devoted to creating online social networks (e.g., FACEBOOK); and those which aim to exchange ideas and information about a given topic (e.g., the STACK OVERFLOW site² for asking and answering questions about programming).

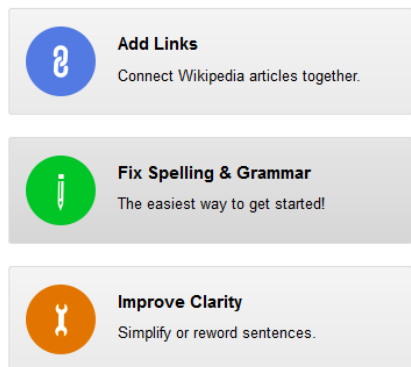
Common choices that a user faces with regard to a particular online community concern: whether they should participate in it at all; if so, whether they should actively contribute to it; and, if so, the form of their contribution, such as its type (e.g., a text or a photo), its

¹The principal authors of this chapter are Federica Cena, Cristina Gena, and Fabiana Venero.

²<http://stackoverflow.com>

Welcome to Wikipedia, *Fabiana veronesi*!

Looking for an easy way to get involved? Just choose one of the three options below, and we'll give you a suggested article to edit.



[← No thanks, return to the page I was reading](#)

Figure 11.1: Explicitly presented choices that new contributors to WIKIPEDIA encounter.

content (e.g., supportive or critical), and its quality (e.g., thoughtful or shallow). Figure 11.1 shows a page that presents new WIKIPEDIA contributors with choices of some of these sorts.

Online communities presuppose that at least some members will contribute media, knowledge, and/or other resources to the community. Since typically no particular amount or type of contribution is required for any individual member, there is a good deal of free choice. According to one way of looking at this type of situation, the optimal behavior of any given member is not to contribute at all but simply to benefit from the contributions made by others, as a “free rider” or “lurker”. But in general a community could not survive if too many members took this attitude. So it is not surprising that a good deal of research in this area has focused on users’ motivations, bringing to light a large variety of reasons why a user may or may not contribute (see, e.g., Ames and Naaman, 2007; Bishop, 2007; Beenen et al., 2004;

Preece et al., 2004; Wang and Fesenmaier, 2003) and often explicitly embracing a persuasion perspective, sometimes resulting in guidelines for community designers who want to encourage contribution (see, e.g., Kraut and Resnick, 2012a) or participation (see, e.g., Kraut et al., 2012).

Taking a choice support perspective, however, shifts our focus from persuading users to supporting them in choosing when, where, and how to participate and contribute. This perspective takes into account the fact that there may be occasions when not participating or contributing is the best option, for the individual member and/or for the community as a whole. But even for people (e.g., organizers of online communities) who are primarily interested in increasing the volume of contribution, this perspective is relevant, because of one of the general relationships between choice support and persuasion that were mentioned in 1.2.2: One way of encouraging someone to do X is to help her successfully make the specific preferential choices that are involved in doing X . For example, if some particular types of contribution are especially rewarding for a particular community member, helping him to identify those types of contribution is a way of making community participation more rewarding for him.

We will illustrate how choices made within online communities can be understood in terms of the ASPECT choice patterns and supported via the ARCADE strategies.

11.2 Choices About Whether to Participate

Before considering individual choice patterns, we will consider a type of community-related choice that can be addressed naturally with all of these patterns: the choice of whether to participate in a particular community at all. Some interesting examples of support for this particular choice are found in efforts of online communities to recruit new members. Even though the goal of recruitment is by definition to acquire new members, the goal is not in general only persuasion but at least in part choice support. After all, a community does not in general want just anyone to join; people whose goals are unlikely to be fulfilled and

who are unwilling to conform to the community's social expectations do not represent particularly desirable members. Kraut et al. [2012] give a number of examples of how communities provide combinations of persuasion and choice support to potential members, including the following:

- Advertisements or FAQ lists that provide information about the community, such as its mission statement, the functionality offered, and data about its members

This method instantiates the strategies *Access Information and Experience*, *Represent the Choice Situation*, and sometimes *Combine and Compute*. When processing this type of communication, a reader can apply the consequence-based pattern, anticipating what it would be like to participate in the community; the socially based pattern, taking into account any social examples that are offered; and maybe even the attribute-based pattern, if they are comparing a number of alternative communities in terms of attributes such as number of members or the cost of participation.

- Showing endorsements from credible sources

A video or text message from a well-known person likewise instantiates the strategy *Access Information and Experience*, though here the main information being offered is that the well-known person endorses the community, which is intended to trigger the advice-taking variant of the socially based pattern.

- Functions that make it easy for existing group members to invite other members whom they know personally to join

Here, the community is encouraging its members to apply the strategy *Evaluate on Behalf of the Chooser* by giving personalized recommendations to their friends and acquaintances. The underlying assumption is that an existing member should be able to do a pretty good job of evaluating the community on behalf of a person whom she knows, on the basis of her familiarity with both the community and the person. The recipient of such a message is being encouraged to apply the subpattern of the socially

based pattern that involves taking advice, though the fact that the inviting member is already a member can be seen as a social example.

- Making it easy for existing group members to share artifacts from the community with nonmembers (e.g., by emailing a link to a blog post)

To the extent to which this method is interpreted as an implicit invitation to join the community, it is similar to the previous method. But primarily it instantiates the strategy *Access Information and Experience* by providing a small sample of what it is like to be a member of the community and thereby supporting consequence-based choice.

- Advice to potential members about how to decide whether to join
For example, the “About” page of the community weblog METAFILTER³ includes the following advice (as of August 2014):

If you’re new to the site, I’d suggest taking a look around, checking out the archives, and getting a feel for the place. You might also consider registering as a member. Members can post comments, customize the look and behavior of the site.

After becoming a member, check out some of the links and think about leaving a comment or two. If you stick around for a while, you’ll get a feel for what types of things are posted as links, and if you find something amazing and/or enlightening, please post it.

The “New User” page of the same site makes it clear that the site owner wants people to think carefully before joining and participating in the community so that inappropriate members and behaviors will be minimized. Barriers such as the one-week waiting period for particular types of contribution can be seen as

³<http://www.metafilter.com/about.mefi>

a forceful form of procedural advice that prevents people from joining and contributing without careful deliberation.

- Provision of previews of what it's like to be a member

As Kraut et al. [2012] point out, the way in which many online communities archive past contributions and interactions among members makes it relatively easy for them to offer potential members a preview of what it is like to be a member: The community merely has to grant nonmembers access to (some of) these records, which may be seen as a low-cost instantiation of the strategy *Access Information and Experience*. The processing supported on the part of the chooser lies somewhere in or near the intersection of the trial-and-error-based and the consequence-based patterns, depending on the extent to which the potential member can be seen as trying out the community or simply acquiring information about the consequences of joining it.

- Requiring a potential member to perform a *separating task* (e.g., editing material written by other community members) that will help him and the community determine how well he fits into the community.

The tactic instantiates the strategy *Advise About Processing* by encouraging (or perhaps even forcing) the potential member to apply the trial-and-error-based pattern in an efficient way.

We will now turn to choices concerning specific contributions to online communities, organizing the discussion in terms of the choice patterns based on consequences, social influence, policies, and trial and error, giving examples of the experience- and attribute-based patterns in passing.

11.3 Consequence-Based Choices in On-Line Communities

In just about every case where a community member can decide whether to make a contribution, we can identify some aspects of the platform or context that can be seen as supporting (or at least influencing) the choice. In this section, we will take as a representative

Views/Day ↕	Quality ↕	Title ↕	Content ↕	Headings ↕	Images ↕	Links ↕	Sources ↕	Tagged with... ↕
23	★★★	Fender Prodigy (talk)	×	×	×	×	×	Add sources
18	★★★	Guitarist (magazine) (talk)	×	×	×	×	×	Add sources
10	★★★	Elite Stratocaster (talk)	×	×	×	×	×	Add sources
21	★★★	Yngwie Malmsteen Stratocaster (talk)	×	×	×		×	Add sources
19	★★★	Compression (physics) (talk)	×		×	×	×	Expand
38	★★★	The Experiment (talk)	×		×	×	×	Expand
113	★★★	Shred guitar (talk)	×		×		×	Unencyclopaedic
111	★★★	Pleonasm (talk)			×		×	Merge
2	★★★	Ibanez GR 220 (talk)	×	×	×	×	×	Merge
5	★★★	TV Jones (talk)	×	×	×	×	×	Wikify
132	★★★	Big School (TV series) (talk)	×		×	×	×	Wikify

Figure 11.2: Example of a table with which SUGGESTBOT recommends editing tasks to a regular WIKIPEDIA contributor. (See <http://en.wikipedia.org/wiki/User:SuggestBot>.)

example the SUGGESTBOT recommender for WIKIPEDIA editors, which grew out of research of Cosley and colleagues (see, e.g., Cosley et al., 2006; Frankowski et al., 2006; Cosley et al., 2007). The basic idea is to present to community members well-organized overviews of possible contributions along with information about the nature of the contributions, the cost of making them, and maybe also the likelihood that other members could make the same contributions.

The table of recommended articles in Figure 11.2 can be seen as facilitating consequence-based choice, though it also has features typical of support for attribute-based choice.

- With regard to the step of identifying options (Table 6.1): For a WIKIPEDIA contributor, the total set of options is in principle the huge set of all WIKIPEDIA pages—and in fact for each page, the contributor can decide which aspect(s) of it to work on. SUGGESTBOT applies the strategy *Evaluate on Behalf of the Chooser* to winnow this vast set of options down to a much smaller consideration set. The recommendation algorithms used to generate the options are described only vaguely on this page: “SuggestBot picks articles in a number of ways based on other articles you’ve edited, including straight text similarity, following wikilinks, and

matching your editing patterns against those of other Wikipedians.”⁴ Publications (cited above) by the group that developed SUGGESTBOT, which has expertise in the recommender systems field, have described a variety of recommendation algorithms that they have experimented with over the years, including one that favored items that the contributor was relatively uniquely able to handle.

- With regard to the anticipation and evaluation of consequences: The strategy *Access Information and Experience* is applied to provide information about two very different types of consequence: (a) what the experience of editing the page will be like (which can be inferred to some extent from the title of the page, the current quality rating, and the types of change that are required); and (b) what the benefits for readers will be if a given page is edited (e.g., to what extent the quality can be improved, which particular aspects of the quality will be improved, and how many viewers are likely to read the improved page).
- With regard to the step of making a final choice, the table instantiates the strategy *Represent the Choice Situation* by representing the information about consequences in a way that makes it relatively easy to apply various of the integration strategies discussed in 5.4 (which, as was noted in 6.9, can sometimes also be applied to the integration of evaluations of consequences). Note that the integration subtask is still not trivial: The representation does not help to distinguish between the two different types of consequence mentioned above; nor does it help the chooser devote more attention to the specific consequences that she considers most important—except that the contributor can sort the table by one of the columns if she considers its content to be especially important. (This example table has been sorted so that all pages requiring improvement to the images are in the upper rows, which might be appropriate for a contributor who either

⁴The contributor to whom the recommendations in this table are being offered has apparently edited guitar-related pages in the past.

enjoys contributing images or thinks that images are especially important parts of an article).

11.4 Socially Based Choices in On-Line Communities

Online communities have an inherently social dimension. So it is not surprising that considerations regarding social examples and expectations are often highlighted in the literature.

11.4.1 Explicit Social Expectations

Online communities provide some of the clearest examples of social expectations that are formulated explicitly. Many communities have standards that determine what counts as desirable and undesirable behavior and publish these as explicit guidelines or etiquette rules. To give just one example: The “Welcome” page of STACK OVERFLOW,⁵ like that of other STACK EXCHANGE communities, explains that “This site is all about getting answers. It’s not a discussion forum. There’s no chit-chat.” It also lists types of question that members should and should not ask. Guidelines like these may at first glance appear to be straightforward attempts at persuasion, but they can also be seen as applications of the ARCADE strategy *Access Information and Experience* with regard to social expectations: A community member who sincerely wants to conform to the expectations of the community would otherwise lack important information.

As Kiesler et al. [2012] note, informing community members about norms is not as straightforward as it may seem. Their “Design Claim 19” reads as follows: “Prominently displayed guidelines may convey a descriptive norm that the guidelines are not always followed.” In our terms, such guidelines, though intended to convey positive information about social *expectations*, can convey negative information about social *examples*. The more general point is that any attempt to apply the strategy *Access Information and Experience* should take into account the fact that the chooser may not take the presented information at face value but rather engage in inference and interpretation.

⁵<http://stackoverflow.com/tour>

We noted in 8.4 that one motivation for people to conform to social expectations is the anticipation that they will suffer negative consequences otherwise. Here again, online communities offer some especially clear examples, in the form of regulatory mechanisms for dealing with those who violate explicit community norms. Kiesler et al. [2012] discuss sanctions that communities impose, ranging from reputation penalties to outright bans. Here again, informing members about specific sanctions can be seen as an application of the strategy *Access Information and Experience*. Note that a community member who refrains from performing an undesirable action *only* because of concern about sanctions can be seen as applying the consequence-based pattern, which here shows some overlap with the socially based pattern. But often, community members see sanctions as just one reason to conform to expectations, alongside a perception of the legitimacy of the norms or a feeling of identification with the community.

11.4.2 Implicit Social Expectations

Of course community members can also perceive and interpret social examples and expectations even when these are not described explicitly. In these cases, it is often hard to know which particular variant(s) of the socially based choice pattern is (or are) being applied.

Representative examples of this range of possible interpretations are provided by two studies reported by Sukumaran et al. [2011], which examined a particular type of choice that contributors to online communities often have to make: whether (a) to provide a contribution that is quick and easy to make or (b) to take the time to work out a potentially more valuable contribution.

In their first study, two variants of a simulated news website were presented to two groups of participants (in a between-participant design). In the first site, each news story was followed by mostly “thoughtful” (i.e., relatively long and carefully formulated) comments provided by fictitious users; the second site was the same except that the user comments were mostly not thoughtful (i.e., short and simple-minded). Participants who used the first site tended to write more thoughtful comments of their own than those who used the second site; and they

were more inclined to predict that other users of the site would provide thoughtful comments.

Considering what was said in Section 8, we can see two possible types of explanation for these results, both of which might be valid at the same time:

1. Users of the first site may have inferred from the prevalence of thoughtful comments that there was a general expectation on the part of other users that comments on this site ought to be thoughtful. If so, they may have been motivated to conform to this expectation for either of the following two reasons discussed in Section 8:

- They could be rewarded or punished for conformance or non-conformance (e.g., in the form of supportive or critical feedback about their comments).
- They considered the expectation to be legitimate (e.g., “If I am going to benefit from using this news site, I ought to respect the expectations that members of the site have concerning my behavior”).

Whether this explanation is valid could have been tested, for example, if the investigators had asked participants in some appropriate way whether they thought they were expected to provide thoughtful comments and whether they were motivated to conform to any such expectation.

2. The participants may simply have viewed the comments of others as examples of how other people behave on the site in question. As was discussed in Section 8, there are several reasons why they might be inclined to follow these social examples even if they did not assume that there was any expectation that they would do so. For example, perhaps people have been making thoughtful comments because they have found it rewarding to do so on this particular site (e.g., because such comments elicited equally thoughtful responses from other users on the site).

Note that these two possible explanations for the effect of the provided social examples suggest different ideas about how to support choice in this situation: For example, applying the strategy *Access In-*

formation and Experience, we could inform users about either the prevailing expectations about behavior or the reasons why people find it worthwhile to provide thoughtful responses on the first site.

The second study by Sukumaran et al. [2011] admits an even wider range of possible interpretations. Here again, two simulated news sites were used that differed in the extent to which they seemed to encourage thoughtful comments. In this second study, no social examples were provided, but the first site differed from the second one in terms of variety of features: how serious the visual style of the site was; whether the words used in the CAPTCHA form were semantically related to thoughtfulness; what specific texts were used to encourage users to comment; and whether the commenting box made it easy to enter a long, thoughtful comment.

As in the first study, users of the first website tended to produce more thoughtful comments; but in contrast to the first study, they did not tend to predict that other users would make relatively thoughtful comments.

In this study, an interpretation in terms of social expectations or examples is less plausible than in the first study, because of the absence of any social information. But there are possible explanations in terms of other choice patterns:

1. The effects of the visual style of the website and of the words used in the interface can be seen as illustrations of the point made in 5.2 about how a person's evaluation criteria in a particular situation can be influenced by subtle aspects of the representation of the choice problem, some of which are logically irrelevant (recall, for example, the study of Mandel and Johnson, 2002, on the influence of the background pattern of a commercial website). In the news site situation, there is probably no compelling consideration that determines how much importance a person should attach to the goal of providing a thoughtful answer relative to the goal of getting through the task quickly, so contextual factors like these can have a large influence. Note that, if a designer manipulates factors like these with the intention of influencing a user's evaluation criteria, they are applying the strategy *Represent the Choice Situation*.

2. The fact that the commenting box in the first website made it easier to enter a thoughtful comment can be understood in terms of the consequence-based pattern: It changed the consequences of attempting to enter a thoughtful comment by reducing the associated cost. Designing the commenting box in this way with the goal of encouraging thoughtful comments can be seen as an instance of the strategy *Design the Domain*, one which is biased in the direction of a particular choice.⁶

Note that neither of these two explanations is specifically related to social contexts; they can be applied even when no other people are involved.

A broad understanding of choice patterns can help us with the interpretation of studies like these by reminding us of the full range of possible explanations for observed phenomena; and it can help us to narrow down the set of possible explanations in advance by designing the studies in such a way that the results allow us distinguish among different explanations.

11.5 Policy-Based Choices in On-Line Communities

We noted in 9.4 that one simple way of acquiring a policy concerning a particular type of choice is to adopt a perceived social expectation as a personal policy. For example, a community member does not need to deliberate about whether to criticize another member in a contribution if he has adopted for himself the community's expectation that personal criticisms are to be avoided.⁷ The explicit formulations of behavioral norms just discussed can be seen as support for not only socially based but also policy-based choice: In particular, if the community norms are designed in such a way that members can easily understand, learn, and apply them, they are more likely to be adopted as individual policies.

⁶The tactic of encouraging people to do something by making the thing simpler or easier to do is referred to in the persuasive technology literature as *reduction*; see Fogg [2003, chap. 3].

⁷Note that he may actually adopt a policy which is related to the norm but is not identical to it, such as: "When you feel like criticizing someone, instead praise them in a recognizably exaggerated way."

A very different type of example of support for policy-based choice can be found in the study of Ames and Naaman [2007], which looked at (among other things) how people used a mobile app called ZONE-TAG for annotating photos with tags before uploading them to FLICKR. Choosing tags for each photo before uploading it can be impractically time-consuming, especially if the user is on the move. The functionality of ZONE-TAG made it possible, when making multiple photos of a given event, to specify some tags for the first photo and then reuse these tags with minimal effort for all subsequent photos of the same event. A user who thinks ahead a bit can decide to proceed in this way even if she would be unwilling or unable to tag each photo individually. This example illustrates again one of the general benefits of policy-based choice: the fact that applying a policy typically requires less time and effort than most other ways of making a choice, especially if the application of the policy has been at least partly automated. In cases where the user would otherwise not bother to tag the photos at all, the support offered for policy-based choice presumably tends to lead to better choices (i.e., one or more meaningful tags vs. none at all), even though tagging of individual photos might yield better choices if users had time for it. From the point of view of design support for policy-based choice, we should note that the “tag stickiness” feature just mentioned was apparently not explicitly designed or presented to users as a form of support for policy-based choice of tags. In consequence, users had to discover this way of using the function themselves, and it was not well suited for all situations (e.g., those in which a user wants to make an occasional photo that does not belong to the main event that he is photographing repeatedly). The design of functionality like this with the explicit goal of supporting policy-based choice could make use of some of the other ideas discussed in Section 9.

One interesting consequence of policy-based choice in the context of online communities is that it can lead to relatively high predictability of an individual’s choices, which can be helpful to other persons (e.g., those who want to identify all of the photos that a user took of a particular event). On the other hand, strict adherence to a policy sometimes yields choices that seem wrong or even incomprehensible in

isolation, at least if the perceiver is not aware of the policy being applied (as when a user employing the strategy sketched above takes a picture during an event that is unrelated to the event while still using the tags for the event).

11.6 Trial-and-Error-Based Choices in On-Line Communities

Especially people who have not yet joined a community or have only recently joined it sometimes have only a vague or inaccurate notion of what it is like to participate in the community and to make specific types of contribution—even if the community offers the sort of preview that was discussed in 11.2. In these cases, applying the trial-and-error-based strategy is a natural approach that can be supported. As was discussed in Section 10, we can distinguish between support for exploration and support for the interpretation of consequences of actions.

11.6.1 Support for Exploration

An example of explicit advice concerning an exploration strategy, from the METAFILTER weblog, was already given in 11.2.

A different tactic is to apply the strategy *Design the Domain* to create an environment that generally facilitates exploration. An example is provided by the “Sandbox” pages of WIKIPEDIA.⁸ One function of these pages is to enable people to acquire skill in using WIKIPEDIA’s markup language and related functions; but these pages also support preferential choice by enabling a user to experience some aspects of what it’s like to edit a WIKIPEDIA page (e.g., whether she finds the use of the markup language cumbersome or distracting). This example illustrates an interesting complication that often arises in connection with exploration: Note that, even as a user is deciding whether he enjoys using the markup language, he is also getting more skilled at using the language, and this increase in skill is likely to influence his level of enjoyment. In cases like this, it is desirable to encourage and facilitate exploration that is extensive enough to enable the chooser to arrive at

⁸<http://en.wikipedia.org/wiki/Wikipedia:Sandbox>

a reliable assessment of what the consequences will be like in the long run.

Another variant of the same tactic is to encourage and facilitate what are sometimes called *microcontributions* (Frankowski et al., 2007) or *peripheral contributions* (Bryant et al., 2005), such as rating, tagging, and commenting on items. Although these are genuine contributions (unlike actions in a sandbox), they are associated with smaller consequences (e.g., in terms of the effort required and the responses of other members). Although microcontributions can benefit a community in various ways, one function is to enable members to experience what it is like to contribute to the community—at relatively low cost and risk to themselves and to the community. A limitation of this function is that the experience of making microcontributions is unlikely to be wholly representative of the more general experience of making contributions (cf. Bryant et al., 2005). This nonrepresentativeness is one reason why encouraging microcontributions has been found to be effective as a compliance tactic for getting people to make larger contributions (see, e.g., the discussion of the *foot-in-the-door technique* by Cialdini and Goldstein, 2004, pp. 602–604), but from the point of view of choice support it is a drawback. Note that the procedural advice given by METAFILTER (quoted in 11.2) encourages the prospective member to progress from small to larger contributions, presumably so as to require more representative experience.

11.6.2 Support for Interpretation of Consequences

Both published research and everyday experience confirm that the contribution choices of contributors to online communities are influenced by the consequences of the contributions that they have made in the past. For example, Preece et al. [2004] found that unpleasant past experience resulting from contributions was named by lurkers as a reason for not actively contributing to a community. The negative experiences named ranged from difficulties with the operation of the platform to unsatisfactory interaction with other group members (e.g., “Received

a rude response to a past post”, p. 215).⁹ Especially in the case of strongly negative experiences, it seems likely that the users not only became aware of negative consequences but also formed negative affective associations with the actions in question, which could lead to experience-based choices not to contribute (cf. 7.5).

Experiencing consequences can also influence the particular form that a contribution takes. For example, in their study on motivations for tagging, Ames and Naaman [2007] report an example of learning from experience that resulted from positive feedback that one user received: After having observed that photos with (particular) tags tended to receive more views, this user decided always to add tags to his/her public photos, as a way to promote their content and gain reputation. In this case, the photo sharing platform was designed so as to make consequences related to feedback from other users (e.g., the number of views) easily noticeable.

Experience with the consequences of contributions can be helpful even to experienced community members in cases where the community environment changes in ways that ought to be taken into account in future choices. Velasquez et al. [2014] found that changes and transformations in the community environment (e.g., concerning content-related policies and shared goals) can cause some users to become “latent”—that is, no longer to contribute content in the same quantity and fashion as they used to. For example, some participants in the user-generated EVERYTHING2 encyclopedia who were interested in building new relationships and getting to know other community members decreased their level of contribution when they saw that their posts containing mainly subjective thoughts and personal discussions were being accepted less frequently and might even be deleted by editors, as a result of a new content policy.

Designers of online communities have a great many opportunities to control both the consequences of members’ choices and the ways

⁹Sometimes the negative experiences observed were those of other community members, which illustrates that people can also learn from the “trial and error” of other persons.

in which the consequences are perceived and interpreted by members, using several of the ARCADE strategies:

- Applying the strategy *Design the Domain*: Since feedback is in general generated by other users, various features of a community platform can influence how much feedback other users provide and what the content of the feedback is. Note that other users who provide feedback are themselves making choices about what feedback to provide. So here we have a case where platform support for one type of choice (“What sorts of contribution should I make?”) takes the form of support for (or persuasive influence of) another type of choice by other persons (i.e., “What sort of feedback should I provide?”).
- Applying one or more of the strategies *Access Information and Experience*, *Represent the Choice Situation*, and *Combine and Compute*: The platform can influence how any feedback provided by other users is perceived by the contributor: Whether the contributor can perceive it at all and how it is represented to the contributor (e.g., in its original raw form; as a histogram showing the distribution of the expressed evaluations, ...).

Tactics based on these strategies can be used in various ways to try to overcome some of the typical pitfalls of feedback interpretation that were discussed in 10.4:

- Ensuring that an adequate amount of feedback is available (e.g., by making it easy and/or motivating for other members to provide feedback).
- Minimizing bias in the feedback (e.g., by aiming to ensure that the interaction design does not introduce bias and maybe even counteracts bias, for example by encouraging members to supply neutral as well as positive or negative feedback).
- Ensuring timeliness of feedback and a clear relationship to the contribution that it refers to (e.g., by ensuring that both the original contribution and any feedback to it are made available without unnecessary delay).

- Ensuring that the feedback given is diagnostic rather than simply evaluative (e.g., by providing not only an evaluation scale but also an opportunity to express more specific ideas).

11.7 Concluding Remarks on Choice in Online Communities

One of the purposes of this section has been to show that designers of online communities can fruitfully take a choice support perspective in addition to the currently more widespread persuasion perspective. The domain of online communities does not merely offer a wealth of examples of interesting opportunities for choice support; it also illustrates the point made in 1.2.2 that persuasion is often most effective when combined with choice support, and vice versa.

As we've mentioned, even a community that uses persuasive techniques to increase its membership will not in general want just anyone to join; a member who fits well is one whose decision to join the community is a decision that she will be satisfied with in the long run. In the other direction, supporting trial-and-error-based choice by ensuring that members receive enough useful feedback on their contributions may require applying persuasive methods to encourage other members to provide feedback.

We hope that this section will encourage community designers to view choice support and persuasion as being best used in combination, like the white and black keys of the piano.

12

Choices Concerning Privacy

12.1 Introduction

Whereas in connection with online communities the persuasion perspective has been much more prominent than the choice support perspective, in our second example domain—privacy-related choices—researchers often explicitly state that they are interested in understanding the choices that people make and helping them to make better choices (see, e.g., the comprehensive survey article by Iachello and Hong, 2007).¹ On the other hand, this research literature on the whole makes little reference to the psychology of choice and choice support. Instead, we often see conceptions of choice that are hard to relate to any well-founded psychological concepts. In particular, a lot of the relevant discussion in this area refers to the concept of *privacy preferences*, which is more misleading than helpful, as we will discuss in 12.5.1. One aim of the present section is to show how a clearer understanding of privacy-related choices can be achieved when concepts like this one are replaced with psychologically grounded ones such as those from the ASPECT model.

¹The principal authors of this chapter are Bettina Berendt and Silvia Gabrielli.

We also wish to call attention to the often underestimated complexity of privacy-related choices relative to most other types of preferential choice in the HCI area. For example, a popular conception is that the quintessential privacy choice is the question of whether to make a photo available to friends or friends-of-friends on a social networking site. We believe that much of the current confusion and concern about poor privacy practices arises from a lack of a principled understanding of privacy-related choices. We therefore dedicate some space to an analysis of this type of choice before showing in subsequent subsections how the ASPECT and ARCADE models can be applied to them.

12.1.1 Three Scenarios for Privacy-Related Choice

We will refer to three scenarios, which can be illustrated as follows with examples:

1. *The social networking site scenario:* A chooser is in possession of a (self-made) photograph of himself, his partner, and their child. The chooser wants to share the photo with some recipient(s) on a social networking site. In the course of this transaction, personal data and information get created and/or transmitted.

2. *The e-commerce scenario:* The chooser wants to purchase some product or service online. The chooser can or has to supply some information *in addition* to the information required for this basic transaction.

3. *The ubiquitous computing scenario:* This scenario covers the main privacy choices of users interacting with ubiquitous and location-enhanced technologies, including sensors (e.g., RFID sensors). The chooser is, for example, an urban traveler who needs to decide whether a particular (group of) people should be allowed to find out her location.

12.1.2 Goals of Choice: What Is Privacy, and What Are “Good” Privacy-Related Choices?

Throughout time and across cultures, people have shown patterns of interacting with others and withdrawing from them. Altman [1976, p. 7] has described these behaviors as *boundary regulation processes*, in

which privacy is “a selective control of access to the self or to one’s group”. These boundary regulation processes have two essential features. First, privacy involves both closedness and openness: The seeking and the avoiding of social interaction are mutually contingent. Second, the “ideal level” (however one might define that concept) of social interaction or privacy changes with time and context. The ideal level also varies because of individual and cultural differences. These control processes for boundary regulation involve the sharing or disclosing of information vs. the withholding of information, as was investigated in detail in the work of Petronio [2002]. Westin [1967] offers a related notion of *information privacy*: “the claim of individuals, groups or institutions to determine for themselves when, how and to what extent information about them is communicated to others” (p. 7). Definitions such as Westin’s (and also their codification in legal regulations on data protection and privacy rights) emphasize that privacy is not only or always about “being let alone” and “disclosing as little data as possible” (see Berendt [2012] for a detailed discussion).

We will distinguish between three types of “others” and the forms of privacy associated with them:²

1. Other people, usually peers. The corresponding form of privacy is called *social privacy*.
2. Commercial entities or other institutions, which raise issues of *institutional privacy*.
3. Governmental entities, which raise issues of *protection from surveillance*.

We define a *privacy-related choice* of a chooser as any choice that results in the granting of access to data about a person to one or more (potential) *knowers*—or withholding access from them. The person whom the data concerns is usually the chooser himself, but it can also be someone else.

²These classifications have been discussed by many authors and with different names; we use the terminology introduced in Raynes-Goldie [2012] and Phillips [2004]. *Information privacy* is often associated with “others” of types 1 and 2.

The central privacy goals, in the sense of a *good outcome* for the person (cf. 3.6), are to achieve the level and/or form of privacy desired by the person. As with other choices, not only considerations of outcome play a role.³

Design choices that have privacy implications are determined not only by end users' goals but also by the goals of site operators. Often, the underlying business model rests on users paying for a free service with their personal data. Our examples will illustrate support for both types of goals.

12.1.3 A Key Characteristic of Privacy-Related Choices: Multiple Interrelated Choices in Several Dimensions

One key characteristic and challenge of privacy-related choices is that they rarely come alone. Instead, a privacy-related choice is often bundled with another choice, often one that is more important to the chooser. Hence a privacy-related choice will often receive less attention than it deserves simply because of limits on attention and time—or it may even be neglected completely. A narrow focus on only one of the bundled choices can also be encouraged by site operators who are mainly interested in obtaining data, as well as by cultural factors.

The first form of bundling concentrates attention on what may be called a *main choice*, which makes at least part of the privacy-related choice a *secondary choice*:

- The main choice in the e-commerce scenario is how to execute a commercial transaction effectively: The user wants to buy a winter jacket, not make decisions about his privacy.
- The main choice in the ubiquitous computing scenario may be how to reach your destination quickly and safely.

³Even this simple definition is faced with a number of challenges. The person's desires may be unknown or unclear, and they may change over time. If the person is not the chooser herself, then the person's desires may conflict with the chooser's wishes. Other people may decide which levels are applicable for the chooser/person, for example when the chooser is below the age of consent or where privacy is regarded as a public good. In the following discussion, we will abstract away from these additional complications.

- The main choice in the social networking site scenario is a choice related to social privacy (e.g., how to show a photo to a particular relative and whom else to show it to). The choices related to institutional privacy—what information should be shared with the site or provider—are often not even perceived.

In all three scenarios, surveillance (i.e., the sharing of information with a governmental entity) generally remains invisible. And although the outcomes of the privacy-related secondary choices may be highly relevant for the chooser, they are often invisible and drawn out over time, and they may be caused by people other than the chooser (see, e.g., Gürses and Diaz, 2013).

A second form of bundling is that one person's privacy-related choices constrain other people's privacy and privacy-related choices, now or in the future:

- In the example given for the social networking site scenario, publishing the photo also discloses information about the other two persons depicted. Interface options for letting the affected persons co-decide are not straightforward and rarely available; and some persons (such as the child in the photo) may not even be able to voice or enforce their decisions.
- Norms of social reciprocity and ease of handling strongly suggest that reactions to a privacy-relevant action should be delivered over the same channel (e.g., social networking site)—which in effect requires people to register and communicate on the same platform.
- Professional norms or concrete activity decisions (such as a decision to teach a class using GOOGLE HANGOUT) can also enforce membership in and activity on certain platforms.
- Choices in e-commerce scenarios may also be constrained by the simple nonexistence of privacy-friendly providers or options, by an excessive price exacted for protecting one's privacy, or by many other factors that make reality differ from an ideal market in which demand for a service can always be met by supply and

no participant is more powerful than any other in determining prices.

- Some constraints on privacy-related choices arise from the joint workings of individual and collective behaviors and technical factors. For example, in order for a system to generate real-time information for ride sharing, multiple parties have to provide location-based data. Another example is the safe encryption of messages with public-key cryptography: This practice requires all participants in an exchange to generate, manage, exchange, and use keys.
- All of these behaviors can also give rise to social examples and expectations (see Section 8) and to habits (7.3) that in turn influence future privacy-related choices in the direction of increased disclosure.

12.1.4 How to Support Privacy-Related Choice? “Nudges” Vs. Awareness Support

In this section on privacy-related choices, we will describe a number of choice support tactics, some embodied in entire tools for supporting privacy-related choices, with reference to the ASPECT and ARCADE models. Many of these tactics can be supported by two categories of tool that currently constitute foci of development and investigation: *privacy nudges* and *data-based privacy awareness tools*. Since tools in both of these categories can support all of the ASPECT choice patterns, we will introduce the basic distinction between them here.

Privacy Nudges

The concept of a *privacy nudge* is nicely illustrated by the three examples presented by Wang et al. [2013]. Although the term *nudge* is commonly associated with influencing choices in a particular direction (cf. the discussion of this concept in 1.2.2), these interface elements can be interpreted naturally within the ARCADE model as forms of choice support. The context is one in which a user of FACEBOOK is about to post some content.

The *picture nudge* shows the user five randomly chosen pictures of members of the set of people who will see the post. This method instantiates the strategy *Represent the Choice Situation* by augmenting an abstract representation of the set of recipients (e.g., “anyone on the internet”) with a more concrete representation in terms of particular individuals. This type of representation presumably encourages thinking about how the various possible viewers of the post will respond to it.

If the five pictures presented were not selected at random by the system but rather chosen so as to be representative of the set of people who were actually likely to view the post, this nudge would also instantiate the strategies *Access Information and Experience* and *Combine and Compute*.

The *sentiment nudge* of Wang et al. [2013] alerts the user if the content of the post seems likely to be perceived as “negative”, as determined by analysis of its words by a sentiment analysis module. This nudge both evokes the consequence-based pattern and applies the strategy *Combine and Compute* to help the chooser predict a particular type of consequence of her action.

Finally, the *timer nudge* explicitly introduces a delay of 10 seconds after the user has submitted the post, during which he can cancel the post. This method is similar to simply advising the user to “think for 10 seconds before sending off your post”, the difference being that the user is almost forced to take the advice (since he has little else to do during the 10 seconds). Hence the nudge implicitly instantiates the strategy *Advise About Processing*. It also illustrates once again (cf. 4.4) that in interactive systems, as opposed to human-human dialog, procedural “advice” often takes the form not of verbal advice but (also) of interaction design that makes it especially convenient or even necessary to take the advice.

Data-Based Privacy Awareness Tools

Data-based awareness tools (see, e.g., Gao and Berendt, 2011; Berendt et al., 2012; Wang et al., 2013) are software tools that apply the strategy *Access Information and Experience*, often in conjunction

with *Combine and Compute* and *Represent the Choice Situation*, in any of various ways to make users aware of potential privacy threats. For example, in support of the consequence-based pattern, the “awareness” can concern the need to make a choice, the assessment of the current situation, and the possible consequences of particular actions (see 12.2 below). In support of the trial-and-error-based pattern, users are made more aware of the consequences of actions once they have performed them (12.6). Data-based awareness tools may—but need not—include privacy nudges.

After this introductory discussion, we will consider each of the ASPECT choice patterns in turn (except for the experience-based pattern), in each case looking at examples of how the concepts associated with that pattern apply to some types of privacy-related choice.

12.2 Consequence-Based Choices About Privacy

It is natural for people to apply the consequence-based pattern to privacy choices, since poor choices can often lead to unfortunate consequences. But this choice pattern raises special challenges for several of the steps in this pattern (see Table 6.1).

12.2.1 Situation Assessment and Recognition of Choice Opportunities

In an ideal world, users would always know when some sort of privacy-related choice was called for and what the relevant features of the current situation were (6.2, 6.3).

In some cases, interface elements indicate straightforwardly where a privacy-related choice does or does not have to be made, as in the following examples, which instantiate the strategies *Access Information and Experience* and *Represent the Choice Situation*:

- Checkboxes for opting out of or into the transfer of information to particular third parties alert the user to a need to choose who will or will not receive particular data.

- A red asterisk in front of a mandatory data field indicates that there is no choice about whether to provide this information.
- The conspicuous provision of telephone contact numbers on the screen signals to the user that she can choose an alternative method of transferring information to the institution in question.

It would in general be infeasible, however, to alert users in a conspicuous way about every privacy-related choice opportunity—even aside from the fact that organizations and individuals who are interested in acquiring personal data often have no interest in increasing users' situation awareness.

For example, in principle, whenever a user visits a website that is going to make use of his personal data in one or more ways, the user might want to choose whether to allow those uses of the data. But often users have no idea what use (if any) is going to be made of their data, and they can hardly take the trouble to find out every time they visit a website. One well-known approach to this problem was the PRIVACY BIRD (Cranor et al., 2006, discussed in 12.5 below), which automatically checked for possibly undesirable consequences and alerted the user to cases where an explicit choice appeared to be worthwhile.

As is discussed by Iachello and Hong [2007, p. 55–57], users of web browsers regularly encounter cues that are provided with the goal of improving situation assessment, such as “lock” icons and warnings that data may be intercepted by third parties. One problem with such cues, which is typical of attempts to support situation assessment via the strategies *Access Information and Experience* and *Represent the Choice Situation*, is that users have only a limited capacity to attend to them, especially since the cues in general do not concern the user's primary task. Even if such a cue is noticed, the user may not be able to make any of the inferences that situation assessment is supposed to support, concerning what will happen if the chooser takes no action and what the consequences of particular actions will be—because of a problem that is mentioned repeatedly in this section: the incomplete mental models that most users have of privacy-relevant factors. For example, if you are informed that data that you enter into a web form might be intercepted by a third party, you probably have little idea of how

likely this event is to occur, who might do the intercepting, or what use they might make of the data. Warnings of this sort are often best seen not as helpful applications of choice support tactics but rather as steps taken to protect a stakeholder from criticism or legal challenges (“Don’t blame us; we warned you!”).

In the ubiquitous computing scenario, we have the additional challenge that it tends to be more difficult than with graphical user interfaces to provide useful cues to support situation assessment (see, e.g., Nguyen and Mynatt, 2002; Iachello and Hong, 2007, pp. 55–57).⁴ In connection with older devices such as audio recorders, warnings such as flashing red lights and regularly spaced beeps have long been used to instantiate the strategy *Access Information and Experience* to support situation assessment. A current challenge is to find analogous methods to provide adequate awareness of the much larger range of sensors that now exist even in an ordinary smartphone (including, e.g., GPS, motion and orientation sensors, and video cameras). The PRIVACY MIRRORS framework of Nguyen and Mynatt [2002] takes into account the broad range of things that users of ubiquitous computing technology in principle ought to be aware of: not only the technical properties of the systems that they are dealing with (cf. Edwards et al., 2001) but also relevant aspects of the physical and social environments.

Many of the specific methods that have been developed for increasing users’ situation awareness can also serve to provide feedback in the context of the trial-and-error-based pattern (see 10.5 above and 12.6 below.)

12.2.2 Deciding When and Where to Make a Choice

The question of when a particular choice is best made (6.4) has been discussed especially often in the privacy domain. One key question is whether users can be expected to make configuration choices at one point in time that will relieve them of the need to make a large number of specific choices later on. For example, Lederer et al. [2004, p. 447–448] argue that “Emphasizing Configuration Over Action” is one of

⁴Bellotti et al. [2002] discuss related issues with “sensing systems” on a more general level.

the typical pitfalls associated with the design of interactive systems that have personal privacy implications. Since having users choose privacy settings is essentially a matter of asking them to select a policy that will be implemented (semi)automatically by the system, this issue is discussed below in the subsection on policy-based privacy choices (12.5).

12.2.3 Making an Appropriate Set of Options Available

One source of difficulty in choosing is when an unnecessarily restricted set of options is available that forces the chooser to deal with negative consequences or difficult tradeoffs that would not arise if a more suitable option were available. An illustration of this point is provided by the efforts made within a ubiquitous computing scenario by Iachello et al. [2005] to determine empirically what set of options would be found suitable by users of an application that enables people to disclose their locations to each other. Options that were considered by the study participants to be useful to have included (a) sending a vague, evasive response such as “I am busy”; and (b) sending an inaccurate indication of the user’s location (e.g., “on the way home” vs. “still at the office”).

On a more general level, Lederer et al. [2004, p. 448] argued that failing to provide “an obvious, top-level mechanism for halting and resuming disclosure” is another of the pitfalls in the design of privacy-relevant systems. They noted as an example that, at the time of their writing, there was typically no simple way for the user of a web browser to block *all* cookies during part of her browsing session. In the intervening years, many web browsers have filled this gap in the set of available options by introducing a “Private Browsing” mode that can be turned on and off easily.

Judiciously expanding the set of options in this way is an application of the strategy *Design the Domain*, which can be favorably compared to the alternative of using other ARCADE strategies to help users to choose among less desirable options.

12.2.4 Holistic Representations of Privacy Choices

One reason for the difficulty of privacy-related choices is that it is often unclear to choosers what the consequences of the various options will be. As we saw in Section 6, there can be many reasons for difficulty in anticipating consequences. One difficulty that seems relatively characteristic of privacy choices in current web-based and ubiquitous systems is due to the way in which options are represented to users—that is, to a problematic application of the strategy *Represent the Choice Situation* (cf. the discussion of framing in 6.7.2): Options are often represented to users in what may be called a *holistic* way that makes it relatively hard to perceive individual consequences clearly. After discussing this type of representation, we will argue for the benefits of more specific representations of individual consequences.

A privacy-related choice situation is often represented by the system's designers as one in which it is desirable to disclose information. This representation can induce users to disclose more information than they otherwise would. In these cases, the strategy *Represent the Choice Situation* is being used with a persuasive intent (cf. 1.2.2) to push the user toward an option that is desirable from the point of view of those who run the system.

One frequent representation is that of *information disclosure as part of an exchange*. This representation draws on real-life characteristics of information disclosure and privacy, especially in e-commerce scenarios. In many commercial and other public transactions, a participant needs to disclose private—or even sensitive—data in order to enable the transaction to occur: In the clothing store, you need to state your size in order to obtain clothes that fit; in the doctor's office, you need to talk about your ailments in order to receive diagnosis and treatment. Many online stores and services promise personalization advantages in return for data. The *perception* that personalization is occurring can lead people to disclose more information about themselves, even if everyone is in fact being given the same recommendations (Kobsa and Teltzrow [2004]). The assumption that questions are asked because they are relevant in the given context may lead people to answer even questions

that would be considered illegitimate outside of the context suggested by the representation (Berendt et al., 2005; Berendt, 2009).

In the social-networking-site or ubiquitous computing scenarios, options that involve information disclosure are often described as involving an exchange of data from the user for a free service from a provider—though the relevant formulations are generally to be found in relatively hard-to-reach documents such as terms-and-conditions pages.

By contrast, the fact that data are being exchanged among peers is made salient by the interface design: Interactive elements for reacting to a new piece of content are placed right next to that content. The fact that such exchanges also feed data to the service provider remains in the background. FACEBOOK typically frames choices made when using the social-networking site as choices about communication with peers and empowering internet users.

Outside of the HCI context, similar framing has been used in the “interfaces” of questionnaires about surveillance choices. For example, in a large-scale survey after the 2013 leaking of details about the PRISM surveillance programme in the U.S., Pew Research⁵ asked questions such as “Should the government be able to monitor everyone’s email to prevent possible terrorism?” and obtained widespread acceptance.

Another type of representation that may have a persuasive intent involves representing multiple privacy-related choices as if they were basically the same—encouraging broad choice bracketing (9.2)—even if the choices differ quite a bit in terms of their consequences for the user. Figure 12.1A shows the beginning of an explanation of choices about providing location information in iOS 7. Note that the choices are framed as whether to provide location information in order to benefit from a service offered via the iPhone. It is only by reading considerably further that the user can learn that some of the services in question benefit not the individual user who is providing the data but rather users in general—for example, by contributing to a crowdsourced road traffic database. Similarly, in the screen that lists system services (Figure 12.1B), there is no visual distinction between services of the first type (such as “Setting Time Zone”) and those of the second type

⁵*Public Says Investigate Terrorism, Even If It Intrudes on Privacy*, 10 June 2013.

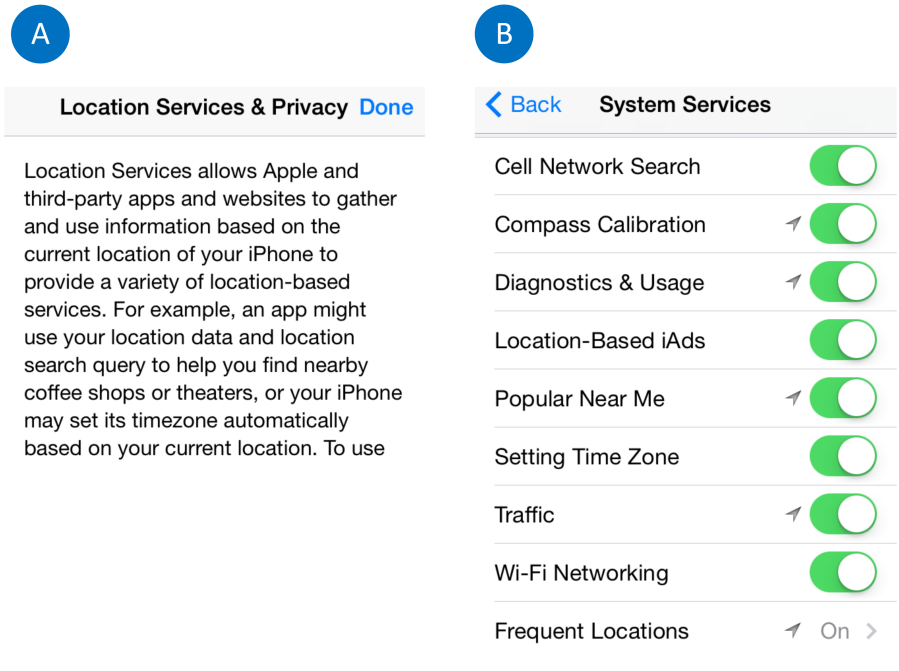


Figure 12.1: Two screens from Apple’s iOS 7 that illustrate how significantly different privacy-related choices can be bundled together in a manner that encourages the chooser to deal with them in the same way.

(e.g., “Traffic”). Finally, there is a switch for turning off *all* location services but no such switch for turning off only services of the second type. If we assume that the distinction between the two types of system service is decision-relevant, a natural form of choice support is to make the difference easily recognizable (*Represent the Choice Situation*) and to provide controls for conveniently handling the two types of service differently (*Design the Domain*).

12.2.5 Challenges and Approaches for Predicting Specific Consequences of Privacy-Related Choices

To put it simply, the consequences of privacy-related choices will be that someone will know something and/or that actions affecting the chooser or others will occur, both of which may inflict harm or bring

benefits (Berendt, 2012). Privacy-related choices face just about all of the general challenges (discussed in Section 6) raised by the prediction and evaluation of possibly uncertain future consequences (e.g., hyperbolic time discounting, as was argued by Acquisti and Grossklags [2004]). They also face a number of specific challenges, most notably the general impossibility of undoing choices (12.6) and the often unpredictable consequences of data processing.

A piece of data that is disclosed does not only get recorded and retrieved; it also gets *processed*. It is hard if not impossible to predict what knowledge (and actions) a piece of data—or even the fact that a piece of data is missing—can contribute to after it has undergone (perhaps repeated) linking to other information and processing such as data mining.⁶ Cumulative risk assessment—understanding how small pieces of data disclosed over time can be combined to contribute to a profile of a person—is even more difficult.

Choices in ubiquitous computing scenarios face a number of specific challenges related to the fact that the ultimate consequences of particular choices tend to be hard to predict without the benefit of specialized knowledge or experience. In a study of people’s responses to five methods for obfuscating GPS trace data collected from them, Brush et al. [2010, p. 7] found that their study participants largely understood the basic operation of the obfuscation methods—but that they had a hard time understanding what the consequences would be of applying a given obfuscation method to data collected over a period of time. For example, participants considered the *subsampling* method of systematically leaving gaps in the collection of data to be basically acceptable, evidently not noticing that its application for an entire day would normally reveal the user’s home location—a consequence which, in other contexts, they judged to be unacceptable.

Nguyen et al. [2008, p. 189] asked a broad variety of shoppers in U.S. shopping malls about the possible dangers of six tracking and recording technologies: credit cards, store loyalty cards, electronic toll collection systems, web server records, store video cameras, and RFID

⁶For a seminal HCI perspective on this issue, see Dix [1990]; for a recent overview of challenges and technical solution approaches, see Berendt [2012].

sensors. Whereas respondents showed a clear understanding of the *benefits* of these technologies, “they had difficulties articulating possible costs or threats of these technologies” (p. 186). For example, though they could easily enough anticipate a straightforward possible consequence of divulging their credit card information—that someone might use their data to make an unauthorized purchase—they did not mention the consequences of allowing organizations to build long-term records of their purchases.

In sum, the interrelatedness of privacy-related choices and the invisibilities, indirections, delays, and nondeterministic consequences of the multiple actions that stakeholders perform with personal information (see also 12.1.3) make it hard for a user to predict what may happen as a result of any given privacy-related choice: A very wide or even unknowable range of knowing and action consequences could result from a privacy-related choice. On the other hand, alerting choosers to this wide range of consequences and the difficulties of predicting them could lead to information overload and also fear. This fear, in turn, could lead to inactivity or otherwise ineffective choice. Support approaches therefore tend to focus on specific consequences.

One general difficulty with anticipating the consequences of configuration choices is that a given configuration option can have a large number of consequences for different situations (e.g., adopting a high velocity for a mouse can work out differently from one type of task to the next). One interesting approach to this difficulty, which can be realized with some types of configuration problem, is to give the chooser an overview of the consequences of a given configuration for a wide variety of situations. PVIZ (Mazzia et al., 2012) helps a user of FACEBOOK to answer questions about the visibility of elements of his profile, such as “With my current privacy settings, which people will be able to see my cell phone number?” PVIZ generalizes the idea realized in FACEBOOK’s *Audience View*, which enables a user to see how her profile will look to a particular individual. To provide an overview of visibility for all possible viewers, PVIZ does some automatic clustering of people and labeling of clusters, as well as graphical visualization that enables the user to see which proportion of members of each cluster will be able

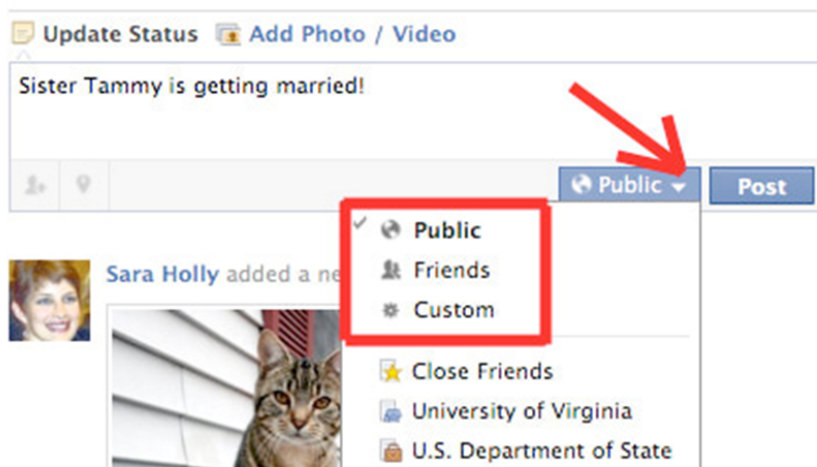


Figure 12.2: Facebook audience selector tool to alert and support *who* choices about possible recipients when posting information on social media.

to see the profile element in question. It therefore instantiates all of the first three ARCADE strategies. The results of user studies can be interpreted broadly as showing the promise of this approach for helping people anticipate the consequences of policy decisions.

12.3 Attribute-Based Choices About Privacy

Privacy-related choices are occasionally represented to users in a way that encourages attribute-based choice.

A frequent type of choice faced by members of social networking sites is the social-privacy choice of whom (among the other members and the public at large) to share some particular content with. The persons in question may be identified in different ways, such as by their names and/or by their attributes. The basic problem being tackled is the difficulty that users have in perceiving and managing the intended and actual “audiences” of posts given their typically long and undifferentiated sets of contacts.

The answer to this problem is sought in various forms of *access control* that have been derived from models established in the security field, as described in the survey by Sayaf and Clarke [2013]. For example, in FACEBOOK today, users are by default alerted to the possibility of making a *who* decision when posting something. This is done via a choice box placed right next to the post-input field (see Figure 12.2). When opened, it provides information about possible recipients and also arranges these options hierarchically, so that the number of choices remains manageable.⁷

In view of the challenging nature of this choice task, which is a type of configuration problem, it is understandable that tools have been developed which aim to make it tractable by applying several of the ARCADE strategies in conjunction. We will use as an example the tool FREEBU (Gao et al., 2012⁸), which helps people create lists of their FACEBOOK contacts.

Figure 12.3 shows one of FREEBU’s several modes that largely supports attribute-based choice: Contacts are grouped on the screen in terms of attributes such as home country, university attended, and whether they know a particular language. The user’s choice task is not, as would be typical of attribute-based choice, to select one or more individuals who are in some respect desirable but rather to compose a new list of friends that is likely to be useful in future choice situations (e.g., when the user is deciding whom to share a particular media item with).

This *column mode* applies ARCADE strategies as follows:

- *Represent the Choice Situation:* The most salient contribution of the interface is the way in which it organizes the relevant items (here: persons) on the screen in a way that facilitates a particular way of generating options (here: possible lists of contacts). Given that members of a useful contact list are often similar with respect to particular attributes, organizing the contacts by at-

⁷The top-level option “custom”, which appears to be rarely used, allows for arbitrarily complex combinations of the basic elements.

⁸<http://people.cs.kuleuven.be/~bo.gao/freebu>

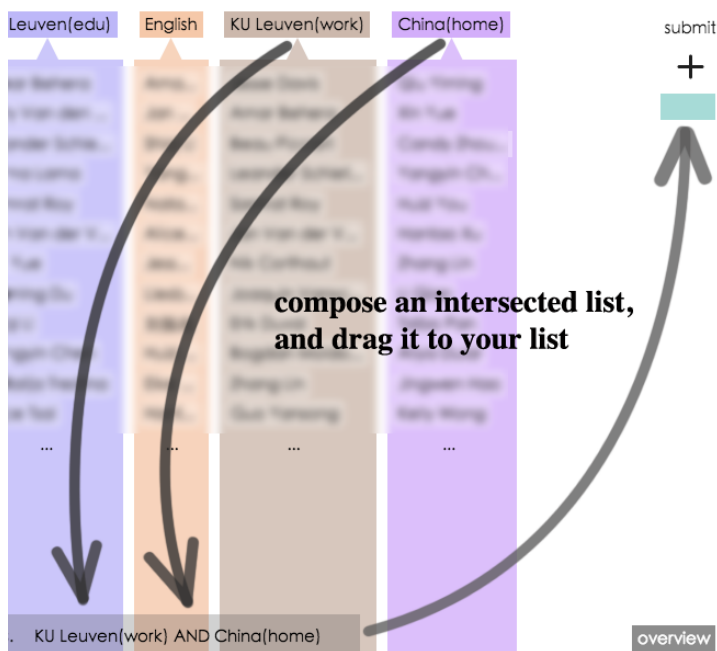


Figure 12.3: An image from the FREEBU interface for organizing FACEBOOK friends into lists. (This image, which obfuscates the friends’ names for privacy reasons, offers users advice about how to use the interface to create a list.)

tribute makes it relatively easy for the user to see promising ways of grouping contacts.

- *Access Information and Experience:* Though it is not the primary purpose of this visualization to convey new information to the user, it can happen that the user becomes aware of relevant facts about his friends that he had not previously noticed (e.g., the fact that one or more of them comes from a particular town).
- *Combine and Compute:* Creating a particular arrangement of items on the screen requires a certain amount of computation—which is quite straightforward in the column mode but more sophisticated in other modes of FREEBU, such as the one that clusters friends on the basis of their relationships to each other.

- *Advise About Processing*: The user interface includes some textual hints as to how to go about creating a friend list. In principle, these hints could be formulated without reference to specific user interface elements (e.g., “Look for sets of friends that resemble each other with respect to two or more attributes ...”). In the FREEBU interface, as in many other such interfaces, procedural advice of this sort is blended with instructions and hints about how to operate the user interface: “You can also first create an intersected column by dragging multiple columns into the ‘intersection’ box at the bottom of the screen; only the friends who share the common attributes of these columns will stay in the intersection box. Then drag the intersection box into the target list.”

12.4 Socially Based Choices About Privacy

Given that privacy-related actions in essence involve more than one person, it is natural that social expectations and examples play an especially important role (see, e.g., Greene et al. [2006]). The importance of social norms is underscored by the often striking differences that exist among cultures with regard to particular types of disclosure (see, e.g., Petronio, 2002, pp. 40–42). Palen and Dourish [2003] note that “privacy management . . . involves combinations of social and technical arrangements that reflect, reproduce and engender social expectations, guide the interpretability of action, and evolve as both technologies and social practices change” (p. 133).

These examples and expectations may concern either the *disclosure* of personal content (e.g., the revelation of intimate personal details in a blog post) or the *accessing* of such content that belongs to others (e.g., “shoulder surfing” to look at the email inbox of a colleague that is displayed on her screen).

In the social networking site scenario, privacy-related social expectations can be expressed explicitly, as through a community policy statement, or implicitly through the behavior of individual members (cf. 11.4). An example of the former case is provided by the “Com-

munity Guidelines” page of the popular diet and fitness community MYFITNESSPAL,⁹ which includes, for example, precise rules (i.e., articulated social expectations) concerning what may and may not be shown in photos that are posted by members who want to show off their progress in improving their physique.

As was discussed in connection with online communities (11.4), when the only available evidence is community members’ behavior, it can be hard for a member to know whether what they are dealing with is an expectation or merely examples that are not necessarily considered appropriate. The behavior of other persons that a chooser interprets can consist of (a) examples that are not causally related to his own behavior or (b) responses to his own choices (in which case the chooser can be seen as applying the trial-and-error-based pattern to acquire knowledge about social examples and expectations).

Both cases are illustrated in a lengthy message board thread on MYFITNESSPAL¹⁰ in which members have contributed more or less revealing and embarrassing sequences of “before and after” pictures of themselves to illustrate their progress in losing weight. These contributions often evoke evaluative comments (e.g., “I love these! It’s amazing how people change”). The pictures and the comments that they evoke can be seen to influence the subsequent contributions of other members (e.g., “Wow these are awesome! Great job everyone! I feel ready to do a new fat face comparison pic!”). Note that it’s not clear in these examples to what extent the evaluative comments refer to (a) the weight-losing progress shown in the photos or (b) the contributor’s willingness to display these photos. This fact illustrates again the point made in 12.1.3 about how privacy-related choices are often “bundled” with other aspects of a person’s behavior.

12.5 Policy-Based Choices About Privacy

Although the term *policy* is relatively seldom applied in the privacy area in the sense of the ASPECT model’s policy-based choice pattern,

⁹<http://www.myfitnesspal.com/welcome/guidelines>

¹⁰Specific reference omitted here for privacy-related reasons.

it is of course possible for a user to have a privacy-related policy in the sense of ASPECT: one or more rules or principles that specify (directly or indirectly) what privacy-related choice will be made in each of a set of situations, possibly as a function of particular parameters of the situation. For example, a user could have one of the following rules: (a) She will never give her birthday to any website. (b) She will never give her birthday when it is possible that it might be used for direct marketing. (c) She will give her birthday to a social networking site but will allow only her friends to access this information.

To be sure, arriving at a set of meaningful policies can be a challenge for the individual user in such a complex domain. This difficulty is presumably one reason for the abundance of privacy-related advice that can be found on the internet—such as the brief article titled *7 Things to Stop Doing Now on Facebook* (Consumer Reports Magazine, June 2010), which in effect suggests seven rules that a FACEBOOK user might adopt as part of his privacy-related policies. Personal policies can also be derived from social expectations of the sort discussed in 12.4.

One advantage of policy-based choice in this domain is that it makes it unnecessary for the user to think continually about all of the various privacy-related problems that might arise—which, as we have seen, can be hard (and possibly unpleasant) to anticipate and evaluate. Instead, she can focus on sticking to her policies, assuming that if she does so everything will probably be all right.

Another advantage of policies is that it is sometimes possible to delegate their execution to some sort of automated agent, which involves applying the strategy *Evaluate on Behalf of the Chooser* (cf. Section 9). This prospect is especially attractive in the privacy domain, since in principle privacy-related choices can arise at the rate of one every few seconds, as when a user is surfing the web or moving around the city with a location-aware app. It is therefore understandable that many interactive systems provide forms in which a user can in effect specify policies for making choices such as those about what personal data should be exposed in particular situations.

An example of such a form in a well-known privacy protection agent is shown in Figure 12.4. This web-based form (to be discussed in more

detail below) differs from most forms of this sort in that the user does not specify what the system is supposed to do autonomously but rather when the system is supposed bring the user into the loop by warning him of the danger of a privacy violation.

Before discussing research on this type of policy specification, we should digress a bit to analyze a frequently used concept that tends to obscure understanding in this context: that of *privacy preferences*.

12.5.1 Note on the Concept of “Privacy Preferences”

As can be seen in Figure 12.4, forms that allow a user to specify a policy of this sort are often labeled with the word “privacy preferences”. This term, which is widely used in the HCI privacy literature in other contexts as well, is more misleading than helpful, because the term *preferences* can be interpreted in at least three ways:

1. As the specific choices that a person makes in particular situations (e.g., “I can make my medical information available to this website now, but I prefer not to”).
2. As a chooser’s more *general predispositions* that influence her choices (as in “My general preference is not to make my medical information available to websites”).
3. As the choices that a person makes about *how to fill in a form* of this type (as in “I could have checked the box asking whether my medical information should be made available to websites, but I preferred not to”).

When the term “privacy preferences” is used in interactive systems or scientific discussions, it is rarely made clear which of these meanings is intended. A consequence is that the general impression is created that these three concepts are tightly correlated:

- A. that a user’s specific privacy-related choices are (largely) determined by general predispositions.
- B. that these predispositions are accurately captured by forms that ask about “privacy preferences”.

With the background of the ASPECT model, we can see that neither of these implicit assumptions is even approximately accurate:

Regarding assumption A: The choices that people make are not always determined primarily by general predispositions of these sorts. For example, within the experience-based pattern the chooser may simply repeat a previous choice that the current situation reminds her of. Within the socially based pattern, she may imitate a momentarily salient social example. When applying the consequence-based pattern, she may anticipate specific consequences that can arise because of unique aspects of the current situation; and her evaluation of these consequences can depend on temporary factors such as the relative momentary salience of her various goals.¹¹

Regarding assumption B: Where more general predispositions are involved, they can take qualitatively different forms, some of which are not naturally captured by preference forms. Here are some examples in terms of three of the ASPECT choice patterns:

- Within the policy-based pattern: a personal policy that applies to some class of choices (e.g., “I will never associate my photo with a social network profile”).
- Within the experience-based pattern: a habit of declining any opportunity to upload a photo to a profile; or a strong negative affective association with this action, which may have been acquired through bad previous experience.
- Within the consequence-based pattern: A general belief that performing this action can lead to serious negative consequences.

¹¹In the relevant psychology literature, the term *construction of preference* is sometimes used to refer to phenomena like these. See Lichtenstein and Slovic [2006] for an influential collection of articles that are relevant to the question discussed in this section. But even in this literature, terms like *preference* and *construction* are rarely explicitly and clearly defined. See also the book by Hausman [2012] for a thorough discussion of the concept of *preferences* from the perspectives of economics, psychology, philosophy, and everyday language. The very existence of a book like this should discourage anyone who employs this term from assuming that readers will have a clear and accurate idea of what he is talking about.

By suggesting a tight correlation among points 1–3 above, the use of the term *privacy preferences* serves as a sort of smoke screen in discussion of privacy-related choices, preventing us from seeing clearly how these choices come about and what we can do to support them. Researchers and practitioners in this area would do well to adopt the policy of never using this term—even though applying this policy may at first require a great deal of self-control and investment of the effort required to figure out what is really being referred to. It might help to remember the following slogan:

People don't "have preferences"; they *make choices*.

12.5.2 Studies of Policy Specification

Returning to the more specific topic of "privacy preferences" forms, we can see that, when a user fills in such a form, she cannot in general be assumed to be reporting straightforwardly on a policy that already exists in her mind. Rather, she is being asked (or required) to choose, from among a set of possible policies that the system can enforce on her behalf, the policy that seems best suited for automatically making choices that she would find appropriate. This task is similar to the subtasks of formulating and evaluating a policy that were discussed in 9.4 and 9.5, the main difference being that in this case a set of possible policies is being provided from which the user has to choose.¹²

Against this background, the question of how natural and useful a given "preferences form" is found by a given user is always an empirical question. A number of studies have looked at questions of this type.

One of the best-known and most realistically tested systems for automatically applying users' privacy policies is the PRIVACY BIRD of Cranor and colleagues (see, e.g., Cranor et al., 2006), which was already mentioned in 12.2.1. This system took advantage of the fact that many websites made their own "privacy policies" concerning their use of personal data available in a formal XML-based language defined by the PLATFORM FOR PRIVACY PREFERENCES (P3P).¹³ The PRIVACY

¹²For more general discussions of the question of the subtleties involved in eliciting people's "preferences" and values, see Fischhoff [2006] and Fischhoff [1991].

¹³<http://www.w3.org/P3P>

Privacy Preference Settings

These settings control when a warning icon will be displayed at the top of your browser window. You can click on the warning icon for more information.

Select Privacy Level: Low Medium High Custom Imported

HEALTH OR MEDICAL INFORMATION

Warn me at web sites that use my health or medical information :

For analysis, marketing, or to make decisions that may affect what content or ads I see, etc.

To share with other companies (other than those helping the web site provide services to me)

FINANCIAL OR PURCHASE INFORMATION

Warn me at web sites that use my financial information or information about my purchases :

For analysis, marketing, or to make decisions that may affect what content or ads I see, etc.

To share with other companies (other than those helping the web site provide services to me)

PERSONALLY IDENTIFIED INFORMATION (name, address, phone number, email address, etc.)

Warn me at web sites that may contact me to interest me in other services or products :

Via telephone

Via other means (email, postal mail, etc.)

And do not allow me to remove myself from marketing/mailling lists

Warn me at web sites that use information that personally identifies me :

To determine my habits, interests, or other characteristics

To share with other companies (other than those helping the website provide services to me)

Warn me at web sites that do not allow me to find out what data they have about me

NON-PERSONALLY IDENTIFIED INFORMATION (demographics, interests, web sites visited, etc.)

Warn me at web sites that use my non-personally identified information :

To determine my habits, interests, or other characteristics

To share with other companies (other than those helping the website provide services to me)

Help Import Settings Export Settings OK Cancel

Figure 12.4: The screen used by the PRIVACY BIRD to elicit a user's policy with regard to privacy-related choices. (From http://www.privacybird.org/tour/1_3_beta/tour.html.)

BIRD created a similarly formal representation of the privacy policy that a user had specified via the form shown in Figure 12.4. The PRIVACY BIRD could then automatically determine, whenever the user visited a website with a formalized privacy policy, whether the user

should be warned (by a red bird) about possibly undesirable disclosure of personal data.¹⁴

P3P is an interesting case for the current discussion because, despite its elegant technical design and high hopes and extensive support in its early stages, it was not widely adopted by websites and users and was ultimately suspended by the relevant W3C working group. A variety of explanations have been offered, some of which serve to remind us of the challenging context that confronts attempts to support privacy-related choices (e.g., insufficient support by browser implementers, the ease of circumventing the PRIVACY BIRD with invalid site policies, the lack of enforcement through legal or self-regulation, and user interface problems).¹⁵ The problem most directly relevant to choice architecture is the difficulty that users have in specifying their personal privacy policies. As Cranor et al. [2006, pp. 7–9] explain, the user may want information disclosure decisions to depend on considerations that cannot be expressed in the policy specification form that is provided (e.g., the branch of industry to which the website belongs or the specific companies with which information will be shared). They are also faced with the more general difficulties of policy formulation and evaluation that were discussed in 9.4 and 9.5, which are especially acute in a domain in which the consequences of actions can be so hard to anticipate.

Lederer et al. [2004] describe some of these difficulties in the privacy context as follows:

The act of configuring preferences is too easily desituated from the contexts in which those preferences apply. Users are challenged to predict their needs under hypothetical circumstances, and they can forget their preferences over time. If they predict wrongly, or remember incorrectly, their configured preferences will differ from their *in situ* needs, creating the conditions for an invasion of privacy. (p. 447)

¹⁴Other P3P-based user agents performed privacy-related actions automatically, such as deciding whether to block cookies or whether to allow access to a user's electronic wallet (cf. Cranor et al., 2006).

¹⁵For a recent brief survey and further references, see Morton et al. [2013].

In a similar vein, Berendt et al. [2005] suggest a number of reasons why people's privacy-related behavior in e-commerce sites can deviate from the policies that they express in questionnaires.

Although the results and considerations just summarized highlight a number of problems with the automation of privacy-related policies, there will presumably always be some efforts to achieve this type of automation at least in connection with choices that arise with high frequency and seem to lend themselves to policy-based choice. One question that arises in this context is whether a policy specification interface should allow the specification of complex rules, referring to a number of variables; or whether only simple rules are needed (cf. Iachello and Hong, 2007, pp. 47–49). As one might expect, there appears to be no single answer that applies to everyone. Using an experience sampling methodology in a field study, Anthony et al. [2007] asked participants on a number of occasions whom they would be willing to share their current location with. The responses of most users could be captured with the simple rules “Never share your location with anyone” or “Always share your location with the people on your white list”. For a handful of participants, however, the choices depended on characteristics of the situation in which they found themselves, which implies that they would require a relatively expressive language to formulate a satisfactory policy for revealing their location.

Benisch et al. [2010] conducted a somewhat similar study in which they noted that many users would require considerably more complex policy specification interfaces than the most commonly available type (i.e., a whitelist of individuals with whom the user is always willing to share) if they wanted to be sure that the system would make roughly the same choices that they themselves would make on a case-by-case basis. The authors also took into account the fact that specifying a complex policy can require considerable effort and that at least some users might therefore choose to adopt a relatively simple policy even if it didn't match their case-by-case choices.

In view of the effort and difficulty associated with manually specifying policies of this sort, it is natural to consider computational support for this process. Fang and LeFevre [2010] presented and evaluated a

PRIVACY WIZARD that uses sophisticated machine learning techniques to construct a policy on the basis of a limited amount of input from a user as to whether he would reveal a particular piece of information to particular individuals or groups of people. They show how the recommended policy can be visualized to the user so that he can better understand and evaluate it. This approach involves applying the strategies *Combine and Compute* and *Evaluate on Behalf of the Chooser* to the subtask of formulating a policy. It will be interesting to see, in future research, to what extent and in what contexts machine learning can help users to arrive at privacy-related policies whose automatic application will yield satisfactory results.

12.6 Trial-and-Error-Based Choices About Privacy

Especially because of the inherent difficulty that users have in anticipating the likely privacy-related consequences of their actions, it is natural—though, as we will see, problematic—for privacy-related choices to be based on trial and error. As was discussed more generally in 10.4, the information and experience that the chooser acquires through trying options out can take various forms.

12.6.1 The Difficulty of Learning From Privacy Violations

Let's consider first the most obvious form: feedback about the significant consequences of the chooser's actions. Among the most significant consequences are violations of the chooser's privacy. Unfortunately, there are several reasons why learning from this type of consequence tends to be problematic.

First, personal data is an information good; once it is out, it cannot be taken back. There can in general be no “undo” button for any specific piece of information (cf. 10.1.2). Even though a “right to be forgotten” has now found its way into the draft of the new European Union privacy directive, this right is associated with many difficulties, including technical ones (see, e.g., Druschel et al. [2012]). The irreversibility of possible consequences is an important property with regard to exploration strategies, especially when the consequences in question are seri-

ous. This fact is taken into account in some user interfaces (e.g., when data are about to be permanently deleted) with warnings like “This action cannot be undone!”, which tell the user, among other things, that simply trying the action out is unlikely to be part of a good exploration strategy. These warnings can be viewed as an instantiation of the strategy *Access Information and Experience* insofar as they inform the user about a property of the consequences of an option; if they are presented emphatically as warnings, they can also serve as advice to think carefully before choosing the option (*Advise About Processing*). These tactics may deserve increased attention where privacy-related choices are concerned. The privacy nudges of Wang et al. [2013] mentioned in can serve, in somewhat different ways, to discourage casual trial and error.

Even when viewed purely as sources of information, privacy violations tend to be relatively hard to process as informative feedback. As was indicated in 12.2, the *prediction* of the results of privacy-related choices tends to be made more difficult by the fact that consequences are often interrelated, invisible, indirect, delayed, and nondeterministic, partly because of the involvement of multiple stakeholders. These same factors also constitute obstacles for the trial-and-error-based pattern by making it hard for a chooser to recognize after the fact what *has* happened as a result of a privacy-related choice. For each of the pitfalls in the interpretation of feedback that were discussed in 10.5, the interested reader should be able to think of examples in the privacy domain. For instance, many privacy violations are undetectable for most users, such as those in the surveillance scenario that involve the reading and analysis of data from servers or transatlantic cables.

12.6.2 Learning From Information About Potentially Problematic Situations

Because of all of these problems with learning from privacy violations, a promising approach is to provide to a user feedback about *potential* privacy threats that have arisen because of her actions but which have (mostly, at least) not resulted in actual problems (cf. the more general discussion in 10.4 of the approach of providing feedback about states

that are likely to be correlated with important outcomes). Specifically, it can be helpful for a user to know who has been able to acquire what information as a result of his actions—and what they may have been able to infer from this information. Even if all of the observation and inference that the user becomes aware of does not lead to any privacy violation, this information can give an indication of the danger of such violations if the user continues to make the same choices.

Providing this type of feedback takes different forms depending on whether the other “knowers” are ordinary individual users (e.g., social network peers) or institutions.

Feedback About Observation by Individuals

System and interface design can help a user to find out about the identities and activities of other persons who have become knowers of that user’s private information. An increasing number of “How To” web pages, apps, and plugins enable users to “observe the observers”.

An example is an app for seeing who has viewed your FACEBOOK profile.¹⁶

Several data-based privacy awareness tools (see the general comments in 12.6.1) summarize, visualize, and score a user’s past publishing behavior. An example is the app PRIVACYCHECK (Figure 12.5), which shows what types of personal information from the user’s FACEBOOK pages are made available by the FACEBOOK API to websites that the user visits.¹⁷

An example of a system that was designed to help users learn by trial and error when to disclose their locations while on the move is the LOCYOUTION system of Tsai et al. [2009]. The user can specify, for each weekday, the times of day during which a particular class of contacts (friends, acquaintances, or strangers) are allowed to see her location. She can also check, via a special feedback page, which persons have tried to see her location at what times and whether that person was allowed to do so by the rules that were defined at the time. In this way, the user can essentially debug the rules over time—though the

¹⁶<https://www.facebook.com/WhoHasSeenYourProfileNewApplication>

¹⁷<http://www.rabidgremlin.com/fbprivacy/>

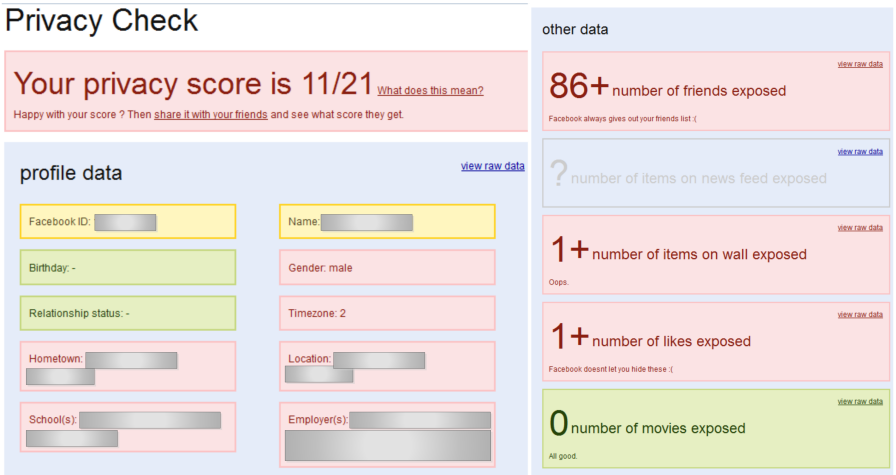


Figure 12.5: Screenshots from the FACEBOOK app PRIVACYCHECK. (The values of some fields are masked in this figure for privacy reasons.)

value of the resulting policy can be limited by the expressiveness of the mechanism for specifying rules (cf. the discussion in 12.5.2). Although the participants in this study by Tsai et al. [2009] expressed general satisfaction with the rule specification mechanism, they also suggested several additional types of rules that they would have liked to have.

A possible drawback of providing even excellent feedback about the responses of one group of persons (e.g., peers such as other users of a social-networking site), is that it may divert the user’s attention away from all those others who are “watching” (cf., the discussion, in earlier subsections, of social privacy vs. institutional privacy and surveillance). This phenomenon can be seen as a case of (maybe inadvertently) biased representation of the choice situation—specifically, the provision of a biased sample of information about who is observing the chooser’s actions.

Feedback About Tracking by Organizations

One way to reduce the bias just mentioned is to make available to ordinary users tools for accessing information about what organizations are

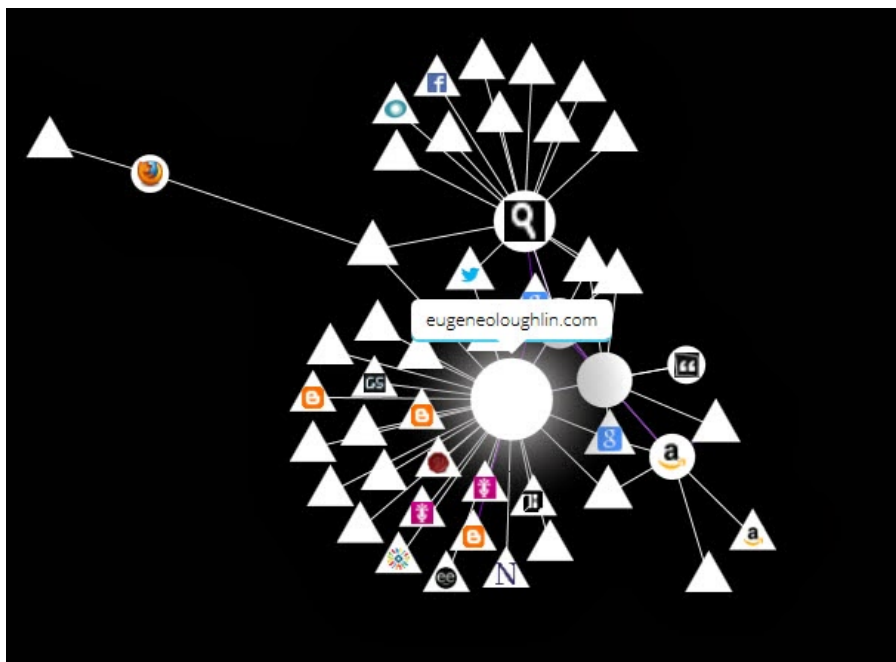


Figure 12.6: Screenshot from the LIGHTBEAM plugin for Mozilla FIREFOX. (From eugeneoloughlin.com.)

monitoring their behavior—and what knowledge they can derive from this monitoring. An example is the plugin LIGHTBEAM¹⁸ (Figure 12.6) for the FIREFOX browser, which enables a user to get an overview of the websites that track his website visits—including cases where one website has recorded his visits to other websites. Typically, users are surprised to see how many sites are tracking their behavior and how much information a single site can acquire.

12.6.3 Characterization in Terms of the ARCADE Model

The feedback provision techniques just discussed are representatives of a large class of choice support techniques that instantiate a combination of the first three ARCADE strategies:

¹⁸<https://addons.mozilla.org/en-US/firefox/addon/lightbeam>

- *Access Information and Experience*, in that information about consequences is being provided;
- *Represent the Choice Situation*, in that this information needs to be represented in a convenient way for the chooser to be able to make use of it; and
- *Combine and Compute*, in that in general some computational processing of the relevant information is required to generate a suitable representation.

12.7 Concluding Remarks on Privacy-Related Choices

In this section, we have given a number of examples of existing or possible tactics for supporting privacy-related choice. But a more important goal of this section has been to shed light on the special challenges that are raised by privacy-related choices, by discussing them in terms of the choice patterns of the ASPECT model. Though we hope that those who design and deploy relevant computing technology will benefit from being able to think about these challenges in terms of the ASPECT choice patterns and the ARCADE strategies, we believe that the challenges are too great to be dealt with completely even by the best-informed interaction and system design. Equally important is the goal of conveying to users of computing technology more realistic mental models of privacy-related choices (including their effects on other people) and a grasp of the options and tools that they have available for dealing with them (cf. 4.7.2).

13

Concluding Remarks

The authors hope that the readers of this publication have acquired a solid foundation for their work in their role as choice architects within the HCI field. With some careful thought, maybe combined with consultation of the cited literature, you should now be in a position to generate and try out ideas about how to help users of computing technology make better choices, whether you are a practitioner working on a specific application or a researcher developing new techniques or principles.

But we also hope and intend that this publication should serve as a basis for further research and publications, and we invite all readers to consider what they can contribute, whether they are PhD students looking for an exciting topic or senior researchers who are able to draw on decades of experience.

13.1 More Focused Analyses

The present publication can be compared with an atlas of an entire continent (for example, Europe) that comprises a large number of diverse countries: It is intended to be just about equally useful to people

in many different areas of HCI, with interest in diverse application domains and particular types of computing technology. A next step is to use this relatively coarse-grained map as a framework within which to work out more specific maps, like those of individual countries. Concretely, we encourage readers to consider creating more specific works titled “Choice Architecture for X ”, where X is a particular type of computing technology (e.g., pervasive computing or multimodal interaction) or a particular application domain (e.g., automotive user interfaces or personal health and well-being). In particular, instead of focusing, as in the present publication, on choices *about* the use of computing technology, these new works can consider everyday choices in any of a wide range of domains. Works like these can focus on particularly relevant tactics derived from the ARCADE support strategies. With this narrower focus, it should be possible not only to discuss existing support solutions and suggest new ones, as we have done in this work, but also to implement and test novel solutions so that they can be adapted on the basis of experience and reused by others, giving rise to guidelines and design patterns.

13.2 Extension to Decision Making by Groups

Many choices are made by a group of people rather than by an individual, as when a group of collaborating authors decides what text processing system to use to prepare their joint article. Everyday group decision making with the support of computing technology appears to be growing in importance with the increasing interconnectedness of users. In the formulation of the ASPECT and ARCADE models, group decision making has so far intentionally been left out of consideration. Extending the models to include groups as choosers is a natural—though challenging—next step. Group decision making in general involves some processes, such as interpersonal negotiation of compromises in cases of conflict of interest, which are not found in individual decision making (see, e.g., Kameda et al., 2003, Sorokin et al., 2004, and Saaty and Peniwati, 2008, for general treatments of group decision making; and Jameson and Smyth, 2007, and Masthoff, 2010, for

discussions of the special characteristics of recommender systems that make recommendations to groups).

13.3 Application to Other Types of Choice in HCI

Everyday choices made by users of computing technology are not the only choices that are of interest in the HCI field. Here are several other types of choice to which it ought to be possible to apply the ASPECT and ARCADE models fruitfully:

1. Choices made—and evaluations expressed—by participants in user studies who answer questions intended to measure their “preferences” (cf. 12.5.1)
2. Choices made by interaction designers
3. Choices made by organizations about whether and how to invest in usability and user experience (cf. Bias and Mayhew, 2005; Schaffer and Lahiri, 2014)
4. Choices made by HCI students and researchers about which research questions to investigate and how to investigate them

For each of these types of choice, we can ask: What (combinations of) choice patterns do people apply? What are the most challenging parts of these patterns for this type of choice? What possible support tactics can we derive from the ARCADE strategies?

13.4 Shouldering Responsibility for the Future of Human Choice

The ways in which people make choices and decisions will continue to evolve rapidly as computing technology provides ever more essential tools for our everyday thinking. Researchers and practitioners in HCI have an exciting opportunity to influence the further evolution of this central aspect of what it means to be human. In particular, they can help to determine whether this evolution will move in the direction of ever-increasing uncritical reliance on computational support, with

increasing vulnerability to bias and manipulation, or whether 21st-century computing technology will ultimately be seen to have increased humans' ability to make good choices for themselves.

No matter how energetically we work to achieve the latter goal, there will inevitably continue to be many interactive systems that are designed with a persuasive intent, the goal being to bias choices in a particular direction that is not necessarily in the interests of their users. It's not just that introducing bias into choice processes is often more lucrative than helping people choose for themselves; as we have seen repeatedly throughout this work, bias is also normally much easier to create and indeed impossible to avoid entirely. Accordingly, we believe it will be necessary to help people learn about the strategies and technologies that are being used to support and/or influence their choices; about the sorts of bias that these tools can introduce, whether intentionally or not; and about what enlightened users can do to retain ultimate control over their choice processes. In short, we should help people learn how to use modern tools for choosing effectively—without being used *by* them.

Acknowledgments

The authors thank Alan Dix and Paul van Schaik for their perceptive comments on an earlier version of this publication and Per Ola Kristensson for many fruitful discussions. Helpful feedback was given by the participants in a number of Anthony Jameson’s courses on this topic.

Anthony Jameson and Silvia Gabrielli thank the Autonomous Province of Trento for funding the research unit PREVOLUTION, in which the foundation for this work was laid, and EIT ICT Labs for funding for the activity “Understanding and Supporting Short-Term Decisions”. The work also benefited from Anthony Jameson’s involvement in the EIT ICT Labs activities APPS FOR YOUR CAR (2013) and 3CIXTY (2014) and in the FP7 EU project GLOCAL (Event-based Retrieval of Networked Media, grant agreement 248984).

Bettina Berendt thanks the Flemish Agency for Innovation through Science and Technology (IWT) and the Fonds Wetenschappelijk Onderzoek–Vlaanderen (FWO) for support through the projects SPION (grant number 100048) and Data Mining for Privacy in Social Networks (grant number G068611N), respectively.

Finally, the authors thank James Finlay and Alet Heezemans of now Publishers and editor-in-chief Ben Bederson for encouraging and enabling us to publish this foundational work even though it required much more time than anyone had expected. If James, Alet, and Ben had insisted on faster completion, this work would have been published

in a premature state and been much less useful as a foundation. If they had been less insistent, the work would never have been published at all.

References

- Alessandro Acquisti and Jens Grossklags. Privacy attitudes and privacy behavior—losses, gains, and hyperbolic discounting. In *Economics of Information Security*, pages 165–178. 2004.
- Irwin Altman. Privacy: A conceptual analysis. *Environment and Behavior*, 8 (1):7–29, 1976.
- Morgan Ames and Mor Naaman. Why we tag: Motivations for annotation in mobile and online media. In Bo Begole, Stephen Payne, Elizabeth Churchill, Rob S. Amant, David Gilmore, and Mary B. Rosson, editors, *Human Factors in Computing Systems: CHI 2007 Conference Proceedings*, pages 971–980. ACM, New York, 2007.
- Elisabeth André, Thomas Rist, and Jochen Müller. Employing AI methods to control the behavior of animated interface agents. *Applied Artificial Intelligence*, 13:415–448, 1999.
- D. Anthony, D. Kotz, and T. Henderson. Privacy in location-aware computing environments. *IEEE Pervasive Computing*, 6, 2007.
- Dan Ariely. *Predictably Irrational*. HarperCollins, New York, 2008.
- Richard P. Bagozzi. The legacy of the technology acceptance model and a proposal for a paradigm shift. *Journal of the Association for Information Systems*, 8(4):243–255, 2007.
- Peter J. Bayley, Jennifer C. Frascino, and Larry R. Squire. Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature*, 436(7050):550–553, 2005.

- Gerard Beenen, Kimberly Ling, Xiaoqing Wang, Klarissa Chang, Dan Frankowski, Paul Resnick, and Robert E. Kraut. Using social psychology to motivate contributions to online communities. In Jim Herbsleb and Gary Olson, editors, *Proceedings of the 2004 Conference on Computer-Supported Cooperative Work*, pages 212–221. ACM, New York, 2004.
- Victoria Bellotti, Maribeth Back, W. K. Edwards, Rebecca E. Grinter, Austin Henderson, and Cristina Lopez. Making sense of sensing systems: Five questions for designers and researchers. In Loren Terveen, Dennis Wixon, Elizabeth Comstock, and Angela Sasse, editors, *Human Factors in Computing Systems: CHI 2002 Conference Proceedings*, pages 415–422. ACM, New York, 2002.
- M. Benisch, P. Kelley, N. Sadeh, and L. Cranor. Capturing location-privacy preferences: Quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 2010.
- B. Berendt, D. Clarke, R. De Wolf, B. Gao, J. Pierson, S. Preibusch, and R. Sayaf. SPION Deliverable 5.1 Report on Research Activities (Identity Management), 2012. COSIC Internal Technical Report, K.U. Leuven, Belgium.
- Bettina Berendt. Ranking – use and usability. *Bulletin of the Belgian Mathematical Society — Simon Stevin*, 16:745–768, 2009.
- Bettina Berendt. More than modelling and hiding: Towards a comprehensive view of web mining and privacy. *Data Mining and Knowledge Discovery*, 24(3):697–737, 2012.
- Bettina Berendt, Oliver Günther, and Sarah Spiekermann. Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, 48(4):101–106, 2005.
- Tilman Betsch and Susanne Haberstroh, editors. *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005a.
- Tilman Betsch and Susanne Haberstroh. Current research on routine decision making: Advances and prospects. In Tilman Betsch and Susanne Haberstroh, editors, *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005b.
- James R. Bettman, Mary F. Luce, and John W. Payne. Constructive consumer choice processes. *Journal of Consumer Research*, 25:187–217, 1998.
- Sudeep Bhatia. Associations and the accumulation of preference. *Psychological Review*, 120(3):522–543, 2013.
- Suresh K. Bhavnani and Bonnie E. John. The strategic use of complex computer systems. *Human-Computer Interaction*, 15(2/3):107–137, 2000.

- Suresh K. Bhavnani, Frederick A. Peck, and Frederick Reif. Strategy-based instruction: Lessons learned in teaching the effective and efficient use of computer applications. *ACM Transactions on Computer-Human Interaction*, 15(1), 2008.
- Randolph G. Bias and Deborah J. Mayhew. *Cost-Justifying Usability: An Update for the Internet Age*. Morgan Kaufmann, San Francisco, 2005.
- Jonathan Bishop. Increasing participation in online communities: A framework for human-computer interaction. *Computers in Human Behavior*, 23: 1881–1893, 2007.
- Gerd Bohner, Hans-Peter Erb, and Frank Siebler. Information processing approaches to persuasion: Integrating assumptions from the dual- and single-processing perspectives. In William D. Crano and Radmila Prislin, editors, *Attitudes and Attitude Change*. Psychology Press, New York, 2008.
- John Brooke. SUS — a quick and dirty usability scale. In Patrick W. Jordan, Bruce Thomas, Bernard A. Weerdmeester, and Ian L. McClelland, editors, *Usability Evaluation in Industry*, pages 189–194. Taylor & Francis, London, 1996.
- Duncan P. Brumby and Andrew Howes. Strategies for guiding interactive search: An empirical investigation into the consequences of label relevance for assessment and selection. *Human-Computer Interaction*, 23:1–46, 2008.
- A. J. Brush, John Krumm, and James Scott. Exploring end user preferences for location obfuscation, location-based services, and the value of location. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, 2010.
- Susan L. Bryant, Andrea Forte, and Amy Bruckman. Becoming Wikipedian: Transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*, 2005.
- Sven Buschbeck, Anthony Jameson, Adrian Spirescu, Tanja Schneeberger, Raphaël Troncy, Houda Khrouf, Osmo Suominen, and Eero Hyvönen. Parallel faceted browsing. In *Extended Abstracts of CHI 2013, the Conference on Human Factors in Computing Systems (Interactivity Track)*, 2013.
- Jerome R. Busemeyer and James T. Townsend. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459, 1993.

- Stuart Card. Information visualization. In Julie A. Jacko, editor, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pages 515–548. CRC Press, Boca Raton, FL, 3rd edition, 2012.
- Stuart K. Card, Thomas P. Moran, and Allen Newell. *The Psychology of Human-Computer Interaction*. Erlbaum, Hillsdale, NJ, 1983.
- J. M. Carroll and M. B. Rosson. The paradox of the active user. In J. M. Carroll, editor, *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*, pages 80–111. MIT Press, Cambridge, MA, 1987.
- Suzanne C. Charman and Andrew Howes. The adaptive user: An investigation into the cognitive and task constraints on the generation of new methods. *Journal of Experimental Psychology: Applied*, 9(4):236–248, 2003.
- Luca Chittaro. Passengers’ safety in aircraft evacuations: Employing serious games to educate and persuade. In *Proceedings of Persuasive Technology 2012*, pages 215–226, 2012.
- Robert B. Cialdini. *Influence: The Psychology of Persuasion*. HarperCollins, New York, 2007.
- Robert B. Cialdini and Noah J. Goldstein. Social influence: Compliance and conformity. *Annual Review of Psychology*, 55:591–621, 2004.
- Robert T. Clemen and Terence Reilly. *Making Hard Decisions With Decision Tools*. Cengage Learning, Australia, 3rd edition, 2013.
- Jonathan D. Cohen, Samuel M. McClure, and Angela J. Yu. Should I stay or should I go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society*, 362: 933–942, 2007.
- Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. Using intelligent task routing and contribution review to help communities build artifacts of lasting value. In Rebecca Grinter, Thomas Rodden, Paul Aoki, Ed Cutrell, Robin Jeffries, and Gary Olson, editors, *Human Factors in Computing Systems: CHI 2006 Conference Proceedings*, pages 1037–1046. ACM, New York, 2006.
- Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. SuggestBot: Using intelligent task routing to find work in Wikipedia. In Tessa Lau and Angel R. Puerta, editors, *IUI 2007: International Conference on Intelligent User Interfaces*, pages 32–41. ACM, New York, 2007.
- Lorrie F. Cranor, Praveen Guduru, and Manjula Arjula. User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction*, 13 (2):135–178, 2006.

- Antonio R. Damasio. *Descartes' Error: Emotion, Reason and the Human Brain*. G. P. Putnam's Sons, New York, 1994.
- Jan De Houwer. Conditioning as a source of liking: There is nothing simple about it. In Michaela Wänke, editor, *Social Psychology of Consumer Behavior*. Psychology Press, New York, 2009.
- Clarisse S. De Souza and Jenny Preece. A framework for analyzing and understanding online communities. *Interacting with Computers*, 16(3):579–610, 2004.
- Ap Dijksterhuis, Maarten W. Bos, Loran F. Nordgren, and Rick B. van Baaren. On making the right choice: The deliberation-without-attention effect. *Science*, 311:1005–1007, 2006.
- Ap Dijksterhuis, Rick B. van Baaren, Karin C. Bongers, Maarten W. Bos, Matthijs L. van Leeuwen, and Andries van der Leij. The rational unconscious: Conscious versus unconscious thought in complex consumer choice. In Michaela Wänke, editor, *Social Psychology of Consumer Behavior*. Psychology Press, New York, 2009.
- Alan J. Dix. Information processing, context and privacy. In *Proceedings of the IFIP TC13 Third International Conference on Human-Computer Interaction*, pages 15–20, Amsterdam, 1990. North-Holland Publishing Co.
- Peter Druschel, Michael Backes, and Rodica Tirtea. The right to be forgotten—Between expectations and practice, 2012. ENISA Report.
- W. K. Edwards, M. W. Newman, and J. Z. Sedivy. The case for recombinant computing. Technical Report Technical Report CSL-01-1, Xerox PARC, 2001.
- Ward Edwards and B. Fasolo. Decision technology. *Annual Review of Psychology*, 52:581–606, 2001.
- Michael D. Ekstrand, John T. Riedl, and Joseph A. Konstan. Collaborative filtering recommender systems. *Foundations and Trends in Human-Computer Interaction*, 4(2):81–173, 2011.
- Seymour Epstein. Demystifying intuition: What it is, what it does, and how it does it. *Psychological Inquiry*, 21:295–312, 2010.
- Lujun Fang and Kristen LeFevre. Privacy wizards for social networking sites. In Michael Rappa, Paul Jones, Juliana Freire, and Soumen Chakrabarti, editors, *Proceedings of the 19th International Conference on World Wide Web*, pages 351–360. ACM, New York, 2010.
- Klaus Fiedler and Peter Juslin, editors. *Information Sampling and Adaptive Cognition*. Cambridge University Press, Cambridge, UK, 2006.

- Baruch Fischhoff. Debiasing. In *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, UK, 1982.
- Baruch Fischhoff. Value elicitation: Is there anything in there? *American Psychologist*, 46(8):835–847, 1991.
- Baruch Fischhoff. Constructing preferences from labile values. In Sarah Lichtenstein and Paul Slovic, editors, *The Construction of Preference*. Cambridge University Press, Cambridge, UK, 2006.
- Martin Fishbein and Icek Ajzen. *Predicting and Changing Behavior: The Reasoned Action Approach*. Taylor & Francis, New York, 2010.
- B. J. Fogg. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann, San Francisco, 2003.
- Craig R. Fox and Amos Tversky. A belief-based account of decision under uncertainty. *Management Science*, 44(7):879–895, 1998.
- D. Frankowski, D. Cosley, S. Sen, L. Terveen, and J. Riedl. You are what you say: Privacy risks of public mentions. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 572–579, 2006.
- Dan Frankowski, Shyong K. Lam, Shilad Sen, F. M. Harper, Scott Yilek, Micheal Cassano, and John Riedl. Recommenders everywhere: The wikiLens community-maintained recommender system. In *International Symposium on Wikis and Open Collaboration*, pages 47–60, 2007.
- Simon French, John Maule, and Nadia Papamichail. *Decision Behaviour, Analysis, and Support*. Cambridge University Press, Cambridge, UK, 2009.
- Wai-Tat Fu and J. R. Anderson. From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, 135(2):184–206, 2006.
- Wai-Tat Fu and Peter Pirolli. SNIF-ACT: A cognitive model of user navigation on the world wide web. *Human-Computer Interaction*, 22:355–412, 2007.
- B. Gao and B. Berendt. Feedback and awareness aspects. In A. Acquisti, E. Balsa, B. Berendt, D. Clarke, R. De Wolf, C. Diaz, B. Gao, S.F. Gürses, A. Kuczerawy, J. Pierson, F. Piessens, R. Sayaf, T. Schellens, F. Stutzman, B. Van Alsenoy, and E. Vanderhoven, editors, *SPION Deliverable 2.1 State of the Art*. 2011.

- Bo Gao, Bettina Berendt, Dave Clarke, Ralf D. Wolf, Thomas Peetz, Jo Pierson, and Rula Sayaf. Interactive grouping of friends in OSN: Towards online context management. In Jilles Vreeken, Charles Ling, Mohammed J. Zaki, Arno Siebes, Jeffrey X. Yu, Bart Goethals, Geoffrey I. Webb, and Xindong Wu, editors, *12th IEEE International Conference on Data Mining Workshops*, pages 555–562. IEEE Computer Society, 2012.
- Margaret M. Gardiner. Episodic and semantic memory. In Margaret M. Gardiner and Bruce Christie, editors, *Applying Cognitive Psychology to User-Interface Design*, pages 135–161. Wiley, Chichester, England, 1987.
- Margaret M. Gardiner and Bruce Christie, editors. *Applying Cognitive Psychology to User-Interface Design*. Wiley, Chichester, England, 1987.
- Gerd Gigerenzer. *Gut Feelings: The Intelligence of the Unconscious*. Penguin, London, 2007.
- Gerd Gigerenzer. *Risk Savvy: How to Make Good Decisions*. Penguin, New York, 2014.
- Gerd Gigerenzer and Peter M. Todd, editors. *Simple Heuristics That Make Us Smart*. Oxford, New York, 1999.
- Daniel T. Gilbert, Elizabeth C. Pinel, Timothy D. Wilson, Stephen J. Blumberg, and Thalia P. Wheatley. Durability bias in affective forecasting. In Thomas Gilovich, Dale Griffin, and Daniel Kahneman, editors, *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, Cambridge, 2002.
- Thomas Gilovich, Dale Griffin, and Daniel Kahneman, editors. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, Cambridge, 2002.
- Paul W. Glimcher and Ernst Fehr. *Neuroeconomics: Decision Making and the Brain*. Elsevier, Amsterdam, 2nd edition, 2014.
- Andreas Glöckner. Does intuition beat fast and frugal heuristics? A systematic empirical analysis. In Henning Plessner, Cornelia Betsch, and Tilmann Betsch, editors, *Intuition in Judgement and Decision Making*. Erlbaum, New York, 2008.
- Mark A. Gluck, Eduardo Mercado, and Catherine Myers. *Learning and Memory: From Brain to Behavior*. Worth, New York, 2nd edition, 2014.
- Wayne D. Gray and Deborah A. Boehm-Davis. Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied*, 6(4): 322–335, 2000.

- K. Greene, V. J. Derlega, and A. Mathews. Self-disclosure in personal relationships. In A. Vangelisti and D. Perlman, editors, *Cambridge Handbook of Personal Relationships*, pages 409–427. Cambridge University Press, Cambridge, UK, 2006.
- Stephen J. Guastello. *Human Factors Engineering and Ergonomics: A Systems Approach*. Taylor & Francis, Boca Raton, FL, 2nd edition, 2014.
- Seda Gürses and Claudia Diaz. Two tales of privacy in online social networks. *IEEE Security & Privacy*, 11(3):29–37, 2013.
- Reid Hastie. Problems for judgment and decision making. *Annual Review of Psychology*, 52:653–683, 2001.
- Reid Hastie and Robyn M. Dawes. *Rational Choice in an Uncertain World*. Sage, Thousand Oaks, CA, 2nd edition, 2010.
- Daniel M. Hausman. *Preference, Value, Choice, and Welfare*. Cambridge University Press, Cambridge, UK, 2012.
- Marti A. Hearst. *Search User Interfaces*. Cambridge University Press, Cambridge, UK, 2009.
- Ralph Hertwig, Greg Barron, Elke U. Weber, and Ido Erev. The role of information sampling in risky choice. In Klaus Fiedler and Peter Juslin, editors, *Information Sampling and Adaptive Cognition*. Cambridge University Press, Cambridge, UK, 2006.
- Judith H. Hibbard and Ellen Peters. Supporting informed consumer health care decisions: Data presentation approaches that facilitate the use of information in choice. *Annual Review of Public Health*, 24:413–433, 2003.
- Thomas T. Hills, Peter M. Todd, and Robert L. Goldstone. The central executive as a search process: Priming exploration and exploitation across domains. *Journal of Experimental Psychology: General*, 139(4):590–609, 2010.
- Robin Hogarth. *Educating Intuition*. University of Chicago Press, Chicago, 2001.
- Robin M. Hogarth. On the learning of intuition. In Henning Plessner, Cornelia Betsch, and Tilmann Betsch, editors, *Intuition in Judgement and Decision Making*. Erlbaum, New York, 2008.
- Christopher K. Hsee. Attribute evaluability: Its implications for joint-separate evaluation reversals and beyond. In Daniel Kahneman and Amos Tversky, editors, *Choices, Values, and Frames*. Cambridge University Press, Cambridge, UK, 2000.

- Darrell Huff. *How to Lie With Statistics*. Norton, New York, 1993. Originally published in 1954.
- G. Iachello, I. Smith, S. Consolovo, M. Chen, and G. Abowd. Developing privacy guidelines for social location disclosure applications and services. In *Proceedings of the 2005 Symposium on Usable Privacy and Security*, New York, 2005. ACM.
- Giovanni Iachello and Jason Hong. End-user privacy in human-computer interaction. *Foundations and Trends in Human-Computer Interaction*, 1 (1):1–137, 2007.
- Walter Isaacson. *Steve Jobs*. Simon & Schuster, New York, 2011.
- Alice M. Isen and Aparna A. Labroo. Some ways in which positive affect facilitates decision making and judgment. In Sandra L. Schneider and James Shanteau, editors, *Emerging Perspectives on Judgment and Decision Research*. Cambridge University Press, Cambridge, UK, 2003.
- Sheena Iyengar. *The Art of Choosing*. Hachette, New York, 2010.
- Sheena S. Iyengar and Mark R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79:995–1006, 2000.
- Anthony Jameson. How can persuasive technology help people choose for themselves? In Shlomo Berkovsky and Jill Freyne, editors, *Proceedings of Persuasive 2013*. Springer, Berlin, 2013. Abstract of a keynote address.
- Anthony Jameson and Krzysztof Z. Gajos. Systems that adapt to their users. In Julie A. Jacko, editor, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*. CRC Press, Boca Raton, FL, 3rd edition, 2012.
- Anthony Jameson and Kerstin Klöckner. User multitasking with mobile multimodal systems. In Wolfgang Minker, Dirk Bühler, and Laila Dybkjær, editors, *Spoken Multimodal Human-Computer Dialogue in Mobile Environments*, pages 349–377. Springer, Dordrecht, 2005.
- Anthony Jameson and Barry Smyth. Recommendation to groups. In Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl, editors, *The Adaptive Web: Methods and Strategies of Web Personalization*, pages 596–627. Springer, Berlin, 2007.
- Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. *Recommender Systems: An Introduction*. Cambridge, Cambridge, UK, 2011.

- Blair T. Johnson, Gregory R. Maio, and Aaron Smith-McLallen. Communication, attitude change: Causes, processes, and effects. In Dolores Albarracín, Blair T. Johnson, and Mark P. Zanna, editors, *The Handbook of Attitudes*. Erlbaum, Mahwah, NJ, 2005.
- Eric J. Johnson and Daniel G. Goldstein. Do defaults save lives? *Science*, 302 (5649):1338–1339, 2003.
- Eric J. Johnson, Suzanne B. Shu, Benedict G. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. Beyond nudges: Tools of a choice architecture. *Marketing Letters*, 23:487–504, 2012.
- Jeff Johnson. *Designing With the Mind in Mind: A Simple Guide to Understanding User Interface Design Rules*. Morgan Kaufmann, Burlington, MA, 2010.
- Joseph G. Johnson and Jerome R. Busemeyer. Rule-based decision field theory: A dynamic computational model of transitions among decision-making strategies. In Tilmann Betsch and Susanne Haberstroh, editors, *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005.
- Helmut Jungermann and Katrin Fischer. Using expertise and experience for giving and taking advice. In Tilmann Betsch and Susanne Haberstroh, editors, *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005.
- Helmut Jungermann, Hans-Rüdiger Pfister, and Katrin Fischer. *Die Psychologie der Entscheidung [The Psychology of Decision Making]*. Spektrum, Heidelberg, 2010.
- Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, 2011.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–295, 1979.
- Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341–350, 1984.
- Daniel Kahneman, Amos Tversky, and Paul Slovic. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, UK, 1982.
- Tatsuya Kameda, R. S. Tindale, and James H. Davis. Cognitions, preferences, and social sharedness: Past, present, and future directions in group decision making. In Sandra L. Schneider and James Shanteau, editors, *Emerging Perspectives on Judgment and Decision Research*. Cambridge University Press, Cambridge, UK, 2003.

- Ralph L. Keeney. *Value-Focused Thinking: A Path to Creative Decisionmaking*. Harvard, Cambridge, MA, 1992.
- David Kieras. Model-based evaluation. In Andrew Sears and Julie A. Jacko, editors, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pages 1191–1208. CRC Press, Boca Raton, FL, 2nd edition, 2008.
- Sara Kiesler, Robert E. Kraut, Paul Resnick, and Aniket Kittur. Regulating behavior in online communities. In Robert E. Kraut and Paul Resnick, editors, *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA, 2012.
- Gary Klein. *Sources of Power: How People Make Decisions*. MIT Press, Cambridge, MA, 1998.
- Alfred Kobsa and Maximilian Teltzrow. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. In *Privacy Enhancing Technologies*, pages 329–343, 2004.
- Derek J. Koehler and Nigel Harvey, editors. *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- Robert E. Kraut and Paul Resnick. Encouraging contribution to online communities. In Robert E. Kraut and Paul Resnick, editors, *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA, 2012a.
- Robert E. Kraut and Paul Resnick, editors. *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA, 2012b.
- Robert E. Kraut, Moira Burke, John Riedl, and Paul Resnick. The challenges of dealing with newcomers. In Robert E. Kraut and Paul Resnick, editors, *Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA, 2012.
- Richard P. Larrick. Debiasing. In Derek J. Koehler and Nigel Harvey, editors, *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Computing*, 8(6):440–454, 2004.

- Jennifer S. Lerner and Philip E. Tetlock. Bridging individual, interpersonal, and institutional approaches to judgment and decision making: The impact of accountability on cognitive bias. In Sandra L. Schneider and James Shanteau, editors, *Emerging Perspectives on Judgment and Decision Research*. Cambridge University Press, Cambridge, UK, 2003.
- Sarah Lichtenstein and Paul Slovic, editors. *The Construction of Preference*. Cambridge University Press, Cambridge, UK, 2006.
- David A. Lieberman. *Human Learning and Memory*. Cambridge University Press, Cambridge, UK, 2012.
- Charles E. Lindblom. The science of muddling through. *Public Administration Review*, 19(2):79–88, 1959.
- Charles E. Lindblom. Still muddling, not yet through. *Public Administration Review*, 39(6):517–526, 1979.
- George Loewenstein and Jon Elster, editors. *Choice Over Time*. Sage, New York, 1992.
- George Loewenstein, Daniel Read, and Roy Baumeister, editors. *Time and Decision*. Sage, New York, 2003.
- Tina Loraas and Michelle C. Diaz. Learning new uses of technology: Situational goal orientation matters. *International Journal of Human-Computer Studies*, 67:50–61, 2009.
- Wendy E. Mackay. Triggers and barriers to customizing software. In Scott P. Robertson, Gary M. Olson, and Judith S. Olson, editors, *Human Factors in Computing Systems: CHI 1991 Conference Proceedings*, pages 153–160. ACM, New York, 1991.
- Naomi Mandel and Eric J. Johnson. When web pages influence choice: Effects of visual primes on experts and novices. *Journal of Consumer Research*, 29:235–245, 2002.
- James G. March. *A Primer on Decision Making: How Decisions Happen*. The Free Press, New York, 1994.
- Judith Masthoff. Group recommender systems: Combining individual models. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 677–702. Springer, Berlin, 2010.
- A. J. Maule. Can computers help overcome limitations in human decision making? *International Journal of Human-Computer Interaction*, 26(2–3): 108–119, 2010.

- Alessandra Mazzaia, Kristen LeFevre, and Eytan Adar. The PViz comprehension tool for social network privacy settings. In *Proceedings of the 2012 Symposium on Usable Privacy and Security*, New York, 2012. ACM. Extended version at <http://alliemazzaia.com/pubs/pviz-tech-report.pdf>.
- Eric McCreath, Judy Kay, and Elisabeth Crawford. i-ems - an approach that combines hand-crafted rules with learnt instance-based rules. *Australian Journal of Intelligent Information Processing Systems*, 9(1):49–63, 2005.
- Luke McDowell, Oren Etzioni, Steven D. Gribble, Alon Halevy, Henry Levy, William Pentney, Deepak Verma, and Stani Vlasheva. Mangrove: Enticing ordinary people onto the semantic web via instant gratification. In Dieter Fensel, Katia Sycara, and John Mylopoulos, editors, *The Semantic Web—ISWC 2003*, pages 754–770. Springer, Berlin, 2003.
- Lorraine McGinty and James Reilly. On the evolution of critiquing recommenders. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 419–453. Springer, Berlin, 2011.
- Joanna McGrenere, Ronald M. Baecker, and Kellogg S. Booth. A field evaluation of an adaptable two-interface design for feature-rich software. *ACM Transactions on Computer-Human Interaction*, 14(1), 2007.
- Joseph T. McGuire and Joseph W. Kable. Rational temporal predictions can underlie apparent failures to delay gratification. *Psychological Review*, 120(2):395–410, 2013.
- Barbara A. Mellers. Choice and the relative pleasure of consequences. *Psychological Bulletin*, 126(6):910–924, 2000.
- Anthony Morton, Bettina Berendt, Seda Gürses, and Jo Pierson. “Tool clinics” – Embracing multiple perspectives in privacy research and privacy-sensitive design. In A. Acquisti, I. Krontiris, M. Langheinrich, and A. Sasse, editors, *My Life, Shared - Trust and Privacy in the Age of Ubiquitous Experience Sharing*. 2013. Dagstuhl Reports 3(7).
- Kanchan Mukherjee. A dual system model of preferences under risk. *Psychological Review*, 177(1):243–255, 2010.
- Benjamin R. Newell, David A. Lagnado, and David R. Shanks. *Straight Choices: The Psychology of Decision Making*. Psychology Press, Hove, UK, 2007.
- D. H. Nguyen and E. Mynatt. Privacy mirrors: Understanding and shaping socio-technical ubiquitous computing. Technical Report GIT-GVU-02-16, Georgia Institute of Technology, USA, 2002.

- D. H. Nguyen, A. Kobsa, and G. R. Hayes. An empirical investigation of concerns of everyday tracking and recording technologies. In *Proceedings of the 10th International Conference on Ubiquitous Computing*, New York, 2008. ACM.
- Donald A. Norman. Cognitive engineering. In Donald A. Norman and Stephen W. Draper, editors, *User Centered System Design: New Perspectives on Human-Computer Interaction*, pages 31–61. Erlbaum, Hillsdale, NJ, 1986.
- Antti Oulasvirta, Janne Hukkinen, and Barry Schwartz. When more is less: The paradox of choice in search engine use. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2009.
- Antti Oulasvirta, Tye Rattenbury, Lingyl Ma, and Eeva Ratta. Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing*, 16: 105–114, 2012.
- Leysia Palen and Paul Dourish. Unpacking “privacy” for a networked world. In Loren Terveen, Dennis Wixon, Elizabeth Comstock, and Angela Sasse, editors, *Human Factors in Computing Systems: CHI 2003 Conference Proceedings*, pages 129–136. ACM, New York, 2003.
- John W. Payne and James R. Bettman. Walking with the scarecrow: The information-processing approach to decision research. In Derek J. Koehler and Nigel Harvey, editors, *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- John W. Payne, James R. Bettman, and Eric J. Johnson. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK, 1993.
- Stephen J. Payne and Andrew Howes. *Adaptive Interaction: A Utility Maximization Approach to Understanding Human Interaction With Technology*. Morgan & Claypool, Palo Alto, CA, 2013.
- Sandra Petronio. *Boundaries of Privacy: Dialectics of Disclosure*. State University of New York Press, Albany, NY, USA, 2002.
- Jella Pfeiffer. *Interactive Decision Aids in E-Commerce*. Springer, Berlin, 2012.
- David J. Phillips. Privacy policy and PETs. *New Media & Society*, 6(6): 691–706, 2004.
- Jennifer K. Phillips, Gary Klein, and Winston R. Sieck. Expertise in judgment and decision making: A case for training intuitive decision skills. In Derek J. Koehler and Nigel Harvey, editors, *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.

- Peter Pirolli. Exploring and finding information. In John M. Carroll, editor, *HCI Models, Theories, and Frameworks*, pages 157–191. Morgan Kaufmann, San Francisco, 2003.
- Peter Pirolli. *Information Foraging Theory: Adaptive Interaction with Information*. Oxford University Press, New York, 2007.
- Carolin Plate, Nathalie Basselin, Alexander Kröner, Michael Schneider, Stephan Baldes, Vania Dimitrova, and Anthony Jameson. Recommendation: New functions for augmented memories. In Vincent Wade, Helen Ashman, and Barry Smyth, editors, *Adaptive Hypermedia and Adaptive Web-Based Systems: Proceedings of AH 2006*, pages 141–150. Springer, Berlin, 2006.
- Henning Plessner, Cornelia Betsch, and Tilmann Betsch, editors. *Intuition in Judgement and Decision Making*. Erlbaum, New York, 2008.
- Jennifer Preece, Blair Nonnecke, and Dorine Andrews. The top five reasons for lurking: Improving community experiences for everyone. *Computers in Human Behavior*, 20(2):201–223, 2004.
- Jenny Preece. *Online Communities: Designing Usability and Supporting Sociability*. Wiley, New York, 2000.
- Pearl Pu and Boi Faltings. Enriching buyers' experiences: The SmartClient approach. In Thea Turner, Gerd Szwillus, Mary Czerwinski, and Fabio Paternò, editors, *Human Factors in Computing Systems: CHI 2000 Conference Proceedings*, pages 289–296. ACM, New York, 2000.
- Howard Rachlin. *The Science of Self-Control*. Harvard, Cambridge, MA, 2000.
- Tim Rakow and Ben R. Newell. Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23:1–14, 2010.
- Kate Raynes-Goldie. *Privacy in the Age of Facebook: Discourse, Architecture, Consequences*. PhD thesis, Curtin University, 2012.
- Daniel Read. Intertemporal choice. In Derek J. Koehler and Nigel Harvey, editors, *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- Daniel Read, George Loewenstein, and Shobana Kalyanaraman. Mixing virtue and vice: Combining the immediacy effect and the diversification heuristic. *Journal of Behavioral Decision Making*, 12:257–273, 1999a.
- Daniel Read, George Loewenstein, and Matthew Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19:171–197, 1999b.

- James Reilly, Jiyong Zhang, Lorraine McGinty, Pearl Pu, and Barry Smyth. Evaluating compound critiquing recommenders: A real-user study. In *Proceedings of the 2007 Conference on Electronic Commerce*, 2007.
- F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors. *Recommender Systems Handbook*. Springer, Berlin, 2011.
- Francesco Ricci and Quang N. Nguyen. Acquiring and revising preferences in a critique-based mobile recommender system. *IEEE Intelligent Systems*, 22(3):22–29, 2007.
- Robert M. Roe, Jerome R. Busemeyer, and James T. Townsend. Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108(2):370–392, 2001.
- Karol G. Ross, James W. Lussier, and Gary Klein. From the recognition primed decision model to training. In Tilmann Betsch and Susanne Haberstroh, editors, *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005.
- Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, 3rd edition, 2010.
- Thomas L. Saaty and Kirti Peniwati. *Group Decision Making: Drawing Out and Reconciling Differences*. RWS Publications, Pittsburgh, 2008.
- Rula Sayaf and Dave Clarke. Access control models for online social networks. In Luca Cavaglione, Mauro Coccoli, and Alessio Merlo, editors, *Social Network Engineering for Secure Web Data and Services*, pages 32–65. IGI Global, 2013.
- Eric Schaffer and Apala Lahiri. *Institutionalization of UX: A Step-by-Step Guide to a User Experience Practice*. Addison-Wesley, Upper Saddle River, NJ, 2nd edition, 2014.
- Benjamin Scheibehenne, Rainer Greifeneder, and Peter M. Todd. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research*, 37(3):409–425, 2010.
- Sandra L. Schneider and James Shanteau, editors. *Emerging Perspectives on Judgment and Decision Research*. Cambridge University Press, Cambridge, UK, 2003.
- P. W. Schultz, Jennifer J. Tabanico, and Tania Rendón. Normative beliefs as agents of influence: Basic processes and real-world applications. In William D. Crano and Radmila Prislin, editors, *Attitudes and Attitude Change*. Psychology Press, New York, 2008.
- Barry Schwartz. *The Paradox of Choice: Why More Is Less*. HarperCollins, New York, 2004.

- Eldar Shafir, Itamar Simonson, and Amos Tversky. Reason-based choice. *Cognition*, 49:11–36, 1993.
- Herbert Simon. Rational choice and the structure of environments. *Psychological Review*, 63(2):129–138, 1956.
- Paul Slovic, Melissa Finucane, Ellen Peters, and Donald G. MacGregor. The affect heuristic. In Thomas Gilovich, Dale Griffin, and Daniel Kahneman, editors, *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, Cambridge, 2002.
- Robert D. Sorkin, Shenghua Luan, and Jesse Itzkowitz. Group decision and deliberation: A distributed detection process. In Derek J. Koehler and Nigel Harvey, editors, *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- Abhay Sukumaran, Stephanie Vezich, Melanie McHugh, and Clifford Nass. Normative influences on thoughtful online participation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3401–3410, New York, NY, USA, 2011. ACM.
- Richard H. Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12:183–206, 1999.
- Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven, 2008.
- Nava Tintarev and Judith Masthoff. Explanation of recommendations. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*. Springer, Berlin, 2010.
- J. Y. Tsai, P. Kelley, P. Drielsma, L. F. Cranor, J. Hong, and N. Sadeh. Who’s viewed you? the impact of feedback in a mobile location-sharing application. In *Proceedings of the 2009 Conference on Human Factors in Computing Systems, CHI 2009, Boston, USA, April 4-9, 2009*. ACM, 2009.
- Endel Tulving. *Elements of Episodic Memory*. Oxford University Press, New York, 1983.
- Daniel Tunkelang. *Faceted Search*. Morgan & Claypool, Palo Alto, CA, 2009.
- Amos Tversky. Elimination by aspects: A theory of choice. *Psychological Review*, 79:281–299, 1972.
- Amos Tversky and Daniel Kahneman. Rational choice and the framing of decisions. *Journal of Business*, 59(4):5251–5278, 1986.
- Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106: 4, 1039–61., 106(4):1039–1061, 1991.

- Amos Tversky and Daniel Kahneman. Cumulative prospect theory: An analysis of decision under uncertainty. *Journal of Risk and Uncertainty*, 5 , 2976 323., 5:297–323, 1992.
- Alcides Velasquez, Rick Wash, Cliff Lampe, and Tor Bjornrud. Latent users in an online user-generated content community. *Journal of Computer Supported Cooperative Work*, 23(1), 2014.
- Viswanath Venkatesh and Fred Davis. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46(2):186–204, 2000.
- Viswanath Venkatesh, Michael G. Morris, Gordon B. Davis, and Fred D. Davis. User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(346):425–478, 2003.
- Bas Verplanken, Vemund Myrbakk, and Erik Rudi. The measurement of habit. In Tilmann Betsch and Susanne Haberstroh, editors, *The Routines of Decision Making*. Erlbaum, Mahwah, NJ, 2005.
- Peter Wakker. *Prospect Theory for Risk and Ambiguity*. Cambridge University Press, Cambridge, UK, 2010.
- Yang Wang, Pedro G. Leon, Kevin Scott, Xiaoxuan Chen, Alessandro Acquisti, and Lorrie F. Cranor. Privacy nudges for social media: An exploratory Facebook study. In *Companion Proceedings of the 22nd International Conference on World Wide Web*, pages 763–770, 2013.
- Youcheng Wang and Daniel R. Fesenmaier. Assessing motivation of contribution in online communities: An empirical investigation of an online travel community. *Electronic Markets*, 13(1):33–45, 2003.
- Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, San Francisco, 2nd edition, 2004.
- Elke U. Weber and Eric J. Johnson. Mindful judgment and decision making. *Annual Review of Psychology*, 60:53–88, 2009.
- A. Westin. *Privacy and Freedom*. Atheneum, New York, NY, USA, 1967.
- Ryen W. White and Resa A. Roth. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool, San Francisco, CA, 2009.
- Max L. Wilson, Bill Kules, m. c. schraefel, and Ben Shneiderman. From keyword search to exploration: Designing future search interfaces for the web. *Foundations and Trends in Web Science*, 2(1), 2010.
- Timothy D. Wilson. *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard, Cambridge, MA, 2002.

- Wendy Wood and David T. Neal. A new look at habits and the habit-goal interface. *Psychological Review*, 114(4):843–863, 2007.
- J. F. Yates, Elizabeth S. Veinott, and Andrea L. Patalano. Hard decisions, bad decisions: On decision quality and decision aiding. In Sandra L. Schneider and James Shanteau, editors, *Emerging Perspectives on Judgment and Decision Research*. Cambridge University Press, Cambridge, UK, 2003.
- Richard M. Young and Allan MacLean. Choosing between methods: Analysing the user's decision space in terms of schemas and linear models. In J. J. O'Hare, editor, *Human Factors in Computing Systems: CHI 1988 Conference Proceedings*, pages 139–143. ACM, New York, 1988.
- Robert M. Young. Rational analysis of exploratory choice. In M. Oaksford and N. Chater, editors, *Rational Models of Cognition*. Oxford University Press, Oxford, 1998.
- Rami Zwick, Amnon Rapoport, Alison K. Lo, and A. V. Muthukrishnan. Consumer sequential search: Not enough or too much? *Marketing Science*, 22(4):503–519, 2003.