HOT THOUGHT

Mechanisms and Applications of Emotional Cognition

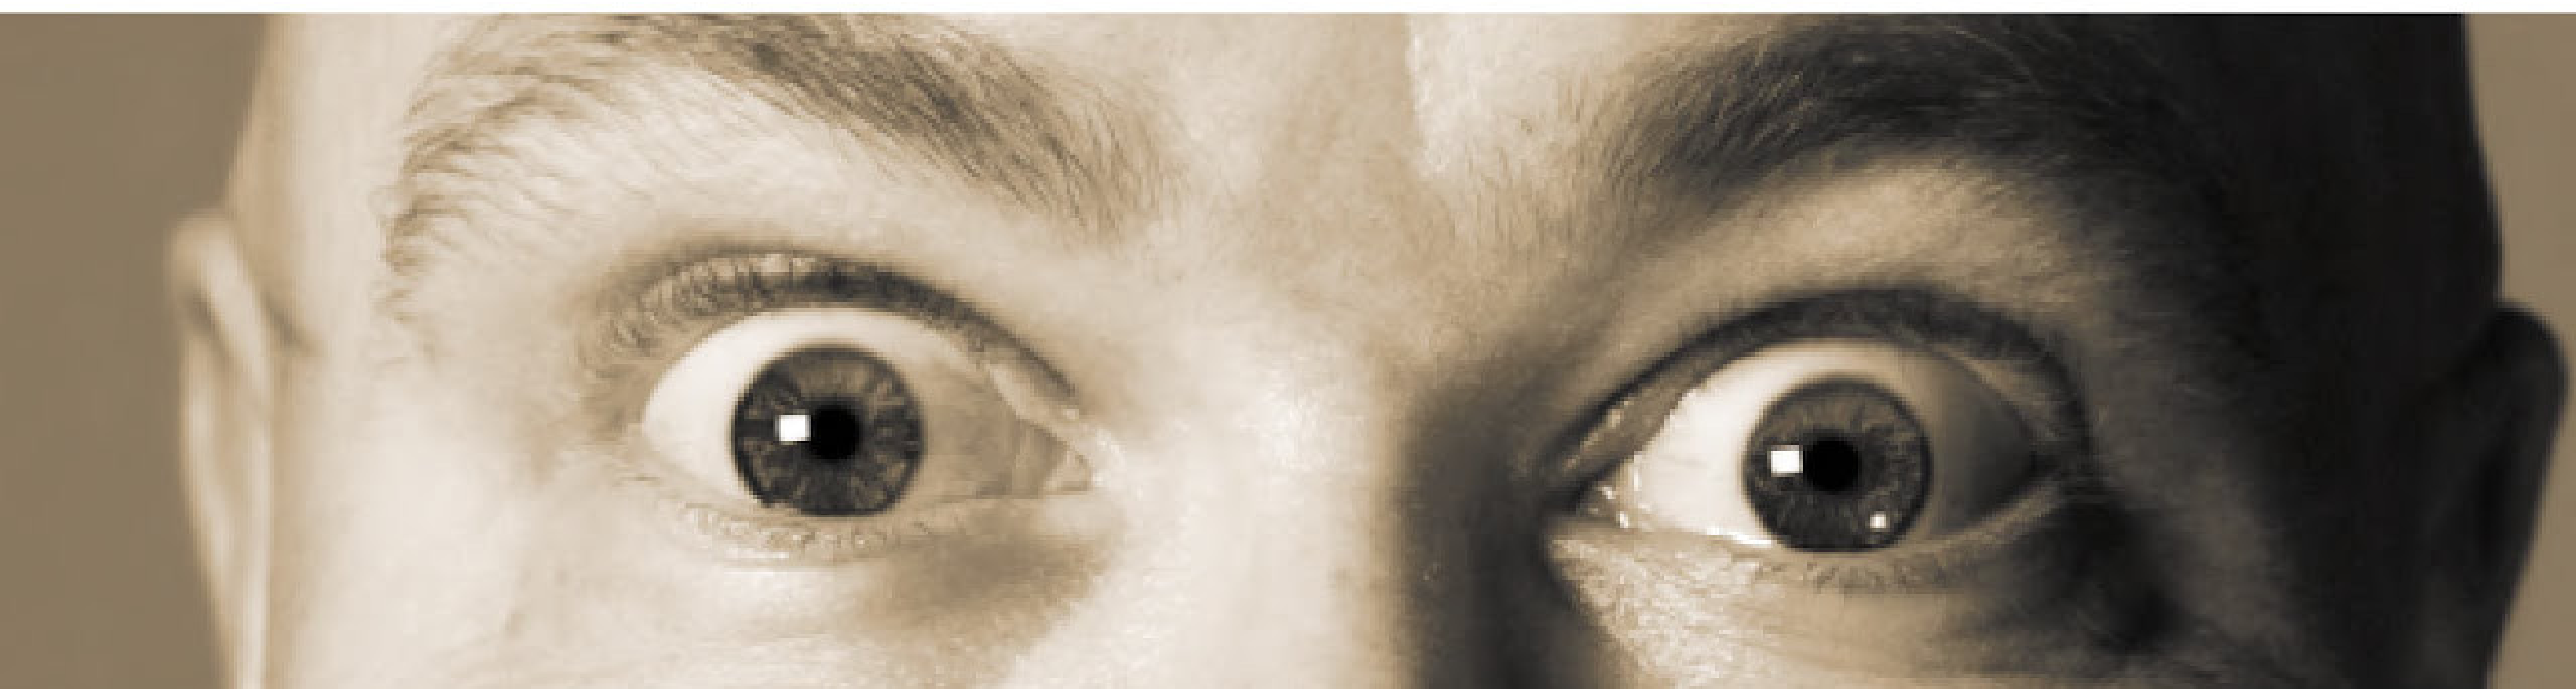

Paul Thagard

Hot Thought

Mechanisms and Applications of Emotional Cognition

Paul Thagard

in collaboration with Fred Kroon, Josef Nerb, Baljinder Sahdra, Cameron Shelley, and Brandon Wagar

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Sabon on 3B2 by Asco Typesetters, Hong Kong, and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Thagard, Paul.
Hot thought : mechanisms and applications of emotional cognition / Paul Thagard ; in collaboration with Fred Kroon ... [et al.].
p. cm.
"A Bradford book."
Includes bibliographical references and index.
ISBN-13: 978-0-262-20164-3 (hc : alk. paper)
ISBN-10: 0-262-20164-X (hc : alk. paper)
1. Emotions and cognition. I. Kroon, Fred. II. Title.

BF311.T416 2006
152.4—dc22 2006041979

10 9 8 7 6 5 4 3 2 1

Contents

Preface

This book is a collection of articles on emotional cognition—hot thought—produced since the 2000 appearance of my book, *Coherence in Thought and Action*. The articles are concerned with mental mechanisms that explain how cognition and emotion interact to produce many kinds of human thinking, from everyday decision making to legal reasoning to scientific discovery to religious belief. Part I of the book describes cognitive, social, neural, and molecular mechanisms that interact in ways that provide explanations of many kinds of human thinking that crucially involve emotion. Part II goes into greater depth in particular domains, showing the application of emotional cognition to the understanding of thinking in law, science, religion, and other areas. The last two chapters address philosophical questions about the normative status of emotional cognition, especially when and how thinking and reasoning *should* be emotional. The final chapter also provides pointers to related current and future work on the nature and significance of emotional cognition.

The chapters reprint the relevant articles largely intact, but I have done some light editing to coordinate references and remove redundancies. Origins of the articles are indicated in the acknowledgments.

Acknowledgments

While I was writing and revising this text, my research was supported by the Natural Sciences and Engineering Research Council of Canada. I have benefited from discussions with many people about the role of emotions in cognition. I am particularly grateful to the coauthors of articles included in this collection: Fred Kroon, Josef Nerb, Baljinder Sahdra, Cameron Shelley, and Brandon Wagar. Please note that Baljinder and Brandon were first authors of their respective articles. For comments on particular chapters, I am indebted to Allison Barnes, Barbara Bulman-Fleming, Peter Carruthers, Mike Dixon, Chris Eliasmith, Christine Freeman, David Gooding, Usha Goswami, Ray Grondin, Tim Kenyon, Patrick Lee, James McAllister, Elijah Millgram, Josef Nerb, Stephen Read, Baljinder Sahdra, Cameron Shelley, Jeff Shrager, Craig Smith, Eliot Smith, Steve Smith, Marcia Sokolowski, Rob Stainton, Robin Vallacher, Chris White, and Zhu Jing.

I am grateful to the respective publishers and to my coauthors for permission to reprint the following articles:

Sahdra, B., and P. Thagard (2003). Self-deception and emotional coherence. *Minds and Machines* 15: 213–231.

Thagard, P. (2001). How to make decisions: Coherence, emotion, and practical inference. In E. Millgram (ed.), *Varieties of Practical Inference*, 355–371. Cambridge, Mass.: MIT Press.

Thagard, P. (2002). Curing cancer? Patrick Lee's path to the reovirus treatment. *International Studies in the Philosophy of Science* 16: 179–193.

Thagard, P. (2002). How molecules matter to mental computation. *Philosophy of Science* 69: 429–446.

Thagard, P. (2002). The passionate scientist: Emotion in scientific cognition. In P. Carruthers, S. Stich, and M. Siegal (eds.), *The Cognitive Basis of Science*, 235–250. Cambridge: Cambridge University Press.

Thagard, P. (2003). Why wasn't O. J. convicted? Emotional coherence in legal inference. *Cognition and Emotion* 17: 361–383.

Thagard, P. (2004). What is doubt and when is it reasonable? In M. Ezcurdia, R. Stainton, and C. Viger (eds.), *New Essays in the Philosophy of Language and Mind. Canadian Journal of Philosophy*, supplementary volume 30, 391–406. Calgary: University of Calgary Press.

Thagard, P. (2005). How to be a successful scientist. In M. E. Gorman, R. D. Tweney, D. C. Gooding, and A. P. Kincannon (eds.), *Scientific and Technological Thinking*, 159–171. Mahwah, N.J.: Lawrence Erlbaum.

Thagard, P. (2005). The emotional coherence of religion. *Journal of Cognition and Culture* 5: 58–74.

Thagard, P. (forthcoming). Critique of emotional reason. In C. de Waal (ed.), *Susan Haack: The Philosopher Replies to Critics*. Buffalo: Prometheus Books.

Thagard, P., and F. W. Kroon (forthcoming). Emotional consensus in group decision making. *Mind and Society*.

Thagard, P., and J. Nerb (2002). Emotional gestalts: Appraisal, change, and emotional coherence. *Personality and Social Psychology Review* 6: 274–282.

Thagard, P., and C. P. Shelley (2001). Emotional analogies and analogical inference. In D. Gentner, K. H. Holyoak, and B. K. Kokinov (eds.), *The Analogical Mind: Perspectives from Cognitive Science*, 335–362. Cambridge, Mass.: MIT Press.

Wagar, B. M., and P. Thagard (2004). Spiking Phineas Gage: A neurocomputational theory of cognitive-affective integration in decision making. *Psychological Review* 111: 67–79.

I am grateful to Todd Nudelman for astute copyediting, and to Abninder Litt for help with the proofreading and index.

1 Mental Mechanisms

Emotional Cognition

Emotional cognition is thinking that is influenced by emotional factors such as particular emotions, moods, or motivations. Here are some examples of situations where people are strongly influenced by their emotions, for good or bad:

A jury member reviews the evidence that an accused man is guilty of murder, but discounts much of it because he seems like a nice guy.
A scientist is told by advisors that her planned research is pointless, but she pursues it because she finds it exciting, and goes on to win a Nobel prize.
An entrepreneur decides to start a novel kind of business because his gut instinct tells him that it will be highly successful, and it is.
A voter wavers over what political candidate to choose, then opts for one with enormous charisma.
A person overcome with anxiety and guilt because of family disasters finds comfort in becoming a born-again Christian.

The primary descriptive aim of this book is to increase understanding of how emotional cognition works in such situations. The secondary normative aim is to suggest ways of improving thinking by appreciating the difference between cases where emotions foster good decisions and those where emotions get in the way. Reasoning is a particular kind of thinking in which decisions are made or beliefs are acquired as the result of the comparative evaluation of different options with respect to various kinds of evidence. Contrary to standard philosophical assumptions, reasoning is often an emotional process, and improving it requires identifying and assessing the impact of emotions.

Accomplishing both the descriptive and normative aims requires an appreciation of the mental mechanisms that underlie mental cognition. Part I of

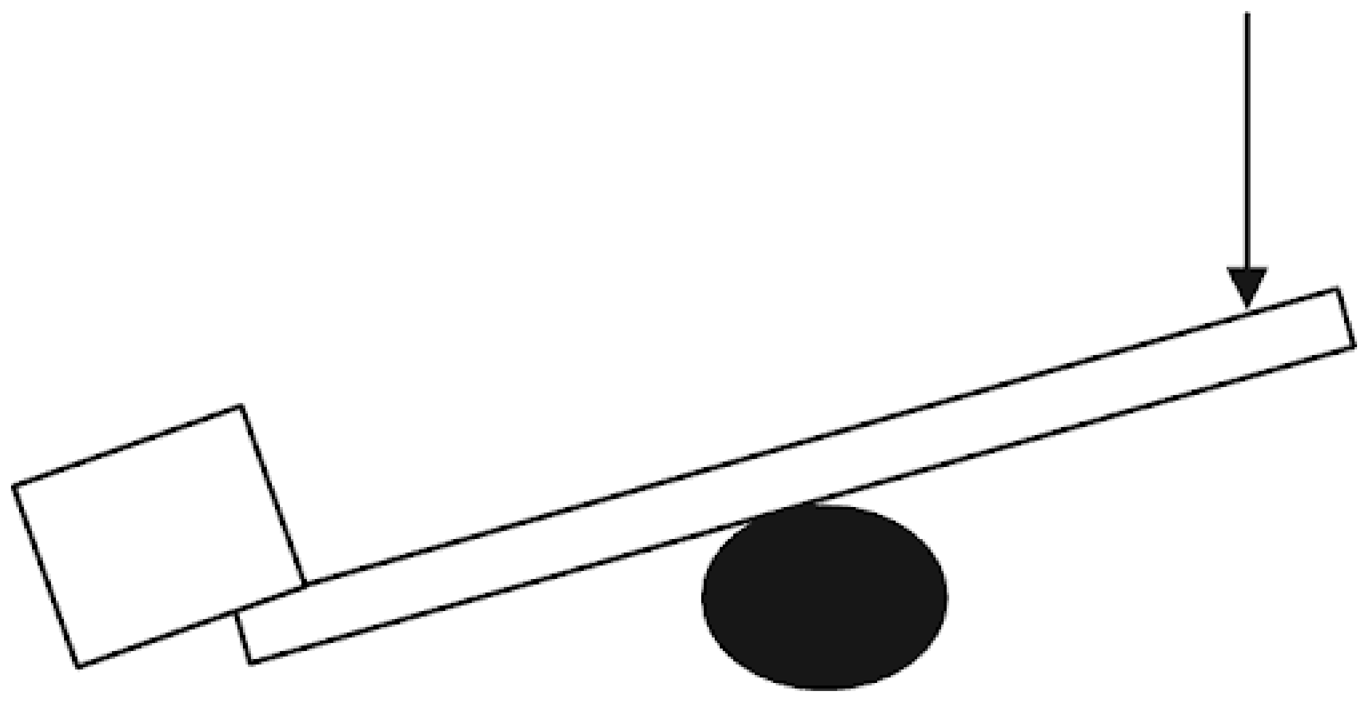

Figure 1.1
A simple machine, the lever.

this book describes some of the cognitive, social, neural, and molecular mechanisms that are crucial for emotional thinking. These mechanisms provide explanations of how emotions often influence people's decisions and other inferences. Part II shows how such mechanisms operate in emotional reasoning in law, science, and religion, and discusses the difference between desirable and undesirable kinds of emotional thinking. Before getting into details, however, it is necessary to say what mechanisms are and how they contribute to explanation.

Machines and Mechanisms

Humans invented simple machines such as levers and wedges long before the beginning of recorded history. Consider the basic lever shown in figure 1.1. It consists of only two parts, a stick and a rock. But levers are very powerful and enabled people to build huge structures such as the Egyptian pyramids. In general, a machine is an assemblage of components that transmit force, motion, and energy to each other to accomplish some task. To describe a machine and explain its operation, we need to specify its components, their properties, and their relations to other parts. Most important, we need to describe how changes to the properties and relations of the parts with respect to force, motion, and energy enable the machine to accomplish its tasks. The lever in figure 1.1 operates by virtue of the fact that the stick is rigid and is on top of the rock, which is solid. Applying force to the top of the stick makes the bottom of the stick move and lift the block, thus accomplishing the machine's task.

Ancient philosophers such as Epicurus and Lucretius realized that natural events can be explained in the same way that people explain the operations of machines. They postulated that all matter consists of atoms whose interactions determine the behavior of the objects that they constitute. According to

Lucretius (1969), even human thought is the result of the interactions of atoms. Natural phenomena are thus explained by *mechanisms*, which are like machines in that their changes result from forces and motions applied to their parts. Unlike machines, however, a natural mechanism is not constructed by humans and has no human-contrived task, although it may have a biological function. Nevertheless, the form of explanation is the same for both natural and constructed mechanisms: specify components, their properties and relations, and describe how changes in force, motion, and energy propagate through the system.

Mechanistic explanation has been fabulously successful in modern science, with triumphs such as Newton's theory of motion, the atomic theory of matter, the germ theory of disease, and the biological theories of evolution and genetics. But many philosophers, from Plato to Descartes, Leibniz, and even some modern thinkers, have resisted the extension of mechanistic explanation to mental phenomena. Plato's forms, Descartes's soul, Leibniz's monads, and twentieth-century phenomenology and hermeneutics are all attempts to understand the human mind in terms other than its material parts. Granted, early mechanistic attempts to explain mental operations were unsatisfactory, as Lucretius' atoms, Hume's associations of ideas, and modern behaviorist stimulus–response theories were much too crude to explain the richness of human thought. But I will now review a range of powerful mechanisms that have become available in the past fifty years for explaining thought in general and emotional reasoning in particular.

Kinds of Mental Mechanisms

Current cognitive science explains human thinking using a confluence of cognitive, neural, molecular, and social mechanisms. Those mechanisms most familiar are cognitive ones, which describe the mind as operating by the application of computational procedures on mental representations (for a survey, see Thagard 2005). Mental representations are cognitive structures, such as concepts, rules, and images, that are manipulated by algorithmic processes such as spreading activation of concepts, chaining of rules, and rotation of images. Connectionist explanations of mental phenomena describe them as arising from representations constructed by the activity of simple neuronlike units and processes that include spreading activation and modification of links between units.

Neural mechanisms are much closer to the operation of the brain in two key respects. First, they involve artificial neurons that are much more neurologically realistic than the highly simplified units found in connectionist

Table 1.1
Constituents of mental mechanisms

Mechanisms	Components	Relations	Interactions	Changes
Social	Persons and social groups	Association, membership	Communication	Influence, group decisions
Cognitive	Mental representations such as concepts	Constituents, associations, implication	Computational processes	Inferences
Neural	Neurons, neural groups	Synaptic connections	Excitation, inhibition	Brain activity
Molecular	Molecules such as neurotransmitters and proteins	Constituents, physical connection	Biochemical reactions	Transformation of molecules

explanations. For example, real neurons have temporal properties such as firing in coordination with each other that are crucial to their ability to support complex mental activity. Second, whereas connectionist models usually employ small numbers of interconnected units, brain mechanisms involve billions of neurons organized into functional areas such as the hippocampus and various parts of the cortex. Neural mechanisms depend on molecular mechanisms, such as the movement of ions within neurons, the flow of neurotransmitters in the synaptic connections between them, and the circulation of hormones through the brain's blood supply.

Finally, because human thought often involves interaction with other people, we need to attend to the social mechanisms that allow one person's thinking to influence another's. Social mechanisms involve verbal and other kinds of communication, including ones that make possible the transfer of emotions as well as other kinds of information. Table 1.1 summarizes some of the ingredients of the four kinds of mechanisms useful for different levels of mental explanation. Chapters 2 through 7 provide much more detailed examples of the mechanisms at each level that are relevant for understanding emotional thought.

The Nature of Complex Mechanisms

My approach to emotions is in keeping with the mechanism-based view of explanation espoused by such philosophers of science as Salmon (1984b), Bechtel and Richardson (1993), and Bechtel and Abrahamsen (2005). Machamer, Darden, and Craver (2000, p. 3) characterize mechanisms as

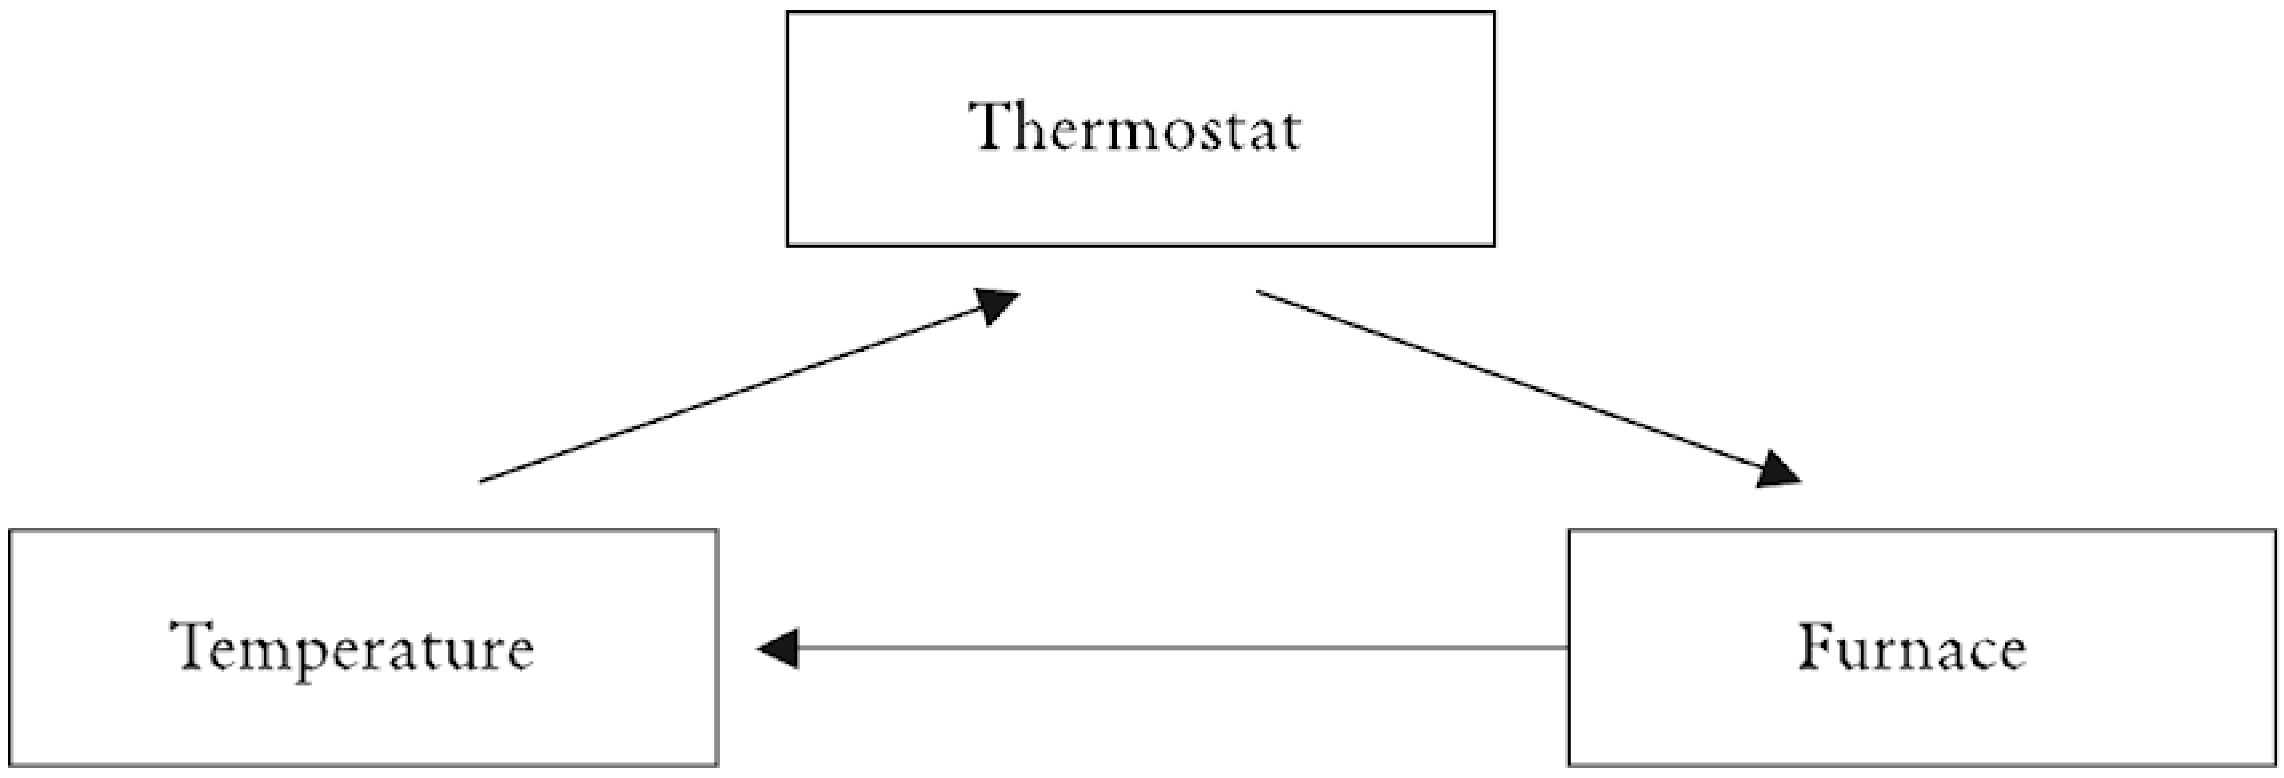

Figure 1.2
Feedback mechanism that regulates temperature.

"entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions." I prefer the term "component" to "entity" because it indicates that the objects in a mechanism are part of a larger system; and I prefer "interaction" and "change" to "activity" because they sound less anthropomorphic. More important, I find the reference to start and finish conditions misleading, because the mechanisms needed to explain emotional thinking involve ongoing feedback processes rather than unidirectional changes.

For a simple example of feedback, consider the system shown in figure 1.2. A thermostat is a device that detects when the temperature drops below a given threshold; it then signals a furnace to start up and heat the air. Taken separately, the thermostat and the furnace are unidirectional machines describable by simple rules: if the temperature is less than *X*, then signal the furnace to start; if a signal to start is received, then start the furnace. Taken together, however, they constitute a feedback mechanism that functions to maintain the temperature within a stable range. There are no start and finish conditions, but rather an ongoing process of temperature regulation. The human body contains many such feedback mechanisms, for example, ones that maintain blood pressure and cholesterol levels. As we will see in the next four chapters, those mechanisms most relevant to explaining emotions—cognitive, neural, molecular, and social—are feedback mechanisms rather than unidirectional ones with start and finish conditions.

Machamer, Darden, and Craver (2000) observe that mechanisms occur as nested hierarchies and that their description is frequently multilevel. They argue that higher-level entities and activities are essential to the intelligibility of those at lower levels and vice versa, and that mental phenomena are best explained by the integration of different levels. I fully agree, and will describe in detail how an understanding of emotional thinking needs to integrate into account mental computations accomplished by molecular mechanisms such as chemical reactions that operate within and between neurons and other cells.

The rest of the book is concerned with applying knowledge of emotional mechanisms to illuminate thinking in the domains of law, science, and religion. Chapter 8 provides an explanation of why the jury in the O. J. Simpson trial did not convict him, arguing that their decision was based on a combination of explanatory and emotional coherence. More generally, chapter 9 discusses the nature of doubt and reasonable doubt in legal and other contexts. Chapters 10 through 12 are concerned with both the positive and negative roles of emotions in scientific thinking. Chapter 10 provides a general overview of the ways in which emotions affect the pursuit, discovery, and evaluation of scientific ideas, and chapter 11 looks at how they contributed to the development of research in a particular case. Chapter 12 provides a distillation of advice on how to be a successful scientist, including valuable emotional habits. Chapters 13 and 14 discuss the application of emotional coherence to religious thinking, the former by looking at a particular case of the self-deception of a minister, the latter considering more generally the emotional content of religious beliefs.

Implicit in all these discussions of the applications of emotional reason is a set of judgments about appropriate standards for when the incursion of emotional elements into cognitive deliberation is legitimate. Chapter 15 begins the discussion of this normative question, which is continued in the concluding chapter 16. I reject the traditional assumption that emotion is inimical to good reason, but also argue against the romantic view that emotion always contributes to good thinking. Chapter 16 also sketches a number of other areas of ongoing investigation of hot thought, including its relevance to conflicts of interest, generation of explanatory hypotheses, and neuroeconomics.

We shall see that many aspects of human thinking and reasoning are influenced by emotions. An understanding of how emotions affect thought both positively and negatively requires the specification of cognitive, social, neural, and molecular mechanisms. Such mechanisms consist of components that undergo interactive changes that may involve feedback processes. Our best strategy for explaining emotional thinking is to look at multiple integrated levels of causal mechanisms.

2 How to Make Decisions

Students face many important decisions: What college or university should I attend? What should I study? What kind of job should I try to get? Which people should I hang out with? Should I continue or break off a relationship? Should I get married? Should I have a baby? What kind of medical treatment should I use? A theory of practical reasoning should have something to say about how students and other people can improve their decision making.

I have often taught a first-year course on critical thinking intended to help students improve their reasoning about what to believe and about what to do. After spending about two-thirds of the course on ways of improving judgments about the truth and falsity of controversial claims in areas such as medicine and pseudoscience, I devote the last third to practical reasoning, with the focus on how people can make better decisions. I discuss both the kinds of erroneous reasoning that decision makers commonly fall into, and some systematic models that have been proposed by psychologists, economists, and philosophers to specify how people *should* make decisions.

Many students in the course dislike these models, and resist the claim that using them is preferable to making decisions simply by intuition. They trust their "gut feelings" more than they trust the analytical methods that require a systematic and mathematical comparative assessment of competing actions that satisfy multiple criteria. The textbooks I use (including Gilovich 1991; Russo and Schoemaker 1989; Schick and Vaughn 1999) encourage people to avoid the use of intuition and instead to base their judgments and decisions on reasoning strategies that are less likely to lead to common errors in reasoning. From this perspective, decision making should be a matter of calculation, not intuition.

While I agree that intuition-based decision making can lead to many problems, I also think that calculation-based decision making of the sort recommended by psychologists and economists has some serious pitfalls. In this chapter, I will try to offer a synthesis and partial reconciliation of intuition and calculation models of decision, using a recently developed theory of emotional coherence (Thagard 2000). This theory builds on a previous coherence-based theory of decision making developed in collaboration with Elijah Millgram. Understanding decision making in terms of emotional coherence enables us to appreciate the merits of both intuition and calculation as contributors to effective practical reasoning.

Decision as Intuition

Suppose you are a student trying to decide whether to study (1) an arts subject such as philosophy or art history in which you have a strong interest or (2) a subject such as economics or computer science which may lead to a more lucrative career. To make this decision intuitively is just to go with the option that is supported by your emotional reactions to the two alternatives. You may have a strongly positive gut feeling toward the more interesting subject along with a strongly negative feeling about the more career-oriented one, or your feelings may be just the opposite. More likely is that you feel positive feelings toward both alternatives, along with accompanying anxiety caused by your inability to see a clearly preferable option. In the end, intuitive decision makers choose an option based on what their emotional reactions tell them is preferable.

There is much to be said for intuitive decision making. One obvious advantage is speed: an emotional reaction can be immediate and lead directly to a decision. If your choice is between chocolate and vanilla ice cream, it would be pointless to spend a lot of time and effort deliberating about the relative advantages and disadvantages of the two flavors. Instead, an emotional reaction such as "chocolate—yum!" can make for a quick and appropriate decision. Another advantage is that basing your decisions on emotions helps to ensure that the decisions take into account what you really care about. If you are pleased and excited about a possible action, that is a good sign that the action promises to accomplish the goals that are genuinely important to you. Finally, decisions based on emotional intuitions lead directly to action: the positive feeling toward an option will motivate you to carry it out.

But emotion-based intuitive decision making can also have some serious disadvantages. An option may seem emotionally appealing because of failure to consider other available options. Intuition may suggest buying chocolate ice cream only because you have failed to consider a lower-fat alternative that would be a healthier choice. Intuition is also subject to the intense craving that drug addicts call "jonesing." If you are jonesing for cocaine, or for a pizza, or for a Mercedes-Benz convertible, your intuition will tell you to choose what you crave, but only because the craving has emotionally swamped other desires that you will be more aware of when the craving is less intense.

Another problem with intuition is that it may be based on inaccurate or irrelevant information. Suppose you need to decide whom to hire for a job. If you are prejudiced against people of a particular sex, race, or ethnicity, then your intuition will tell you not to hire them, even if they have better qualifications for doing the job well. It is difficult to determine introspectively whether your intuitions derive from reliable and relevant information.

Finally, intuitive reasoning is problematic in group situations where decisions need to be made collectively. If other people disagree with your choices, you cannot simply contend that your intuitions are stronger or better than the intuitions of others. Defending your emotional reactions and attempting to reach a consensus with other people requires a more analytical approach than simply expressing your gut feelings.

Decision as Calculation

Experts on decision making recommend a more systematic and calculating approach. For example, Bazerman (1994, p. 4) says that rational decision making should include the following six steps:

1. Define the problem, characterizing the general purpose of your decision.
2. Identify the criteria, specifying the goals or objectives that you want to be able to accomplish.
3. Weight the criteria, deciding the relative importance of the goals.
4. Generate alternatives, identifying possible courses of action that might accomplish your various goals.
5. Rate each alternative on each criterion, assessing the extent to which each action would accomplish each goal.
6. Compute the optimal decision, evaluating each alternative by multiplying the expected effectiveness of each alternative with respect to a criterion times the weight of the criterion, then adding up the expected value of the alternative with respect to all criteria.

We can then pick the alternative with the highest expected value and make a decision based on calculation, not on subjective emotional reactions. Using slightly different terminology, Russo and Schoemaker (1989, chap. 6) recommend essentially the same kind of decision making process based on multiple weighted factors.

Some students dismiss this kind of process as robot-like, and find it offensive that important decisions in their lives might be made mathematically. A cartoon in the *New Yorker* (Jan. 10, 2000, p. 74) shows a man sitting at a computer and saying to a woman: "I've done the numbers, and I will marry you." Some decisions, at least, seem inappropriately based on doing the numbers. But is the emotional dismissal of Bazerman's six-step calculation method justified? We can certainly see some notable advantages of the calculation method over the intuition method. First, it is set up to avoid neglecting relevant alternatives and goals. Second, it makes explicit the consideration of how the various alternatives contribute to the various goals. Third, it puts the decision making process out in the open, enabling it to be carefully reviewed by a particular decision maker and also by others involved in a group decision process.

However, the calculation method of decision making may be more difficult and less effective than decision experts claim. Suppose you are trying to decide between two courses of study, say philosophy versus computer science, and you systematically list all the relevant criteria such as how interesting you find the subjects and how well they fit with your career plans. You then weight the criteria and estimate the extent to which each option satisfies them, and proceed to a calculation of the expected value of the competing choices. Having done this, you find that the expected value of one option, say philosophy, exceeds that of the other. But what if you then have the reaction—"I don't want to do that!" Your emotional reaction need not be crazy, because it may be that the numerical weights that you put on your criteria do not reflect what you really care about. Moreover, your estimates about the extent to which different actions accomplish your goals may be very subjective and fluid, so that your unconscious estimation is at least as good as your conscious one. I once knew someone who told me that she made decisions by first flipping a coin, with heads for one option and tails for another. When the coin came up heads, she would note her emotional reaction, which gave her a better idea of whether she really wanted the option associated with heads. She then used this emotional information to help her make a choice between the two options.

There is empirical evidence that calculation may sometimes be inferior to intuition in making good judgments. Damasio (1994) describes people with injuries that have disconnected the parts of their brains that perform verbal reasoning and numerical calculation from emotional centers such as the amygdala. With their abstract reasoning abilities intact, you might think that the patients become paragons of rationality, like Spock or Data but there are effective algorithms for approximating the maximization of coherence construed as constraint satisfaction (Thagard and Verbeurgt 1998).

I will now make this account of coherence more concrete by showing how it applies to inference about what to do. Elijah Millgram and I have argued that practical inference involves coherence judgments about how to fit together various possible actions and goals (Millgram and Thagard 1996; Thagard and Millgram 1995). On our account, the elements are actions and goals, the positive constraints are based on facilitation relations (the action of going to Paris facilitates my goal of having fun), and the negative constraints are based on incompatibility relations (you cannot go to Paris and London at the same time). Deciding what to do is based on inference to the most coherent plan, where coherence involves evaluating goals as well as deciding what to do.

More exactly, deliberative coherence can be specified by the following principles:

Principle 1: Symmetry Coherence and incoherence are symmetrical relations: If a factor (action or goal) F_1 coheres with a factor F_2, then F_2 coheres with F_1.

Principle 2: Facilitation Consider actions $A_1 \ldots A_n$ that together facilitate the accomplishment of goal G. Then

(a) each A_i coheres with G,
(b) each A_i coheres with each other A_j, and
(c) the greater the number of actions required, the less the coherence among actions and goals.

Principle 3: Incompatibility

(a) If two factors cannot both be performed or achieved, then they are strongly incoherent.
(b) If two factors are difficult to perform or achieve together, then they are weakly incoherent.

Principle 4: Goal priority Some goals are desirable for intrinsic or other noncoherence reasons.

Principle 5: Judgment Facilitation and competition relations can depend on coherence with judgments about the acceptability of factual beliefs.

Principle 6: Decision Decisions are made on the basis of an assessment of the overall coherence of a set of actions and goals.

In order to assess overall coherence, we can use the computer program DECO (short for "deliberative coherence"). DECO represents each element (goal or action) by a neuron-like unit in an artificial neural network and then spreads activation through the network in a way that activates some units and deactivates others. At the end of the spread of activation, the active units represent elements that are accepted, while the deactivated ones represent elements that are rejected. DECO provides an efficient and usable way to compute the most coherent set of actions and goals.

At first glance, deliberative coherence might seem like a variant of the calculation model of decision making. Figuring out which action best coheres with your goals sounds like Bazerman's calculation of the expected value of alternatives based on the extent to which they satisfy weighted criteria. But there are some crucial differences. Unlike Bazerman's proposal, the deliberative coherence model of decision does not take the weights of the goals as fixed. In DECO, units representing some of the goals get initial activation in accord with principle 4, goal priority, but their impact depends on their relation to other goals: even a basic goal can be deactivated, at least partially, by other goals. The impact of goals on decision making depends on their activation, which depends on their relation to other goals and to various actions. For example, students trying to decide what to do on the weekend might start off thinking that what they most want to do is to have fun, but realize that having fun is not so important because it conflicts with other goals such as studying for an important exam or saving money to pay next term's tuition.

Psychologically, decision as coherence is very different from decision as calculation. Calculations are conscious and explicit, displayable to everyone on pencil and paper. In contrast, if coherence maximization in human brains is similar to what happens in the artificial neural networks used in DECO, then assessment of coherence is a process not accessible to consciousness. What comes to consciousness is only the result of the process of coherence maximization: the realization that a particular action is the one I want to perform. Thus, as an account of how decisions are made by people, deliberative coherence is closer to the intuition model of decision than to the calculation model. Coherence is not maximized by an explicit, consciously accessible calculation, but by an unconscious process whose output is the intuition that one action is preferable to others. There is, however, a major difference between the deliberative coherence account of decision making and the intuition account: intuitions about what to do are usually emotional, involving feelings that one action is a good thing to do and that alternatives are bad things to do. Fortunately, coherence theory can naturally be extended to encompass emotional judgments.

Emotional Coherence

In the theory of coherence stated above, elements have the epistemic status of being accepted or rejected. We can also speak of degree of acceptability, which in artificial neural network models of coherence is interpreted as the degree of activation of the unit that represents the element. I propose that elements in coherence systems have, in addition to acceptability, an emotional *valence*, which can be positive or negative. Depending on the nature of what the element represents, the valence of an element can indicate likability, desirability, or other positive or negative attitudes. For example, the valence of Mother Theresa for most people is highly positive, while the valence of Adolf Hitler is highly negative. Many other researchers have previously proposed introducing emotion into cognitive models by adding valences or affective tags (Bower 1981, 1991; Fiske and Pavelchak 1986; Lodge and Stroh 1993; Ortony, Clore, and Collins 1988; Sears, Huddy, and Schaffer 1986). Kahneman (1999) reviews experimental evidence that evaluation on the good/bad dimension is a ubiquitous component of human thinking.

Just as elements are related to each other by the positive and negative deliberative constraints described in the last section, they also can be related by positive and negative valence constraints. Some elements have intrinsic positive and negative valences, for example *pleasure* and *pain*. Other elements can acquire valences by virtue of their connections with elements that have intrinsic valences. These connections can be special valence constraints, or they can be any of the constraints posited by the theory of deliberative coherence. For example, if someone has a positive association between the concepts of *dentist* and *pain*, where *pain* has an intrinsic negative valence, then *dentist* can acquire a negative valence. However, just as the acceptability of an element depends on the acceptability of all the elements that constrain it, so the valence of an element depends on the valences of all the elements that constrain it.

The basic theory of emotional coherence can be summarized in three principles analogous to the qualitative principles of coherence above:

1. Elements have positive or negative valences.
2. Elements can have positive or negative emotional connections to other elements.

3. The valence of an element is determined by the valences and acceptability of all the elements to which it is connected.

As already mentioned, coherence can be computed by a variety of algorithms, but the most psychologically appealing model, and the model that first inspired the theory of coherence as constraint satisfaction, employs artificial neural networks. In this connectionist model, elements are represented by units, which are roughly analogous to neurons or neuronal groups. Positive constraints between elements are represented by symmetric excitatory links between units, and negative constraints between elements are represented by symmetric inhibitory links between units. The degree of acceptability of an element is represented by the activation of a unit, which is determined by the activation of all the units linked to it, taking into account the strength of the various excitatory and inhibitory links.

It is straightforward to expand this kind of model into one that incorporates emotional coherence. In the expanded model, called "HOTCO" for "hot coherence," units have valences as well as activations, and units can have input valences to represent their intrinsic valences. Moreover, valences can spread through the system in a way very similar to the spread of activation, except that valence spread depends in part on activation spread. An emotional decision emerges from the spread of activation and valences through the system because nodes representing some actions receive positive valence while nodes representing other actions receive negative valence. The gut feeling that comes to consciousness is the end result of a complex process of cognitive and emotional constraint satisfaction. Emotional reactions such as happiness, anger, and fear are much more complex than positive and negative valences, so HOTCO is by no means a general model of emotional cognition. But it does capture the general production by emotional inference of positive and negative attitudes toward objects, situations, and choices.

It might seem that we can now abandon the cognitive theory of deliberative coherence for the psychologically richer theory of emotional coherence, but that would be a mistake for two reasons. First, emotional coherence must interconnect with other kinds of coherence that involve inferences about what is acceptable as well as about what is emotionally desirable. The valence of an element does not depend just on the valences of the elements that constrain it, but also on their acceptability. Attaching a negative valence to the concept *dentist*, if it does not already have a negative valence from previous experience, depends both on the negative valence for *causes-pain* and the acceptability (confidence) of *causes-pain* in the current context. The inferential situation here is analogous to expected utility theory, in which the expected utility of an action is calculated by summing, for various outcomes, the result of multiplying the probability of the outcome times the utility of the outcome. The calculated valence of an element is like the expected utility of an action, with degrees of acceptability analogous to probabilities and valences analogous to utilities. There is no reason, however, to expect degrees of acceptability and valences to have the mathematical properties that define probabilities and utilities. Because the valence calculation depends on the acceptability of all the relevant elements, it can be affected by other kinds of coherence. For example, the inference concerning whether to trust someone depends largely on the valence attached to them based on all the information you have about them, where this information derives in part from inferences based on explanatory, analogical, and conceptual coherence (Thagard 2000).

The second reason for not completely replacing the cold (nonemotional) theory of deliberative coherence with the hot theory of emotional coherence is that people can sometimes come up with incompatible hot and cold judgments about what to do. Unconsciously using deliberative coherence may produce the judgment that you should not do something, while emotional coherence leads you in a different direction. For example, students seeing the first nice spring day at the end of a long Canadian winter might decide emotionally to go outside and enjoy it, while at the same time reasoning that the alternative of finishing up overdue end-of-term projects is more coherent with their central goals such as graduating from university. I am not the only person capable of thinking: "The best thing for me to do is *X*, but I'm going to do *Y*." Jonesing in reaction to vivid stimuli can make emotional coherence swamp deliberative coherence.

The theory of emotional coherence provides a psychologically realistic way of understanding the role of intuition in decisions. My gut feeling that I should go to Paris is the result of an unconscious mental process in which various actions and goals are balanced against each other. The coherence process involves both inferences about what I think is true (e.g., I'll have fun in Paris) and inferences about the extent to which my goals will be accomplished. But the coherence computation determines not only what elements will be accepted and rejected, but also an emotional reaction to the element. It is not just "go to Paris—yes" or "go to Paris—no," but "go to Paris—yeah!" or "go to Paris—yuck!"

As we just saw, however, emotional coherence may be better as a descriptive theory of how people make decisions than as a normative theory of how people should make decisions. Judgments based on emotional coherence may be subject to the same criticisms that I made against intuitive decisions:

unconscious mental process in which many factors are balanced against each other until a judgment is reached that accepts some beliefs and rejects others in a way that approximately maximizes coherence.

This does not mean that practical and theoretical reasoning should be sneered at. Reasoning is a verbal, conscious process that is easily communicated to other people. People are rarely convinced by an argument directly, but the fact that reasoning does not immediately translate into inference does not make it pointless. Making reasoning explicit in decisions helps to communicate to all the people involved what the relevant goals, actions, and facilitation relations might be. If communication is effective, then the desired result will be that each decision maker will make a better informed intuitive decision about what to do.

Improving inference is both a matter of recognizing good inference procedures such as Informed Intuition and watching out for errors that people commonly make. Such errors are usually called fallacies by philosophers and biases by psychologists. Psychologists, economists, and philosophers have identified a variety of error tendencies in decision making, such as considering sunk costs, using bad analogies, and being overconfident in judgments. Noticing the role of emotional coherence in decision making enables us to expand this list to include emotional determinants of bad decision making such as jonesing and failing to perceive the emotional attitudes of other people. In this paper I have emphasized the positive strategy of making decisions using a recommended procedure, Informed Intuition, but a fuller account would also develop the negative strategy of avoiding various tendencies that are natural to human thinking and that often lead to poor decisions.

The coherence model of decision making allows goals to be adjusted in importance while evaluating a decision, but it does not address the question of how we adopt new goals. Millgram's (1997) account of practical induction is useful for describing how people in novel situations can develop new interests that provide them with new goals. A full theory of decision making would have to include an account of where human goals come from and how they can be evaluated. People who base their decisions only on the goals of sex, drugs, and rock and roll may achieve local coherence, but they have much to learn about the full range of pursuits that enrich human lives.

Conclusion

I have tried in this chapter to provide students and other people with a model of decision making that is both natural and effective. Practical inference is not simply produced by practical syllogisms or cost-benefit calculations, but requires assessment of the coherence of positively and negatively interconnected goals and actions. This assessment is an unconscious process based in part on emotional valences attached to the various goals to be taken into consideration, and yields a conscious judgment that is not just a belief about what is the best action to perform but also a positive emotional attitude toward that action. Reason and emotion need not be in conflict with each other, if the emotional judgment that arises from a coherence assessment takes into account the relevant actions and goals and the relations between them. The procedure I recommend, Informed Intuition, shows how decisions can be both intuitive and reasonable. (See chaps. 15 and 16 for further discussion of normative issues.)

3 Emotional Analogies and Analogical Inference

Paul Thagard and Cameron Shelley

Introduction

Despite the growing appreciation of the relevance of affect to cognition, analogy researchers have paid remarkably little attention to emotion. This chapter discusses three general classes of analogy that involve emotions. The most straightforward are analogies and metaphors *about* emotions, for example "Love is a rose and you better not pick it." Much more interesting are analogies that involve the transfer of emotions, for example in empathy in which people understand the emotions of others by imagining their own emotional reactions in similar situations. Finally, there are analogies that generate emotions, for example analogical jokes that generate emotions such as surprise and amusement.

Understanding emotional analogies requires a more complex theory of analogical inference than has been available. The next section presents a new account that shows how analogical inference can be defeasible, holistic, multiple, and emotional, in ways to be described. Analogies about emotions can to some extent be explained using the standard models such as ACME and SME, but analogies that transfer emotions require an extended treatment that takes into account the special character of emotional states. We describe HOTCO, a new model of emotional coherence, that simulates transfer of emotions, and show how HOTCO models the generation of emotions such as reactions to humorous analogies. Finally, we supplement our anecdotal collection of emotional analogies by discussing a more comprehensive sample culled from Internet wire services.

Analogical Inference: Current Models

In logic books, analogical inference is usually presented by a schema such as the following (Salmon 1984a, p. 105):

Objects of type *X* have properties *G*, *H*, etc.
Objects of type *Y* have properties *G*, *H*, etc.
Objects of type *X* have property *F*.
Therefore: Objects of type *Y* have property *F*.

For example, when experiments determined that large quantities of the artificial sweetener saccharine caused bladder cancer in rats, scientists analogized that it might also be carcinogenic in humans. Logicians additionally point out that analogical arguments may be strong or weak depending on the extent to which the properties in the premises are relevant to the property in the conclusion.

This characterization of analogical inference, which dates back at least to John Stuart Mill's nineteenth-century *System of Logic*, is flawed in several respects. First, logicians rarely spell out what "relevant" means, so the schema provides little help in distinguishing strong analogies from weak. Second, the schema is stated in terms of objects and their properties, obscuring the fact that the strongest and most useful analogies involve relations, in particular causal relations (Gentner 1983; Holyoak and Thagard 1995). Such causal relations are usually the key to determining relevance: if, in the above schema, *G* and *H* together cause *F* in *X*, then analogically they may cause *F* in *Y*, producing a much stronger inference than just counting properties. Third, logicians typically discuss analogical arguments and tend to ignore the complexity of analogical inference, which requires a more holistic assessment of a potential conclusion with respect to other information. There is no point in inferring that objects of type *Y* have property *F* if you already know of many such objects that lack *F*, or if a different analogy suggests that they do not have *F*. Analogical inference must be defeasible, in that the potential conclusion can be overturned by other information, and it must be holistic in that everything the inference maker knows is potentially relevant to overturning or enhancing the inference.

Compared to the logician's schema, much richer accounts of the structure of analogies have been provided by computational models of analogical mapping such as SME (Falkenhainer, Forbus, and Gentner 1989) and ACME (Holyoak and Thagard 1989). SME uses relational structure to generate candidate inferences, and ACME transfers information from a source analog to a target analog using a process that Holyoak, Novick, and Melz (1994) called copying with substitution and generation (CWSG). Similar processes are used in case-based reasoning (Kolodner 1993), and in many other computational models of analogy.

But all of these computational models are inadequate for understanding analogical inference in general and emotional analogies in particular. They

(see-in (centroscymnus deep-ocean-water) see-in-1)
(inhabit (centroscymnus deep-ocean-water) inhabit-1)
(enable (have-1 see-in-1) enable-1)
(because (absorb-1 penetrate-1) because-1)
(adapt (see-in-1 inhabit-1) adapt-1)

Target: coelacanth-3
(have (coelacanth rod-pigment-3) have-3)
(absorb (rod-pigment-3 473nm-light) absorb-3)
(penetrate (473nm-light deep-ocean-water) penetrate-3)
(see-in (coelacanth deep-ocean-water) see-in-3)
(enable (have-3 see-in-3) enable-3)
(because (absorb-3 penetrate-3) because-3)

Operating in specific mode, HOTCO is asked what depth the coelacanth inhabits, and uses the proposition INHABIT-1 in the source to construct for the target the proposition

(inhabit (coelacanth deep-ocean-water) inhabit-new)

Operating in broad mode and doing general copying with substitution and generation, HOTCO can analogically transfer everything about the source to the target, in this case generating the proposition that coelacanths inhabit deep water as a candidate to be inferred.

However, HOTCO does *not* actually infer the new proposition, because analogical inference is defeasible. Rather, it simply establishes an excitatory link between the unit representing the source proposition INHABIT-1 and the target proposition INHABIT-NEW. This link represents a positive constraint between the two propositions, so that coherence maximization will encourage them to be accepted together or rejected together. The source proposition INHABIT-1 is presumably accepted, so in the HOTCO model it will have positive activation which will spread to provide positive activation to INHABIT-NEW, unless INHABIT-NEW is incompatible with other accepted propositions that will tend to suppress its activation. Thus analogical inference is defeasible, because all HOTCO does is to create a link representing a new constraint for overall coherence judgment, and it is holistic, because the entire constraint network can potentially contribute to the final acceptance or rejection of the inferred proposition.

Within this framework, it is easy to see how analogical inference can employ multiple analogies, because more than one source can be used to create new constraints. Shelley (1999) describes how biologists do not simply use the centroscymnus analog as a source to infer that coelacanths inhabit deep water, but also use the following different source:

Source 2: ruvettus-2
(have (ruvettus rod-pigment-2) have-2)
(absorb (rod-pigment-2 474nm-light) absorb-2)
(penetrate (474nm-light deep-ocean-water) penetrate-2)
(see-in (ruvettus deep-ocean-water) see-in-2)
(inhabit (ruvettus deep-ocean-water) inhabit-2)
(enable (have-2 see-in-2) enable-2)
(because (absorb-2 penetrate-2) because-2)
(adapt (see-in-2 inhabit-2) adapt-2)

The overall inference is that coelacanths inhabit deep water because they are like the centroscysmus and the ruvettus sources in having rod pigments that adapt them to life in deep water. Notice that these are deep, systematic analogies, because the theory of natural selection suggests that the two source fishes have the rod pigments because the pigments adapt them to life in deep ocean water environments. When HOTCO maps the ruvettus source to the coelecanth target after mapping the centroscysmus source, it creates excitatory links from the inferred proposition INHABIT-NEW with both INHABIT-1 in the first source and INHABIT-2 in the second source. Hence activation can flow from both these propositions to INHABIT-NEW, so that the inference is supported by multiple analogies. If another analog or other information suggests a contradictory inference, then INHABIT-NEW will be both excited and inhibited. Thus multiple analogies can contribute to the defeasible and holistic character of analogical inference.

The new links created between the target proposition and the source proposition can also make possible emotional transfer. The coelacanth example is emotionally neutral for most people, but if an emotional valence were attached to INHABIT-1 and INHABIT-2, then the excitatory links between them and INHABIT-NEW would make possible spread of that valence as well as spread of activation representing acceptance. In complex analogies, in which multiple new excitatory links are created between aspects of one or more sources and the target, valences can spread over all the created links, contributing to the general emotional reaction to the target. The section after next provides detailed examples of this kind of emotional analogical inference.

Analogies about Emotions

The *Columbia Dictionary of Quotations* (available electronically as part of the Microsoft Bookshelf) contains many metaphors and analogies concerning love and other emotions. For example, love is compared to religion, a master, a pilgrimage, an angel/bird, gluttony, war, disease, drunkenness, insanity, market exchange, light, ghosts, and smoke. It is not surprising that writers discuss emotions nonliterally, because it is very difficult to describe the experience of emotions in words. In analogies about emotions, verbal sources help to illuminate the emotional target, which may be verbally described but which also has an elusive, nonverbal, phenomenological aspect. Analogies are also used about negative emotions: anger is like a volcano, jealousy is a green-eyed monster, and so on.

In order to handle the complexities of emotion, poets often resort to multiple analogies, as in the following examples:

(1) John Donne:
Love was as subtly catched, as a disease;
But being got it is a treasure sweet.

(2) Robert Burns:
O, my love is like a red, red rose,
That's newly sprung in June:
My love is like a melodie,
That's sweetly play'd in tune.

(3) William Shakespeare:
Love is a smoke made with the fume of sighs,
Being purged, a fire sparkling in lovers' eyes,
Being vexed, a sea nourished with lovers' tears.
What is it else? A madness most discreet,
A choking gall and a preserving sweet.

In each of these examples, the poet uses more than one analogy or metaphor to bring out different aspects of love. The use of multiple analogies is different from the scientific example described in the last section, in which the point of using two marine sources was to support the same conclusion about the depths inhabited by coelacanths. In these poetic examples, different source analogs bring out different aspects of the target emotion, love.

Analogies about emotions may be general, as in the above examples about love, or particular, used to describe the emotional state of an individual. For example, in the movie *Marvin's Room*, the character Lee played by Meryl Streep describes her reluctance to discuss her emotions by saying that her feelings are like fishhooks—you can't pick up just one. Just as it is hard to verbalize the general character of an emotion, it is often difficult to describe verbally one's own emotional state. Victims of posttraumatic stress disorder frequently use analogies and metaphors to describe their own situations (Meichenbaum 1994, pp. 112–113):

- I am a time bomb ticking, ready to explode.
- I feel like I am caught up in a tornado.
- I am a rabbit stuck in the glare of headlights who can't move.
- My life is like a rerun of a movie that won't stop.
- I feel like I'm in a cave and can't get out.
- Home is like a pressure cooker.
- I am a robot with no feelings.

In these particular emotional analogies, the target to be understood is the emotional state of an individual, and the verbal source describes roughly what the person feels like.

The purpose of analogies about emotions is often explanatory, describing the nature of a general emotion or a particular person's emotional state. But analogy can also be used to help deal with emotions, as in the following quote from Nathaniel Hawthorne: "Happiness is a butterfly, which, when pursued, is always just beyond your grasp, but which, if you will sit down quietly, may alight upon you." People are also given advice on how to deal with negative emotions, being told for example to "vent" their anger, or to "put a lid on it."

In principle, analogies about emotions could be simulated by the standard models such as ACME and SME, with a verbal representation of the source being used to generate inferences about the emotional target. However, even in some of the above examples, the point of the analogy is not just to transfer verbal information, but also to transfer an emotional attitude. When someone says "I feel like I am caught up in a tornado," he or she may be saying something like "My feelings are like the feelings you would have if you were caught in a tornado." To handle the transfer of emotions, we need to go beyond verbal analogy.

Analogies That Transfer Emotions

As already mentioned, not all analogies are verbal: some involve transfer of visual representations (Holyoak and Thagard 1995). In addition, analogies can involve transfer of emotions from a source to a target. There are at least three such kinds of emotional transfer, involved in persuasion, empathy, and reverse empathy. In persuasion, I may use an analogy to convince you to adopt an emotional attitude. In empathy, I try to understand your emotional reaction to a situation by transferring to you my emotional reaction to a similar situation. In reverse empathy, I try to get you to understand my emotion by comparing my situation and emotional response to it with situations and responses familiar to you.

The purpose of many persuasive analogies is to produce an emotional attitude, for example when an attempt is made to convince someone that abortion is abominable or that capital punishment is highly desirable. If I want to get someone to adopt positive emotions toward something, I can compare it to something else toward which he or she already has a positive attitude. Conversely, I can try to produce a negative attitude by comparison with something already viewed negatively. The structure of persuasive emotional analogies is:

You have an emotional appraisal of the source *S*.
The target *T* is like *S* in relevant respects.
So you should have a similar emotional appraisal of *T*.

Of course, the emotional appraisal could be represented verbally by terms such as "wonderful," "awful," and so on, but for persuasive purposes it is much more effective if the particular gut feeling that is attached to something can itself be transferred over to the target. For example, emotionally intense subjects such as the Holocaust or infanticide are commonly used to transfer negative emotions.

Blanchette and Dunbar (2001) thoroughly documented the use of persuasive analogies in a political context, the 1995 referendum in which the people of Quebec voted whether to separate from Canada. In three Montreal newspapers, they found a total of 234 different analogies, drawn from many diverse source domains: politics, sports, business, and so on. Many of these analogies were emotional: 66 were coded by Blanchette and Dunbar as emotionally negative, and 75 were judged to be emotionally positive. Thus more than half of the analogies used in the referendum had an identifiable emotional dimension. For example, the side opposed to Quebec separation said "It's like parents getting a divorce, and maybe the parent you don't like getting custody." Here the negative emotional connotation of divorce is transferred over to Quebec separation. In contrast, the yes side used positive emotional analogs for separation: "A win from the YES side would be like a magic wand for the economy."

HOTCO can naturally model the use of emotional persuasive analogies. The separation-divorce analogy can be represented as follows:

Source: divorce
(married (spouse-1 spouse-2) married-1)
(have (spouse-1 spouse-2 child) have-1)

Source: you	*Target: Hamlet*
(fire (boss, you) s1-fire)	(kill (uncle, father) t1-kill)
(lose (you, job) s2-lose)	(lose (Hamlet, father) t2-lose)
	(marry (uncle, mother) t2a-marry)
(cause (s1-fire, s2-lose) s3)	(cause (t1-kill, t2-lose) t3)
(angry (you) s4-angry)	(angry (Hamlet) t4-angry)
(depressed (you) s5-depressed)	(depressed (Hamlet) t5-depressed)
(cause (s2-lose, s4-angry) s6)	(cause (t2-lose, t4-angry) t6)
(cause (s2-lose, s5-depressed) s7)	(cause (t2-lose, t5-depressed) t7)
(indecisive (you) s8-indecisive)	
(cause (s5-depressed, s8-indecisive) s9)	

The purpose of this analogy is not simply to draw the obvious correspondences between the source and the target, but to transfer over your remembered image of depression to Hamlet.

Unlike persuasive analogies, whose main function is to transfer positive or negative valence, empathy requires transfer of the full range of emotional responses. Depending on his or her situation, I need to imagine someone being angry, fearful, disdainful, ecstatic, enraptured and so on. As currently implemented, HOTCO transfers only positive or negative valences associated with a proposition or object, but it can easily be expanded so that transfer involves an *emotional vector* which represents a pattern of activation of numerous units, each of whose activation represents different components of emotion. This expanded representation would also make possible the transfer of "mixed" emotions.

Empathy is only one kind of explanatory emotional analogy. We already saw examples of analogies whose function is to explain one's own emotional state to another, a kind of reverse empathy in that it enables others to have an empathic understanding of oneself. Here is the structure of reverse empathy:

I am in situation *T* (target).
When you were in a similar situation *S*, you felt emotion *E* (source).
So I am feeling an emotion similar to *E*.

Here is a final example of analogical transfer of emotion: "Psychologists would rather use each other's toothbrushes than each other's terminology." This analogy is complex, because at one level it is projecting the emotional reaction of disgust from use of toothbrushes to use of terminology, but it is also generating amusement. A similar dual role is also found in the following remark in *The Globe and Mail*: "Starbucks coffee shops are spreading through Toronto faster than head lice through a kindergarten class." Both these examples convey an attitude as does the remark of country music star Garth Brooks: "My job is like what people say about pizza and sex: When it's good, it's great; and even when it's bad, it's still pretty good." Note that this is a multiple analogy, for what its worth. The writer Flaubert also used an analogy to convey his attitude toward his work: "I love my work with a love that is frantic and perverted, as an ascetic loves the hairshirt that scratches his belly." Let us now consider analogies that go beyond analogical transfer of emotions and actually generate new emotions.

Analogies That Generate Emotions

A third class of emotional analogies involves ones that are not about emotions and do not transfer emotional states, but rather serve to generate new emotional states. There are at least four subclasses of emotion-generating analogies, involving humor, irony, discovery, and motivation.

One of the most enjoyable uses of analogy is to make people laugh, generating the emotional state of mirth or amusement. The University of Michigan recently ran an informational campaign to get people to guard their computer passwords more carefully. Posters warn students to treat their computer passwords like underwear: make them long and mysterious, don't leave them lying around, and change them often. The point of the analogy is not to persuade anyone based on the similarity between passwords and underwear, but rather to generate amusement that focuses attention on the problem of password security.

A major part of what makes an analogy funny is a surprising combination of congruity and incongruity. Passwords do not fit semantically with underwear, so it is surprising when a good relational fit is presented (change them often). Other emotions can also feed into making an analogy funny, for example when the analogy is directed against a person or group one dislikes:

Why do psychologists prefer lawyers to rats for their experiments?

1. There are now more lawyers than rats;
2. The psychologists found they were getting attached to the rats;
3. And there are some things that rats won't do.

This joke depends on a surprising analogical mapping between rats in psychological experiments and lawyers in their practices, and on negative emotions attached to lawyers. Further surprise comes from the addendum that psychologists have stopped using lawyers in their experiments because the results did not transfer to humans. Another humorous analogy is implicit in rate, rich, and suggestive that it has the emotional appeal of an excellent scientific theory or mathematical theorem. Holyoak and Thagard (1995, chap. 8) describe important scientific analogies such as the connection with Malthusian population growth that inspired Darwin's theory of natural selection. Thus scientific and other elegant analogies can generate positive emotions such as excitement and joy without being funny.

Not all analogies generate positive emotions, however. Ironies are sometimes based on analogy, and they are sometimes amusing, but they can also produce negative emotions such as despair:

> HONG KONG (January 11, 1998 AF-P)—Staff of Hong Kong's ailing Peregrine Investments Holdings will turn up for work Monday still in the dark over the fate of the firm and their jobs.... Other Peregrine staff members at the brokerage were quoted as saying Sunday they were pessimistic over the future of the firm, saddled with an estimated 400 million dollars in debts. "I'm going to see the Titanic movie... that will be quite ironic, another big thing going down," the South China Morning Post quoted one broker as saying.

Shelley (2001) argues that irony is a matter of "bicoherence," with two situations being perceived as both coherent and incoherent with each other. The Peregrine Investments–Titanic analogy is partly a matter of transferring the emotion of despair from the Titanic situation to the company, but the irony generates an additional emotion of depressing appropriateness.

The final category of emotion-generating analogies we want to discuss is motivational ones, in which an analogy generates positive emotions involved in inspiration and self-confidence. Lockwood and Kunda (1997) have described how people use role models as analogs to themselves, in order to suggest new possibilities for what they can accomplish. For example, an athletic African American boy might see Michael Jordan as someone who used his athletic ability to achieve great success. By analogically comparing himself to Michael Jordan, the boy can feel good about his chances to accomplish his athletic goals. Adopting a role model in part involves transferring emotions, e.g., transferring the positive valence of the role model's success to one's own anticipated success, but it also generates new emotions accompanying the drive and inspiration to pursue the course of action that the analogy suggests. The general structure of the analogical inference is:

My role model accomplished the goal *G* by doing the action *A*.
I am like my role model in various respects.
So maybe I can do *A* to accomplish *G*.

The inference that I may have the ability to do *A* can generate great excitement about the prospect of such an accomplishment.