Cambridge University Press & Assessment
978-1-108-42164-5 — Social Inquiry and Bayesian Inference
Tasha Fairfield , Andrew E. Charman
Excerpt
More Information

Part I

Foundations

Cambridge University Press & Assessment
978-1-108-42164-5 — Social Inquiry and Bayesian Inference
Tasha Fairfield , Andrew E. Charman
Excerpt
More Information

1 Introduction

Bayesian Reasoning for Qualitative Research

The way we intuitively approach qualitative case research is similar to how we read detective novels or murder mysteries. We want to explain what happened – who killed Samuel Ratchett on the Orient Express, or how democracy emerged in South Africa. We consider different hypotheses along the way, drawing on our own ingenuity and the literature we have read – whether other Agatha Christie mysteries, or theories of regime change. As we gather evidence and discover new clues, we revise our assessment about which hypothesis provides the best explanation. Bayesianism provides a natural framework to govern how we should adjust our degree of belief in the truth of a hypothesis – for example, *A lone gangster slipped onboard the train and killed Samuel Ratchett as revenge for being swindled*, or *Mobilization from below drove democratization in South Africa by altering economic elites' regime preferences* (Wood 2001), given our previous knowledge and the new information that we learn during our research.

To illustrate how Bayesian updating works intuitively, consider the following more extended example. During the Latin American debt crisis in the 1980s, an intriguing phenomenon arose in which new presidents who had explicitly promised voters that they would *not* implement austerity measures nevertheless imposed harsh neoliberal reforms after taking office. Stokes (2001) entertains two alternative explanations for why these presidents violated their policy mandates. The first hypothesis proposes that presidents sought to represent the people's best interests and realized that austerity was the only way to fix the economy, despite voters' trepidations. The second hypothesis proposes instead that presidents were motived by opportunities for their own private financial gain, for example, kickbacks and bribes from business. Now consider the case of Venezuela's President Pérez, one of Latin America's policy-mandate violators – why did he opt to impose "neoliberalism by surprise"? Given whatever relevant information we know, we might strongly favor one of Stokes' (2001) hypotheses over the other, we might weakly favor one of them, or we might simply be indifferent.

In Bayesian terminology, our initial view about the plausibility of a hypothesis is called the *prior probability*.

Now suppose we learn the following information. According to multiple sources, Pérez learned how bad Venezuela's economic situation was only after taking office. And advisors reported that Pérez was paying close attention to developments in Peru, where he saw that the heterodox economic stabilization policies that his friend President García had implemented were not working (Stokes 2001:68–69). In light of this new information,[1] we might change our view about which hypothesis – representation, or rent-seeking – is more plausible. The Bayesian language for this revised view, which takes into account both our previous knowledge and the new evidence, is the *posterior probability*.

Figure 1.1 displays the results of this intuitive Bayesian updating exercise as conducted with roughly 80 participants at the Syracuse Institute for Qualitative and Multi-Method Research. Not surprisingly, given that the students came from a wide range of backgrounds and subfields, prior views were distributed across the spectrum, from strongly favoring the representation hypothesis to strongly favoring the rent-seeking hypothesis. But after considering the Venezuelan case evidence, beliefs tended to shift in favor of the representation hypothesis, with most participants converging on a weak preference for that explanation over the rent-seeking alternative.

This book aims to improve the way we evaluate the import of qualitative evidence and conduct case study research by drawing on insights from logical Bayesianism, an inferential framework originally developed for the natural sciences. Logical Bayesianism conceptualizes probability as the rational degree of belief that we should hold in the truth of a proposition – for example, a causal hypothesis – given the information we possess, which is inevitably limited. Bayesian probability, and Bayes' rule in particular, provides a rigorous framework for reasoning under conditions of uncertainty and incomplete information. In principle, Bayesian probability provides a unified approach for all scientific inquiry. From this perspective, we reexamine central debates on the logic of inference, research design, and analytic transparency in qualitative research.

Bayesianism is enjoying a revival across many fields, from astronomy and data science to macroeconomics and political polling, and it has much to offer qualitative social science. First and foremost, Bayesianism provides a solid methodological foundation for causal inference in case studies and

[1] We assume for the sake of illustration that readers are not already familiar with these particulars of the Venezuelan case. If these details are already known, they should inform the prior probability of the hypotheses, and coming across this information again would not change our initial views.

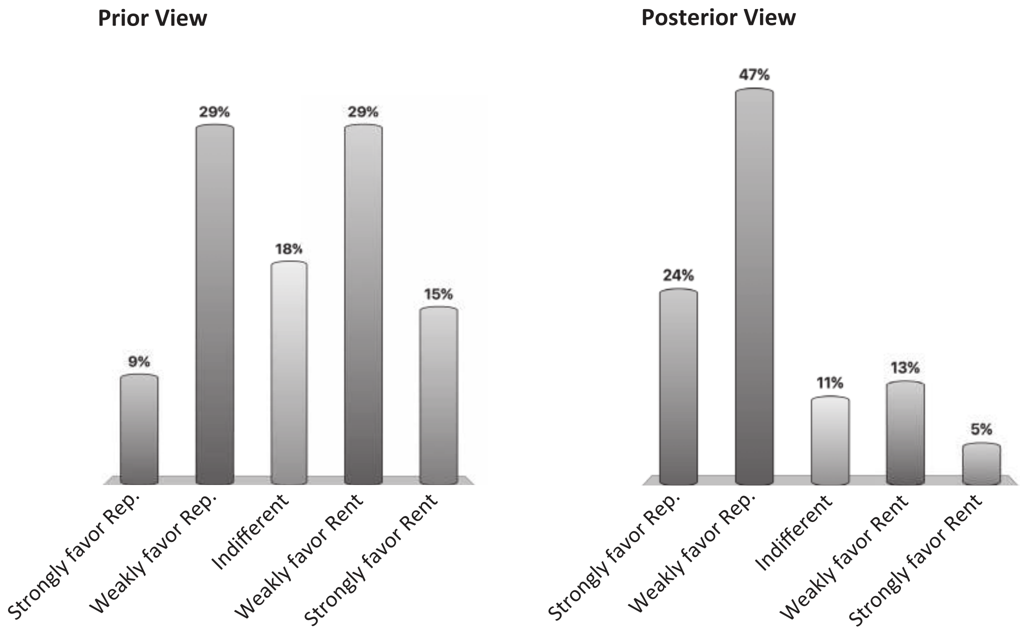

Figure 1.1 **Intuitive Bayesian updating.** Representation vs. rent-seeking hypotheses, drawing on Stokes 2001.

qualitative research more broadly. In fact, we will argue that Bayesianism is the only sound approach for causal analysis with qualitative, nonstochastic data. Second, learning the basic principles of Bayesian reasoning can help leverage and improve intuition. Bayesianism mirrors the way we naturally approach inference and implicitly underpins much of our common sense, but it also helps us avoid cognitive biases that can lead to sloppy reasoning. Even if scholars choose not to explicitly apply some of the more formal Bayesian techniques that we will introduce, learning and practicing these techniques can nevertheless help improve inferential judgements when conducting traditional, narrative-form case study analysis.

Third, Bayesianism facilitates consensus building and promotes knowledge accumulation, by providing a clear framework for scrutinizing inferences and pinpointing sources of disagreement. Bayesianism directs us to ask if disagreements stem from different initial views regarding which explanations are most plausible, given that everyone brings very different background knowledge to the discussion, and/or if scholars are assessing the inferential weight of the evidence differently, why their assessments diverge, and whether one line of reasoning about the evidence is more justifiable and compelling than another. Moreover, Bayesianism forces us to think very carefully about our hypotheses, making us aware of shortcoming in how they have been specified and pushing us to refine and articulate theories more precisely – which can in itself make an important contribution to resolving debates and promoting knowledge accumulation.

Finally, following Bayesian principles more explicitly enhances research transparency, which is a growing concern in light of the “replication crisis” in social science. We contend that the overarching concern in all scientific inquiry should be *reliability* of inference, which encompasses but extends beyond the notion of replication. Assessing reliability entails asking how much confidence we can justifiably hold in our conclusions. Bayesian probability is ideally suited to this task, because it provides the natural language for evaluating uncertainty.

1.1 Placing Our Approach in Perspective

Bayesian probability holds the potential to clarify and remap methodological debates within political science and to fully substantiate the importance of qualitative research vis-à-vis large-N statistical analysis and other quantitative research traditions. Despite recent innovations in process tracing (e.g., Bennett 2015), applying Bayesian probability in qualitative research remains a frontier that has not been definitively addressed. Qualitative methodologists have not yet recognized the full potential and ramifications of Bayesian probability; in fact, many misconceptions persist in efforts to apply Bayesianism in case study research. Nor have quantitative social scientists who view Bayesian analysis as a powerful technical tool (e.g., Iversen 1984, Jackman 2009, Gelman et al. 2013) fully realized the broader implications of Bayesian probability for inference and research design, although we build here on pioneering works by Western and Jackman (1994) and Western (2001) that set out some key foundational principles. Humphreys and Jacobs (2015) break new ground by exploring implications for research designs that combine within-case clues and cross-case datasets within a Bayesian framework, yet the logical Bayesian perspective we espouse leads to distinct recommendations regarding how best to operationalize test strength, select cases, and iterate between theory building and theory testing.

Accordingly, we seek to contribute to three areas of methodology: process tracing, qualitative methods more broadly, and multi-method research. While Bayesian justifies many common-sense practices, our perspective also introduces ideas that go against some established ways of thinking in many realms of social science. However, in the spirit of *Rethinking Social Inquiry* (Brady and Collier 2010), we hope to broaden the scope of debate on how to make inference more rigorous and to help bridge the gap between qualitative and quantitative methods in a way that more firmly establishes the value of qualitative research.

1.1.1 Process Tracing

A growing movement within political science has identified Bayesian probability as the methodological foundation of process tracing, which Bennett and Checkel (2015:4) define broadly as "the use of evidence from within a case to make inferences about causal explanations of that case." As part of an initiative to establish process tracing as a rigorous method, the literature has moved from informal analogies to Bayesianism (McKeown 1999, Bennett 2008, Beach and Pedersen 2013) toward efforts to more formally apply Bayesian analysis in qualitative research (Rohlfing 2012, Bennett 2015, Humphreys and Jacobs 2015, Fairfield and Charman 2017). But whereas Bayesian statistical techniques have been successfully elaborated for large-*N* quantitative research, efforts to apply Bayesian probability in qualitative case research are still under development. Most efforts to formalize Bayesian process tracing have examined only a few illustrative pieces of evidence (Rohlfing 2013, Bennett 2015) and/or have included only highly simplified process-tracing clues (Humphreys and Jacobs 2015). Moreover, we have found that a number of important conceptual and technical points have been overlooked, incorrectly handled, or inadequately addressed within the literature that is innovating in this terrain.[2]

We aim to advance the process-tracing literature by providing a careful exposition of the foundations of Bayesian "probability as extended logic" from the physical sciences (Jeffreys 1939, Cox 1961, Jaynes 2003, Gregory 2005), which contrasts with the more "subjective" treatments of Bayesianism in most philosophy of science and medical testing literature that inform much of the existing work on Bayesian process tracing.[3] We elaborate concrete guidelines to help scholars avoid potential pitfalls when endeavoring to apply Bayesian reasoning in process tracing, and we illustrate how to proceed when working

[2] Problems of this sort are also pervasive in qualitative methods literature that invokes Bayesianism more informally (e.g., Beach and Pedersen 2016).

[3] See Section 1.3.1.

with the kinds of complex, diverse, and nuanced real-world evidence that scholars gather during in-depth fieldwork and archival research.

1.1.2 Qualitative Methods

From a broader qualitative methods perspective, Bayesianism yields two major payoffs. First, it provides the rigorous foundation that is needed to definitively explicate and legitimate qualitative research. Earlier efforts to understand qualitative case studies and their relation to large-N quantitative research from a frequentist statistical perspective (e.g., King, Keohane, and Verba 1994) inevitably attributed a subordinate role to the former; there simply is no principled rationalization for small-N qualitative research within frequentism. If applied in strict accordance with its foundational tenets, frequentist techniques can only be used to analyze stochastic data, and large, independent samples are often considered critical for accurate inference. In contrast, Bayesianism narrows the divide between qualitative and quantitative research, because all inference in principle proceeds in the same manner, by applying Bayes' theorem and the other rules of probability theory. Furthermore, in the words of Pierre Simon Laplace (1814), Bayesian probability is essentially "common sense reduced to calculation." As such, it validates many intuitively sensible practices that have long characterized qualitative research but are often discouraged by frequentist-oriented disciplinary norms, including nonrandom case selection and iteration between theory development and data analysis.

Second and relatedly, Bayesianism provides a simple, intuitive alternative to the wide range of inferential approaches that are advocated and debated within the qualitative methods literature. On the one hand, our framework makes it unnecessary to distinguish between approaches such as pattern matching (Campbell 1975), congruence analysis (George and Bennett 2005), process tracing (George and McKeown 1985), causal narrative (Sewell 1996), or the comparative sequential method (Falleti and Mahoney 2015). While these represent important initiatives to better understand the logic of qualitative research, a Bayesian perspective effectively subsumes them all, by revealing that inference always entails reasoning from evidence – whether within-case, cross-case, or a combination of both – to identify the best available explanation from a concrete set of alternatives. We emphasize that whereas Bayesianism in qualitative social science is strongly associated with within-case analysis, the same logic applies to comparative case studies that aim to assess theories with scope conditions that extend beyond a single case. On the other hand, our Bayesian framework elucidates fundamental problems with alternative approaches including qualitative comparative analysis (QCA) and fuzzy-set analysis (e.g., Ragin 1987, 2000). When hypotheses depart from strict necessity or sufficiency, crisp set theory in effect introduces ad hoc procedures to relax the rules of logic, instead of applying probabilistic reasoning. Fuzzy set theory in contrast builds on vagueness as a central conceptual principle, whereas we argue that uncertainty – codified in Bayesian probability – is the appropriate framework for scientific inference.

1.1.3 Multi-Method Research

Our vision of Bayesian probability as a unified framework for scientific inference contrasts with the majority of literature on multi-method research, which at least implicitly operates within a frequentist framework (e.g., Lieberman 2005, 2015, Gerring 2012, Weller and Barnes 2014, Seawright 2016, Goertz 2017). Authors maintain that combining different techniques – often statistical analysis and case studies – harnesses complementary sources of causal leverage. Statistical regression is generally conducted to establish a correlation, and case studies are included to illustrate causal mechanisms. However, many authors overlook the fact that the techniques they combine are associated with distinct epistemological principles and are designed to address different questions; the former aim to estimate population-level parameters, and the latter often aim to explain particular outcomes in specific cases (Mahoney 2010:141, Goertz and Mahoney (2012)). It is difficult to see how a case study would strengthen an inference about the average causal effect in the eyes of a committed frequentist, or how the numerical value of a regression coefficient would bring much insight to understanding a particular case or set of cases. Seawright (2016:5) provides a cogent discussion of these problems, noting that "because qualitative and statistical approaches produce results that are different in kind, it is only possible to assess … convergence very abstractly."

We contend that Bayesian probability is the only sound option for integrating qualitative and quantitative information on a more or less equal footing, because it is the only rigorous framework for which the same logic of inference applies across all types of data – whether stochastic or nonstochastic, experimental or observational, quantitative or qualitative. Efforts to bring both qualitative evidence and quantitative data to bear on a single research question without Bayesian probability (the approach pursued by Seawright 2016) inevitably prioritize either a quantitative method or a qualitative method as the main engine of inference.

Our perspective shares common ground with Humphreys and Jacobs (2015), who illustrate that multiple different goals, including estimating average effects, assessing case-level explanations, and comparing theories, can all be accommodated within a Bayesian model. However, we suggest moving away from the emphasis on average causal effects and distributions of unobservable causal types in a population – notions which remain rooted in frequentism – in favor of what we regard as a simpler and more direct Bayesian approach in which causal hypotheses for explaining known outcomes of given cases are the primary propositions of interest. Here we agree with Mahoney (2015:202) that thinking about average causal effects is not very useful for many research agendas. We also caution that prospects for formally combining inferences from quantitative and qualitative research will be constrained by the inevitable difficulty of nonarbitrarily quantifying the inferential weight of evidence in social science contexts. Yet despite the challenges, we believe that both quantitative and qualitative components of research can benefit from applying Bayesian insights.

1.2 A Guide for Readers

This book aims to reach a broad readership. Qualitative research practitioners and scholars who include case studies within multi-method designs are of course central audiences. Our Bayesian framework and practical recommendations for inference and research design apply not just to single case studies, but also to small-N and medium-N comparative research, ranging from historical analysis to reconstruction of more contemporary policymaking processes, as well as studies that aim to combine qualitative information with quantitative data. Yet we also aim to foster greater understanding of Bayesian inference among scholars whose primary research involves large-N analysis, experimental designs, and/or formal theory. All research draws on insights from qualitative information – regardless of the core analytical approach employed – and we believe that Bayesian probability can serve as an important bridge between qualitative and quantitative methods.

With these diverse audiences in mind, we have sought to make the book accessible and engaging for readers with a wide range of technical backgrounds. For those who do not have any previous exposure to Bayesian inference or classical statistics, we stress that no mathematical skills are required beyond basic arithmetic and algebra. Indeed, no previous methodological