

How Can Yes-No Questions Be Informative?

Emmanuel J. Genot (Université Lille III, UMR 8163 STL)
Justine Jacot (Université Paris I, IHPST)*

July 9, 2009

1 Introductory Remarks

In this paper, we propose to analyze the informativeness of questions, in the context of ‘inquiry games’, where a player, *Inquirer*, intends to establish a certain conclusion by means of a combination of reasoning and interrogation, using information (answers) from other players, *Sources*. This paper is one step toward a precise game-theoretic foundation for these games, which were suggested by Hintikka to be a model of rational inquiry.¹ Sources may be partially unaware of Inquirer’s goal, and becoming aware of her goal may make them unwilling to achieve cooperation (i.e. answer truthfully whenever they can). This introduces a discrepancy with situations amenable to standard game theory, in which the space of strategies is common knowledge among players.

Section 2 introduces the idea of inquiry games, and presents an example borrowed to Holmesian lore, in order to show how questions introduce information in inquiry contexts. We focus on ‘yes-no’ questions, because of their semantic properties which in turn guarantee their availability in questioning strategies. Section 3 presents the difficulties standard game theory face to represent the kind of strategies an inquirer has to rely on to prevent her question to inform her sources when this information could affect their answers. We conclude with some remarks about the relations between the present work and issues in epistemology, and the application of game-theory.

2 Inquiry As A Game

2.1 A Game-Theoretic Analogy

Hintikka suggested that empirical inquiry can be modeled as a game in which a player, *Inquirer*, tries to establish a given conclusion C from a set of (theoretical) premises T (a subset of her background knowledge), using both reasoning steps (logical moves) and interrogative steps (interrogative moves), in a given context (an underlying model

*The authors wish to thank Sebastian Enqvist, Frédéric Jouneau, Gabriel Sandu, and Bernard Walliser, who offered valuable comments on previous versions of this paper. They are also greatly indebted to three anonymous referees, whose insightful remarks led to major modifications of the initial paper, and hopefully some clarifications and improvements.

¹In Hintikka’s information-seeking games, Inquirer plays against Nature, trying to prove a conclusion C from premises T , together with answers from a unique source, *Oracle* (see essays in [7] and [8]), which is usually assumed to be truthful (but see [6]).

M). Hintikka calls *Interrogative Model of Inquiry* (IMI) the epistemological theory stemming from this analogy with games, and he argues that this model is able to bring under a new light the idea that (deductive) logic is a norm of good reasoning.

Inquiry games are indeed a game-theoretic interpretation of the kind of reasoning formalized in tableaux systems, i.e. building a model of *T* which is also a countermodel of *C*, once generalized to apply to cases where a given conclusion does not follow from premises alone. In these cases, subtableaux which remain open (compatible with the premises and the negation of the attempted conclusion) represent possibilities left so far uneliminated by the premises. The role of questions – and answers – is to narrow down the range of alternatives left open, i.e. the various ways *M* might be, given the information assumed (premises) and obtained (answers) by Inquirer.

The analogy is easy to understand in the case of *semantic trees* (as introduced in [13]), to which rules for interrogative moves are added.² One can interpret the successive nodes as positions in an extensive game. A strategy in such a game is thus a sequence of logical and interrogative moves calculated to close the tree – if the necessary answers are available – or obtain a countermodel – if they are not – as quickly as possible. It can be shown that the ‘best’ interrogative strategy parallels the best deductive strategy (*Strategy Theorem*, [9]). Since first-order logic is not decidable, the ‘best’ interrogative strategy cannot be specified recursively. Yet this parallelism helps understand the role of logic in reasoning, and vindicates the ‘Sherlock Holmes sense of logic’, i.e. the conception that rules for *deductive* inference, play a role in problem-solving and inquiry.³

Hintikka and his collaborators have been mostly interested in making this parallelism explicit. Hence most of the discussion has focused on cases where a unique source – Nature, or Oracle – provides answers, and in which this source is assumed to be truthful. Rules for excepting answers are nonetheless discussed (see [9]), and it has been shown that they are enough to generate contractions and revisions in the sense of Belief Revision Theory (see [2]). Yet, in both cases, available answers are a ‘static’ parameter.

Nature, or multiple (and possibly conflicting) sources, give the answer whenever it is available in a particular context (model). Sources are not treated as strategic players, and are insensitive to some of Inquirer’s strategic choices.⁴ Some sources may be unreliable, and there can be several strategies (from Inquirer’s part) to cope with uncertainty. But this uncertainty is not the result of choices made by sources as players in a strategic game. All of this could be viewed as the result of Nature’s moves, but Nature is non-strategic player, i.e. a player with a ‘flat’ preference ordering over outcome.

Taking the game-theoretic analogy seriously means that one should consider ‘strategic’ sources, i.e. sources as agents, and evaluate the consequences on inquiry games.

²For propositional questions, interrogative rules state that any formula to which a *branching rule* would be applied (β -type formulas, in [13]) as the result of a logical move, can serve as the presupposition of a question. If an answer (input from the model) is available, Inquirer is dispensed from reasoning by cases, and no branching occurs. With existential or negated universal formulas (δ -type) an answer introduces an individual identifiable in the model, and thus dispenses to reason with an ‘arbitrary’ object or a dummy name (see [9]). See [2] for such a system for propositional logic.

³This sense is used by Holmes himself, e.g. in “A Study in Scarlet”: “From 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.” ([1], p. 23).

⁴Explicit definitions are creative in interrogative games. Hence, two questions which are, given *T*, equivalent by definition, can be such that no answer to the first is available from any source, while an answer to the second is. Likewise, some measurement may be exclusive (when resources are limited, or when measurement affects the system to be measured enough to introduce uncertainty). Hence, even ‘non-strategic’ sources may exhibit sensitivity to choice of questions, order, etc.

Games against Nature, or with passive (non-strategic) sources, should be considered as ‘degenerated’ inquiry games, and one should try to identify the features of more general inquiry games. The next section will propose to identify some of them, starting from an example.

2.2 An Example of Holmesian Inquiry

Holmesian inquiry combines deduction and questions, and Holmes regularly keeps his sources of answers in the dark as to the motives of his questioning. Sometimes, disguising the purpose of his question is crucial to the usefulness of the information obtained, as illustrated by the following example (from the short story “Silver Blaze”, in Conan Doyle’s *Memoirs of Sherlock Holmes*).

Example 1. *Investigating the disappearance of a race horse, Silver Blaze, and the death of his trainer John Straker, Holmes suspects that Straker accidentally died while attempting to wound the horse, in order to rig a forthcoming race, and make a profit betting against Silver Blaze. Holmes has found in Straker’s pocket a bill for an expensive dress, and suspects that Straker bought it for his mistress. The expenses incurred by a double life would explain Straker’s scheme. Instead of asking the widow whether she owns the dress, Holmes asks her whether they formerly met, feigning to remember her wearing a similar dress. She answers him that it is unlikely, since she does not own such a dress. Holmes then concludes that his hypothesis is correct.*

The basic ideas of the MI can be illustrated informally from this example. Before addressing his question to the widow, Holmes considers equally possible two scenarios (or rather, two classes of scenarios, everything else being equal). In one one, Straker bought an expensive dress for his wife; while in the other, he bought it for his mistress. In the first case (scenario) Straker would have had no need for the extra money he would have won betting against his horse in a rigged race (or at least, it wouldn’t have been a sufficient motive). In the second case, since this expense would not have been the only one incurred by living a double life, he would have needed the money. Obtaining evidence that Straker bought the dress for his mistress, given Holmes’ background assumptions, would be evidence enough to ascribe Straker a powerful motive for rigging the race.⁵

Holmes is attempting to eliminate the first scenario, using an answer from the widow. Clearly, his reasons for asking the question: “Didn’t we meet last year, while you were wearing such-and-such a dress?” rather than: “Did your husband buy you such-and-such a dress (or not)?” are a consequence of his expectation about the widow’s willingness to answer. She is less likely to guess his purpose (the conclusion he’s aiming at) in the first case, than in the second. Assuming that she would not want to be known that her husband had a mistress, she’s likely to refuse to answer (or lie outright) if the first question is asked.

Sherlock Holmes possesses two different types of private information, in the sense that it is not shared with Straker’s widow. The first one is factual: Straker bought such-and-such dress. The second is about his own intention: the reason why he wants to

⁵However the two scenarios are not equally likely, in Holmes’ opinion. The bill was addressed to someone else, and Holmes, at some point, states that no gentleman would carry in his pocket someone else’s bills. Hence Straker must have a motive for having the bill addressed to him under an alias, which makes the hypothesis that the second scenario is the actual one more likely. This kind of consideration is particularly important when one lacks evidence (or doubts one’s evidence) to eliminate all but one scenarios. In the present case, given Holmes’ background assumptions, the initial ‘educated guess’ is corroborated by the answers obtained from sources.

know whether the widow possesses such-and-such dress. Were either of these information revealed to the widow, she would be likely to lie. Hence Holmes has to choose his questions in such a way as not to give away either.

How can the choice of a question to ask make a difference as to what Holmes reveals? Consider, for simplicity sake, that Holmes has as only options the two questions above – or rather, only two strategies, one in which he uses the first, and the other in which he uses the second. Both are ‘yes-no’ questions, i.e. their presupposition is a tautology, hence no prior reasoning step is needed to ask those questions.⁶ Their introduction does not, as a consequence, depend on prior moves, and both strategies could be identical up to the choice of either questions. This choice is made on the basis of whatever information is likely to be given away. The next section discusses how epistemic logic can represent this informativeness of ‘yes-no’ questions.

2.3 Informativeness of Information-Seeking Questions

Consider a propositional question Q of the form: “Is it the case that ϕ or ψ ?” (assuming that ϕ and ψ are pairwise exclusive, but not necessarily jointly exhaustive) asked by an agent X , and overheard by an agent Y . Assume furthermore that Y believes that X has asked Q as a request for information.⁷ From this latter assumption, Y will likely infer that: (i) X believes that either ϕ holds or ψ holds; (ii) X does not know whether ϕ holds or ψ holds; and (iii) X wants to know which of ϕ and ψ holds.

If ϕ and ψ are exhaustive given X ’s background assumptions, but not according to Y , and if Y considers X to be a trustworthy source, then Y learns something from X ’s asking the question from (i). If moreover Y ’s background assumptions include some information tying X ’s ignorance as to whether ϕ or ψ holds to some other facts, then Y is likely to learn even more from (ii). What is learned from (iii) is trickier to represent. Let’s discuss information obtained from the two first conclusions, in order to show the contrast with information obtained from the third.

The difference between the set of scenarios to which Y ’s attention is restricted, given Y ’s background assumptions, before and after hearing X ’s question, can be represented as an update of a S5 Kripke structure, since indiscernibility is an equivalence relation.⁸ The ‘old’ equivalence class of epistemically indiscernible scenarios (compatible with Y ’s prior information) is updated, after hearing X ’s question, i.e. restricted to only those scenarios compatible with (i) and (ii). In the new structure, the related scenarios are those which were so related before the new information has been obtained, *and* are compatible with this information.⁹

In the special case of ‘yes-no’ questions, no information comes from (i). The presupposition of the question is a tautology, and updating the Kripke structure results in a trivial update. Hence they are informative, in the sense of elimination of possibilities, in the above case only insofar as Y ’s background assumptions link X ’s ignorance to

⁶What is needed, however, is a rule allowing to introduce a ‘yes-no’ question at any step of an argument, which is equivalent to the cut rule. See [9], and [2] for an implementation in semantic trees.

⁷Notice that Y could mistake the purpose of X ’s question, since Q could be an ‘exam’ question aimed at establishing whether Y knows the answer, while X already knows it. X could also use Q to check whether Y answers truthfully, if X knows that Y believes that one of the answers hold), or whether Y has some information about the answer (and whether it corresponds to X ’s information).

⁸We assume that information resulting from X ’s question is compatible with Y ’s background knowledge.

⁹*Dynamic Epistemic Logic* (DEL, see [14]) is designed to describe the information flow, and adds so-called ‘dynamic operators’ (i.e. operators build out of formulas of the underlying epistemic language) which can express the result of updating an epistemic structure in the above mentioned way.

other facts, i.e. as a result of (ii).¹⁰

Notice that, in the example of the last section, the two questions Holmes may choose do not differ in terms of the conclusions warranted by the ignorance they reveal, with respect to the information Holmes is privy too, and would like the widow to keep on ignoring – namely, that her husband bought a dress. Asking whether Straker bought his wife such and such dress, or not, is compatible with the possibility that Straker bought no dress at all, in the sense that the answer ‘no’ is compatible with the scenario in which Straker bought no dress. Asking the question does not give away Holmes’ (private) information that he did. And asking whether the the widow remembers meeting him or not (while wearing such and such dress) does not give away this information, either. Hence, the difference in questioning strategies must lie in the what they reveal of Holmes’ intended use of the information obtained through these questions.

The kind of information obtained as a result of type (iii) inferences is however rather different. One important feature of inquiry games is the distinction between the *principal* question – to which the hypothesis under investigation is one of the potential answers – and *instrumental* questions, i.e. questions the function of which is to help eliminate certain scenarios in which the hypothesis under investigation does not hold. In inquiry games with strategic sources, sources do not always know which hypothesis is being investigated, and sometimes cooperation can only be achieved provided that this is ignored. Information of type (i) and (ii) is, from a ‘source’ player’s point of view, useful to reduce uncertainty as to the *actual situation* the game is played in. By contrast, information of type (iii) is useful to reduce the (source) player’s uncertainty not only as to which strategy is being followed by the ‘inquirer’ player, but also as to *which game is played*.

If Straker’s widow were to realize that, when inquiring (indirectly) about the dress, Holmes intends to use this information to conclude that the dress was bought for Straker’s mistress, she would be reluctant to answer. From Holmes’ point of view, it seems important that she does not learn that her husband has indeed bought a dress, if she does not know that already. She could then infer that, since he did not buy it for her, he must have bought it for a mistress. Subsequently, she might want to prevent her husband’s adultery to be revealed, especially if she is shrewd enough to anticipate Holmes’ others ‘deductions’. Hence Holmes has to keep his motive as opaque as possible. He has to disguise the purpose of his question, i.e. the aim of the game.

Modeling situations of this kind seems also to require a sophisticated apparatus to represent the distinction between what one knows, and what one is aware of. Anticipating conclusions is sometimes conditional to paying attention to possibilities (scenarios) implicitly known, rather than reducing one’s uncertainty through learning and updating.¹¹ Representing such patterns of reasoning requires subtle distinctions between ‘knowledge’ and ‘awareness’.

¹⁰A simple case is as follows. *X* has submitted a paper to a conference, and meets *Y*, who is in the Program Committee. *X* then asks *Y* whether his paper has been accepted, since the deadline for acceptance notification is passed. Prior to *X*’s question, *Y* has no specific information about *Y*’s paper, but *Y* knows that, presently, all accepted contributors have been notified. Hence *Y*’s background assumptions include the following conditional: ‘if *X*’s paper is accepted, then *X* knows it’. After *X*’s question, *Y* can conclude that *X*’s paper has not been accepted, and answer *X*’s question, though *Y* could not have known that *X*’s paper had been rejected, had not *X* asked the question. (Thanks to an anonymous referee for the example.)

¹¹Indeed, Straker’s widow knows that if her husband bought a dress, and this dress is not for her, then he bought it for his mistress, in the sense that she needs not any outside source of information to answer the question whether the conditional holds, or not. Learning that her husband has bought a dress would allow her to conclude that he had a mistress, provided that she becomes aware of the above conditional. Holmes’ precautions not to reveal the fact that Straker did buy a dress is meant to maintain this unawareness.

Modeling inquiry *games* has to give anticipation of possible inferential moves (availability of some inferential strategy) a decisive role. This requires a departure from standard game theory, insofar as such anticipations, from the part of ‘source’ players, correspond to guessing which game is played. Moving to nonstandard game theory provides in turn an elegant solution to the representation of the distinction between knowledge and awareness, as will be established in the next sections.

3 An Implementation in GT

We propose to model the way an inquirer obtains information by questioning as information-seeking games, closely related to extensive games with unawareness developed in [3]. This modeling provides an interesting explanation of informativeness of ‘yes-no’ questions. In this kind of games, one or several players may not be aware of other players’ possible moves at some nodes of the game. As a consequence, they may not be aware of possible moves of their own. Players can also update their strategy when they become aware of other players’ available strategies.

In our example, this lack of awareness is not only worthy but may be taken advantage of by Holmes, helping him to conceal some of his intentions. This game involves only two players, Holmes, acting as *Inquirer*, and Straker’s widow, the *Source*. Holmes wants to reach a conclusion without his the widow’s knowing it. He thus wants to conceal the purpose of the questions instrumental in establishing this conclusion.

3.1 Information-Seeking And Standard Game Theory

In this section, we propose to model a situation, slightly different to that from **Ex. 1**, by a ‘standard’ extensive game, showing which role considerations about knowledge and awareness (yet in a non-technical sense) play in the player’s choices. The situation differs from the example in that we consider Holmes’ options to be: (i) to ask the widow about her husband’s being unfaithful; and (ii) to ask about her owning a certain dress. Both questions are ‘yes-no’ questions, hence need no prior information to be raised. Let us use the following notational convention:

ϕ = Straker bought a dress, ψ = Straker’s widow possesses such a dress,
 χ = Straker had a mistress

Holmes (*I*) considers premises $P_1 = \phi$ and $P_2 = (\phi \wedge \neg\psi) \rightarrow \chi$: he knows both, and is aware of them. ‘Knowing’ here means simply that *I* has answer to the question whether P_1 holds, or not, and likewise to the question whether P_2 holds. Being aware means that he has considered the question, and knows that he knows the answer.

Straker’s widow (*S*) knows that P_2 , in the above sense: considering the question whether P_2 holds or not, she would be able to answer without any need for further information, that it does. However, she has not considered the question, and hence does not know that she knows. She has not either considered the question whether P_1 holds or not, and does not know that she does not know.

Finally, Holmes wants to know whether or not Straker had a mistress. Obviously, asking the blunt question with presupposition $(\chi \vee \neg\chi)$ could hurt the widow’s feelings, incurring a risk to make her unwilling to answer this question whether she knows the answer or not, or to lie outright, or to refuse to answer further questions, etc. Moreover, the widow is not a reliable source with respect to this question: even if she would remain cooperative, she could perfectly fail to know that her husband had a mistress,

or genuinely believe that he had none. Hence Holmes, as a rational inquirer, will not ask this question.

A more sensible strategy is the following: ask the question with presupposition $(\psi \vee \neg\psi)$; and (ii) if answer $\neg\psi$ is obtained, infer from P_1, P_2 and $\neg\psi$, that χ holds. The widow, in this case, is a reliable source: if she owns the dress, she can't fail to know it, nor could she fail to know that she does not if she does not.¹²

Fig. 1 represents the game corresponding to the situation described above. It requires that I plays first, followed by S , hence the game is sequential, and will be described in extensive form, showing its different possible courses. Since player I and player S do not know the same premises, the information in the game is asymmetric. However, once a move has been made, it is known and recalled by all players. I 's goal is to get an answer to his 'yes-no' question in order to conclude whether or not χ is the case, while S 's goal is to give an answer and guess what I really wants to know. I 's best chance to get an answer seems obviously to ask the question with presupposition $(\psi \vee \neg\psi)$ – denoted $?(\psi \vee \neg\psi)$.

We assume that a true answer yields a payoff of +10 for I , whatever the question is, while other actions of S yields a null payoff for I . S gets a payoff of +10 when she lies, in all cases, because she minimizes the risk given her uncertainty about I 's private information (see below). When she refuses to answer, she gets a lesser payoff, since she can always fear that I might ask another question, and so on.




Figure 1:

As can be seen in the game tree, I gets a positive payoff only in the case of S 's saying the truth, whatever the question asked by him. But, as previously noted, S

¹²Were answer ψ to be obtained, Holmes would have several options. Because of the explanatory value of conclusion χ , he could try to dismiss the answer, i.e. not accept it, and push further his investigation (e.g. by finding another source). This line would be supported by the reasons Holmes has to expect the answer $\neg\chi$ to be more likely than answer χ (Straker has bought the dress using an alias). Or he could try to find another hypothesis to explain Straker's need for money.

has no interest in saying the truth when $?(χ \vee \neg χ)$ is asked. Things are a bit more complicated when I asks $?(ψ \vee \neg ψ)$. In this case, S has no reason to be suspicious about her husband's unfaithfulness, but this question draws S 's attention to the dress.

Assume that asking $?(ψ \vee \neg ψ)$ has the following effect: in addition to learning Holmes's knowledge and ignorance (see 2.3), the widow also realizes that P_2 holds.¹³ Payoffs in **Fig. 1** are given on the assumption that the widow would prefer to lie about her husband unfaithfulness if she knew it, and that she is ready to lie in order to prevent anyone from learning, or concluding, that he had a mistress. Answering – truthfully – that $\neg ψ$ has a different effect, depending on Holmes' knowing or not that $φ$. If the widow's preference not to lie in general, and about the dress in particular, is outweighed by the reasons she may have to prefer keeping her husband's (possible) unfaithfulness secret, then she will prefer to lie about the dress. Intuitively speaking, it means that if the widow is reasonably suspicious of Holmes' private information (as a result of her attention being drawn to the importance of the dress) her best interest will be to lie.¹⁴

If the above representation is all that there is to say about inquiry games, then the problem I faces is to find reliable sources with particular preferences, i.e. preferences that will not interfere with their willingness to answer truthfully. This conclusion is however totally uninteresting. For a proficient inquirer is more likely to try to find a way to 'bypass' a given source's preferences, as Holmes does. And there is no way in the above modeling allowing to understand how he does. So far, we have considered a situation in which: (i) the widow does know which game is played, but (ii) does not know which is the state of Nature. Holmes' strategy, in **Ex. 1**, is indeed to prevent her from realizing *which game she is playing*. Let us now show how to represent this.

3.2 Preventing Strategic Inference

In **Ex. 1**, Holmes asks a question seemingly unrelated to the purpose of his inquiry, as far as the widow knows – for at the moment, she only thinks that he is investigating her husband's murder, rather than his involvement in a criminal activity. Holmes asks the widow whether they already met before. Let us denote this question $?(θ \vee \neg θ)$. The intuitive explanation is straightforward: Holmes must 'disguise' his question about $ψ$ asking whether $θ$ instead, which would seem *irrelevant to the inquiry*. But this question allows Holmes to ask $?(ψ \vee \neg ψ)$, in such a way that the widow's reasoning of the preceding section does not occur. Since the widow will think the question unrelated to $χ$ (and only incidentally related to $ψ$), she has no reason to lie. The task of the present section is to explain how this is possible, and to model the corresponding games.

Holmes' goal is not to obtain an answer to the question $?(θ \vee \neg θ)$, for he already knows the answer ($\neg θ$).¹⁵ After asking his question, and getting the expected (and true) answer $\neg φ$,¹⁶ Holmes follows with a shrewd move: he refuses the widow's answer,

¹³More precisely, we consider the conditional probability that the widow realizes that P_2 holds, given that question $?(ψ \vee \neg ψ)$ is asked, to be 1 (or arbitrarily close to 1).

¹⁴The widow's information does not include $φ$, which can be represented as the result of a move from Nature, occurring prior to the beginning of the game represented in **Fig. 1**. Hence, there is a 'mirror' (sub)game which follows Nature's (possible) move $\neg φ$, i.e. the widow does not know whether she is in one subgame or the other (the game is both with imperfect and asymmetric information). Holmes' question about the dress raises her personal probability that she is in the $φ$ -side of the game rather than the $\neg φ$ side. The widow has nonetheless, in the conditions we assumed, a uniform strategy in this game (i.e. lying about the dress in both subgames).

¹⁵This would make question $?(θ \vee \neg θ)$ an 'exam' question, to be mistaken for a request for information (see n. 7).

¹⁶What if the widow answers $θ$ instead? Holmes could follow with the very same move planned to follow

asserting θ (as if indeed his first question had been a ‘exam’ question) following with the ‘yes-no’ question $?(\psi \vee \neg \psi)$. The rationale for this strategy is as follows: he draws the widow’s attention on a fact of least importance, insisting on it, namely θ , and he asks the question he is really interested (about ϕ) in as a follow-up of the first, *as if it were an exam question* meant to show her that she was mistaken about θ .¹⁷

We must therefore add a third option at the outset, for **Fig. 1** to be a more faithful representation of the ‘game’ of **Ex. 1**. Besides $?(\chi \vee \neg \chi)$ and $?(\psi \vee \neg \psi)$, I has the option $?(\theta \vee \neg \theta)$, where I draws S ’s attention on something completely different from what she might expect, given what she knows about I ’s interests. But this question has another virtue: not only will S ’s attention be diverted, but S will also be unaware of other strategies I could have followed, preventing her from guessing in what game she is. How to represent this unawareness is the subject of the remainder of this section. The new option does not simply add a new subgame, and a complete representation of the game will have to be delayed until the next section. Let us explain why.

In classical game theory, players’ rationality implies that they are all aware of the structure of the game, and such an awareness is common knowledge. Common knowledge of all possible moves in the game is not a reasonable assumption in games of inquiry, precisely because some possible moves need to be hidden until they are performed (thus remaining unattended if they are not), to prevent the source from planning in advance her answers, in case it would impair the conduct of inquiry. Following some of the developments presented by Halpern and Rêgo in [3], we propose to introduce the notion of *awareness* – and *unawareness* – as central to information-seeking games, making them a special case of *extensive games with unawareness*.

An extensive game with unawareness represents what players are aware of at each node of the game, resulting in what Halpern and Rêgo call an *augmented game*. An augmented game shows the way a player views the game at a given stage, that is what she believes to be the true game in that situation. An ‘augmented game’ is thus a *trimmed* game tree, compared to the complete (standard) representation of the corresponding extensive game. Thus, the complete description of a game *with unawareness* is constituted by a set of augmented games, one for the modeler and one for each game other players think they might be in at some stage of the game.

Since an augmented game is built after a standard extensive game, let us briefly recall the definition of such a game¹⁸: an extensive game G with perfect information is a structure $G = (N, H, Z, P, (U_i)_{i \in N})$ where:

1. $N = \{1, \dots, n\}$ is a set of players;
2. H is a set of sequences (histories) of the game (whith $Z \subset H$ the set of terminal histories of the game);

answer $\neg \theta$, for the widow could be mistaken, rather than lying. In the former case, he could obtain (as expected) $\neg \phi$ (for she could not be mistaken about the dress). Would Holmes obtain answer ϕ to the second question, he would take it as an indication that the first answer, θ , was equally likely to be mistaken, as to be deliberate. (remember that Holmes thinks unlikely that the widow owns the dress). In this case, he would certainly try to find another source, everything else having failed. But this would have less to do with the widow’s willingness to keep her husband’s memory spotless, than with the widow’s being an absent-minded person, or a pathological liar.

¹⁷Actually, in the short story, Holmes describes his ‘move’ as a question, when exposing his solution to the owner of Silver Blaze, and employer of Straker. But what he has done was to insist on ψ , just as if asking the question and giving the answer himself, only to be ‘corrected’ by the widow. This is however equivalent to asking the question about ψ and using the widow’s answer, though it certainly lowers even more the risk of the widow’s second-guessing Holmes’ intent.

¹⁸We follow the definition of [12].

3. $P : H \rightarrow N$ is a function which assigns to each non-terminal history a member of N ;
4. $U_i : Z \rightarrow \mathbb{R}$ is a utility function for each player $i \in N$ which associates a payoff to each terminal history of the game.

An augmented game G' is based on a game G and augments G by “describing each agent’s awareness level at each node, where [a player’s] awareness level at a [given node] is essentially the set of runs (complete histories) in G that [the player] is aware of at [that] node.”¹⁹ What essentially distinguishes an augmented game from a standard extensive game is that at each node, it is not only specified which player’s turn it is, but also what is that player’s awareness level. Intuitively, this specifies the game the player believes, at that node, to be playing. ▽
△

3.3 A Partial Model

If we return to our first representation of Holmes’ questioning strategy, we have seen that we could add to I ’s first move a third branch, representing his option to ask the question $?(\theta \vee \neg \theta)$. However, before introducing considerations of awareness, it was not sufficient to model **Ex. 1**. **Fig. 2** represents the information-seeking games *with unawareness*, where unawareness of (possible) moves is represented by indexing each move player S is unaware of with $*_S$. We have thus the description of the game from the modeler’s viewpoint. (For sake of simplicity, we assume that modeler’s representation of the game and I ’s one is the same.)

After I ’s added move $?(\theta \vee \neg \theta)$, S has, as always, three options,,: saying the truth, lying or saying nothing. But as shown in the game tree **Fig. 2**, this does not affect I ’s next move – since Holmes’ strategy in **Ex. 1** can be carried even in case of other answers. Indeed, I feigns to takes answer θ for granted, and follows with a new ‘yes-no’ question $?(\psi \vee \neg \psi)$, as if it were an ‘exam’ question (stressing answer ψ as expected).²⁰

As we said previously, the utility function associated with each player’s preference ordering is not the same. Concerning the preferences ordering of I , we assume that he always prefers to get the truth, then to have no answer, a lie being thus what he dislikes most. Besides, the cost for a ‘yes-no’ question is negligible, unless the ‘yes-no’ question is a disguised one, since in this case I needs to find the best way to hide his intention, possibly by adding a ‘yes-no’ question to the first he asked. This cost is however balanced by the increased value of the information obtained, and we simply assume that they cancel out.

Representing S ’s preferences a trifle more complicated. Indeed, S may be assumed to have a preference for telling the truth over not answering, and for the latter over lying. But she certainly prefers to lie, and then not to answer, if she believes that truth might upset her, i.e. if she has a reason to think that I has concealed his intention or if the question draws her attention on an unpleasant fact (e.g. the fact that her husband might have had a mistress), as we already said. We therefore represent in the game-tree payoffs values that reflect these preference orderings.

The game in **Fig. 2** represents the description of the game from the point of view of the modeler (and I), after Nature’s first move. Adding the branch $?(\theta \vee \neg \theta)$ cuts off

¹⁹[3].

²⁰In Conan Doyle’s short story, this corresponds to Holmes’ waiving the widow’s first answer, insisting that he perfectly remembers the dress. As we mentioned in n. 17, Holmes does not ask a question, but his strategy is equivalent to putting a ‘yes-no’ question to the widow.




Figure 2: The modeler's game

S 's possibility to become aware of other moves available to I at first stage of the game. Asking $?(\theta \vee \neg\theta)$ seems irrelevant to the inquiry.²¹

Informally speaking, asking this question makes S believe she is in another situation, hence playing another game, unrelated to I 's inquiry (given what she knows about the latter). Without going into the detail of representing S 's beliefs, **Fig. 3** shows what game she believes to be playing, after I 's first move at $?(\theta \vee \neg\theta)$.²²

Because S thinks she plays the game of **Fig. 3**, and following her (unaltered) preference ordering, S has no incentive to lie or to refuse to answer. The 'yes-no' question $?(\theta \vee \neg\theta)$ 'signals' indeed (for S) a situation unrelated to the current inquiry (as determined by the background knowledge of S). The strategy S then follows is only related to this local state (local situation) she believes to be the actual situation. Therefore we can say that, in the augmented game G^S , S follows a *local strategy*.

In a standard extensive game, a strategy for a player $i \in N$ is a function that assigns an action a in $A(h) = \{a : (h, a) \in H\}$ to each non terminal history $H \setminus Z$ for which $P(h) = i$. The important point here is that a strategy specifies an action for each non

²¹We implicitly assume that the game between I and S is played in a broader context, where S has some expectations about I 's demand for information, just as Straker's widow has expectations about Holmes' interests, in relation to the case of her husband's death. Asking about a former meeting, or, in the game of **Fig. 2**, choosing option $?(\theta \vee \neg\theta)$, gives the impression that Holmes (I) has shifted to another set of (unrelated) interests.

²²Recall that the awareness level of a player at a given node depends on the information available to this player at that node, including information about other player's move. This is readily represented by excluding from S 's awareness the possible histories following other possible initial moves of I .




Figure 3: The game from S 's viewpoint

terminal history after which it is i 's turn to move, which implies that i is aware of all possible moves in the game, even if some terminal histories are not effectively reached. This no longer holds in a game with awareness. In this case, the strategy of a player is a collection of local strategies, one for each augmented game depending of his awareness level.²³

As a last remark, notice that in the game shown in **Fig. 1**, there was no Nash equilibrium with pure strategies. On the contrary, in the game shown in **Fig. 3**, there are several Nash equilibria, one of which is Pareto-optimal ($\langle ?(\theta \vee \neg\theta), T, ?(\theta \wedge (\psi \vee \neg\psi)), T \rangle$). In the end of sec. 3.1, we suggested that an inquirer's problem is not to find appropriate sources, but to adequately use the available sources. This last remark suggests how this translates in game-theoretic terms. One of the most useful skills of a proficient inquirer seems to be the ability to *design games in which the source's anticipations are limited to local strategies*.

4 Concluding Remarks

So far, logical modeling of epistemic states has focused on dynamics of information, as affecting the accessibility relation over a space of scenarios (worlds, alternatives, etc.). The effect of demand for information expressed by questions, as part of reasoning strategies of agents, and the way it shapes expectations and strategies of other

²³For technical details, see [3].

agents has not received proper attention. We have argued that, in some cases, demands for information relative to strategic aspects of reasoning are important enough to be integral part of models of interaction. This supports the relevance of epistemological models representing explicitly research interests.

Several formal frameworks have been proposed in which demand for information plays a role in representation of the epistemic or doxastic state of an agent. In Hintikka's IM , this can be represented by the Inquirer's *Range of Attention*, i.e. the additional ('yes-no') questions she is prepared to introduce in an interrogative argument (see [5]). In I. Levi's model, an inquirer's demand for information is represented by an *ultimate partition*, which specifies the range of relevant expansions of the current 'belief state' of the inquirer (see [10], ch. 2). Levi discusses cases where several inquirers can agree on a common ground (basic corpus) and a common demand for information (basic partition). A related proposal is E. J. Olsson's introduction of *research agendas* to the representation of doxastic states (see [11]). The two latter proposals are related to Belief revision Theory and the AGM framework (albeit both Levi and Olsson adopt a critical stance towards it). Relations between the three approaches are briefly discussed in [2]

When inquiry is such that it should be modeled as a *non-cooperative game*, inferences of agents about other agent's demand for information matters as much as inference from their ignorance to state of Nature. The latter has received deserved attention in the literature, while the former has been almost totally overlooked. We have proposed to identify this information as strategic information, to treat it through extensive games with unawareness. An interesting open problem is to represent in epistemic languages this type of information, since the dynamic logics for belief and knowledge update which have recently been flourishing are not designed to accommodate this phenomenon.

Inquiry games, or information-seeking games, as developed by Hintikka, are described as games against Nature, hence are degenerate games, game-theoretic representations of decision problems. Though the possibility of studying game-related aspects, such as equilibria concepts, is offered by such a modeling, this possibility has hardly been explored. An exception is [4]: providing a game-theoretic account of interrogative tableau systems, the paper discusses Nature's 'cooperative' ('non-cooperative') strategies – as maximizing (minimizing) the likelihood of tableau closure. Though this is a first step toward a more-than-analogical relation between inquiry games and game theory, the scope of this approach is not clear, given that Nature is a non-strategic player.

Our example shows that inquiry games can be generalized to more than one players (vs. Nature), but also that standard game-theory is ill-suited to model such games. Our treatment yields however an unexpected vindication of Hintikka's contention that inquiry is a game against Nature. Since Nature is a non-strategic player, Nature's 'strategy' does not depend on the Inquirer's goal. Any strategic player who cannot anticipate an Inquirer's demand for information in a way that would affect the player's disposition to answer, is in this respect just as non-strategic as Nature is.

A competent inquirer has often to devise clever means to maintain a 'veil of ignorance'. Keeping the anticipations of an agent, to whom questions are asked, as 'local' as possible is the daily bread of inquirers of all ilks, be they pollsters, sociologists, or criminal investigators.²⁴ Likely, in proper (non-degenerate, multi-player) inquiry

²⁴Keeping *investigators'* expectations is sometimes crucial also, as illustrated by double-blind protocols.

games, the aim of Inquirer is to play a game as close as possible to a game against Nature.

References

- [1] Conan Doyle. *The Complete Sherlock Holmes*. Doubleday, 1960.
- [2] E. J. Genot. The Game of Inquiry. Unpublished manuscript (Preliminary version presented at the LOCCOL'08 Conference (Prague), 10–14 September 2008), 2009.
- [3] J. Y. Halpern and L. Chaves Rego . Extensive Games With Possibly Unaware Players. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, Hakodate, Japan*, pp. 744–751, 2006.
- [4] S. Harris. GTS and Interrogative Tableaux. *Synthese*, 99:329–343, 1994.
- [5] J. Hintikka. Reasoning About Knowledge in Philosophy: the Paradigm of Epistemic Logic. In *TARK '86*, pp. 63–80. Morgan Kaufmann, 1986.
- [6] J. Hintikka. The Interrogative Approach to Inquiry and Probabilistic Inference. *Erkenntnis*, 26:3:429–442, 1987.
- [7] J. Hintikka. *Inquiry as Inquiry: A Logic of Scientific Discovery*. Kluwer, 1999.
- [8] J. Hintikka. *Socratic Epistemology*. Cambridge University Press, 2007.
- [9] J. Hintikka, I. Halonen, and A. Mutanen. Interrogative Logic as a General Theory of Reasoning. [7], pp. 47–90.
- [10] I. Levi. *Mild Contraction*. Oxford University Press, 2004.
- [11] E. J. Olsson and D. Westlund. On the Role of Research Agenda in Epistemic Change. *Erkenntnis*, 65:165–183, 2006.
- [12] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [13] R. Smullyan. *First-Order Logic*. Springer, 1968.
- [14] H. P. van Ditmarsch, W. van der Hoek, and B. P. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.

But in these cases, 'field' investigators are themselves sources of answers.