

The Anatomy of Bias

How Neural Circuits Weigh the Options

Jan Lauwereyns

**The MIT Press
Cambridge, Massachusetts
London, England**

© 2010 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

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 Syntax and Times by SNP Best-set Typesetter Ltd., Hong Kong. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Lauwereyns, Jan, 1969–

The anatomy of bias : how neural circuits weigh the options/Jan Lauwereyns.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-262-12310-5 (hardcover : alk. paper) 1. Neural circuitry—Mathematical models. 2. Decision making—Physiological aspects. 3. Bayesian statistical decision theory. 4. Neural networks (Neurobiology)—Mathematical models. I. Title.

QP363.3.L38 2010

153.8'3—dc22

2009021061

10 9 8 7 6 5 4 3 2 1

Prelude: A Double Inheritance, a Single Question

From day one of my shadowy life in science there loomed large the inescapable figure of William James. Though my educational lineage as an experimental psychologist traces back most easily to Albert Michotte (of Leuven, Belgium) and then Wilhelm Wundt (of Leipzig, Germany), it was the hypnotic writing style of the American pioneer of psychology that first captured my attention, or started my imagination. It was not that I had read much of William James, or much of anything, really, but there were quotes floating around, little gems that kept surfacing time and again, in undergraduate courses and textbooks, in research articles, in papers presented at conferences, and even in texts far beyond safe academic waters. The quote that seemed most pervasive was this one, in which James (1950, pp. 403–404) claims everyone knows what attention is:

Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others, and is a condition which has a real opposite in the confused, dazed, scatterbrained state which in French is called *distraction*, and *Zerstreutheit* in German.

I cannot remember where or when I first ran into this curious impromptu definition. I vaguely recall having a bit of a pedantic dispute with one of my teachers, Tom Carr (then at Michigan State University), about the grammar, particularly in the first sentence, with the “objects or trains of thought.” Tom thought the mind took possession of either an actual “object” or else “a train of thought,” but I insisted that such would be an odd pair of alternatives, a categorical error, almost, moving from a physical thing to a mental representation. Instead I suggested that the choice had to be between “an object of thought” and “a train of thought,” with emphasis either on a static or a dynamic mental representation. (I’m not sure I explained my concern very clearly at

the time; in any case, my subtle linguistic criticism would not have sounded very convincing in my thick Dutch accent.) Tom, of course, soon regretted inviting me to sit in on his graduate course on cognitive processes.

The Beautiful but Ill-Defined Concept of Attention

Today the little passage from *The Principles of Psychology* still retains some of its mysterious beauty for me, even though (or perhaps even more so because) I have now recited it dozens of times in my own undergraduate and graduate courses on brain and cognition at Victoria University of Wellington, New Zealand. (Dozens of times, yes, but no more than once or twice in each course, to be sure; it may be my mantra, but I don't hum it in class.) I like pointing out how funny it is that there should follow something after the rather confident observation that everyone knows what attention is. If we all know what it is, then why do we keep talking about it? I then usually go on to unpack the quote and briefly explore the ramifications of each of the mentioned characteristics of attention, with the opponent mechanisms of facilitation versus inhibition of information processing, the similarity to the spotlight metaphor of attention (often attributed to Posner, 1980, though Posner did not actually use the word "spotlight" in that paper), the notion that attention can be seen as the gateway to consciousness, and so on.

James sounds so hypnotic thanks to the leisurely cadence (and slightly idiosyncratic punctuation) with which he advances through the dense brush of concepts and abstract ideas, indulging, for instance, in a triplet of adjectives to specify a state ("confused, dazed, scatterbrained") which he then proceeds to push into the unknown again by associating it with exotic words, borrowed from French ("*distraction*," which, italicized, apparently does not simply mean "distraction") and German ("*Zerstreuung*," which, roughly translated to an English neologism, would be "scatteredness" or, indeed, the scatterbrained state). The text is so lucid and lovable that it forgets to teach and becomes confusing, providing the ultimate image of what it was trying to say all along, namely, that we can know, and yet not really know, what attention is. This is precisely the point of departure for a science that wishes to address the core experience of being (or life) as we know it. William James called it "psychology," a term that no longer carries quite the same meaning as it did when he had its principles in mind. Psychology has become fractionated, and the various objects or trains of thought that James brought together in his *opus magnum* would in the present day find their offspring in the divergent fields of cognitive neuroscience, psychoanalysis of the post-Lacan variety, literary theory, philosophy of mind, evolutionary psychology, experimental economics, and several other lines of inquiry. Being a fan of James and the breadth

of his vision, I will recklessly endeavor to scavenge materials from these various fields with the single aim of producing a coherent but open-minded account of attention, or bias versus sensitivity, or how the activities of neurons allow us to decide one way or another that, with a faint echo of Hamlet in the background, something appears to be or not to be.

In the meantime, it may be useful to situate the inspiring quote in its original context. It is buried deep in the first volume of the *Principles*, in chapter XI on the topic of, indeed, attention. James started the chapter with the observation that “psychologists of the English empiricist school” (by which he refers to Hume and Locke among others) had failed to give the concept of selective attention its fair due in their investigations; they had too easily made experience “tantamount to the mere presence to the senses of an outward order” (1950, p. 402), as if attention were merely a slave to its objects. This could not be the entire story, James reasoned (still on p. 402):

Millions of items of the outward order are present to my senses which never properly enter into my experience. Why? Because they have no interest for me. My experience is what I agree to attend to. Only those items which I notice shape my mind—without selective interest, experience is an utter chaos.

Here, we are dangerously close to circular reasoning, paradoxical thinking, or a form of thought that would have no place in science, other than as a boundary condition, or the ultimate void of meaninglessness. My attention is supposed to go to items that hold some interest for me, but how can I select such items unless I already know what they are? What use would I have for a selection mechanism that operates only on things that are known to me already? James seems to suggest that it is primarily a matter of mental hygiene; selective interest is needed to shape my mind and to prevent utter chaos. It would be a cognitive economy of sorts, and thanks to the organizational power of my attention, and the space it frees up in the toy model of the world inside my head, I would perhaps be able to think deeper, weigh different options and courses of action, and make better decisions. If this is true, the study of attention should receive an urgent and central place in the human effort to understand all things human; if it is indeed the core factor that stands between experience and chaos, then all areas of human life, from the most mundane to the most esoteric, would stand to benefit from a detailed examination of its characteristics, its conditions and implications.

The Problematic Idea of a Book

For better or worse, beautiful and ill-defined as it was, the topic of attention seemed an excellent place for an eager mind to enter the world of science.

I was eager enough, and so my mind did just that (enter the world of science) exactly there (on the topic of attention), though it took me a while, and a string of great teachers, before I realized it myself. In part, the delay in my vocation was due to something I like to think of as a benign multiple personality disorder, which is probably not so very uncommon, given the number of highly talented scientists I have met who are engaged in some other intellectual passion, artistic or musical. Two of my favorite examples are the expressionist painter–engineer–computational neuroscientist Minoru Tsukada and the baritone–psychophysicist–physiologist R. H. S. (Roger) Carpenter. One not-so-very-hidden hidden agenda of this book will be to present an example of radical cross-disciplinary interests that actually might be good for each other. Given that I have two left hands and a tendency to spill liquids, and knowing that I couldn’t sing two notes to save my life, my own creative impulses had to be directed to the art of words, to poetry, prose, any form of writing that cared about how it said what it said, or even professed that saying something meant saying something in a certain way, and if you wanted to say it right, you had better consider how you were going to say what you wanted to say in the first place. Thus, I was condemned to literature. And in the literary arena I stumbled across William James again:

As we take . . . a general view of the wonderful stream of our consciousness, what strikes us first is [the] different pace of its parts. Like a bird’s life, it seems to be made of an alternation of flights and perchings. The rhythm of language expresses this, where every thought is expressed in a sentence, and every sentence closed by a period. The resting-places are usually occupied by sensorial imaginations of some sort, whose peculiarity is that they can be held before the mind for an indefinite time, and contemplated without changing; the places of flight are filled with thoughts of relations, static or dynamic, that for the most part obtain between the matters contemplated in the periods of comparative rest.

I encountered this wonderful, lyrical passage from *The Principles of Psychology* (James, 1950, p. 243) in a collection of essays, entitled *The Language of Inquiry*, by the American poet Lyn Hejinian (2000, p. 121). Hejinian’s own work is notorious—celebrated by some, including me, cast aside by others—for its relentless commitment to experimentation and literary invention; she was on the forefront of the so-called L=A=N=G=U=A=G=E movement in poetry, driven by a very diverse group of poets who emerged as leading new voices in California’s Bay Area in the 1970s. Their one common feature was a shared meta-analytic interest in the medium of their art. Of these poets, Hejinian is probably the poet with the most “scientific” attitude. In fact, the William James quote that Hejinian used in her book is very telling about her

prerogatives, demonstrating exactly what James said about attention and how we only take note of what holds an interest for us. For Hejinian, the sentence is the principle unit of language, or, stated in a formula,

POETRY = SENTENCE ART.

Hejinian's William James quote hit me like a boomerang. A few years ago I was sure I had had my fill of the intentionality of visual selective attention (the title of my PhD thesis). But if I had thought to escape the legacy of William James, the hard-core cognitive psychologists, the reductionist spotlight theorists of attention, then here was James with a vengeance, right in my beloved territory of experimental poetry. Hejinian's quote was on the mark in several ways, not just with respect to her own poetics and interests, or as an illustration of the near-paradoxical mechanism of attention, which gives us only what we had already, but also as a final vindication of my interpretation of the other quote, with the static versus dynamic modes of thought (the grammatical issue with the objects and trains). The new passage proved me right and Tom Carr wrong, I concluded, and so maybe I had a future as a William James kind of scholar after all.

Slowly, then, the idea of writing about my field of study started taking shape; perhaps to really approach the topic of attention and decision making in the way that best exploited my skills, feeding from my strengths and acknowledging my weaknesses, I would have to not only process and compile technical reports of scientific experiments in line with the culture of the profession but also seek to understand and be able to talk about the findings and observations in a language that would invite readers from outside the core circle of cognitive scientists to take part in the inquiry. Again, William James (1950, pp. 252–253), as (not entirely accurately) quoted by Lyn Hejinian (2000, p. 122), would be my guide:

The truth is that large tracts of human speech are nothing but signs of direction in thought.... These bare images of logical movement... are psychic transmissions, always on the wing, so to speak, and not to be glimpsed except in flight.

The original version by James had "psychic transitions," though I think I prefer Hejinian's "transmissions." But especially the last phrase invariably reverberates in my mind's ear whenever I reread the quote. The echo must needs be heard, "always on the wing, so to speak, and not to be glimpsed except in flight." I take it as an image of the extraordinary power of human speech, even when all that talking rather feebly seems to give "nothing but signs of direction in thought." These directions of thought are exactly the stuff that scientific progress is made on, and it is in the exercise of speech that we become aware

of the psychic transactions, once again, “always on the wing, so to speak, and not to be glimpsed except in flight.”

Writing about my field of study, then, might achieve something like this, giving a glimpse of the logical movements in flight, showing how the cognitive neuroscience of attention and decision making got to where it is at present, and looking ahead to where it might be going next. If James is right, and if my speech is to the point, this exercise—or the bird’s eye view, to continue the metaphor—would be useful even to the experts who are defining the research programs. Which trends are important? Which kinds of experiments are viable? Which approaches are likely to yield new information? I think we need to think about such things every once in a while, before rushing off to more data collection:

The role of the thinker is not so much to utter truths or establish facts, but to distinguish among the large population of true facts those that are important and relevant from those that are not.

The words are from Manuel DeLanda (2004, p. 7) in his characterization of Gilles Deleuze’s epistemology. They are appropriate in any context, I believe. DeLanda, by the way, is another fine example of a very productive cross-disciplinary spirit. He started out as an experimental filmmaker but found his way to philosophy through the work of the idiosyncratic French philosopher, who happened to share DeLanda’s passion for cinema. DeLanda proceeded to make it his mission to save the useful bits of Deleuze’s work for the science-oriented analytic philosophers who generally do not have much time for what is rather pejoratively termed “continental philosophy.” It turned out Deleuze had said useful things about dynamics, chaotic systems, and so forth for even the most mathematically inclined biologist with a phobia for full sentences.

DeLanda’s, and Deleuze’s, point is that we sometimes need to take one step back and survey the journey traveled so far to find the best direction to move forward. One experiment easily leads to another, but does it really tell us something new, does it actually move to a more precise level of description, or does it merely provide another paraphrase of some already-established correspondence between one thing and another?

At first glance, the current publication culture in science does not seem to be very conducive to the kind of speech act that I wish to engage in with the present book. Researchers are under extreme pressure to publish boatloads of papers as quickly as possible, preferably but not necessarily in top journals. To accommodate this problem, over the last ten or twenty years new journals have sprung up like mushrooms in a damp forest. The number of new online journals in the last two or three years alone is simply staggering. The principle

benefit is obvious; researchers can now easily archive massive amounts of data. The costs, however, are indirect, often ignored, and potentially disastrous for the entire scientific enterprise. The archived data are often of substandard quality, or even if they are good-quality data, the analyses have been rushed, the discussion is too summary, and the paper in general lacks sharpness and precision. Worse, the sheer volume of published reports makes it impossible for researchers to keep track of all relevant publications; there are vastly more papers than before, and virtually all of these papers pass to oblivion in a few years or even a few months. The net effect of publishing extensively is that you are expected to publish even more—funding agencies, administrators, and people on the sidelines expect more of the same. All of these factors combined create the clear and present danger that labor-intensive, long-term research programs become an economic liability for the individual researcher. Who would ever have the luxury of taking a half year off to write a book-length critical review? One monograph equals four or five technical reports that will forever remain unpublished.

Perhaps we simply need to stop worrying and do what we wish to do, or, as a fourteenth-century Buddhist priest in Kyoto suggested (Kenkō, 2001, p. 53), “A man who has determined to take the Great Step should leave unresolved all plans for disposing of urgent or worrisome business.” Luckily, science is hosted by institutes and organizations that do provide researchers with some protection against immediate pressures that run counter to “the Great Step” of sustained, rational thinking. Throughout my work on this book, I have had the good fortune of being kept out of the rain by four such umbrellas. I gratefully acknowledge them here: Victoria University of Wellington, New Zealand, my academic home base, which gave me the opportunity to take an eleven-month research and study leave from June 2008 to April 2009; the Brain Science Research Center of Tamagawa University, Tokyo, Japan, the host institute that took me in as a Visiting Scholar during that time; the Japanese Society for the Promotion of Science, which awarded me a very generous Long-Term Invitation Fellowship for Research in Japan (award number L-08507); and The MIT Press, a strong early supporter of the project.

I would strongly like to encourage other researchers to take advantage of similar opportunities to step out of the usual routines of daily life in the lab and reflect on the big picture of our fields of study. It is only through such efforts that we can hope to escape the pitfall of hastily publishing too many papers containing too little information. Let us take time and think things through. In the words of Lyn Hejinian’s wonderfully uncompromising poem “Happily” (2000, p. 385):

Constantly I write this happily
 Hazards that hope may break open my lips
 What I feel is taking place, a large context, long yielding, and to doubt it would be a
 crime against it
 I sense that in stating “this is happening”
 Waiting for us?

We are waiting for us, indeed, to make this take place, “a large context, long yielding,” providing “hazards that hope may break open my lips”—odd words that force some kind of shift: of perspective, if not of paradigm.

Of course, now the burden is on me to provide something substantial and new. How will I do that? Like Hejinian, I will use odd words in the hope of forcing some kind of shift, not of paradigm but of perspective. Though my topic stays very close to what William James thought we all knew, I will talk about “bias” and “weighing the options,” or concepts that are not exactly synonymous with the more common terms “attention” and “decision making.” Indeed, I will spend a good amount of time and effort pleading for a language that is as ergonomic and simple as possible to move from brain to behavior, and back again, in ways that allow researchers to be more precise than before about the correspondence between the things we do and the things that happen inside our head.

I may be a William James kind of scholar, but the problem with “attention” is that it is too vague—so there will be a desire for a touch of good old behaviorist thinking in the fabric of this book as well. The aim is to combine breadth of vision with conceptual minimalism. For this, “bias” will be my key concept.

The Architecture of My Great Step

What is bias exactly? Does it simply mean preference for one thing rather than another? What is the role of bias when we make judgments about objects and events, from simple, sensory categorization to complex, social assessment? How does bias emerge in neural circuits? My great step, the book I must now write, should give an integrative account of the structure and function of bias as a core brain mechanism that attaches different weights to various information sources, prioritizing some cognitive representations at the expense of others. The principal objective is to provide a critical analysis of this core brain mechanism in a single volume, offering an engaging and comprehensive narrative. The project is not to produce a textbook in the classical sense—it does not purport to replace, for instance, the authoritative 800-page reference work on *Memory, Attention, and Decision-Making* by Edmund T. Rolls (2008). Instead, my hubris will be to think I should seek to open up the research field to new or other modes of thought, believing that there may be ideas and

knowledge to be gained from explicit efforts to rethink the data, applying a different perspective—here, with emphasis on the pervasive concept of bias.

The book is primarily aimed at the kinds of minds that would be ready to begin a postgraduate program in psychology or neuroscience, or perhaps the third or fourth year in the more demanding undergraduate programs, but I hope it will also attract more advanced readers from the same and neighboring research fields. As an authored book, speaking in a single voice and providing an integrative account, it promises to be an attractive point of entry to an otherwise amorphous and vast literature—a literature which may be quite daunting for anyone who is willing to invest time in reading one “content-heavy” book but simply does not have the time or the means to survey the various expansive bodies of literature on the topic of bias in decision making.

The book endeavors to be innovative in two major ways. First, I identify “bias” (rather than more noncommittal terms such as “selection” or “attention”) as a core concept to study the neural underpinnings of action control and various forms of decision making. This allows researchers to apply existing computational tools and be relatively precise about how neurons contribute to the control of behavior. Second, I aim to provide explicit linking propositions with other contemporary analyses of decision making in an effort to stimulate mutual interest (and multidisciplinary discussion) among students and professionals from within and outside the area of cognitive neuroscience.

Chapter 1, “Bayes and Bias,” develops the concept of bias on the basis of notions from Bayesian probability, signal detection theory, and current neural models of decision making. “Bias” here refers to the unequal weight given to different sources of information—as a result of preferences, prejudices, or other forms of selective processing. The aim is to translate the computational properties of bias into easily recognizable neural signatures. These signatures are introduced by means of concrete examples from the experimental literature. The chapter also includes a brief overview of the types of experimental paradigms that provide relevant data—this stands as a short “discourse on method” (inviting a brief detour on epistemology).

Chapter 2, “Wish Come True,” focuses on how the reward factor is implemented in neural mechanisms for information processing. A distinction is made between anticipatory mechanisms that bias the neural system on the basis of reward probability and synergistic mechanisms that influence the quality and resolution of information processing (from perception to memory). This chapter also builds a bridge to the new field of “neuroeconomics” and perspectives from evolutionary psychology, as well as literary, philosophical, and psychoanalytic approaches to the concept of desire.

Chapter 3, “Fear Materialized,” explores the flip side of the coin—negative motivation as a determinant of the priority given to different sensory inputs.

Important questions include the extent of overlap with the reward-oriented mechanisms, the possibility of opponent dynamics, and the similarity of the underlying (anticipatory versus synergistic) algorithms. I carefully examine previous reports of “automatic attention” evoked by potent emotional stimuli and suggest a reinterpretation with the concept of bias. The mechanisms of fear conditioning turn out to stimulate neural plasticity and produce structural forms of bias. Again, the discussions afford comparison with notions from neuroeconomics and evolutionary psychology, as well as more psychodynamic explorations of fear and aversion.

In chapter 4, “The Familiarity Fallacy,” I examine the neural underpinnings of what is perhaps the most stereotypical image of bias in society—namely, racism, sexism, or any form of prejudice in which familiarity goes hand in hand with priority and beneficial treatment. Do such effects emerge inevitably from the structure of neural circuits, or can we divorce familiarity bias from prejudice and social inequality? If information propagation depends on learning and preestablished neural connections, then what are the behavioral implications? The ability to learn is an obvious prerequisite for flexible and well-adapted behavior, and this, in turn, raises the question of when the familiarity fallacy is really a *fallacy* (or where do benefits end and costs begin?).

Chapter 5, “The Proximity Trap,” departs from the classic observation, first offered in the work of the German Gestalt psychologists, that spatial and temporal properties of information, such as proximity, similarity, and continuation, can generate perceptual grouping. In this chapter I revisit these phenomena from the perspective of bias in neural processing, along with other “history” effects that have typically not been associated with the literature on perceptual grouping (e.g., repetition priming). A dominant theme here is the autonomous nature of the perceptual organization, giving rise to “pop-out” phenomena and exogenous attention capture. A comparison is made with the physics of attractors and self-organization in chaotic environments.

In chapter 6, “Less Is More,” I discuss the implications of the simple fact that individuals or organisms operate in a constrained “answer space”—they can only do so much. Bottleneck and other metaphors of selection have abounded in psychology and neuroscience alike, but I propose that the *reinforcing* aspect of selection has been underinvestigated. Horse race and other competition models are congruent with the notion that competition raises the stakes and thus could lead to an intensification of information processing, while at the same time the winner of the competition “takes all” (or will become the object of more extensive processing, such as being maintained in working memory or transferred to long-term memory). This might lead to a systems bias for bias, following a slogan such as “Less is more,” so that

selection and bias are applied even in situations in which it would in principle be possible to be noncommittal, weighing all information equally. Relevant discussions in this regard are those on the human mind's penchant for simplicity and its potentially damaging implications.

Finally, in chapter 7, "Utopia—A World without Bias," I conduct a thought experiment, trying to imagine what a world without bias would be like. From the wealth of evidence presented in the previous chapters it is quite clear that bias in neural processing is so pervasive as not to be eradicated by any means, nor would we necessarily benefit from the absence of bias (chapters 2 and 3 clearly connect bias with inclusive fitness and evolutionary success). Instead, we must focus on how to *control* bias—knowing when it can be used advantageously but also being alert to conditions that require us to work against bias. The opposition to bias, and the self-control that it implies, may be a crucial component in adaptive behavior of a "higher order," in terms of delay of gratification, long-term investment, and sacrifice for group or social benefit.

Before and After

If all goes well and my great step delivers on its promises, the book that has now been written will die a slow death in the best scientific tradition, being superseded by the work of a student or, better yet, the works (plural!) of students (plural!). The scientific investigation of bias in decision making is still in its infancy and will continue to benefit greatly from the immense rate of technological development. Only a half century ago, the British psychologist Donald Broadbent (1961, p. 302) wrote the following:

It seems likely that no conceivable observation can overcome the uncertainty about the position and velocity of an electron, since determining one disturbs the other. Events within the skull are not inherently unobservable: they are merely inaccessible to our present techniques. The proper analogy for the most speculative statements in psychology is not with, say, a proposition about velocity through the ether, where there is some doubt whether the proposition is meaningful. It is rather with propositions about mountains on the other side of the moon, which are perfectly meaningful but rather difficult to test. The mountains are as difficult to observe as the velocity, but the reasons for the difficulty are purely technological.

The statement is very much true today. In defiance of skeptics who have time and again declared the death of science in general, or the inherent inaccessibility of the relation between brain and mind in particular, the enterprise of cognitive neuroscience has all but made the most arcane phenomena of human consciousness tractable in the laboratory. The present book is a part of that enterprise and invites others to join in.

Of course, I did not just randomly pick a quote from Broadbent to broadcast the invitation. Donald Broadbent was one of a handful of scientists who helped shape the literature on selective attention during the 1950s and 1960s in cognitive psychology; his well-known filter model was, in part, inspired by the work of William James. Broadbent was the most frequently cited scientist in my own master's thesis, and I have a vivid memory of a talk he gave, less than a year before he passed away, at the XXVth International Congress of Psychology in Brussels in the summer of 1992 (when I was recruited as one of the little people in yellow T-shirts, the equivalent of ballkids at Wimbledon).

In one way, Donald Broadbent brought me to William James. In another way, Lyn Hejinian, via the great modernist poet Gertrude Stein (sometimes simplistically but not entirely ineffectively designated “the Picasso of poetry”), brought me to William James. My double inheritance revolves around the one question that has occupied much of my thinking in any domain: How do we choose what we see? Or what we hear, or notice, or perceive, or... Do we have a choice in the matter, or has the choice been made for us by some kind of inflexible cognitive apparatus? The one question soon breaks up in many slightly different versions in the Garden of Forking Questions, to corrupt the title of a short story by Jorge Luis Borges. *The Anatomy of Bias* presents my account of seventeen years' gardening.

Many people have helped me, encouraged me, mentored me, or challenged me. I would like to thank my colleagues, friends, teachers, and students, Arnoud van Adrichem, Hugo Bousset, Marc Brysbaert, Joseph Bulbulia, Roger Carpenter, Tom Carr, Brian Coe, Doug Davidson, Peter De Graef, Dirk Derom, Arkadii Dragomoshchenko, Géry d'Ydewalle, Fernanda Ferreira, Wim Fias, Hiroki Fujiwara, Dave Harper, John Henderson, Andrew Hollingworth, Yoshinori Ide, Shoichiro Iwakiri, Todd Jones, Reiko Kawagoe, Shunsuke Kobayashi, Masashi Koizumi, Jan Lammertyn, Bill Manhire, Kae Nakamura, Kiwao Nomura, Greg O'Brien, Michael Palmer, Masamichi Sakagami, Sue Schenk, Gerd Segers, Muneyoshi Takahashi, Yoriko Takikawa, Jan Theeuwes, Minoru Tsukada, Ken-Ichiro Tsutsui, Johan Velter, Leo Vroman, Johan Wagemans, Katsumi Watanabe, Matt Weaver, Jeff Wickens, Carolyn Wilshire, Regan Wisnewski, and Bob Wurtz.

Most of all, with this book I would like to pay tribute to the one teacher who taught me more than any other, and it will be obvious that that must be Okihide Hikosaka.

The final words of gratitude, from a helpless heart and a human brain, belong to the mysterious world of the unsayable, yet they wish to be said, and I wish to say them here, for Shizuka, Nanami, and Shinsei.