INVESTIGATION · FOUNDATIONS

The Epistemology of Criminal Investigation: Generating Knowledge Beyond a Reasonable Doubt

What does it mean for a criminal investigation to know that someone is guilty? This article examines how investigations turn traces into knowledge, how courts define the standard of proof, and why probability, stories and explanation all play a part.

ModQu research article · 2,174 words · 22 references (APA 7th edition)

A criminal investigation is an attempt to know something about the past that nobody can now observe directly. A death has occurred; was it murder, and if so, who did it? The investigator cannot go back and watch. They must reason from traces that remain: bodies, documents, digital records, physical marks and the memories of witnesses. The court that eventually judges the case must then decide whether that reasoning is strong enough to justify punishment. Epistemology, the branch of philosophy concerned with knowledge and justified belief, offers a useful way to understand what investigators and courts are doing, and why they sometimes get it wrong.

This article looks at criminal investigation as a way of generating knowledge. It considers what counts as knowing in a legal context, how the standard of proof "beyond a reasonable doubt" is understood, and three models of how evidence is turned into a conclusion: probabilistic reasoning, story construction and inference to the best explanation.

Knowledge, belief and proof

Philosophers have traditionally defined knowledge as justified true belief: to know something, a person must believe it, it must be true, and they must have good reasons for believing it. Gettier (1963) showed that this definition is incomplete, because a person can have a justified true belief that is true only by luck. Imagine an investigator who believes a suspect is guilty because of a fingerprint that was, unknown to anyone, planted; if the suspect is in fact guilty for other reasons, the investigator's belief is true and justified, but few would say the investigator knew it.

The legal system cannot wait for philosophers to resolve such puzzles. It works with a more practical concept: proof to a standard. A court does not claim to know with certainty that a defendant is guilty. It claims that the evidence presented meets a defined threshold. Laudan (2006) argues that the criminal trial should be understood as an epistemic engine designed to find the truth, but one that also deliberately distributes the risk of error. Because convicting an innocent person is regarded as far worse than acquitting a guilty one, the threshold for conviction is set high. Blackstone's (1769) maxim that it is better that ten guilty persons escape than that one innocent suffer expresses this asymmetry.

Beyond a reasonable doubt

The criminal standard in common law countries is proof beyond a reasonable doubt. In the United States, the Supreme Court held in *In re Winship* (1970) that due process requires this standard for every element of a criminal offence. In Canada, *R. v. Lifchus* (1997) set out how judges should explain it: a reasonable doubt is not imaginary or frivolous, nor based on sympathy or prejudice, but on reason and common sense, and it is logically connected to the evidence or lack of evidence. Proof beyond a reasonable doubt does not require absolute certainty, but it requires more than proof that the accused is probably guilty.

Courts have generally refused to express the standard as a number, and research suggests that people interpret it differently. Simon and Mahan (1971) asked judges, jurors and students to state the probability of guilt they associated with different standards of proof and found substantial variation both between and within groups. Laudan (2006) argues that this vagueness is a serious problem for a system that aims to control the rate of wrongful convictions.

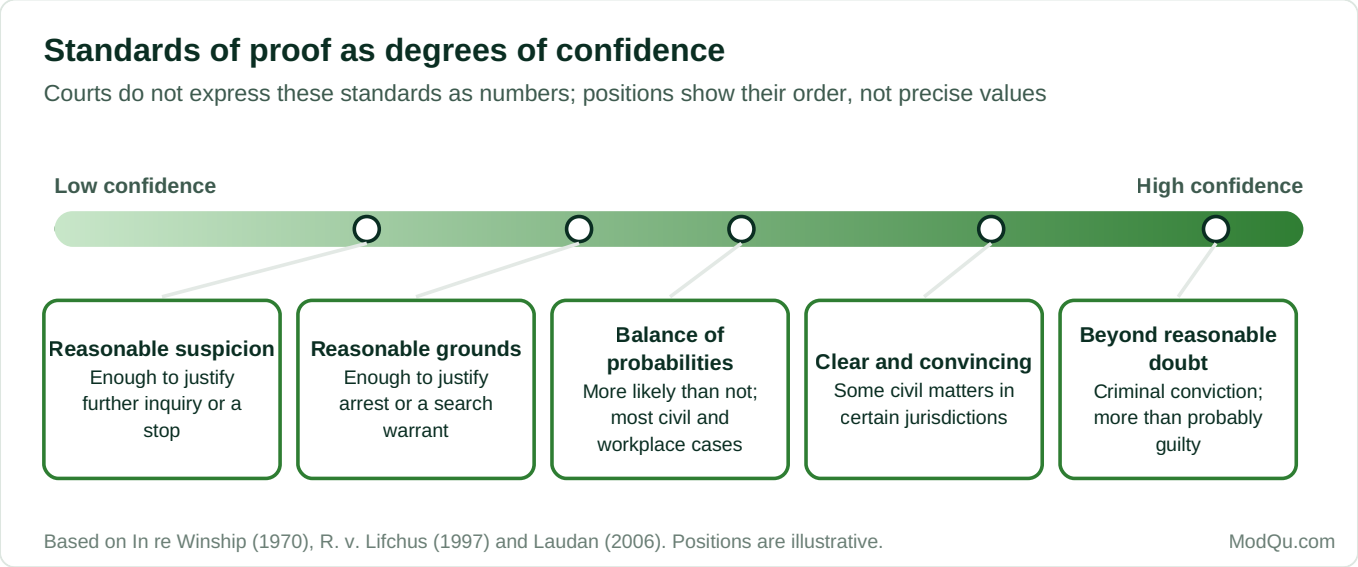


Figure 1. The main standards of proof arranged by the confidence they require. An investigation moves along this scale as evidence accumulates, and each stage authorises different actions. Diagram by ModQu.

For investigators, the standard matters long before trial. An investigation progresses through thresholds: reasonable suspicion to begin certain inquiries, reasonable grounds to arrest or search, and finally sufficient evidence to charge. At each stage the investigator must ask whether the evidence would eventually be capable of meeting the criminal standard, and whether the alternatives a court might consider have been addressed.

Three models of reasoning from evidence

Probabilistic reasoning

One way to think about proof is in terms of probability. Each item of evidence changes the probability that a hypothesis is true, and the strength of the change depends on how much more likely the evidence is if the hypothesis is true than if it is false. Bayesian approaches formalise this, and scholars have shown how complex legal arguments can be represented as networks of linked probabilities (Fenton et al., 2013, 2016).

Probabilistic thinking exposes common errors. The best known is the prosecutor's fallacy: confusing the probability of the evidence given innocence with the probability of innocence given the evidence. Thompson and Schumann (1987) showed that mock jurors often made this error, and also its mirror image, the defence attorney's fallacy, which dismisses strong evidence because many people in a large population could share a characteristic.

The case of Sally Clark in England shows the stakes. Clark was convicted in 1999 of murdering her two infant sons. An expert witness told the jury that the chance of two sudden infant deaths in a family like hers was about one in 73 million, a figure obtained by squaring the chance of a single death. The calculation wrongly assumed that the deaths were independent, and the figure was presented as if it were the probability of innocence. Hill (2004) showed that when the dependence between deaths in the same family and the rarity of double infant homicide were taken into account, the statistics did not favour murder. Clark's conviction was quashed on appeal in 2003 for other reasons, but the case became a standard illustration of statistical error in court.

Courts in England have been cautious about allowing juries to be taken through formal Bayesian calculations. In *R v Adams* (1996), the Court of Appeal held that introducing Bayes' theorem to the jury risked confusing them and usurping their task. The tension remains: probability provides the correct logic for weighing evidence, but legal decisions are made by people who reason more naturally in other ways.

Story construction

Psychological research on jurors suggests that they reason less by calculating probabilities than by building stories. Pennington and Hastie (1992) developed the story model, in which jurors organise the evidence into a narrative of what happened, judge stories by their coverage of the evidence, their coherence and their uniqueness, and choose the verdict that best fits the story they accept. Their experiments showed that evidence presented in story order was more persuasive than the same evidence presented in the order of witnesses.

Investigators construct stories too. Innes (2003), in his study of homicide inquiries, describes how detectives build an account of what happened, who did it and why, and how this account then guides the search for further evidence. Stories are indispensable, because they connect isolated facts into an explanation. They are also dangerous, because a coherent story can feel convincing even when parts of it are weak, and because once a story is accepted, evidence that does not fit tends to be discounted. Simon (2012) reviews research showing how this coherence effect can lead investigators, prosecutors and jurors to become more confident than the evidence warrants.

Inference to the best explanation

A third model combines elements of both. On this view, legal proof is a form of inference to the best explanation: the fact-finder compares the competing explanations offered by each side and asks which best accounts for the evidence (Lipton, 2004; Pardo & Allen, 2008). In a criminal case, conviction is justified when the prosecution's explanation is plausible and there is no plausible explanation consistent with innocence. Pardo and Allen argue that this matches how courts actually reason better than a purely numerical model does.

This view has direct implications for investigation. If proof is a contest between explanations, then the investigator's job is not only to build the strongest case for one explanation but to identify and test the alternatives a court will consider. An investigation that has never examined how the evidence could be explained if the suspect were innocent has not produced knowledge sufficient for proof, however strong its own story seems.

Three models of reasoning from evidence to proof

	Probabilistic	Story model	Best explanation
Core idea	Each item updates the probability of guilt	Evidence is organised into a narrative of events	Competing explanations are compared
Strength	Correct logic for weighing evidence	Matches how people actually reason	Focuses attention on alternatives
Main risk	Numbers misused or misunderstood	Coherent stories feel stronger than they are	Plausible alternatives may never be generated
Key sources	Fenton et al. (2016); Thompson & Schumann (1987)	Pennington & Hastie (1992); Simon (2012)	Lipton (2004); Pardo & Allen (2008)

Summary by ModQu ModQu.com

Figure 2. Three complementary ways of understanding how evidence becomes proof. In practice, careful investigators use all three. Diagram by ModQu.

How investigations produce knowledge

Bringing these models together suggests what a sound criminal investigation does epistemically.

1. **It gathers evidence with known credibility.** Each item has a documented source and a known risk of error. Anderson et al. (2005) and Schum (1994) stress that the credibility of each source, and the chain of inference from evidence to conclusion, must be analysed explicitly.
2. **It generates and tests rival explanations.** Following the logic of falsification (Popper, 1959), it seeks evidence that would show its preferred account to be wrong.
3. **It attends to diagnosticity.** It distinguishes evidence that merely fits the case from evidence that discriminates between explanations, and avoids treating the former as proof.

4. **It reasons abductively but checks deductively.** It forms explanations from the evidence (Josephson & Josephson, 1994) and then asks what else must be true if the explanation is right, and whether it is.
5. **It reports its uncertainty.** It distinguishes what has been established from what is probable and what remains unknown.

Haack (2014) argues that legal proof is best understood in terms of how well evidence supports a conclusion, considering how far the evidence is supportive, how secure it is independently and how comprehensive it is. An investigation that is strong on all three, with evidence that fits, is reliable on its own terms and covers the relevant questions, produces conclusions that deserve confidence.

Investigative knowledge and trial proof

Investigation and trial are often described as stages of the same process, but they have different epistemic roles. The investigation is open-ended: it gathers information, much of which will never be used, and it must remain willing to change direction. The trial is closed: it considers only admissible evidence, presented by two opposing parties, and must reach a verdict on a defined charge. The investigator therefore knows more than the court in one sense, because they have seen everything, and less in another, because their view has not been tested by an opponent.

This difference creates a particular responsibility. Information that points away from the suspect may never reach the court unless the investigator records it and passes it on. Disclosure rules in many jurisdictions require the prosecution to provide such material to the defence, but those rules depend on the investigation having recorded and retained it. The epistemic quality of the trial thus depends on the epistemic honesty of the investigation that preceded it.

The difference also explains why investigators should think like a sceptical court long before trial. An explanation that satisfies an investigation team may fail when the defence offers an alternative that the team never considered. The best protection is to anticipate that challenge: to ask, at each stage, what a well-informed opponent would say about the evidence, and to gather the evidence that would answer it.

Certainty, finality and new evidence

Criminal verdicts aim at finality, but knowledge is always open to revision. New scientific techniques, newly discovered witnesses or the confession of another person can show that a conviction or an acquittal was mistaken. Legal systems handle this tension through appeals, review bodies and, in some jurisdictions, limited exceptions to the rule against retrial. From an epistemological point of view, every verdict is a judgment made on the evidence then available, not a permanent truth. Investigators who keep complete records, preserve evidence and document the alternatives they considered make it possible for later reviewers to test that judgment if new evidence emerges.

When the system fails

Wrongful convictions show what happens when these conditions are not met. Garrett (2011) found that many of the first 250 DNA exonerations in the United States involved evidence that appeared strong at trial but was flawed: mistaken eyewitnesses, unreliable forensic claims, false confessions and informants with incentives to lie. In most, the jury had been persuaded beyond a reasonable doubt. The failures were epistemic: the evidence was less credible, and the alternatives less thoroughly excluded, than the people deciding the case believed.

These cases show that the standard of proof cannot by itself guarantee accurate outcomes. A high threshold protects the innocent only if the evidence presented is reliable and the alternatives have been properly explored. That responsibility falls mainly on the investigation, because by the time a case reaches court the opportunities to find exculpatory evidence, test alternative suspects or preserve fragile material have often passed. Courts can only weigh what investigators put before them.

Lessons beyond criminal justice

The epistemology of criminal investigation has lessons for any organisation that draws conclusions from evidence. Decisions should be matched to a stated standard of proof. Evidence should be assessed for credibility as well as content. Stories should be tested, not merely told. Rival explanations should be listed and examined. And conclusions should carry an honest statement of what remains uncertain. A research report, a market analysis and a criminal case are all claims to knowledge, and all deserve the same discipline.

References

Anderson, T., Schum, D., & Twining, W. (2005). *Analysis of evidence* (2nd ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511610585

Blackstone, W. (1769). *Commentaries on the laws of England* (Vol. 4). Clarendon Press.

Fenton, N., Neil, M., & Berger, D. (2016). Bayes and the law. *Annual Review of Statistics and Its Application*, *3*, 51–77. https://doi.org/10.1146/annurev-statistics-041715-033428

Fenton, N., Neil, M., & Lagnado, D. A. (2013). A general structure for legal arguments about evidence using Bayesian networks. *Cognitive Science*, *37*(1), 61–102. https://doi.org/10.1111/cogs.12004

Garrett, B. L. (2011). *Convicting the innocent: Where criminal prosecutions go wrong*. Harvard University Press.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–123. https://doi.org/10.1093/analys/23.6.121

Haack, S. (2014). *Evidence matters: Science, proof, and truth in the law*. Cambridge University Press. https://doi.org/10.1017/CBO9781139626866

Hill, R. (2004). Multiple sudden infant deaths—coincidence or beyond coincidence? *Paediatric and Perinatal Epidemiology, 18*(5), 320–326. https://doi.org/10.1111/j.1365-3016.2004.00560.x

Innes, M. (2003). *Investigating murder: Detective work and the police response to criminal homicide*. Oxford University Press.

In re Winship, 397 U.S. 358 (1970).

Josephson, J. R., & Josephson, S. G. (Eds.). (1994). *Abductive inference: Computation, philosophy, technology*. Cambridge University Press. https://doi.org/10.1017/CBO9780511530128

Laudan, L. (2006). *Truth, error, and criminal law: An essay in legal epistemology*. Cambridge University Press. https://doi.org/10.1017/CBO9780511617515

Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). Routledge.

Pardo, M. S., & Allen, R. J. (2008). Juridical proof and the best explanation. *Law and Philosophy, 27*(3), 223–268. https://doi.org/10.1007/s10982-007-9016-4

Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology, 62*(2), 189–206. https://doi.org/10.1037/0022-3514.62.2.189

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.

R v Adams [1996] 2 Cr App R 467 (Court of Appeal of England and Wales).

R. v. Lifchus, [1997] 3 S.C.R. 320 (Supreme Court of Canada).

Schum, D. A. (1994). *The evidential foundations of probabilistic reasoning*. Wiley.

Simon, D. (2012). *In doubt: The psychology of the criminal justice process*. Harvard University Press.

Simon, R. J., & Mahan, L. (1971). Quantifying burdens of proof: A view from the bench, the jury, and the classroom. *Law & Society Review, 5*(3), 319–330. https://doi.org/10.2307/3052837

Thompson, W. C., & Schumann, E. L. (1987). Interpretation of statistical evidence in criminal trials: The prosecutor's fallacy and the defense attorney's fallacy. *Law and Human Behavior, 11*(3), 167–187. https://doi.org/10.1007/BF01044641