
How Can We Support Users' Preferential Choice?

Anthony Jameson

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

Silvia Gabrielli

CREATE-NET
Trento, Italy

Per Ola Kristensson

Computer Laboratory, University of Cambridge
Cambridge, UK

Katharina Reinecke

Harvard University
Cambridge, MA, USA

Federica Cena, Cristina Gena, and Fabiana Vernero

University of Turin
Turin, Italy

Copyright is held by the author/owner(s).
CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract

Users of computing technology are constantly making choices about how to use the technology which are *preferential* in the sense that there is no correct or incorrect option. We argue that HCI should devote more attention to helping users to make better preferential choices, tapping into the vast pool of relevant psychological research. After offering a quick high-level overview of this research, we introduce four general strategies for exploiting it in interaction design and illustrate these strategies with reference to examples. Looking at selected other paradigms that involve influencing preferential choice, we explain how our framework can lead to greater coverage and conceptual clarity.

Keywords

Preferential choice, decision making, psychology, interaction design, choice architecture, persuasive computing, recommender systems

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User interfaces – Evaluation / methodology

General Terms

Design, Human Factors, Theory

Introduction

We want to raise the question of how HCI can do more to support users' *preferential choices* about their use of computing technology. With this term, we refer to the broad class of cases in which a user can choose among two or more options, none of which is correct or incorrect but one of which can be preferred to the others. Here are some examples, related to a situation in which a driver of a car is considering whether and how to compose text messages while driving:¹

- Which of these recently received text messages should I read or listen to now?
- Shall I respond to this particular message now or put it off until later?
- If I respond, shall I use speech recognition or tapping on a virtual keyboard?
- ... and shall I send a quick generic answer or create an interesting text?
- How should I configure my in-car text message processing system?
- Should I take the time to practice dictating and correcting text messages before I try to do so while driving?

Preferential choice can be contrasted with *nonpreferential* choice, which is a matter of choosing the single correct action for achieving a given goal (e.g., "Which of these unfamiliar buttons do I have to push if I want to start dictating a message?").

Many HCI design principles, guidelines, and practices can be seen as supporting both preferential and nonpreferential choice; but preferential choice gives rise to characteristic cognitive processes that deserve special attention. To date, particular types of preferential choice have often been

treated separately, using in each case ideas that seem especially relevant to that category of choice (e.g., research about the forces that influence the behavior of mobile email users, discussed in [15]).

By contrast, we propose a generally applicable, domain-independent approach to supporting preferential choice that is based on a comprehensive overview of the psychological processes that are involved in such choices. In essence, we identify general phenomena and design tactics of which domain-specific phenomena and design ideas can be seen as specific instantiations. The intended benefits are that (a) we can leverage a vast amount of psychological research that is not specific to any domain; and (b) we can see commonalities among superficially different attempts to support users' preferential choice.

We begin by presenting a compact overview of cognitive processes that can occur when a person makes a preferential choice. We then consider corresponding design tactics aimed at supporting these processes.

A Compact Overview of Processes Involved in Preferential Choice

Preferential choice has been studied in a number of areas of psychological research, including: judgment and decision making (see, e.g., [12]), naturalistic decision making ([11]), the Reasoned Action approach ([5]), research on habitual behavior ([24]), behavioral economics ([21]), self-control ([18]), and compliance tactics ([3]).

Providing a coherent overview of the HCI-relevant processes using terminology taken directly from the original research literature would result in an opaque mosaic of concepts such as *recognition-primed decision making*, *choice bracketing*, *coherent arbitrariness*, *complex ambivalence*, and

¹This example scenario (which is not the subject of research of any of the authors of this paper) was suggested by recent research in the Automotive IUI group at DFKI (<http://automotive.dfi.de/index.php/en/home>). We thank Christian Müller for helpful discussions of these examples.

internalities, each of which would require its own explanation. As an alternative, we offer the overview shown in Table 1: It lists a large number of “questions” that a chooser might conceivably “ask” him- or herself while making a preferential choice. Each of these questions serves as an easily understandable *label* for a set of more specialized research concepts and associated empirical and theoretical results. Though in some cases such questions may be consciously asked and addressed by a computer user, the processing represented in the table by a question often occurs without any verbal formulation or conscious deliberation—whatever particular definition of the elusive concept of *consciousness* one may prefer to use (see, e.g., [23]). Moreover, in any single case where a user makes a choice, only some (often very small) subset of the considerations represented in the table will actually be involved.

Comments on the Overview

The group of questions starting with “How good would the consequences be if I chose this option?” in the middle column of the table belongs to one familiar view of decision making: The chooser anticipates and evaluates the (possibly uncertain and temporally remote) consequences of his or her possible actions and bases his or her choice in some way on these evaluations. Since exhaustive consideration of consequences is in general infeasible and unnecessary, many simplifying shortcuts and heuristics are employed (see, e.g., [17], [8]).

Even stronger simplifications of the choice process are represented by the questions in the lower left-hand corner of the table, which reflect a contrasting view of decision making as being often determined by considerations that do not directly concern the consequences of options (see, e.g., [13] for a deep discussion of the relationship between such

nonconsequentialist and *consequentialist* ways of thinking about choices).

Whereas the social factors (“What does the social context suggest?”) can be seen as reflecting the results of learning on the part of other persons in the near and distant past, learning on the individual level can also be important, as can be seen in the last group of questions on the right: In cases where a user makes a given choice repeatedly (as often happens with choices about the use of computing technology), the user can form habits, acquire policies, and refine his or her understanding of the system in question and of the typical consequences of using it in particular ways. Learning from experience can also give rise to changes in the user’s tastes or skills (see the bottom of the middle column), which will have implications for future choices.

Using the Overview to Generate Design Ideas

Let us assume now that we are working on interaction design intended to help users perform particular tasks (e.g., dealing with incoming text messages while driving) and have noticed a case where a user needs to make a nontrivial preferential choice (e.g., when and how to respond to a given incoming message). Suppose we are concerned that users may make choices that are not in their best interests—or not in the best interests of other legitimate stakeholders (e.g., other drivers on the road). We want to consider what we might do to help users make “better” choices (taking into account the fact that what counts as a “good” choice is partly up to the user him- or herself, as is indicated near the beginning of Table 1).

In a nutshell, our high-level strategy is to consider, for each of the questions in the table that is relevant to the choice of interest, what we as designers can do to influence the treatment of that question for the better. Given the large

Table 1. Overview of questions that a chooser may (consciously or not) consider when making a preferential choice.
(Explanation in text.)

How am I going to approach this choice?

What resources do I have available for making this choice?

How much time and attention do I have?

Shall I think about the choice in advance, instead of deciding on the spot?

Shall I take a broader view of this choice?

Shall I think in terms of a general policy instead of in terms of a single choice?

Shall I engage in trial and error before making the choice?

What does it mean for me to make a "good choice" in this case?

Do I want to focus on ...

... getting good results?

... saving time?

... avoiding difficult or unpleasant thinking?

... making a choice that I can justify to others?

What's going on here?

What, if anything, do I need to make a choice about?

What are the important things that are happening now?

What's going to happen next?

What does my own experience with this type of choice suggest?

What have I done in the past in this type of situation?

What (good or bad) habits of mine apply here?

What policy (if any) have I adopted for this type of choice?

What option(s) does my self-image suggest?

What does the social context suggest?

What have other people recommended?

What do people like me do in this situation?

What rules, norms, and commitments do I need to take into account?

What options should I consider?

What options are available to me?

(If there are a lot of options:) Which of these many options should I consider carefully?

How good would the consequences be if I chose this option?

(If the option involves performing an action:)

How is it going to feel to perform the action?

How good am I at executing this type of action?

How good or bad will the things that happen immediately be?

(If the action has consequences that will occur later:)

How good or bad will the things that happen at some later time be?

How much should I discount the value of the later consequences?

(If it is a matter of selecting one of a number of objects to use for some purpose:)

How desirable are the attributes of this object?

(If it is a matter of implementing a general policy:)

What will the *steady-state situation* be like once I have been implementing this policy for a while?

How might choosing this option change the conditions for future choices that I make?

(If it is a matter of configuring a system:)

How will the new configuration change my immediate experience?

How will the new configuration change the consequences of my future actions?

How might choosing this option change my tastes (i.e., the way I evaluate consequences of future choices)?

How might choosing this option change my knowledge or skills (i.e., my ability to execute particular actions in the future)?

How shall I deal with the complexity of the possible consequences?

What are the relative importances to me of the various possible consequences?

How shall I deal with my uncertainty about the consequences?

What simplifying shortcuts should I use to arrive at a choice in the face of all these possible consequences?

How could I justify the choice of this option ...

... to other people?

... to myself?

(After the choice has been made:) What can I learn from my experience with this choice?

What were the actual consequences of the option that I chose?

How good were these actual consequences (relative to my expectation)?

What have I learned about ...

... the consequences that I can expect from this option?

... how satisfied I am likely to be with this option in the future?

... what option I should choose in similar situations in the future?

number of questions in the table, this strategy gives us quite a number of potential leverage points; and these are further multiplied by the fact that, for any given question, there are four general ways in which we might in principle try to influence its treatment by the user. These are visualized in an abstract graphical way in Figure 1.

1. The simplest strategy is just to call the user's attention to the question—that is, to encourage a sort of processing that might not otherwise occur.

2. We can also try to *support* whatever effort the user may make to answer the question accurately, in two basic ways:

- Provide information that is relevant to the answering of the question.

Although information can sometimes be provided in documentation and training material, it can often be incorporated more subtly and effectively into the interaction design, as we will see in the next section.

- Perform operations that support the answering of the question.

Examples include measures like (a) arranging objects on a display in a way that makes it easier to answer a given question and (b) performing relevant internal computations.

3. We can try to *push* the user in the direction of a particular answer to the question, in any of the following ways:

- Change the salience or accessibility of information in such a way as to bias the answer in one direction or another.
- Explicitly suggest a particular answer to a specific question—or to the entire choice question itself.

The most obvious cases in which “pushing” seems applicable are those in which the designers for some reason want to increase the likelihood of a choice that may not be in the user's own best interest—for example, trying to induce a user to subscribe to an advertising mailing list. But pushing

tactics are often appropriate even when we have the user's own interests at heart, in several types of situation:

- The user may have previously decided on a behavior policy and want to be “pushed” to stick to that policy.
- Pushing the user is often a way of saving the user the time and effort required to deal with a particular question: The user may be quite willing to delegate work to the system at the price of giving the system some influence on the choice that will be made.

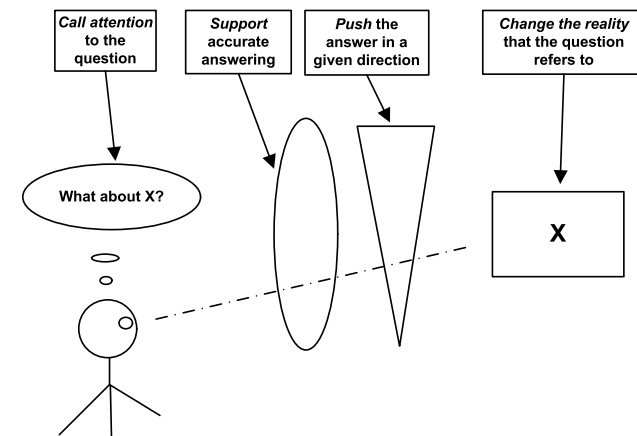


Figure 1. Visualization of four high-level strategies for influencing the way in which a user deals with a question that can arise in the course of making a preferential choice.

(The two shapes in the middle represent, from left to right, a magnifying lens and a distorting lens.)

4. Finally, we may alter the reality that the question refers to. For example, with regard to the question “What options are available to me?”, instead of simply improving the user's awareness of the available options, we can create (or eliminate) one or more options.

This fourth strategy is harder to grasp than the other three, because it is not uniquely associated with influencing choice processes as such. For example, we might choose to introduce a new option because it has some attractive benefits even when viewed in isolation. On the other hand, changing the underlying reality is also a way of influencing choice processes as such. For example, it has often been found that supplementing two existing options *A* and *B* with a third option *C* changes the user's relative preference between *A* and *B* (see, e.g., [20]). As this phenomenon illustrates, altering the underlying reality is often a way of "pushing" the chooser in a given direction, though in a way that is usually more indirect and subtle than the ways discussed above.

Examples of Design Tactics Derived From This Framework

So far in this paper, we have identified (a) several dozen "questions" that users can be thought of as possibly addressing when they make preferential choices and (b) four high-level strategies for influencing how a user deals with a given question. Multiplying these numbers together yields a large number of theoretically possible design tactics (of which, to be sure, most are inapplicable in any given concrete situation). Moreover, the application of any single tactic can take different forms in different specific contexts. We offer in this section several examples of choice-influencing tactics, indicating in each case how it is suggested by our framework. These examples concern the texting-while-driving scenario introduced at the beginning of this paper. We expect that experts in this particular domain can improve on the ideas suggested here—and that they may reject some of them completely. But this discussion should illustrate how our conceptual framework helps to (a) generate nonobvious new design ideas and (b) understand

better the psychological justification of even familiar design tactics.

Example of Calling Attention to an Aspect of the Choice

Maybe the most obvious questions in Table 1 that we might want to call users' attention to concern the *consequences* of choosing to text while driving. On the other hand, it would be a questionable design goal to induce drivers to consider, every time they are tempted to create a text message, the possible benefits and costs of doing so in that particular case.

It may make more sense to call users' attention in some way to the question of *choice bracketing* ([19]): whether they want to consider each "to text or not to text" choice in isolation or rather come up with a general policy for dealing with such choices. Research has revealed various benefits of *broad bracketing*, including these: 1. On a practical level, the chooser may be able to deliberate about the larger choice in a more favorable context (e.g., away from the automobile). 2. Less obviously, the chooser can think in terms of consequences that follow not from an individual action but rather, in the long run, from a sustained pattern of actions. For example, not responding to a single unimportant message may have the immediate consequence of offending the sender of that message; but if the driver consistently maintains a policy of not responding to such messages, in time his or her correspondents are likely to notice the pattern and no longer respond negatively when they do not receive an answer.

A broad view of this type of choice could be encouraged, for example, by verbal exhortations in documentation material, but it is usually more interesting to consider how it can be encouraged by interaction design. For example, merely offering the user a configuration option that determines what texting options will be available when the car is in

motion may cause the user to think in broader terms, since the configuration option raises a *general* question about texting behavior.

Examples of Informational Supporting and Pushing

As is well known, examples provided by other persons (see the question “What do people like me do in this situation?” in Table 1) can exert a powerful influence on preferential choice; but since the choices made by other persons are often not directly perceivable, the user’s answering of this question can often be “supported” or “pushed” considerably by the interaction design. The display of information about other people’s evaluations and choices is already a ubiquitous feature of Web 2.0 platforms. Somewhat similar information is found in automatically generated message signatures like “Composed on my iPhone with ShapeWriter”, and designers of mobile texting systems might consider how such information could best be aggregated and represented to influence users’ choices in an appropriate way.

Another set of questions where supporting and pushing could be applied concerns learning from experience (see the right-hand column of Table 1). It is not easy for drivers to learn from experience about the consequences or the desirability of particular forms of texting while driving. The consequences of the most important type—accidents—constitute low-probability events, and every text message that is sent without causing an accident can be seen as support for the hypothesis that texting while driving is safe. One approach here would be to provide feedback about more subtle direct consequences of the behavior, such as deviations from the normal driving path or the diminished proportion and range of eye gazes directed at the road.

Example of Changing the Underlying Reality

Perhaps the simplest examples of influencing choice by changing the underlying reality involve the addition or subtraction of options. Suppose, for example, that we add a facility (cf., e.g., [4]) which, when invoked with a simple command, automatically generates a brief status message about the user’s current situation that bears some resemblance to human-generated tweets and messages (e.g., “I’m caught in a traffic jam outside of Paris”). The obvious way in which this new option can influence choice is by offering drivers an alternative way of creating a text message. More subtly, the mere presence of the possibility of generating a typical message without any use of human intelligence may cause the user to see in a new light the option of manually writing such messages—and hence influence the user’s choices even when the new option is not chosen.

Comparison With Other Paradigms

As the reader will have noticed, many of the ideas in our framework can also be found in—or at least be derived from—other paradigms that are well known in the HCI field. This overlap is inevitable and desirable, since, after all, HCI has not ignored users’ preferential choice; it has simply not treated it as a single category of tightly interrelated processes in the way proposed here. We comment here on a selection of paradigms that overlap with ours in terms of their concepts and implications. One of our goals is to encourage cross-fertilization and conceptual clarification.

“Nudging” and Choice Architecture

In their book *Nudge: Improving Decisions About Health, Wealth, and Happiness* ([21]), Thaler and Sunstein apply to various areas of life the suggestive concept of a *choice architecture*. Unlike almost all other research on judgment and decision making, this work makes occasional use of

concepts and application examples from HCI, and accordingly uses of these concepts have been springing up in the HCI literature. On close inspection, though, the notions of *choice architecture* and *nudging* turn out to stand for an inspiring but loosely organized collection of ideas about how to influence people's (preferential or nonpreferential) choices, with a focus on means other than incentives and laws.² With our framework, by systematically distinguishing among (a) the psychological processes being targeted and (b) the high-level strategies for targeting them, we aim to help channel some of the enthusiasm about these ideas into the generation of effective design solutions.

Persuasive Computing

Persuasive computing (see, e.g., [6]) concerns the use of computing technology to influence (preferential) choices that people make (e.g., concerning health- or environment-related behavior). A major difference is that the choices themselves do not in general concern the use of computing technology. A consequence of this difference is that many of the design tactics that can be derived from our framework are not applicable. For example, since a person is not in general using a computational device to access yet another serving of chocolate cake at a party, improved interface design cannot help them at the moment of choice in many of the ways discussed above. Moreover, the emphasis in this area is on *persuasion*, which in the terminology used here involves *pushing* rather than *supporting*. We hope that our conceptual framework will help to show how this unnecessarily limited focus can be expanded.

²The high-water mark of conceptual coherence comes in chap. 5 when the authors show how their principles of choice architecture can be captured with the mnemonic acronym "NUDGES".

³For example: "What have I done in the past in this type of situation?" (for straightforward content-based recommendation); "What do people who like me do in this situation?" (for collaborative filtering); and various questions about the evaluation of attributes and consequences (for knowledge-based recommendation). See, e.g., [2] for a discussion of the relationships among the approaches to recommendation mentioned here.

Recommender Systems

Recommender systems can be seen as influencing users' preferential choices concerning products to buy, documents to read, and a variety of other types of item. As with persuasive computing, the choices normally do not concern the use of computing technology, though there are exceptions, such as systems that recommend commands to use (e.g., [14]) or interface configurations to adopt (e.g., [1]). Most recommenders can be viewed as taking over the task of *winning* a large set of options down to a small number that the user can contemplate with reasonable effort (cf. the question "Which of these many options should I consider carefully?" in Table 1). The processing that recommender systems do to perform this winnowing can in most cases be seen as the automated answering of other questions in Table 1.³ The *explanations* that are often provided by recommenders (see, e.g., [22]) can be seen as influencing choices in various ways that can be characterized in terms of the framework that we are proposing. For example, they sometimes offer information about the likely consequences of choosing an option, and they may provide a ready-made justification that the user can present to other people. As this last example suggests, the roles that explanations of recommendations can play in supporting preferential choice can be diverse and subtle; we suggest that our framework can lead to a better understanding of them.

Other Overlapping Paradigms

HCI literature and standard practices include many guidelines and principles that can be seen as helping users to make good choices, such as ensuring the visibility of

options and giving a preview of the consequences of actions. Most of these ideas are applicable to both preferential and nonpreferential choice, though often in different ways. They fit into our framework alongside many tactics that are characteristic of preferential choice.

Similarly, well-known HCI paradigms such as distributed cognition ([9]) and Norman's model of action ([16]) include some of the same processes and tactics that were discussed above, but they are mostly applied to phenomena other than preferential choice, such as performing a complex collaborative task or figuring out the correct way to operate a system so as to achieve a particular goal.

Research Context and Summary of Contributions

This paper is part of a larger research effort that comprises: 1) a comprehensive synthesis of research on preferential choice from psychology, HCI, and other fields; and 2) efforts to bridge the gap between this research synthesis and the needs of HCI researchers and practitioners who wish to deal with methodological and substantive issues related to users' preferential choice. The comprehensive synthesis requires more space than is available even in a normal journal article; the authors are currently working on extended publications, and the first author is presenting a half-day course at CHI 2011 titled *Choice and Decision Making for HCI*.

The second type of effort is more controversial: As has often been noted, making research results from psychology and other fields useful in HCI is usually not straightforward, and it is sometimes even questioned whether it makes sense to try (cf. [10]). We are therefore offering this alt.chi paper as a separate proposal.

We do not claim that the framework summarized in this paper is currently ready to be put into the hands of busy

HCI practitioners who need to make good design decisions on the fly. Because of the multifaceted nature of preferential choice, we assume that, for such a scenario, it will often be necessary to create more practically oriented instruments such as training programs and analysis methods that may be specific to a particular type of system or a particular class of preferential choice (see, e.g., [7], for an early effort along these lines). We do claim that HCI *researchers* should be able to use this framework to organize in an effective way their thinking about users' preferential choice processes and ways of supporting them.

Acknowledgments

The contributions of the first two authors were supported by the Autonomous Province of Trento in the context of the research unit Prevolution. We thank the commentators on the original alt.chi submission—Ryan Kelly, Alan Dix, John Riedl, Ben Bederson, Maria Karam, and two anonymous commentators—for thoughts that will influence our future work in this area.

References

- [1] A. Bunt, C. Conati, and J. McGrenere. Supporting interface customization using a mixed-initiative approach. In T. Lau and A. R. Puerta, editors, *IUI 2007: International Conference on Intelligent User Interfaces*. ACM, New York, 2007.
- [2] R. Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, 2002.
- [3] R. B. Cialdini. *Influence: The Psychology of Persuasion*. Harper Collins, New York, 2007.
- [4] C. Endres and D. Braun. Pleopatras: A semi-automatic status-posting prototype for future in-car use. In *Adjunct Proceedings of the Second International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 2010.

- [5] M. Fishbein and I. Ajzen. *Predicting and Changing Behavior: The Reasoned Action Approach*. Taylor & Francis, New York, 2010.
- [6] B. J. Fogg. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann, San Francisco, 2002.
- [7] S. Gabrielli and A. Jameson. Obstacles to option setting: Initial results with a heuristic walkthrough method. In T. Gross, J. Gulliksen, P. Kotzé, L. Oestreicher, P. Palanque, R. Prates, and M. Winckler, editors, *Human-computer interaction - INTERACT 2009, 12th IFIP TC 13 International Conference*, pages 400–403. Springer, Berlin, 2009.
- [8] G. Gigerenzer and P. M. Todd, editors. *Simple Heuristics That Make Us Smart*. Oxford, New York, 1999.
- [9] J. Hollan, E. Hutchins, and D. Kirsh. Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2):174–196, 2000.
- [10] A. Jameson. How is an experimental subject like a computer user? A critical notice of “Applying Cognitive Psychology to User-Interface Design”, edited by M. M. Gardiner and B. Christie. *Acta Psychologica*, 69:279–298, 1988.
- [11] G. Klein. *Sources of Power: How People Make Decisions*. MIT Press, Cambridge, MA, 1998.
- [12] D. J. Koehler and N. Harvey, editors. *Blackwell Handbook of Judgment and Decision Making*. Blackwell, Malden, MA, 2004.
- [13] J. G. March. *A Primer on Decision Making: How Decisions Happen*. The Free Press, New York, 1994.
- [14] J. Matejka, W. Li, T. Grossman, and G. Fitzmaurice. CommunityCommands: Command recommendations for software applications. In *UIST 2009 Conference Proceedings: ACM Symposium on User Interface Software and Technology*, pages 193–202, 2009.
- [15] C. A. Middleton and W. Cukier. Is mobile email functional or dysfunctional? Two perspectives on mobile email usage. *European Journal of Information Systems*, 15(3):252–260, 2006.
- [16] D. A. Norman. Cognitive engineering. In D. A. Norman and S. W. Draper, editors, *User Centered System Design: New Perspectives on Human-Computer Interaction*, pages 31–61. Erlbaum, Hillsdale, NJ, 1986.
- [17] J. W. Payne, J. R. Bettman, and E. J. Johnson. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK, 1993.
- [18] H. Rachlin. *The Science of Self-Control*. Harvard, Cambridge, MA, 2000.
- [19] D. Read, G. Loewenstein, and M. Rabin. Choice bracketing. In S. Lichtenstein and P. Slovic, editors, *The Construction of Preference*. Cambridge University Press, Cambridge, UK, 2006.
- [20] E. Shafir, I. Simonson, and A. Tversky. Reason-based choice. In S. Lichtenstein and P. Slovic, editors, *The Construction of Preference*. Cambridge University Press, Cambridge, UK, 2006.
- [21] R. H. Thaler and C. R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven, 2008.
- [22] N. Tintarev and J. Masthoff. Explanation of recommendations. In P. B. Kantor, F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*. Springer, Berlin, 2010.
- [23] T. D. Wilson. *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard, Cambridge, MA, 2002.
- [24] W. Wood and D. T. Neal. A new look at habits and the habit-goal interface. *Psychological Review*, 114(4):843–863, 2007.