

INFORMATION FORAGING THEORY

Adaptive
Interaction
with Information



PETER PIROLI

Information Foraging Theory
Adaptive Interaction with Information

Peter Pirolli



OXFORD
UNIVERSITY PRESS

2009

OXFORD

UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2007 by Peter Pirolli

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016

www.oup.com

First issued as an Oxford University Press paperback, 2009

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Pirolli, Peter.

Information foraging theory : adaptive interaction with information / Peter Pirolli.
p. cm. — (Oxford series in human-technology interaction)

ISBN 978-0-19-538779-7

1. Information behavior. 2. Human-computer interaction. I. Title. II. Series.
ZA3075.P57 2007

025.5'24—dc22 2006021795

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper

To my lifelong partner Jacqui, and my parents Remo and Betty. This could not have happened without your love.

Preface

To understand the evolution of things, one must understand something about their history as well as the environmental forces that had shaping influences upon them. Information Foraging Theory evolved through a series of fortuitous historical accidents, as well as a number of enduring shaping forces. A critical event was my move to the Palo Alto Research Center (PARC). Soon after I came to PARC at the beginning of 1992, I became involved in trying to develop studies and models around a set of projects that were collectively called intelligent information access. This included the novel information visualization systems investigated in the User Interface Research Area (see, e.g., Card et al., 1999) as well as the new techniques for browsing and searching being created in the Quantitative Content Area (e.g., Rao et al., 1995). As part of this effort, a group of us (including Stu Card, Dan Russell, Mark Stefik, and John van Gigh from California State University—Sacramento) were running some quick-and-dirty studies of people such as business intelligence analysts and MBA students. Our

studies of people doing information-intensive work started to give me some sense of the range of phenomena that we would need to address. Our study participants clearly were faced with massive volumes of information, often under deadline conditions, and making complex search decisions based on assessments that were enveloped in a great deal of uncertainty.

These information-intensive tasks seemed to be different than the human-computer interaction tasks that were being addressed by cognitive engineering models in the early 1990s, or the science, math, and programming tasks addressed by intelligent tutoring systems of that same period. Such cognitive models addressed tasks that tended to occur in task environments that (although large and complex) were well defined by a circumscribed domain of possible goals, elements of domain knowledge (e.g., about Lisp programming, algebra, word processing), and potential actions (e.g., in a formal language, or in a user interface). In contrast, the behavior of people seeking

information appeared to be largely shaped by the structure or architecture of the content—the *information environment*—and only minimally shaped by the user’s knowledge of user interface. In addition, the structure of the information environment was fundamentally probabilistic. Consequently, behavior was also dominated by choices made in the face of uncertainty and the continual evaluation of the expected costs and benefits of various actions in the information environment, in contrast to the near-certain costs and benefits of actions taken in traditional cognitive modeling domains of the time.

It was clear that it was going to be a challenge to develop theories for information-intensive tasks. Mulling about this issue, I was drawn to work in two areas in which I had done some reading. The first was the work in the late 1980s of John R. Anderson (e.g., Anderson, 1990), who was putting forth the argument that to understand mechanisms of the mind, one must first try to figure out the environmental problems that it solves. John developed the method of rational analysis and applied this approach to memory, categorization, and other areas of cognition with considerable success. I wondered if the approach could be applied to the analysis of the information environment and how it shapes information seeking behavior.¹ The second area of interest was behavioral ecology (e.g., Smith, 1987), which suggested that very diverse strategies adopted by people could be systematically predicted from optimization analysis that focused first on scrutiny of the environment. This particular interest of mine originated as an undergraduate at Trent University, where physiological psychology included coverage of ethology (the precursor to behavioral ecology) and anthropology included what is known as cultural materialism (the precursor to current evolutionary-ecological approaches to anthropology). Working through the literature in these areas, I was led to optimal foraging theory, and particularly to the book by Stephens and Krebs (1986) that is the source of the conventional models discussed in chapter 2. I quite literally had an “ah-ha” experience in the middle of a late-night conversation with Jacqui LeBlanc in which I laid out the basic analogies between information foraging and optimal foraging theory.

In July 1992, I wrote a working paper titled “Notes on Adaptive Sense Making in Information Ecologies,” which discussed the possible application of conventional foraging models and the core mathematics of Stephens and Krebs to idealized information foraging tasks. The working paper got two kinds

of reactions. The first was one of disbelief in the analogy, for a variety of relatively good reasons (e.g., humans are not rational, information is not food). The second was that the ideas were “audacious” (to quote Jock Mackinlay). Fortunately, Stu Card (my manager and colleague in the User Interface Research Area) pushed me to pursue this approach, and he has been my main sounding board for the development of the theory over the years. By the fall of 1993, I had enough material to present a seminar at the University of California—Berkeley called “Sense Making in Complex Information Ecologies.”

In the decade that followed, the fruitfulness of Information Foraging Theory was apparent from the way that it could be used to bring messy data into crystal clear focus. The first time this happened was in application to the Scatter/Gather study presented in chapter 6. Simple analyses of the logs of users interacting with the system seemed to indicate that users were behaving in a nonsystematic way in their allocation of time or in their choices of interface actions. The application of optimal foraging models resulted in another of those “ah-ha” experiences in which suddenly the data plots all fell neatly on lines predicted by theory. Like catching a perfect wave in surfing, the feeling one gets from that moment when one gains power over a small portion of the universe is hard to recount without the skill of poetry (which I do not have), and it is the reward that keeps you coming back.

Acknowledgments

In writing this book, I was fortunate to have input from a great panel of reviewers: Marc Mangel, Julie Heiser, John R. Anderson, and Jakob Nielsen. Jacqui LeBlanc read the earliest versions of the manuscript. Each provided a unique perspective from their respective fields. I am particularly grateful to Marc, who made many suggestions about the math and about connections to a richer history of work in “traditional” foraging theory. I have also been fortunate to work with Alex Kirlik, an editor who provided much-needed collegial advice throughout this project.

PARC has been an especially fertile and supportive environment and I must especially thank my managers Stu Card, Kris Halvorsen, and Mark Stefik for their continued interest in this work. I also have been the beneficiary of funding support from the

most enlightened government funding agencies. I must thank the Office of Naval Research for continued funding over many years and intellectual support from three great program managers, Helen Gigley, Astrid Schmidt-Nielsen, and Susan Chipman. The Advanced Development Research Activity has provided substantial funding support and sustained passion and interest from Lucy Nowell and Heather McMonagle. The Spencer Foundation provided discretionary funding as part of my National Academy of Education Fellowship that was used to support my dilettante ventures into behavioral ecology and evolutionary-ecological anthropology.

Many people have collaborated and contributed to this project over the years. Stu Card helped me shape many of the ideas by asking the right questions. Bernardo Huberman provided a wealth of innovative ideas from a physicist's perspective, most notably about ultradiffusion, random graph processes, and cooperative computational processes. Jim Pitkow was the force that got many of us interested in analyzing the emergent dynamical properties of large aggregates of content, users, and topology on the Web. Special thanks go to Dan Russell, who got all of us interested in sense making in the first place. Pamela E. Sandstrom independently discovered the idea of Information Foraging Theory in her work on scholarly communication (chapter 8), and she has graciously shared her insights and results over the years. Lada Adamic provided me with unpublished data used to calculate the correlation between inlinks and outlinks in chapter 3. Ed Chi took up the notion of information scent and cashed the idea into real usability analysis systems, and more recently has taken on the idea of social information foraging as the backbone for new information access techniques. Wai-Tat Fu helped move Information Foraging Theory into the realm of the Web with his leadership on the SNIF-ACT project. Ayman Farahat, Christiaan Royer, and Raluca Budiu helped develop a hardened system for generating spreading activation networks from online collections to replace the ad hoc code I initially started with. Hinrich Schuetze is credited with providing me with the first statistics from a large document corpus that were used to demonstrate that spreading activation nets could be used to predict information scent. Sara Kiesler provided useful recommendations on the literature on the relation of cooperative processes and innovation. Julie

Morrison, Rob Reeder, Pam Schraedley, Mija Van Der Wege, and Vikram Jiswal contributed enormously to the efforts to study Web users. Marti Hearst provided data that proved to be crucial to the analysis of Scatter/Gather and collaborated with me on the Scatter/Gather studies along with Patti Schank and Chris Diehl. I also thank Jan Pedersen for inviting me to work on Scatter/Gather in the first place. Jakob Nielsen has been my guide in understanding how information foraging theory relates to concrete Web usability issues. Over the years, Jared Spool has developed the notion of information scent as a conceptual tool for practitioners, and he has always been generous in sharing the specific guidelines that have been developed as a result (many of these are presented in chapter 9).

Finally, I thank Jacqui LeBlanc for being there when the lightning first struck and for her support during the first chapter drafts written during an idyllic stretch of time on the porch of the Dolphin Inn in Cayucos after morning surf sessions. A more loving and lovely muse I could not ask for.

Note

1. As a graduate student working with Anderson, I have notes and working papers from 1982 in which Anderson was already beginning to suggest that function with respect to the environment would be a crucial to developing and evaluating theories of cognition.

References

- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Card, S. K., Mackinlay, J. D., & Schneiderman, B. (1999). *Information visualization: Using vision to think*. San Francisco: Morgan-Kaufman.
- Rao, R., Pedersen, J. O., Hearst, M. A., Mackinlay, J. D., Card, S. K., Masinter, L., et al. (1995). Rich interaction in the digital library. *Communications of the ACM*, 38(4), 29–39.
- Smith, E. A. (1987). Optimization theory in anthropology: Applications and critiques. In J. Dupré (Ed.), *The latest on the best* (pp. 201–249). Cambridge, MA: MIT Press.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

Contents

1. Information Foraging Theory: Framework and Method, 3
 2. Elementary Foraging Models, 30
 3. The Ecology of Information Foraging on the World Wide Web, 48
 4. Rational Analyses of Information Scent and Web Foraging, 68
 5. A Cognitive Model of Information Foraging on the Web, 89
 6. A Rational Analysis and Computational Cognitive Model of the Scatter/Gather Document Cluster Browser, 110
 7. Stochastic Models of Information Foraging by Information Scent, 137
 8. Social Information Foraging, 149
 9. Design Heuristics, Engineering Models, and Applications, 166
 10. Future Directions: Upward, Downward, Inward, and Outward, 183
- Index, 193

INFORMATION FORAGING THEORY

Information Foraging Theory

Framework and Method

Knowledge is power.

—Sir Francis Bacon,
Meditationes Sacrae.
De Hæresibus (1597)

Modern mankind forages in a world awash in information, of our own creation, that can be transformed into knowledge that shapes and powers our engagement with nature. This information environment has coevolved with the epistemic drives and strategies that are the essence of our adaptive toolkit. The result of this coevolution is a staggering volume of content that can be transmitted at the speed of light. This wealth of information provides resources for adapting to the problems posed by our increasingly complex world. However, this information environment poses its own complex problems that require adaptive strategies for information foraging. This book is about *Information Foraging Theory*, which aims to explain and predict how people will best shape themselves for their information environments and how information environments can best be shaped for people.

Information Foraging Theory is driven by three maxims attributable in spirit, if not direct quotation, to Allen Newell's (1990) program of Unified Theories of Cognition:¹

1. *Good science responds to real phenomena or real problems.* Human psychology has evolved as an adaptation to the real world. Information foraging theory is concerned with understanding representative problems posed by the real-world information environment and adaptive cognitive solutions to those problems.
2. *Good science makes a difference.* Information Foraging Theory is intended to provide the basis for application to the design and evaluation of new technologies for human interaction with information, such as better ways to forage for information on the World Wide Web.
3. *Good science is in the details.* The aim is to produce working formal models for the analysis and prediction of observable behavior.

Like much of Newell's work, the superficial elegance and simplicity of these maxims unfurls into complex sets of entailments. In this book I argue that the best approach to studying real information

foraging problems is to adopt *methodological adaptationism*, which directs our scientific attention to the ultimate forces driving adaptation and to the proximate psychological mechanisms that are marshaled to produce adaptive solutions. Thus, the methodology of Information Foraging Theory is more akin to the methodology of biology than that of physics, in contrast with the historical bulk of experimental psychology. To some extent, this choice of methodology is a consequence of the success with which Information Foraging Theory has been able to draw upon metaphors, models, and techniques from optimal foraging theory in biology (Stephens & Krebs, 1986). The concern with application (Newell & Card, 1985) drives the theory to be relevant to technological design and evaluation, which requires that models be truly predictive a priori (even if approximately so) rather than a “good fit” explanation of the data a posteriori, as is the case with many current psychological models. Being concerned with the details drives the theory to marshal a variety of concepts, tools, and techniques that allow us to build quantitative, predictive models that span many levels of interrelated phenomena and interrelated levels of explanation. This includes the techniques of task analysis through state-space and problem-space representations, rational analysis and optimization analysis of adaptive solutions, and production system models of the cognitive systems that implement those adaptive solutions.

Audience

The intent of this book is to provide a comprehensive presentation of Information Foraging Theory, the details of empirical investigations of its predictions, and applications of the theory to the engineering and design of user interfaces. This book aims primarily at an interdisciplinary audience with backgrounds and interests in the basic and applied science aspects of cognitive science, computer science, and the information and library sciences. The theory and methodology have been developed by drawing upon work on the rational analysis of cognition, computational cognitive modeling, behavioral ecology, and microeconomics. The crucible of empirical research that has shaped Information Foraging Theory has been application problems in human-information interaction, which is emerging as a new branch in the

field traditionally known as human-computer interaction. Although the emphasis of this book is on theory and research, the insights and results are intended to be relevant to the practitioner interested in a deeper understanding of information-seeking behavior and guidance on new designs. Chapter 9 is devoted entirely to practical applications of the theory.

By its nature, Information Foraging Theory involves the use of technical material such as mathematical models and computational models that may not be familiar to a broad audience. Generally, the technical aspects of the theory and models are presented along with succinct discussion of the key concepts, insights, and principles that emerge from the technical parts, along with illustrative examples, metaphors, and graphical methods for understanding the key points. The aim of this presentation is to provide intuitive understanding along with technical precision and insight.

Frameworks, Theories, and Models

Like other programs of research in the behavioral and cognitive sciences, Information Foraging Theory can be discussed in terms of the underlying framework, the theory itself, and the models that specify predictions in specific situations. *Frameworks* are the general pools of concepts, assumptions, claims, heuristics, and so forth, that are drawn from to develop theories, as well the methods for using them to understand and predict the world. Often, frameworks will overlap. For instance, information processing psychology is a broad framework that assumes that theories about human behavior can be constructed out of information processing concepts, such as processes that transduce physical sensations into sensory information, elements storing various kinds of information, and computational processes operating over those elements. A related framework, connectionism, shares these assumptions but makes additional ones about the nature of information processing being neuronlike. Although bold claims may be made by frameworks, these are typically not testable in and of themselves. For instance, whether the mind is mostly a general purpose learning machine or mostly a collection of exquisitely evolved computational modules are not testable claims in and of themselves.

Theories can be constructed within frameworks by providing additional assumptions that allow one to

make predictions that can be falsified. Typically, this is achieved by specifying a *model* for a specific situation or class of situations that makes precise predictions that can be fit to observation and measurement. For instance, a model of information seeking on the Web (SNIF-ACT) is presented in chapter 5 that predicts the observed choice of Web links in given tasks. It includes theoretical specifications of the information processing model of the user, as well as assumptions about the conditions under which it applies (e.g., English-speaking adults seeking information about unfamiliar topics). The bulk of this book is about Information Foraging Theory and specific models. The aim of this introductory chapter is to provide an outline of the underlying framework and methodology in which Information Foraging Theory is embedded. However, before presenting such abstractions, a simple example is offered in order to illustrate the basic elements and approach of Information Foraging Theory.

Illustration

The basic approach of Information Foraging Theory can be illustrated with a simple example that I hope is familiar to many, involving the task of finding a good, reasonably priced hotel using the World Wide Web (Pemberton, 2003). A typical hotel Web site (see figure 1.1) will allow a user to search for available hotels in some specified location (e.g., “Paris”) and then allows the user to sort the results by the hotel star rating (an indicator of quality) or by price (but not both). The user must then click-select each result to read it, because often the price, location, and features summaries are inaccurate. Lamenting the often poor quality of such hotel Web sites, Pemberton (2003) suggested that improved “usability is about *optimizing the time* you take to achieve your purpose, how well you achieve it, and the satisfaction in doing it... How fast can you find the perfect hotel?” This notion of usability is at the core of Information Foraging Theory.

For illustration, consider the somewhat simplified and idealized task of finding a low-priced, two-star hotel in Paris.² This example shows (in much simplified form) the key steps to developing a model of information foraging: (a) a rational analysis of the task and information environment that draws on optimal foraging theory from biology and (b) a production system model of the cognitive structure of task.

Rational Analysis of the Task and Information Environment

Figure 1.2 presents an analysis of results of search for two-star Paris hotels that I conducted on a popular hotel Web site. The Paris hotel descriptions and prices were returned as a vertical list presented over several Web pages. I sorted the list by star rating and went to the page that began to list two-star hotels. In figure 1.2, the x-axis indicates the order of two-star hotel listings in the search result list when sorted by star rating, beginning at the first two-star hotel through the last two-star hotel, and the y-axis indicates price. Prices fluctuate as one proceeds down the list of Paris hotels. As noted above, this particular hotel Web site, like many others, does not allow the user to sort by both quality (star rating) and price—one must choose one or the other sorting. Assume a rational (and perhaps somewhat boring) hotel shopper who was concerned only with being frugal and sleeping in a two-star hotel. If that shopper methodically scanned the two-star hotel listings, keeping track of only the lowest priced hotel found so far, the lowest price encountered would decrease as plotted in figure 1.3. That is, the shopper would at first find a relatively rapid decrease in lowest price, followed by fewer improvements as the scan progressed. Figure 1.4 shows the savings attained (compared with the very first hotel price found on the list) by continuing to scan down the list. Figure 1.4 is a typical *diminishing returns* curve in which additional benefits (returns) diminish as one invests more resources (in this case, scan time).

A diminishing returns curve such as figure 1.4 implies that the expected value of continuing to scan diminishes with each additional listing scanned. If the list of search results were very long—as is often the case with the results produced by Web search engines—there is usually a point at which the information forager faces the decision of whether it is worth the effort of continuing to search for a better result than anything encountered so far. In the particular example plotted in figure 1.2, there were no additional savings for the last 18 items scanned. Figure 1.3 includes a plot of the expected minimum price encountered as a function of scanning a search result list, and figure 1.4 includes a plot of the expected savings as a function of scanning. These expectations were computed assuming that observed hotel prices in figure 1.2 come from a standard



FIGURE 1.1 A typical Web page from a hotel search site.

distribution of commodity prices (see the appendix for details). Assuming that our hypothetical rational hotel shopper valued time (time is money), the question would be whether the savings expected to be gained by additional scanning of hotel results was worth the time expected to be expended.

In contrast to this simple illustration, typical information problems solved on the Web are more

complicated (Morrison, Pirolli, & Card, 2001), and the assessments of the utility of encountered items in information foraging depend on more subtle cues than just prices. However, the basic problem of judging whether continued foraging will be useful or a waste of valuable time is surely familiar to Web users. It turns out that this problem is very similar to one class of problems dealt with in optimal foraging theory.

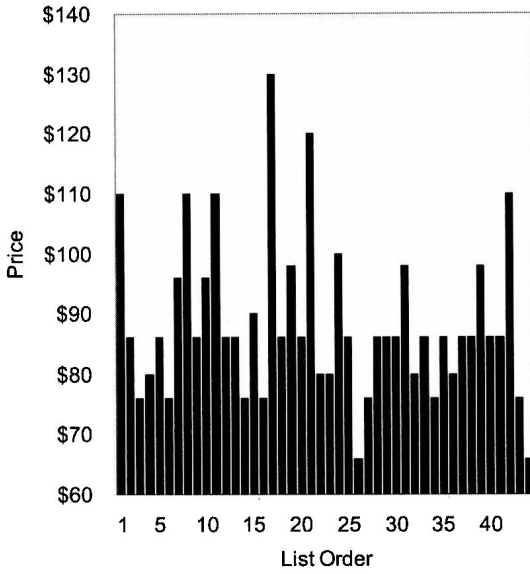


FIGURE 1.2 Prices of two-star Paris hotels in the order encountered in the results of a search of a hotel Web site.




FIGURE 1.4 Diminishing returns of savings as a function of list order. The observed savings is the difference between the observed minimum price found so far and the first price encountered (\$110), presented in figure 1.3. The expected savings is the difference between the expected minimum price and first price encountered.




FIGURE 1.3 The minimum two-star Paris hotel price as a function of order of encounter. The observed prices are the same as those in figure 1.2. The observed minimum is the least expensive hotel price found so far in a process that proceeds through the prices in the order listed. The expected minimum is a prediction based on the assumption that prices are being sequentially and randomly sampled from a fixed distribution of prices (see the appendix for details).

An Optimal Foraging Analogy

Many animals forage in patchy environments, with food arranged into clumps. For instance, a bird that feeds on berries in bushes will spend part of its time searching for the next bush and part of its time berry picking after having found a bush. Often, as an animal forages in a patch, it becomes harder to find food items. In other words, foraging within a food patch often exhibits a diminishing returns curve similar to the one in figure 1.5. Such diminishing returns may occur, for instance, because prey actively avoid the forager as they become aware of the threat of predation. Diminishing returns may also occur because the forager has a strategy of picking off the more highly profitable items first (e.g., bigger berries for the hypothetical bird) from a patch with finite resources. Like the hypothetical Web shopper discussed above, the problem for a food forager facing diminishing returns in a patch is whether to continue investing efforts in getting more out of the patch, or to go look for another patch.

Figure 1.5 is a graphical version of a simple *conventional patch model* (Stephens & Krebs, 1986) based on *Charnov's Marginal Value Theorem* (Charnov,




FIGURE 1.5 Charnov's Marginal Value Theorem states that the rate-maximizing time to spend in patch, t^* occurs when the slope of the within-patch gain function g is equal to the average rate of gain, which is the slope of the tangent line R^* .

1976). The model depicted in figure 1.5 assumes that an animal foraging for food encounters only one kind of food patch at random that is never reencountered. When searching for the next food patch, it takes an average of t_B amount of time to find the next patch (between-patch time). Once a patch is encountered, foraging within the patch returns some amount of energy (e.g., as measured by calories) that increases as a function, g , of the time, t_w , spent foraging within the patch. Figure 1.5 shows a diminishing returns function, g , for within-patch foraging. The problem for the forager is how much time, t_w , to spend within each patch before leaving to find the next patch.

The conventional patch model assumes that the animal forager optimizes the overall rate of gain, R , that characterizes the amount of energy gained per unit time of foraging:

$$R = \frac{g(t_w)}{t_B + t_w}, \quad (1.1)$$

or the amount of energy (calories) gained from an average patch divided by the time spent traveling from one patch to the next (t_B) plus the time spent foraging within a patch (t_w). The optimal amount of time, t^* , to spend in a patch is the one that yields the maximum rate of gain, R^* ,

$$R^* = \frac{g(t^*)}{t_B + t^*}. \quad (1.2)$$

Charnov's Marginal Value Theorem (Charnov, 1976) is a mathematical solution to this problem of determining t^* . It basically says that a forager should leave a patch when the rate of gain within the patch [as measured by the slope of $g(t_w)$ or more specifically the derivative $g'(t_w)$] drops below the rate of gain that could be achieved by traveling to, and foraging in, a new patch. That is, the optimal forager obeys the rule,

if $g'(t_w) \geq R^*$, then continue foraging in the patch; otherwise,
when $g'(t_w) < R^*$, then start looking for a new patch.

Charnov's Marginal Value Theorem can be illustrated graphically in figure 1.5 for this simple problem (one kind of patch, randomly distributed in the world). First, note that the gain function g begins to climb only after t_B , which captures the fact that it takes t_B time to go from the last patch to a new patch. If we draw a line beginning at the origin to any point on the gain function, g , then the slope of that line will be the overall rate of gain R , as specified in equation 1.1. Figure 1.5 shows such a line drawn from the origin to a point just tangent to the function g . The slope of this line is the optimal rate of gain R^* as computed in equation 1.2. This can be verified graphically by imagining other lines drawn from the origin to points on the function g . None of those lines will have a steeper slope than the line plotted in figure 1.5. The point at which the line is tangent to g will be the point at which the rate of gain, $g'(t_w)$ within the patch is equal to R^* . This point also determines t^* , the optimum time to spend within the average patch.

Production System Models

The rational analyses in Information Foraging Theory, which often draw from optimal foraging theory, are used to inform the development of production system models. These rational analyses make minimal assumptions about the capabilities of foragers. Herbert Simon (1955) argued that organisms are not optimal, rational agents having perfect information and unlimited computational resources. Rather, organisms exhibit *bounded rationality*. That is, agents are rational and adaptive, within the constraints of the environment and the psychological machinery

available to them biologically. Production system models provide a way of specifying the mechanistic structures and processes that implement bounded rationality. On the one hand, production systems have been used in psychology as a particular kind of computer simulation formalism for specifying the information processing that theorists believe people are performing. On the other hand, production systems have evolved into something more than just a class of computer simulation languages: They have become theories about the basic information processing architecture of cognition that is implemented in human brains (Anderson, 1983; Anderson & Lebiere, 1998; Newell, 1990).

In general, as used in psychology,³ production systems are composed of a set of *production rules* that specify the dynamics of information processing performed by cognition (*how* we think). Production rules operate over memories (or databases) that contain symbolic structures that represent aspects of the external environment and internal thought (*what* we think about). The system operates in a cyclical fashion in which production rules are selected based on the contents of the data memories and then executed. The execution of a production rule typically results in some change to the memories.

The production system models presented in this book are extensions of ACT theory (Anderson et al., 2004; Anderson & Lebiere, 1998). ACT (Adaptive Control of Thought) theory assumes that there are two kinds of knowledge, *declarative* and *procedural* (Ryle, 1949). Declarative knowledge is the kind of knowledge that a person can attend to, reflect upon, and usually articulate in some way (e.g., by declaring it verbally or by gesture). Declarative knowledge includes the kinds of factual knowledge that users can verbalize, such as “The ‘open’ item on the ‘file’ menu will open a file.” Procedural knowledge is the know-how we display in our behavior, without conscious awareness. For instance, knowledge of how to ride a bike and knowledge of how to point a mouse to a menu item are examples of procedural knowledge. Procedural knowledge specifies how declarative knowledge is transformed into active behavior.

ACT-R (the most recent of the ACT theories) has a memory for each kind of knowledge (i.e., a *declarative memory* and a *procedural memory*) plus a special *goal memory*. At any point in time, there may be a number of goals in goal memory, but the system behavior is focused to achieve just one goal at a time.

Complex arrangements of goals and subgoals (e.g., for developing and executing plans to find and use information) can be implemented by manipulating goals in goal memory.

Production rules (or *productions*) are used to represent procedural knowledge in ACT-R. That is, they specify how to apply cognitive skill (know-how) and how to retrieve and use declarative knowledge. Table 1.1 presents an example of a production system for the task of finding a low-cost hotel using a Web site. The example in table 1.1 is not intended to be a psychologically plausible model, but rather it illustrates key aspects of production system models and how they are used in this book. The productions in table 1.1 are English glosses of productions written in ACT-R 5.0, which is discussed in greater detail below.⁴ Each production rule is of the form

IF ⟨condition⟩, THEN ⟨actions⟩.

The condition of a rule specifies a pattern. When the contents of declarative working memory match the pattern, the rule may be selected for application. The actions of the rule specify additions and deletions of content in declarative working memory, as well as motor commands. These actions are executed if the rule is selected to apply. In ACT-R, each production rule has conditions that specify which goal information must be matched and which declarative memory must be retrieved. Each production rule has actions that specify behavioral actions and possibly the setting of subgoals. Typically, ACT-R goal memory is operated on as what is known in computer science as a push-down stack: a kind of memory in which the last item stored will be the first item retrieved. Hence, storing a new goal is referred to as “pushing a goal on the stack,” and retrieval is referred to as “popping a goal from the stack.”

The production rules in table 1.1 assume that declarative memory contains knowledge encoded from the external world about the location and content of links on a Web page. The productions also assume that an initial goal is set to find a hotel price, and the productions accomplish the task by “scanning” through the links keeping track of the lowest price found so far. This involves setting a subgoal to judge the minimum of the current best price and the price just attended when each link is scanned. Table 1.2 presents a trace of the productions in table 1.1