EXPLAINING THE EVIDENCE

How do we make sense of complex evidence? What are the cognitive principles that allow detectives to solve crimes and lay people to puzzle out everyday problems? To address these questions, David Lagnado presents a novel perspective on human reasoning. At heart, we are causal thinkers driven to explain the myriad ways in which people behave and interact. We build mental models of the world, enabling us to infer patterns of cause and effect, linking words to deeds, actions to effects, and crimes to evidence. But building models is not enough; we need to evaluate these models against evidence, and we often struggle with this task. We have a knack for explaining, but less skill at evaluating. Fortunately, we can improve our reasoning by reflecting on inferential practices and using formal tools. This book presents a system of rational inference that helps us evaluate our models and make sounder judgments.

DAVID LAGNADO is Professor of Cognitive and Decision Sciences in the Department of Experimental Psychology, University College London (UCL). He has written more than 100 articles and co-authored a textbook on the psychology of decision-making. He has worked with US intelligence, the UK government and various legal and financial institutions, looking at methods to improve reasoning and decision-making.

EXPLAINING THE EVIDENCE

How the Mind Investigates the World

DAVID A. LAGNADO

University College London

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781107006003
DOI: 10.1017/9780511794520

First published 2022

A catalogue record for this publication is available from the British Library.

Library of Congress Cataloging-in-Publication Data
NAMES: Lagnado, David A., 1962– author.
TITLE: Explaining the evidence : how the mind investigates the world / David A. Lagnado.
DESCRIPTION: Cambridge, United Kingdom ; New York, NY : Cambridge University Press, 2022. | Includes bibliographical references.
IDENTIFIERS: LCCN 2021019468 (print) | LCCN 2021019469 (ebook) | ISBN 9781107006003 (hardback) | ISBN 9780521184816 (paperback) | ISBN 9780511794520 (epub)
SUBJECTS: LCSH: Evidence–Psychological aspects. | Reasoning (Psychology) | Judgment. | Cognition. | BISAC: PSYCHOLOGY / Cognitive Psychology & Cognition | PSYCHOLOGY / Cognitive Psychology & Cognition
CLASSIFICATION: LCC BF761 .L34 2022 (print) | LCC BF761 (ebook) | DDC 153.4/3–DC23
LC record available at https://lccn.loc.gov/2021019468
LC ebook record available at https://lccn.loc.gov/2021019469

ISBN 978-1-107-00600-3 Hardback
ISBN 978-0-521-18481-6 Paperback

For Tracy

Contents

Figures

Tables

Preface

> *The apartment was in the wildest disorder – the furniture broken and thrown about in all directions . . . On a chair lay a razor, besmeared with blood. On the hearth were two or three long and thick tresses of grey human hair, also dabbled in blood, and seeming to have been pulled out by the roots. Upon the floor were found . . . an earring of topaz, three large silver spoons . . . and two bags containing nearly four thousand francs in gold . . .*
>
> *Of Madame L'Espanaye no traces were here seen; but an unusual quantity of soot being observed in the fire-place, a search was made in the chimney, and (horrible to relate!) the corpse of the daughter, head downward, was dragged therefrom; it having been forced up the narrow aperture for a considerable distance.*
>
> 'The Murders in the Rue Morgue', Edgar Allan Poe (1841)

Who can resist a murder mystery? Whether it's a classic detective novel or a psychological thriller, we rise to the challenge of the whodunit, revelling in the twists and turns, the false leads and the surprise denouement. And despite their tragedy, we are fascinated by real-life crime stories, when evil characters do bad things to good people. In this book I will explore another mystery: how do we – with our limited-capacity minds – manage to tackle such puzzles? This mystery is remarkable but often passes unnoticed. We take our inferential skills for granted. But, if you stop to think, it's amazing that we can formulate – let alone solve – such complex problems. How do we generate hypotheses, weigh up evidence and draw sensible conclusions? How do we leap from a smattering of clues – a bloody razor, uprooted hair, a body shoved up a chimney – to conjure possible causes of a crime scene we have never witnessed before? Even the ability to come up with the *wrong* answer is a feat worthy of awe.

To see how the mind works its magic, we must look beneath the surface, at the cognitive machinery that impels us. Firstly, we are causal thinkers, driven to explain the myriad ways in which people (and things)

behave and interact. Who killed the daughter? Why was her body rammed up the chimney? Why was no money taken? To answer such questions we build *mental* models – crude models 'in the mind' that we can probe and manipulate, allowing us to explain what we have observed, to predict what happens next and to imagine how events might have been different. These models enable us to reason about intricate patterns of cause and effect, linking words to deeds, actions to effects and crimes to evidence. They allow us to dig deeper into human behaviour, helping us to understand rather than just describe what we witness.

But building models is not enough, especially if we are trying to solve a crime or find someone guilty. We need to support our models with evidence, and this adds another layer of complexity. We must evaluate how well the evidence fits our favoured account as compared to other possibilities and also assess the quality of the evidence itself. Who was the source of the blood? How reliable was the blood test? Can we trust what the witnesses say? How do we factor these concerns into our overall judgment? To make things even harder, those closely involved in a crime often have motive to deceive us. Is the accused giving an honest account or lying through his teeth?

Despite being adept at creating mental models to explain events, we find evaluating these models much harder. We can readily devise an elaborate story of how a suspect might have committed the crime and yet struggle to produce convincing evidence for this version of events. We have a knack for *explaining*, but less skill at *evaluating*. And sometimes a compelling narrative can outstrip the evidence, with disastrous consequences.

Fortunately, all is not lost. We can improve our reasoning, learning to evaluate our hypotheses with consistency and balance. And there are tools to support us. I will present a system of rational inference that can help us build better models and make sounder judgments.

Structure of the Book

I use crime investigation as a lens through which to explore the principles of evidential reasoning. My premise is that we use the same forms of reasoning, whether we are solving crimes, diagnosing diseases or dealing with everyday concerns. But I focus on legal cases because they provide engaging examples – where bad reasoning can have serious costs.

Each chapter starts with a different crime case, showing how investigators, legal experts or jurors have grappled with difficult and often contentious questions. The cases highlight important principles of reasoning

which are not restricted to legal settings but apply to any sphere where people must use evidence to draw conclusions.

Chapter 1 traces a crime from the discovery of a body to the verdict of a court, exposing the challenges faced by decision-makers at each step of the process. How do we generate initial hypotheses from sparse information? How do we develop our hypotheses as we gather new evidence? How do we make sense of a large body of conflicting evidence to reach a final decision?

The next two chapters make the case for the critical role of causal models in reasoning. Chapter 2 focuses on the models we manipulate in our heads and on our impressive capacity to reason about new situations (even the terrible scene at Madame L'Espanaye's apartment!). But how do we know when our causal reasoning is good or bad? Chapter 3 introduces the causal modelling framework, showing us how to draw rational inferences from evidence. In Chapters 4 and 5 we look at how people reason and make judgments under uncertainty, with Chapter 4 focusing on people tackling small-scale problems in the lab and Chapter 5 switching to the expert reasoning of detectives solving murder cases.

Our focus then shifts to evidence evaluation. Chapter 6 looks at the legal concept of evidence, while Chapter 7 explores the subtleties that arise when people assess competing causes of the same evidence. Chapter 8 examines the strategies that people use to gather and interpret evidence, focusing on classic confirmation bias.[1] Chapter 9 then looks at how people construct stories to make sense of complex evidence, using the well-known Sally Clark case.

In Chapter 10, I present a Bayes net approach to modelling legal cases and suggest that this method could be used as a training tool to improve inferential reasoning in the legal domain. Note that this is a more technical chapter than the others and can be skipped by those not wanting the nitty-gritty details. But I do recommend reading the Agatha Christie story and looking at the proposed Bayes net.

Chapter 11 summarizes the key messages in the book by applying them to our thinking during the coronavirus crisis. I show how many of the concepts – from psychology and formal modelling – recur in expert and lay reasoning during the pandemic. I argue that the danger of jumping to

[1] Confirmation bias occurs when we gather and interpret evidence according to our prior expectations or preferences. Sometimes this is a reasonable strategy but in other contexts it can lead to biases. Again the murders at Rue Morgue provide a neat example. The police collect and evaluate evidence under the (very reasonable) assumption that the murderer is human. But in this peculiar case this assumption is misleading.

premature causal claims, without careful evidence evaluation, is particularly rife in situations of radical uncertainty. But causal thinking is also an essential tool for dealing with such crises.

Who Is This Book for?

This book is for anyone interested in evidential reasoning – how people reason about evidence and how this reasoning can be improved using formal models. Several disciplines are straddled in the book, so I hope it appeals to students from psychology, crime science, law, computer science and beyond. Ideally, I'd like to encourage the crossing of discipline boundaries: for example, introducing psychologists to legal decision-making, law students to causal modelling, and computer scientists to modelling of legal cases. More ambitiously, I hope the book might be of interest to expert decision-makers outside academia, such as crime investigators, forensic scientists and criminal lawyers.

The book is aimed at students and researchers, and I introduce all necessary technical concepts as the book progresses. It only presupposes some basic probability theory. I actually think that an introduction to causal modelling is a nice route into probabilistic thinking because (as I claim in the book) we are better attuned to causal rather than statistical reasoning. To accompany the book the website www.explainingtheevidence.com provides additional learning materials, including an introduction to probability theory and models of all the Bayes net examples to download and explore.

What This Book Does Not Cover

Human reasoning admits of many levels of explanation. In this book I focus on the cognitive level, how individual minds represent and reason about evidence, with little discussion of the broader social context of reasoning, and nothing on the lower-level neuroscience. Partly this is due to lack of space, but I also feel that the cognitive level is the natural starting point, indispensable for seeing how all the other levels fit together. Neither do I cover the broad literature on the non-cognitive factors – such as emotions and prejudice – that undoubtedly influence legal decision-making. Again, I believe that understanding the cognitive perspective is a precursor to seeing how these other factors might distort and bias our reasoning.

Two Senses of Why

A central theme in the book is that we are better at explaining than evaluating. This can be a slippery distinction to grasp because the two activities are so closely linked – we use evidence to build and revise our models, and we use our models to look for evidence. One way to think about the distinction is to see these activities as responses to two different types of why-question. Carl Hempel (1965) distinguishes between *explanation-seeking* and *reason-seeking* why-questions. The former question asks us why things in the world happen as they do, whereas the latter asks for the justifying grounds to support our claims. For example, when a husband is accused of killing his wife, we can ask why he would do it – what would be his motive for such a crime? To answer this question we point to his violent nature and the fact that his wife wanted to leave him. But we can also ask why we should believe that he did it – how strong is the evidence? Here we point to forensics, CCTV footage or witness testimony.

The two activities are closely interwoven, especially if we don't know the truth of the things we are trying to explain. We use evidence to generate our explanations, and our explanations in turn guide our search for evidence. But they answer distinct questions – *explaining* is about what happens (or happened) in the world, whereas *evaluating* is about our reasons to believe these claims. And our cognition seems more geared towards the former than the latter – or so I shall argue in this book.

Acknowledgements

The seeds of this book were planted a long time ago, when I worked on the project 'Evidence, Inference and Enquiry' under the expert guidance of Philip Dawid, David Schum and William Twining. Around the same time Judea Pearl's book on causality was stirring up cognitive science, and I was fortunate to work with Steven Sloman, applying these new ideas to human reasoning. My next big step came when I started collaborating with Norman Fenton and Martin Neil, who taught me the art and science of modelling legal cases. This led to a six-month stint at the Newton Institute in Cambridge, where I benefitted from discussions with many people, including Richard Gill, Bill Thompson, Leila Schnepps, Julia Mortera, Anne Ruth Mackor, Christian Dahlman, Marjan Sjerps, Charles Berger, Jacob de Zoete, Floris Bex, Henry Prakken, David Bentley, Pat Wiltshire, David Balding and Paul Roberts.

Another key step was working with the Bayesian Argumentation via Delphi (BARD) team on an IARPA-funded project on collaborative Bayesian reasoning; I learned a huge amount from Ulrike Hahn, Kevin Korb, Ann Nicholson, Erik Nyberg, Toby Pilditch, Stephen Dewitt, Marko Tesic, Nicole Cruz, Kirsty Philips, Saoirse Connor Desai and Alice Liefgreen. We even had some laughs along the way.

My colleagues at UCL have (often wittingly) been sounding boards for many of the ideas in this book; in particular, I would like to thank Christos Bechlivanidis, Adam Harris, Nigel Harvey, Ruth Morgan, David Shanks, Marteen Speenkenbrink, Cheryl Thomas, and all members of the Causal Cognition lab, including special guests Marius Usher, Samantha Kleinberg, Jonathan Livengood and Katya Tentori. Special thanks also to my PhD students, past and present, who have helped shape my thinking and my tea-drinking: Neil Bramley, Tobias Gerstenberg, Alice Liefgreen, Milena Nikolic, Nadine Smit, Erica Yu, Paul Troop, Lara Kirfel, Tamara Shengelia, Claire Machan, Matija Franklin and Cristina Leone.

Many colleagues went beyond the call of duty and gave me feedback on the book itself. I'm extremely grateful to Neil Bramley, Sam Dupret, Toby Pilditch, Nichola Raihani, Tracy Ray and William Twining for commenting on various chapters, and to Christos Bechlivanidis, Nancy Cartwright, Philip Dawid, Norman Fenton, Martin Neil, Magda Osman and Marion Vorms for enduring the whole book! I shudder to think how things would have turned out without their help and advice. Thanks also to Cambridge University Press – Janka Romero, Emily Watton and Hetty Marx – for their help and infinite patience, to Cara Chamberlain for precision copy-editing, Melanie Gee for forensic indexing and Vinithan Sethumadhavan for putting it all together.

Many thanks to my family and friends for encouragement and moral support, including my nephew Isaac (aged 11¾) who read the first chapter and said it was very boring.

Finally, this book is dedicated to my wife Tracy, for making it all worthwhile.