The Holmesian logician: Sherlock Holmes’ “Science of Deduction and Analysis” and the logic of discovery

Emmanuel J. Genot¹

Received: 16 March 2020 / Accepted: 16 May 2020 / Published online: 29 May 2020
© The Author(s) 2020

Abstract

This paper examines whether Sherlock Holmes’ “Science of Deduction and Analysis,” as reconstructed by Hintikka and Hintikka (in: Eco U, Sebeok TA (eds) *The sign of three*: Peirce, Dupin, Holmes, Indiana University Press, Bloomington, 1983), exemplifies a *logic of discovery*. While the Hintikkas claimed it does, their approach remained largely programmatic, and ultimately unsuccessful. Their reconstruction must thus be expanded, in particular to account for the role of memory in inquiry. Pending this expansion, the Hintikkas’ claim is vindicated. However, a tension between the naturalistic aspirations of their model and the formal apparatus they built it on is identified. The paper concludes on suggestions for easing this tension without losing the normative component of the Hintikkas’ epistemological model.

Keywords Sherlock Holmes · Merrill B. and Jaakko Hintikka · Logic of discovery

1 Introduction

In the fictional universe Sherlock Holmes inhabits, logicians are capable of astounding feats that seem well beyond the abilities of their real-world counterparts. According to Holmes himself,

[f]rom a drop of water [...] a logician could infer the possibility of an Atlantic or a Niagara without having seen or heard of one or the other [...]. By a man’s finger-nails, by his coat-sleeve, by his boot, by his trouser-knees, by the callosities of his forefinger and thumb, by his expression, by his shirt-cuffs—by each of these things a man’s calling is plainly revealed. That all united should fail to enlighten the competent inquirer in any case is almost inconceivable. (*A Study in Scarlet*, I, 2)

✉ Emmanuel J. Genot
emmanuel.genot@fil.lu.se

¹ Department of Philosophy, University of Lund, Lund, Sweden

Dismissing Holmes' claims as fiction may be tempting but would be a mistake for he describes real-world skills. In 1877, doctor-in-training Arthur Conan Doyle served as an assistant to Scottish surgeon Joseph Bell (1837–1911) at the Royal Edinburgh Infirmary. Bell encouraged his students to draw inferences from observations of their patients and was keen on schooling them on how to do so. Ten years later, an episode of Bell's life, recounted by one Doctor Charles Watson McGillivray and reported in E. Liebow's *Joe Bell, Model of Sherlock Holmes* (1982), would inspire Holmes' claims that one can deduce "a man's calling" from (*inter alia*) the calluses on his hands:

Well do I remember the gasping astonishment of an outpatient to whom [Bell] suddenly remarked, 'Of course I know you are a beadle and ring the bells on Sundays at a Church in Northumberland somewhere near the Tweed.' 'I am all that,' said the man, 'but how do you know? I never told you.' 'Ah,' said Bell, when the outpatient had left bewildered, 'you all know about that as well as I did. What! You didn't make that out! Did you not notice the Northumbrian burr in his speech, too soft for the south of Northumberland? One only finds it near the Tweed. And then his hands. Did you not notice the callosities on them caused by the ropes? Also, this is Saturday, and when I asked him if he could not come back on Monday, he said he must be getting home tonight. Then I knew he had to ring the bells to-morrow. Quite easy, gentlemen, if you will only observe and put two and two together.' (Liebow 1982, p. 135)

Not only did Conan Doyle model Holmes' abilities (and attitude) on Bell's, he also formed his own picture of the "Science of Deduction and Analysis" (as per the title of the second chapter of *A Study in Scarlet*) that underlies them. On occasion, Conan Doyle lets Holmes depict this picture, explaining the principles of logic and probability that support his conclusions, and echoing in the fiction the real-world epistemological views of early proponents of George Boole's 'new logic', in particular, Stanley William Jevons (1835–1882).¹ Jevons' ideas predate by 50 years the discussion of scientific discovery from the 1930s, that came to an abrupt end when Popper (1992) and Reichenbach (1938) asserted the impossibility of a rational reconstruction of discovery processes. This assertion (hereafter, the *Popper–Reichenbach thesis*), which rallied hypothetico-deductivists and probabilistic inductivists in spite of otherwise opposed views on scientific methodology, resulted in a long-lasting consensus that there can be no 'logic' of discovery.

Consequently, Holmes' "Science of Deduction and Analysis" reflects epistemological and methodological insights suppressed by the Popper–Reichenbach thesis. In the early 1980s, it caught the attention of Merrill B. and Jaakko Hintikka, who articulated a partial reconstruction of Sherlock Holmes' methodology as a prelude to a full-blown

¹ Jevons was first identified as the origin of Holmes' conception of logic by Japanese philosopher of science Uchii Soshichi in a popular book (*Sherlock Holmes' Theory of Reasoning* [1988], in Japanese, no translation currently available, English title provided by Pr. Uchii in personal communication). Pr. Uchii later expanded his arguments in two electronic essays (Uchii 1999, 2010). Jevons' main contribution to formal logic was his "logical alphabet", a reformulation of Boole's algebra of logic that eliminated uninterpreted numerical expressions (see Lewis et al. 1959, pp. 14–15) and became the starting point of logical algebraists (particularly Peirce and Schroeder). Jevons implemented his logical alphabet in a mechanical computer in the 1870s, and used it to formalize reasoning by elimination of possibilities, a method often used by Sherlock Holmes.

logic of discovery. Under the title “Sherlock Holmes confronts modern logic” (Hintikka and Hintikka 1983), the Hintikkas proposed that “[t]he crucial part of the task of the Holmesian “logician” [...] is not so much to carry out logical deductions as *to elicit or to make explicit tacit information*.” (Hintikka and Hintikka 1983, p. 156, authors emphasis).

Building upon J. Hintikka’s work on the formal semantics of questions carried in the 1960s and 1970s, the Hintikkas defined “tacit information” as information available as answers to *questions*. Then they were able to show that deductive anticipations can guide questioning strategies in problem-solving, even for non-deductive problems.²

In this paper, we examine how much the Hintikkas captured of Holmes’ “Science of Deduction and Analysis” and whether their reconstruction qualifies as a *logic* of discovery. The Hintikkas’ seminal paper was however largely programmatic and their approach proved eventually unsuccessful. This would lead J. Hintikka to adopt (after M. B. Hintikka’s death in 1987) a somewhat different take on interrogative inquiry in the late 1980s and 1990s, abandoning some of the ideas introduced in their joint article. Therefore, while sympathetic to the Hintikkas’ original project and to J. Hintikka’s interrogative model of inquiry (IMI), our examination must be a critical one. Our criticism is nevertheless constructive and its outcome is essentially a vindication of the Hintikkas’ view that Holmes’ “Science of Deduction and Analysis” is a logic of discovery.

We first present an overview of the Hintikkas’ reconstruction of Holmes’ “Science of Deduction and Analysis” (Sect. 2) and their discussion of memory (Sect. 3), upon which we expand with Holmes’ account of memory (Sect. 4) that we support with some neuroscience (Sect. 5). We then return to inferences Holmes calls deductive but do not appear so (Sect. 6) and evaluate whether their rational reconstruction satisfies criteria for a *logic* of discovery (Sect. 7). We conclude on the prospect of unifying the deductive and associative components of Holmesian method. This paper builds on earlier technical work on the Hintikkas’ model and the IMI (in particular Genot 2018; Genot and Gulz 2016; Genot and Jacot 2012, 2018) while avoiding obfuscation by formalism, although we included summaries in footnotes for assertions backed by a formal construction, rather than asking our reader to take our word for them.

2 The “Science of Deduction and Analysis”, logical proofs, and strategies

Sherlock Holmes’ goal is typically answering a ‘big’ question, for instance, *Who stole the race horse Silver Blaze and murdered his trainer?* or *Where did Irene Adler hide the*

² In this paper, and unless stated otherwise, we understand a *problem* as a pair $\langle K, Q \rangle$ where K is a body of background knowledge and Q a question (specifiable as a set of possible answers) both expressed in some language (formal or not) $\mathcal{L}$. A problem is *deductive* if, and only if, some answer to Q is entailed by K , that is, if it is not possible to interpret all the elements of K as true and at least one of the answers to Q as false. Hence, a problem is *non-deductive* iff it requires an extension of K . A non-deductive problem is solved (non-trivially) if there is a (consistent) extension $K' \supset K$ such that an answer to Q is entailed by K' . Consequently, solving a non-deductive problem $\langle K, Q \rangle$ so construed amounts to transforming $\langle K, Q \rangle$ into a deductive problem $\langle K', Q \rangle$. For examples of deductive and non-deductive problems (including one from Sherlock Holmes’ *Adventure of Silver Blaze*), see the Appendix of Genot and Jacot (2018).

photograph of herself in the company of the King of Bohemia? That goal may in turn be characterized as the identification of the individual, object, or location satisfying the definite description occurring in the ‘big’ question. The inquiry terminates when this ‘big’ question receives a conclusive answer, that is, when Holmes and the other parties involved all agree that a given individual, object, or location, satisfies the definite description. In order to reach that endpoint, Holmes collects answers to ‘small’ questions, for instance: *Did the watchdog in the stables bark at Silver Blaze’s thief?* or *Where would Irene Adler first look at if she believed her house on fire?* Holmes selects these small questions based on anticipations of what their answers would entail in conjunction with his other background assumptions, and on how closer he would then get to a conclusive answer to the ‘big’ question.

The Hintikkas’ suggestion in Hintikka and Hintikka (1983) amounts to the following: in the above description of Holmesian inquiry, *individual* and *satisfaction* can be taken with the sense they have in logical theory and formal semantics. Holmes’ background assumptions and the definite description can sometimes be paraphrased in some first-order language, but first-order languages are not expressive enough to formalize all lines of inquiry, in particular when they involve reasoning about someone else’s beliefs (see Genot and Jacot 2012). Still, the first-order case suffices to establish the relation between Holmesian inquiry and deduction in the strict sense of logical theory. Namely, Sherlock Holmes’ ‘big’ question is answered when there is a deductive proof that such and such individual satisfies the definite description occurring in the ‘big’ question, based on whatever premises end up being accepted by the parties involved in the investigation at the end of the inquiry. Holmes’ background assumptions are typically not strong enough to warrant a deductive proof of the appropriate sort at the outset of the inquiry. Hence, the premises accepted at the end of the inquiry must often include answers to ‘small’ questions asked in the course of the inquiry, that strengthen Holmes’ background knowledge and transform an initial non-deductive problem into a deductive one (cf n. 2, p. 3). Thus, when a line of inquiry can be formalized in a first-order language, first-order metatheory suffices to explain how deduction weighs upon question selection.

While the Hintikkas recognized the importance of modeling less-than-perfect rationality, they also acknowledged that logical models incur unrealistic idealizations. Subsequently, they sketched a decision-theoretic analysis of strategy selection based on a utility function favoring lower cognitive loads on working memory (J. Hintikka added constraints on attention in Hintikka 1989). Unfortunately, utilities cannot in general support strategic reasoning in first-order proofs.³ After M. B. Hintikka’s death, J. Hintikka and his collaborators abandoned utilities and characterized the relation between deduction and interrogation with the resources of first-order semantics and proof theory alone (see Hintikka et al. 1999 for formal results and Hintikka 2007 for

³ For the formally minded reader: some questioning strategies can fall short of providing enough information to entail an answer to the big question. Given the semi-decidability of first-order logic, some of these unconclusive strategies can be infinite (i.e. have infinitely many steps). Thus, the standard means for strategic reasoning from utilities (elimination of dominated strategies) would require the elimination of infinite strategies. This would be tantamount to solving the Halting Problem, which is uncomputable by Turing machines, and formally equivalent to the decision problem for first-order logic (for a detailed formal argument, see Genot 2018; Genot and Jacot 2018).

an informal exposition). Their main result, the *Strategy Theorem* (Hintikka et al. 1999, p. 59), assumes an inquirer who reduces cognitive load by minimizing the number of cases to reason about and the number of individuals to consider in those cases, and thus has the same preferences as a theorem-prover. The theorem itself establishes that the inquirer's best strategy for selecting 'small' questions, assuming her background knowledge, mirrors the shortest proof of the conclusion she is trying to reach assuming the same background knowledge extended by answers to her 'small' questions. Resource-sensitive deductive anticipations (about the shortest proof from an extended background knowledge) therefore guide information-seeking (the selection of questions that provide the extension of the background knowledge, allowing to realize the proof) without need for explicit utilities.

Let us conclude this section with two illustrations of how the Hintikkas' model (and the IMI) clarifies the role of deduction in Holmes' inquiries. In *The Adventure of Silver Blaze*, Holmes investigates the theft of a racehorse, and knows that a watchdog was in the stables at the time of the theft. Holmes assumes that a well-trained watchdog would not fail to bark at a stranger. Together with this assumption, one potential answer to the question whether the dog barked (that it did) entails that the thief was not known to the dog, while the other (that it did not) entails that the thief must have been someone the dog knew well. Entailment is the semantic counterpart of deduction, and so the above amounts to saying that Holmes selects his question based on what he expects to deduce from its answer, and once obtained, that he *deduces* that the thief was someone the dog knew well. Similarly, in *A Scandal in Bohemia*, Holmes makes two assumptions: that, upon believing that one's house is on fire, one first checks the location of one's most prized possession; and that Irene Adler's photograph with the King of Bohemia is her most prized possession. Those two background assumptions together entail that the spot where Irene Adler's gaze lands immediately after a fire alarm will be the hiding place of the photograph. Again, one could say that Holmes *deduced* the hiding place of the photograph. Let us now have a closer look at how the Hintikkas integrate questions within proofs.

3 The limits of analysis and the role of memory

Deductive proofs comprise two types of steps: *analytic* steps, which decompose the premises in subformulas; and *non-analytic* steps, which introduce formulas that cannot be obtained by analytic steps. In the Hintikkas' model, a line of inquiry is represented as the construction of a semantic proof where analytic steps make explicit the information introduced in the proof at earlier stages. At first, this information is given by background assumptions alone. Later, as inquiry proceeds, answers to 'small' questions are added. When the information available can be expressed in a first-order language, analytic steps decompose the premises according to standard first-order compositional semantics, while non-analytic steps are used to generate more questions.

The Hintikkas made the important addition of *interrogative* steps, that are used to elicit information about the state of Nature. No additional notation or operators are needed however: disjunctions serve as preconditions for *whether*-questions, and existential statements for *who*-, *what*- and *where*-questions (*why*-questions do not have

object-language preconditions in the model). Analytic steps may suffice to obtain those formulas from formulas occurring at earlier stages, but cannot introduce relevant questions that involve new notions and concepts (represented, when a first-order reconstruction is possible, by new predicates), that is, concepts (*ditto*, predicates) that do not appear in the premises accepted at that step. In those cases, non-analytic steps are necessary.

The “curious incident of the dog in the night-time” from *The Adventure of Silver Blaze* is one of those cases whose reconstruction requires non-analytic steps. The police holds a suspect (Fitzroy Simpson), known to bet (and lose) on horse races, for the theft of a racehorse (Silver Blaze) and the murder of the horse’s trainer (John Straker). The evening prior to the theft, Simpson visited the stables but was forced to leave when a lad let a watchdog loose on him. Simpson claims that he was after insider information ahead of an upcoming race. But Inspector Gregory, in charge of the inquiry, believes that Simpson’s real motive was to scout the place and later steal the horse and sell it to Gypsies (who were camping nearby and left the morning after the theft) to clear his debts. Gregory’s working theory is that Straker expected a second visit, armed himself with a knife found on his nightstand (more on this knife later), went to the stables and interrupted Simpson, who killed him before stealing the horse. Sherlock Holmes has doubts and asks Gregory in a roundabout way whether the dog barked during the night:

“Is there any point to which you would wish to draw my attention?”

“To the curious incident of the dog in the night-time.”

“The dog did nothing in the night-time.”

“That was the curious incident,” remarked Sherlock Holmes. (*The Adventure of Silver Blaze*)

We may reconstruct Holmes’ reasoning about Simpson as a decision tree summarizing the construction of his deductive–interrogative argument, depicted in Fig. 1. Question $Q1$ is the ‘big’ question and the premises sum up the (explicit) background information that both Gregory and Holmes are attending to. Note that the premises do not yet mention the dog’s behavior, since Holmes considers it via questions $Q2$ and $Q3$, which Gregory overlooked. Holmes obtains $A2$ from Inspector Gregory (cf. above), and $A3$ and $A4$ from his own memory (cf. below) enabling the *reductio* of Simpson’s guilt. The question–answer pairs ($Q3$, $A3$) and ($Q4$, $A4$) illustrate the role of memory in Holmesian inquiry, which the Hintikkas characterize as follows:

In some cases, the great detective has to carry out an observation or even an experiment to answer the question. More frequently, all he has to do is to perform an *anamnesis* and recall certain items of information which he already had been given and which typically had been recorded in the story or novel for the use of the readers, too, or which are so elementary that any intelligent reader is expected to know them. (Hintikka and Hintikka 1983, p. 159)

Paraphrasing the Hintikkas, ($Q3$, $A3$) represents the *anamnesis* of elementary items of information (general knowledge about well-trained watchdogs) that anyone (including Gregory) should know. The pair ($Q4$, $A4$) similarly represents how Holmes recalls the newspapers account of Simpson’s evening visit (discussed between Holmes and

Premises: {*Simpson has motive, opportunity, owns a weapon, and can be placed at the crime scene.*
A dog was kept in the stables the night of the theft.}

Q1: Did Simpson steal Silver Blaze and killed John Straker?

Q2: Did the dog bark at the thief?

A2: No

Q3: To whom the dog would not bark at?

A3: Only someone the dog knows well

Lemma: *Either the dog knows Simpson well, or Simpson is not guilty (A2,A3)*

The dog knows Simpson well

Simpson is not guilty

Q4: Does the dog know Simpson well?

(A1: No)

A4: No

Fig. 1 The curious incident in the night-time

Watson for the reader's benefit at the very beginning of the short story). Holmes then deduces that Simpson cannot be guilty, and in the process, obtains a new descriptor for the thief of Silver Blaze (*someone the dog knew well*).

In the Hintikkas' model, questions *Q2* and *Q4* are supported by the disjunctions *either the dog kept in the stable barked at the thief, or it didn't* and *either the dog knows Simpson well, or it doesn't*. Question *Q3* is supported by the assumption that *if a trained watchdog does not bark at someone, then it knows that person?* The latter can in turn be construed as an answer to the question *is there a person to which a trained watchdog would not bark, or is there no such person?* which is also supported by a disjunction (*either there is a person to which a trained watchdog would not bark, or there isn't*). All these disjunctions are instances of the Excluded Middle, but they cannot be obtained from analysis of the premises, as they involve terms that do not appear in them.

The Excluded Middle is a (classical) tautology, and thus introducing an instance of it is always permitted in a (classical) deductive proof. The same holds for the Hintikkas' interrogative extensions of deductive proofs. But if the model allowed for unrestricted introduction of instances of the Excluded Middle, it could not discriminate between Holmes and Gregory. Subsequently, the set of the instances of the Excluded Middle an inquirer could introduce in a given deductive–interrogative proof is restricted by a parameter (the *range of attention*, Hintikka 1989). However, after abandoning decision theory and utilities (see Sect. 2, especially n. 3, p. 4) J. Hintikka never made this restriction explicit again nor addressed its relation to memory. This resulted in a loss of explanatory power for the model in complex cases, and limited its empirical applications.⁴

⁴ For the formally-minded: results presented in Hintikka et al. (1999) require select introductions of instances of the Excluded Middle, that recombine vocabulary from either the premises or the conclusion. More complex cases (like the “curious incident of the dog in the night-time”) and empirical

4 Holmes’ “brain attic” as a strategic account of memory

Perhaps surprisingly, one need not look any further than the Holmesian canon to find the sketch of a strategic account of memory in inquiry. This account is sketched in *A Study In Scarlet*, when Watson casually mentions that the Earth revolves around the Sun, only to find out that Holmes didn’t know that fact. Watson’s surprise turns to discombobulation when Holmes equally casually proclaims that he will do his best to forget everything about it. Holmes then offers the following clarification:

I consider that a man’s brain originally is like a little empty attic, and you have to stock it with such furniture as you choose. A fool takes in all the lumber of every sort that he comes across, so that the knowledge which might be useful to him gets crowded out, or at best is jumbled up with a lot of other things, so that he has a difficulty in laying his hands upon it. Now the skillful workman is very careful indeed as to what he takes into his brain-attic. He will have nothing but the tools which may help him in doing his work, but of these he has a large assortment, and all in the most perfect order. It is a mistake to think that that little room has elastic walls and can distend to any extent. Depend upon it there comes a time when for every addition of knowledge you forget something that you knew before. It is of the highest importance, therefore, not to have useless facts elbowing out the useful ones. (*A Study in Scarlet*, I, 2)

Holmes’ notion of a ‘brain attic memory’ (hereafter, BAM) is a fictional-world explanation for plot devices that would otherwise look like lucky guesses, but it is partially vindicated by current neurocognitive models. According to these models, human memory is a *content-addressable memory* (hereafter, CAM) that takes data as input and i.e., activates cliques of neurons (see below) where similar data is stored (see e.g., Hebb 2002; Marr 1969). By contrast, a *random access memory* (RAM), typical of artificial computers, takes addresses as input and returns the data stored at those addresses. Since grouping similar data in neighboring location (addresses or cliques of neurons) improves the performance of a CAM (but not that of a RAM) Holmes’ recommendation to keep “a large assortment [of tools], and all in the most perfect order” is a cogent recommendation to improve the performance of human CAM, which his BAM approximates (at least in that respect). This is particularly true relative to domains of expertise. Holmes’ is indeed a “skillful workman” specializing in criminal inquiry, and maintaining his BAM/CAM in “the most perfect order” supports relevant associations that often escape other criminal investigators.

The Holmesian canon offers more than a few examples of how a well-maintained “brain attic” supports Holmes’ investigations. Holmes has made a special study of tobacco ashes and even written a monograph on the topic (“Upon the Distinction Between the Ashes of the Various Tobaccos”, mentioned in *A Study in Scarlet*, *The*

Footnote 4 continued

application of the IMI (e.g. to collaborative inquiry learning, in Hakkarainen and Sintonen 2002) require questions featuring vocabulary absent from the premises. The absence of an explicit parameter for the range of attention is shown (in Genot and Gulz 2016) to result in the ultimate failure of the IMI to account for empirical subjects’ question selection. For alternative proofs for the results in Hintikka et al. (1999) with explicit restrictions on the range of attention, see Genot (2018).

Boscombe Valley Mystery and *The Hound of the Baskervilles*). He has also gathered very specific knowledge from special sciences, such as physical anthropology (that people tend to write graffiti at or just below eye level). This expert knowledge makes Holmes sensitive to details at a crime scene at Lauriston Garden (visited in *A Study in Scarlet*) that inspectors Lestrade and Gregson have overlooked or misinterpreted. Specifically, Holmes spots Trichinopoly cigar ashes on the floorboards and estimates the height of a murderer from an inscription (“Rache”) written on a wall. From this and other clues, he concludes that the murderer is a man, contrary to Lestrade and Gregson who suspect a woman (perhaps named Rachel). Ruling out a woman even suggests a motive and an ethnic origin (*rache* means revenge, in German), both later confirmed by Holmes’ investigation. Eventually, Holmes positively identifies the culprit in part thanks to Trichinopoly ashes in his hotel room.

In *A Study in Scarlet*, Holmes approaches the crime scene without preconceptions. By contrast, in *The Adventure of Silver Blaze*, Holmes initially accepted the police theory, and the narration begins with Holmes explaining to Watson that he must revise his initial assessment. The silence of the dog and its consequences is, in this context, more of a serendipitous discovery than the outcome of a systematic attention to details leading to equally systematic recall as in Lauriston Garden.⁵ Another serendipitous discovery is the odd shape of Straker’s knife, that makes a poor self-defence weapon (first explained away as an emergency choice absent a better option). Holmes learns from Watson that the knife is a veterinary scalpel, out-of-place on Straker’s nightstand, as he had no veterinary training. But Holmes knows of surgical procedures favored by unscrupulous racehorse handlers, which suggests an entirely new direction for his inquiry. Gregory and Watson fail to associate the silence of the dog with generalities about watchdog behavior, and thus to rule out Simpson’s guilt. Similarly, Gregory fails to identify Straker’s knife as a veterinary instrument and overlooks a possible association with surgical procedures used to rig a race. Watson acknowledges that something is amiss with the knife, but lacks knowledge of veterinary procedures to pursue any further.

5 The neuroscience of the brain attic

Neurological models of human CAM are more sophisticated than Holmes’ BAM, but their increased sophistication actually adds weight to Sherlock Holmes’ notion that a well-organized memory assists with inquiry and to the Hintikkas’ suggestion that memory plays a role in “intelligent inquiry.” To see this, let us first remark that Holmes assumes a one–one correspondence between memory items and *loci* in the brain, while theories of human CAM do not. The appropriate space metaphor would be that of a warehouse where similar parts of different items would be stored in the same location.

⁵ The notions of *serendipity* and *serendipity pattern* were coined by Merton (1948) to characterize “the fairly common experience of observing an unanticipated, anomalous, and strategic *datum* which becomes the occasion for developing a new or extending an existing theory” (p. 506). Merton credited Charles S. Peirce and his notion of abduction for influencing his definition (for a detailed history of the term before and after Merton’s paper, see Barber and Merton 2004). For “the curious incident of the dog in the night-time” as a serendipity pattern, see Genot and Gulz (2016); Genot and Jacot (2018).

However, the spacial metaphor then breaks down, as “the same location” must be understood as *literally the same*. Let us turn to that property for a moment.

Items stored in memory are *distributed* across cliques of neurons according to features those neurons are sensitive to. For instance, the memory of the shape of the diamond of a particular deck of playing cards would be distributed across neurons coding its particular rhombus shape and neurons coding its particular shade of red. Subsequently, recalling this particular diamond would require co-activation of (at least) two cliques of neurons. And since those cliques of neurons would code rhombuses and shades of red associated with other memories, there would be a nonzero probability that recalling the particular red rhombus of that particular deck of cards would also trigger recall of other (non-red) rhombuses or other (non-rhombus-shaped) red things. A distributed CAM is thus an economical solution to the problem of limited storage space that Holmes complained about. A more appropriate metaphor would thus be that of memory as a library of blueprints for modular items for a 3D printer, with any given item being stored as an index of blueprints for the modules it is made of.

A distributed CAM thus supports all manners of associations that may be relevant or not depending on the circumstances. In order to avoid association overload, recall is supported by an attentional mechanism called *contextual focus* (see Gabora 2010). In a nutshell, focused attention favors associations of perceived, imagined, or recalled objects or circumstances, with memory items that share their *typical* features. For instance, in a state of focused attention, a chair (perceived or imagined) would be associated with other items of furniture with legs and seats (stools, benches) and backrest (armchairs). By contrast, a state of de-focused attention opens to associations with less-than-typical features, where the same chair might be associated with other objects whose shape adapts to that of some body parts. Contextual focus is in part under volitional control, which opens the possibility to modulate strategically one’s attention to foster out-of-the-box associations, like the association between a comfortable padded chair and a bean bag (possibly suggesting a beanbag chair). Furthermore, the relevant notion of strategy can be made mathematically precise within game theory.⁶

⁶ For the formally-minded: assuming the game-theoretic form of sequential decision problems (as a sequential 1-player game vs Nature), contextual focus under volitional control can be mathematically described as strategic in the game-theoretic sense. Specifically, recall from a distributed CAM is a probabilistic process (recalling an item carries a probability for recalling others). In a sequential decision problem where some steps are ‘recall steps’, the outcome of such steps would be characterized by a lottery over possible associations. A strategy assigning to positions where a decision-maker has to make a move, a lottery between possible moves, is called a *behavioral strategy*. Since switching contextual focus modifies the probability of associating the current content of one’s attention or imagination with items stored in memory, it affects the lottery over the possible outcomes of a recall step. And since different lotteries characterize different behavioral strategies, switching contextual focus is tantamount to a change in strategy selection. A more detailed strategic analysis (complete with an explicit description of the probabilistic model of recall) is presented in Genot and Jacot (2018), and argued to support the notion of “strategic use of memory” in problem-solving consistent with the Hintikka’s project. The role of contextual focus in creative thinking and problem-solving was first identified by Liane Gabora, presented with its supporting empirical evidence in Gabora (2010), and shown to better organize the data than Darwinian theories of creative thinking (like the “blind variation and selective retention” theory) in Gabora (2011). With J. Jacot, we adapted Gabora’s rejoinder to Darwinian theories of creative thought in Genot and Jacot (2018) to show that the strategic analysis of recall offers a better account than Darwinian theories of serendipity (such as P. van Andel’s widely accepted account exposed in Van Andel 1994).

A strict, mathematical characterization puts strategic recall on a par with strategic deductive reasoning in non-deductive inquiry (in the sense of the *Strategy Theorem* discussed in Sect. 2). More precisely, strategic recall can account for the opening of new lines of inquiry via non-analytic steps, including serendipitous ones. Interestingly, Conan Doyle often describes Holmes as ‘daydreaming’, perhaps an indication that he suspected the role of contextual focus. Non-analytic steps of Holmes’ reasoning about the dog (and the knife) in *The Adventure of Silver Blaze* can be reconstructed as the outcome of mnemonic processes, and this reconstruction opens to fruitful comparisons with cases of serendipitous discoveries in the history of science (for instance, the serendipitous discovery of the microwave background radiation of the universe, cf. Genot and Jacot 2018). Moreover, at times, Holmes’ reasoning, involves question–answer steps that follow from associations and similarity relations, without obvious deductive inferential steps. Strategic recall can account for this type of reasoning, which would otherwise present a challenge for the Hintikka’s model, for instance the very first display of Holmes’ deductive skills in the opening chapter of *A Study in Scarlet*.

6 The reconstruction of association-based reasoning

At the beginning of *A Study in Scarlet*, John Watson just returned from Afghanistan, where he spent a few months convalescing from a wound sustained in combat. Living on a modest pension of which he intends to make the most, Watson is looking for some fellow to share a rent with. An acquaintance introduces him to Sherlock Holmes, who is looking for the same, in the chemical laboratory of a London hospital where the following interaction ensues:

“Dr. Watson, Mr. Sherlock Holmes,” said Stamford, introducing us.

“How are you?” he said cordially, gripping my hand with a strength for which I should hardly have given him credit. “You have been in Afghanistan, I perceive.”

“How on earth did you know that?” I asked in astonishment.

“Never mind,” said he, chuckling to himself. (*A Study in Scarlet*, I, 1)

Watson only gets an explanation for Holmes’ conclusion after moving in the 221B Baker Street residence. Watson has just read an essay titled “The Book of Life” (from which our introductory quote is borrowed) and expresses his skepticism towards the claims of its anonymous author (“What an ineffable twaddle!”). Holmes reveals himself as the author and endeavors to defend his claims with an example:

“Those rules of deduction laid down in that article which aroused your scorn are invaluable to me in practical work. Observation with me is second nature. You appeared to be surprised when I told you, on our first meeting, that you had come from Afghanistan.”

“You were told, no doubt.”

“Nothing of the sort. I knew you came from Afghanistan. From long habit the train of thoughts ran so swiftly through my mind that I arrived at the conclusion without being conscious of intermediate steps. There were such steps, however.

The train of reasoning ran, 'Here is a gentleman of a medical type, but with the air of a military man. Clearly an army doctor, then. He has just come from the tropics, for his face is dark, and that is not the natural tint of his skin, for his wrists are fair. He has undergone hardship and sickness, as his haggard face says clearly. His left arm has been injured. He holds it in a stiff and unnatural manner. Where in the tropics could an English army doctor have seen much hardship and got his arm wounded? Clearly in Afghanistan.' The whole train of thought did not occupy a second. I then remarked that you came from Afghanistan, and you were astonished."

"It is simple enough as you explain it," I said, smiling." (*Ibid.*, 1)

Holmes' claim that he is relying on the "rules of deduction" is somewhat puzzling. Holmes reasons from associations between Watson's title (Doctor), demeanor (posture, stiff arm) and appearance (skin tone, gauntness) to possible occupation (army doctor) and circumstances associated with it (foreign deployment, injury, sickness). As noted by the Hintikkas (1983, p. 163), the intermediate conclusion of Holmes' reasoning is literally expressed as the answer to a question (*Where in the tropics...?*) and, with Watson identified as an army doctor, only requires an *anamnesis* (that there is a war in Afghanistan, where the sun shines bright and the sanitary situation is hazardous) to obtain a unique location for Watson's last deployment.

We may reconstruct the above steps as a sequence of deductive steps. To see this, let us introduce the Simon–Newell model of problem-solving (see Newell et al. 1972), which captures information-gathering situations as a special case, and is thus more general than the Hintikkas' model of Holmesian deduction. In the Simon–Newell model, a *problem* is specified through a class of possible solutions, typically expressed in some regimented notation (sequences of first-order formulas for theorem-proving applications, sequences of steps coded in algebraic chess notation for chess problems, etc.). Those possible solutions are ordered by some preference relation (sometimes expressed as a utility function). A problem-solver explores that solution space, using domain-specific knowledge and heuristics to pick a starting point and select a course to navigate. The problem is solved when a candidate is found that satisfies the constraints for a solution, that is, is high enough in the preference ordering (or has a high enough utility), with "high enough" being contextually defined by the heuristics.

Couching Holmes' reasoning in the Simon–Newell model is rather straightforward. Holmes' first (and implicit) question is that of Watson's occupation, to which corresponds the space of all possible professional occupations. Holmes uses his background knowledge (Watson's title) to pick a starting point and his observations (tanned skin, gaunt silhouette, haggard face and peculiar posture) to further navigate that space. A likely enough solution (*Watson is an army doctor who spent time in the tropics*) triggers the next (explicit) question, that of Watson's last deployment. The space of possible solutions is now of geographical regions around the globe divided by countries. Again, Holmes uses his knowledge of current affairs (countries in which the British Empire maintains an active military presence in 1887) and of Watson's characteristics (same as above) to navigate that space, leading to a satisfying candidate solution (Afghanistan).

Holmes' reasoning can be reconstructed using the "rules of deduction" with some computer implementations of the Simon–Newell model. Associations that support the elimination of possibilities must be translated in conditional clauses (for instance "if someone is a medical doctor and has a military posture, then they are an army doctor"). Heuristic rules seek for sets of conditions ("John Watson is a doctor", "John Watson has a military posture") to which the *Modus Ponens* may be applied in conjunction with those clauses to obtain an intermediate conclusion ("John Watson is a medical doctor with a military posture"). The non-deductive part of this process is how the set of conditions is obtained in the first place, which is typically some pattern-recognition algorithm (see Simon 1973; Genot 2018 for the corresponding heuristics in the Hintikkas' model).

Clearly, a detailed reconstruction of Holmes' reasoning requires assumptions about Holmes' distributed memory. Still, it would exemplify the exploration of a solution space and subsequent elimination of candidate solutions captured by the Simon–Newell model. Elimination of epistemic possibilities is precisely the ancillary role that Jevons assigned to logic and probability calculus in rational inquiry. Since Jevons was most likely Conan Doyle's source (cf. our Introduction, and n. 1, p. 2), qualifying Holmes' reasoning as the product of the "rules of deduction" could be the result of Conan Doyle's lumping deduction and probability together which, although literally incorrect, is not without cogency, given that they share the same role in inquiry.

7 Holmesian inquiry and normative epistemology

So far, we have seen that psychological (associative) processes involved in the discovery of solutions for a large class of problems, are amenable to formal analysis and rational reconstruction. This already undermines the Popper–Reichenbach thesis, but not any formal analysis can qualify as a 'logic.' The formal analyses of justification processes proposed by Popper and Reichenbach support normative judgments about justification methods, and endeavor to order those methods as better or worse. They qualify as 'logic of justification' by virtue of this normative component. Similarly, a formal analysis of discovery processes must support normative judgements, in particular rankings of discovery methods as better or worse relative to how fast they converge to a solution for a given problem. The criterion that a logic of discovery must be normative (and in particular, support rankings) is implicit in the Hintikkas' article, in particular in the following:

The crucial part of the task of the Holmesian "logician", we are suggesting, is not so much to carry out logical deductions as *to elicit or to make explicit tacit information*. This task is left unacknowledged in virtually all philosophical expositions of logical reasoning, of deductive heuristics, and of the methodology of logic and mathematics. For this neglect the excuse is often offered that such processes of elucidations and explication cannot be systematized and subjected to rules. It may indeed be true that we cannot usually give effective rules for heuristic processes. It does not follow, however, that they cannot be discussed

and evaluated, given a suitable conceptual framework. (Hintikka and Hintikka 1983, pp. 156–157)

An example of “deductive heuristics” would be the appeal to auxiliary constructions (lemmas), typically to shorten deductive proofs. Assuming a preference for shorter proofs, a strategy with lemmas is in general preferable to a strategy without lemmas. In formal proofs, lemmas are handled via the Cut Rule, which is, in the proof system that underlies the Hintikkas’ model, equivalent to introducing an instance of the Excluded Middle. Any language as expressive as a first-order language capable of enumerating objects would generate countably many candidate lemmas for any proof in that language. Furthermore, some candidates not only would not shorten the proof, but might delay its completion indefinitely. Thus, there can be no effective method to eliminate those candidates. By the same token, there is no general effective method for selecting the best lemma. Since the Cut Rule is also the device by which yes-or-no questions are introduced in the Hintikkas’ model, they are correct to point that heuristics in their “logic of discovery” are no more problematic than in logic *tout court*. That “they [can] be discussed and evaluated, given a suitable conceptual framework” is illustrated *inter alia* by the *Strategy Theorem*, which requires select applications of the Cut Rule, for which there is no effective method, while retaining a normative content.⁷

Herbert Simon argued in much the same way as the Hintikkas, that a normative framework for discovery must be able to rank discovery strategies as better or worse and to provide with recommendations for choosing discovery strategies (see Simon 1973). Taking stock in computer implementations of his model of problem-solving (in Zytkow and Simon 1988), Simon further illustrated his point with domain-specific algorithms for solving particular classes of problems, namely the discovery of empirical laws from qualitative data (descriptions of chemical reactions) or quantitative data (descriptions of physical systems). In both cases, the discovery algorithms can be ranked as better or worse based on how fast they converge to known empirical laws. Simon argued that the ability to support normative claims qualifies his framework as a ‘logic of discovery’ for the class of problem they address, in the same sense that deductive logic (Popper) or probability theory (Reichenbach) would qualify as ‘logic(s) of justification’ for empirical hypotheses.

Simon’s heuristics discovery algorithms are parasitic on heuristics for automated theorem-proving, that may also be tweaked to generate interrogative–deductive proofs *à la* Hintikkas. Algorithms of that sort can solve non-deductive problems where the yes-or-no questions cannot be obtained analytically from the premises but are formulated in the same vocabulary as the premises (see Genot 2018). In more complex cases similar to “the curious incident of the dog in the night-time,” altogether new terms must be introduced (such as *barking* or *being the master of* in association with *watchdog*).

⁷ For the formally-minded: the *Strategy Theorem* expresses a normative claim, namely, that the best discovery strategies for a class of non-deductive problems (Holmesian problems expressible in a first-order language) parallel the best deductive strategies for a class of deductive problems (the deductive transformation of the non-deductive problems, cf. n. 2, p. 3). However, due to the semi-decidability of first-order logic, there can be no effective method to transform a non-deductive problem into a deductive one (some candidate transformations where the premise set is still too weak to entail the desired conclusion may generate infinite proof attempts, and eliminating them would be tantamount to solving the Halting Problem). Hence, we are in a situation where normative claims are possible even in the absence of effective rules.

The search space for questions grows exponentially, and heuristics that do more than re-composing the vocabulary of the premises become necessary. We have suggested in previous sections that human memory can constrain heuristics search in the desired way, thanks to its distributive structure and how it supports associations.

In the light of the above, we can understand the Hintikkas' insistence on the role of memory as a move towards a natural epistemology that would still qualify as a 'logic of discovery' *à la* Simon. Indeed, the conceptual import of the *Strategy Theorem* is preserved in contexts where re-composition of vocabulary is insufficient, provided that heuristics for associations are available (cf. Genot 2018; Genot and Jacot 2018). Now, and as noted earlier, associations supervene on the memory of a particular agent (or, in the multi-agent scenarios discussed in Genot and Jacot 2018, the memory of different agents). Hence, case studies may require a detailed reconstruction of the context of discovery, including individual characteristics of the discoverer's (or discoverers') memory, and this may be viewed by some as too big a step into psychological territory for a normative epistemological project.

8 Concluding remarks

In the late 1980s, Zytkow and Simon (1988) took stock of the advances of automated discovery algorithms to express their consequences *vis-à-vis* the Popper–Reichenbach thesis:

The traditional distinction that places studies of scientific discovery outside the philosophy of science, in psychology, sociology, or history, is no longer valid in view of the existence of computer systems of discovery. It becomes both reasonable and attractive to study the schemes of discovery in the same way as the criteria of justification were studied: empirically as facts, and logically as norms. (Zytkow and Simon 1988, p. 65)

Arguably, the psychological components are part of the facts of discovery, not its logic. Establishing them for the purpose of case studies and rational reconstructions (that are the bread and butter of philosophers of science) may require some speculation. Nevertheless, we can find support in Sherlock Holmes again here, for as long as we stay within the confines of "the scientific use of imagination [where] we have always some material basis on which to start our speculation" (*The Hound of the Baskervilles*, 4), we can look at (domain-specific) norms. The next issue is how much mileage we can get from those norms.

The Hintikkas demonstrated how far deductive norms can carry us. They successfully captured the strategic role of deduction in problem-solving. Their decision-theoretic approach to strategy selection with bounded rationality collapsed eventually, but foreshadowed J. Hintikka's proof-theoretic account and the *Strategy Theorem*, that still stand. Finally, they recognized that entailment-based inferential norms cannot carry very far when it comes to introducing new concepts through questions, and they pointed to the role of memory, beginning with Sherlock Holmes' "brain attic," a suggestion that we followed here (and in greater details in Genot and Jacot 2018).

If we are looking to semantics capable of generating non-trivial (non-flat) preferences over questions, some cognitive semantics is most likely our best bet. For instance, Charles J. Fillmore's deep case grammar (Fillmore 1968) and frame semantics Fillmore (1976) (see also Gamerschlag et al. 2013 for recent developments) capture how certain words—interpreted by concepts—evoke others within a lexical domain—interpreted by a conceptual frame—thereby constraining discourse construction. Such an account could be co-opted to model how, in *The Adventure of Silver Blaze*, *dog* evokes *bark* and triggers Holmes' (covert) question. Similarly, construction grammars, which evolved from Fillmore's work—with Lakoff (1987) and Langacker (1987a,b)—impose pragmatic and semantic constraints on syntactic constructions, reversing the Hintikkas' approach which imposed syntactic restrictions (the *range of attention*, which is just a set of sentences) to mitigate semantic permissions (from the truth-values of tautologies). Also, and quite obviously, cognitive semantics is by design sensitive to bounds on cognitive resources (and has the means to express them). Switching to cognitive semantics would have definite advantages for capturing interrogative strategies. Let us briefly suggest two of them (which are really two sides of the same coin).

The first advantage of cognitive semantics over model-theoretic semantics is its fine-grained interpretation for parts of discourse (substantives, adjectives and verbs) that accounts for relations of neighborhood and evocation (for instance, how *watch-dog* and *thief* could evoke *barking*). Contrast with model-theoretic semantics, with its coarse-grained semantic structures lumping parts of discourse as predicates, and its need for *ad hoc* conditional clauses to express their relations. The second and closely related advantage is the straightforward relation between evocation and recall, and thus between conceptual and neurological levels of description. This relation supports the unification of semantic and associative components of interrogation strategy selection. Contrast with any extension of the Hintikkas' model (including ours) where the associative component acts as a syntactic filter over semantic permissions.

A cognitive approach to Holmes' "Science of Deduction and Analysis" would however require a level of formalization comparable to that of the Hintikkas' (and the IMI). Formalization is not *per se* desirable in cognitive semantics (and is seldom a priority) but a cognitive *aggiornamento* of the Hintikkas' model would require to map the *Strategy Theorem*, that orders strategies based on proof-theoretic criteria, to an equivalent construction that would order strategies based on cognitive semantics criteria. We hypothesized that this mapping would yield further constraints on strategies. But in order to test this hypothesis, a cognitive semantics must be able to capture the same inferential operations as the Hintikkas' semantic and proof-theoretic approach does, and thus, to interpret a logic (via some translation scheme).

Acknowledgements Open access funding provided by Lund University. I owe to J. Jacot my understanding of the relation between the neuroscience of BAM/CAM and question evocation without which (as pointed by an anonymous reviewer) this paper would only provide a marginal contribution to the discussion of erotetic frameworks for discovery. I am indebted to Pr. Uchii for drawing my attention to Jevons' logic (on a par with his ideas about probability) to interpret Holmes' method. I also wish to thank an anonymous referee for pointing out that J. Hintikka's 1990s IMI burgeoned in the 1960s and 1970s, making the 1983 model an unsuccessful decision-theoretic graft on his earlier logical work. I owe to the same reviewer the most enjoyable compliment received in my academic career, namely that this paper was "written in reader-friendly style" and "informative but also nice to read" and will let the reader decide whether it is deserved.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Barber, E., & Merton, R. K. (2004). *The travels and adventures of serendipity*. Princeton: Princeton University Press.
- Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1–68). New York: Holt, Rinehart, and Winston.
- Fillmore, C. J. (1976). Frame semantics and the nature of language. *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, 280, 20–32.
- Gabora, L. (2010). Revenge of the “Neurds”: Characterizing creative thought in terms of the structure and dynamics of memory. *Creativity Research Journal*, 22(1), 1–13.
- Gabora, L. (2011). An analysis of the blind variation and selective retention theory of creativity. *Creativity Research Journal*, 23(2), 155–165.
- Gamerschlag, T., Gerland, D., Osswald, R., & Petersen, W. (2013). *Frames and concept types: applications in language and philosophy* (Vol. 94). Berlin: Springer.
- Genot, E. J. (2018). Strategies of inquiry. *Synthese*, 195(5), 2065–2088. <https://doi.org/10.1007/s11229-017-1319-x>.
- Genot, E. J., & Gulz, A. (2016). The interrogative model of inquiry and inquiry learning. In C. Bashkent (Ed.), *Perspectives on interrogative models of inquiry* (pp. 15–33). Berlin: Springer.
- Genot, E. J., & Jacot, J. (2012). How can yes-or-no questions be informative before they are answered? *Episteme*, 9(2), 189–204.
- Genot, E. J., & Jacot, J. (2020). The brain attics: The strategic role of memory in single and multi-agent inquiry. *Synthese*, 197, 1203–1224. <https://doi.org/10.1007/s11229-018-1743-6>.
- Hakkarainen, K., & Sintonen, M. (2002). The interrogative model of inquiry and computer-supported collaborative learning. *Science and Education*, 11(1), 25–43.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory* (1st ed.). Hove: Psychology Press.
- Hintikka, J. (2007). *Socratic epistemology*. Cambridge: Cambridge University Press.
- Hintikka, J., & Hintikka, M. B. (1983). Sherlock holmes confronts modern logic: Toward a theory of information-seeking through questioning. In U. Eco & T. A. Sebeok (Eds.), *The sign of three: Peirce, Dupin, Holmes*. Bloomington: Indiana University Press.
- Hintikka, J., & Hintikka, M. B. (1989). Reasoning about knowledge in philosophy: The paradigm of epistemic logic. *The logic of epistemology and the epistemology of logic* (pp. 17–35). Dordrecht: Kluwer.
- Hintikka, J., Halonen, I., & Mutanen, A. (1999). Interrogative logic as a general theory of reasoning. *Inquiry as inquiry: A logic of scientific discovery* (pp. 47–90). Dordrecht: Kluwer.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago press.
- Langacker, R. W. (1987a). *Foundations of cognitive grammar: theoretical prerequisites* (Vol. 1). Redwood: Stanford University Press.
- Langacker, R. W. (1987b). *Foundations of cognitive grammar: Descriptive application* (Vol. 2). Redwood: Stanford University Press.
- Lewis, C. I., Langford, C. H., & Lamprecht, P. (1959). *Symbolic logic*. New York: Dover Publications.
- Liebow, E. (1982). *Dr. Joe Bell, Model for sherlock holmes*. Bowling Green: Bowling Green University Popular Press.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology*, 202, 437–470.

- Merton, R. K. (1948). The bearing of empirical research upon the development of social theory. *American Sociological Review*, 13, 505–515.
- Newell, A., Simon, H. A., et al. (1972). *Human problem solving* (Vol. 104). Englewood Cliffs: Prentice-Hall.
- Popper, K. R. (1992). *The logic of scientific discovery* (1st ed.). London: Routledge. 1935 (in German).
- Reichenbach, H. (1938). On probability and induction. *Philosophy of Science*, 5(1), 21–45.
- Simon, H. A. (1973). Does scientific discovery have a logic? *Philosophy of Science*, 40(4), 471–480.
- Uchii, S. (1999). Sherlock holmes and probabilistic induction. <http://philsci-archive.pitt.edu/167/>.
- Uchii, S. (2010). Sherlock holmes on reasoning. <http://philsci-archive.pitt.edu/5306/>.
- Van Andel, P. (1994). Anatomy of the unsought finding, serendipity: Origin, history, domains, traditions, appearances, patterns and programmability. *The British Journal for the Philosophy of Science*, 45(2), 631–648.
- Zytkow, J. M., & Simon, H. A. (1988). Normative systems of discovery and logic of search. *Synthese*, 74(1), 65–90.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.