Rational Ignorance

Author(s): Sylvain Bromberger
Source: *Synthese*, Vol. 74, No. 1, Knowledge-Seeking by Questioning, Part I (Jan., 1988), pp. 47-64
Published by: Springer
Stable URL: https://www.jstor.org/stable/20116485
Accessed: 19-09-2018 12:11 UTC

SYLVAIN BROMBERGER

RATIONAL IGNORANCE

1. INTRODUCTION

Different people know different things and are ignorant about different things. For instance, much that inhabitants of Massachusetts know and don't know is different from what inhabitants of Provence know and don't know, and much of what physicists know and don't know is different from what linguists know and don't know. These differences are to a large extent traceable to differences in circumstances and opportunities. But some are traceable to deliberate choices: people can and do choose to find out certain things and to remain ignorant about others. And they don't all make the same choices, even when the same choices are available to them.

In this paper, I want to explore some aspects of these choices, and of the constraints that limit them. My ultimate reasons for being interested in these choices have to do with concrete issues in epistemology and the study of cognitive development. I will say a few words about those before concluding. But they will not be my major focus. For the time being, I want to concentrate on preliminary conceptual matters that need to be appreciated before we can make real progress with the more concrete problems.

Let me set the stage with a few stipulative definitions.

(1) *IGNORANCE* is the relationship between a person P and a set of questions Q when P does not know the correct answer to any of the members of Q and has no strong views as to what the correct answer to any of them is.[1]

(2) A *PERSON'S* (*P'S*) *IGNORANCE AT TIME* t is the set of questions that is the maximal Q_t to which P stands in the ignorance relation at time t. (I use 'Q_t' as a variable ranging over sets of questions that constitute someone's ignorance at time t.)

(3) A person is *LESS IGNORANT AT TIME* t_2 *THAN AT TIME* t_1 if and only if $Q_2 < Q_1$, i.e., if and only if there is at least one question that is a member of Q_1 but not of Q_2 but not vice versa.

Synthese **74** (1988) 47–64

A decrease in P's ignorance between t_1 and t_2 can come about by chance or it can come about as a result of deliberate action. Deliberate action requires two stages: first, the selection of a question from Q_t, and then, secondly, the carrying out of whatever activities are required to get the answer to that question.

(4) The first stage of a decrease of ignorance is *RATIONAL* if and only if it consists in selecting a question in a rational way, i.e., with a view to maximize values and to minimize costs.

(5) A person P is a *RATIONAL IGNORAMUS AT TIME t* if and only if at time t, P has and uses a rational policy for selecting the next question to be eliminated from P's ignorance.

In what follows,[2] I will explore implications of these definitions. How much and what sorts of things must one know to be rationally ignorant? How do what values come into play in the choices of a rational ignoramus? What problems must a rational ignoramus solve?

2. KNOWING WHAT ONE DOES NOT KNOW

Since a rational ignoramus must deliberately select from his ignorance at time t one question for elimination, he must be able to survey the membership of his ignorance. How else can he establish his preferences? But to be able to do this, our ignoramus must be quite knowledgeable!

2.1. *Representation and Understanding of Questions*

A rational ignoramus must be able to represent questions, i.e., be able to formulate and to understand interrogative sentences in some language.[3] Offhand, that may not seem to be a very constraining requirement. Most people we know will easily meet it. Even so, we should be mindful that it is a very complex requirement. In fact, at present, no semanticist can fully describe it.

To get an inkling of the difficulty – and without even considering what is required to be able to formulate questions – compare what is involved in understanding a declarative sentence with what is involved in understanding an interrogative sentence. To understand a declarative sentence one has to be able to determine what semanticists call its *truth conditions*, i.e., what has to be the case if the sentence is to be true. What that involves is not understood completely either, but at least there are a number of theories that give us a very good inkling. But what constitutes understanding an interrogative sentence? What constitutes understanding

(1) What is the distance between London and Paris?

or Heidegger's beloved question

(2) Why is there something rather than nothing?

It can't be knowing truth-conditions, since questions do not have truth conditions.

Could it be knowing (or being at least in a position to determine) the truth conditions of every possible[4] answer to the question? For instance – in the case of the first example above – could it be knowing the truth conditions of

(3) The distance between London and Paris is one inch.
(4) The distance between London and Paris is two inches.
(5) ....

and so on *ad infinitum*?

That hardly seems plausible: we do understand many questions whose possible answers we are in no position to state. Think of the second example above. Or for that matter think of

(6) What does one need to know in order to know that one does not know something?

How many of us could state all the possible answers to that one, even if we had all the time needed to utter them? And in any case, knowing these truth conditions would not be enough. We would also have to know that each possible answer *is* a possible answer. And that is very different from knowing truth-conditions. Thus understanding an interrogative sentence cannot come down to knowing the truth conditions of each of its possible answers.

2.2. *Knowing of Each Item of Ignorance That It Is a Bona Fide Item of Ignorance*

Our ignoramus must not only be able to formulate and understand questions, he must know that some of them represent genuine items of ignorance, i.e., he must know or have good reasons to believe that he does not know their answer. To know that one does not know that, one has to know quite a bit!

First one has to know, or at least believe, that the *presupposition* of the question is true. Let me explain with an example:

(7) What is the age of the King of France?

That question does not represent a genuine item of ignorance. No one does not know what the age of the King of France is. Not because everyone knows: no one knows what the age of the King of France is either, and not because it is a universal mystery. But because there is no King of France. The question presupposes that there is a King of France, and because there isn't, the question has no answer.[5]

Contrast that with

(8) What binds neutrons and protons in an atomic nucleus?

It presupposes that there are atoms, protons, and atomic nuclei. Since that is true, the question can – in my case does – represent a genuine item of ignorance.

Put somewhat roughly, the presupposition of a question is a proposition to the effect that some thing or things of a certain character exist that the question is about.[6] To know a question represents a *bona fide* item of ignorance, one must thus not only know that one cannot answer it, one must know that its presupposition is true.

Secondly one must know that *the question arises*. What I have in mind here is – like so many of the other topics in this discussion – complicated to analyze, but easy to illustrate. Consider

(9) How long is the number 5?

or

(10) What is the main verb of an electron?

or

(11) What is the square root of the Empire State Building?

None of us, I presume, can give the answer to any of these questions, not because we do not know the answer, but because there are no right or correct answers. On the other hand there is presumably an answer to the question

(12) With how much force does the earth attract the moon?

though once upon a time that question would have seemed crazy to people.

Knowing that a question arises – as I use the term – requires knowing general principles of a peculiarly abstract kind, e.g., that all objects attract each other with a force that has magnitude and direction (hence the question "What is the force with which the earth attracts the moon?" arises); that atoms have mass (hence the question "What is the mass of a lithium atom?" arises); that every event has a cause (hence the question "What was the cause of the fire?" arises for every fire); that objects have inertia (hence "What force is responsible for change of direction or velocity?" arises in connection with every deviation from straight uniform travel). Let me call such principles *question raising principles*.

Knowing that a question arises is very similar to knowing that the presupposition of the question is true. The difference lies in the fact that knowing that a question arises requires knowing general principles of a relatively high level of abstraction, whereas knowing that the presuppositions hold requires knowing that a certain thing or certain things exist (or that certain events have occurred, or that certain facts hold).

Thirdly, one has to know that one does not know the answer. That covers many different kinds of situations. There is the kind of situation in which one is able to cite many possible and plausible answers, but has absolutely no reason to prefer any of them (has no evidence that favors one above the others). That is probably how most of us stand with regard to

(13) What is the distance between London and Paris?

or

(14) What will be the name of the last person admitted to next year's freshman class of M.I.T.?

To recognize instances of it one must – among other things – be able to evaluate and to compare evidence.

Next, there is the kind of situation in which one does have evidence that makes some answer more likely than others, but the evidence is not conclusive, e.g.,

(15) What will the weather be like tomorrow?

To recognize that sort of situation one must be able to determine when evidence is conclusive, or at least strong enough to warrant the belief that one knows the answer.

Next still, there is the kind of situation in which one can come up with only answers one knows to be false and yet has reason to believe that the question is sound, i.e., arises and has true presuppositions. Until recently no one knew why the moon looked larger when near the horizon than when right overhead. A number of theories existed, but each was refuted by good evidence. To recognize that type of situation one must have all the knowledge mentioned above and furthermore be able to determine that a certain set of possible answers is not complete even though one cannot extend it.

2.3. *Knowing Enough About What One Does Not Know*

We are examining what an ignoramus must know in order to be able to function as a rational ignoramus. We have said that he must be able to represent items in his ignorance. We have said that he must know that these are items of ignorance. But of *how many* items of his ignorance should he know this? How large a list should he have? Which questions from his ignorance should it contain?

Unless rationality can require the impossible, it cannot be a complete list. It cannot contain every member of his ignorance. That follows from fairly straightforward considerations. There are an infinite number of objects, events, and states of affairs whose existence any individual is doomed not to know about. Each of these gives rise to an indefinitely large number of questions. And he is thus doomed to remain ignorant about the fact that he does not know the answer to those questions.

It also follows straightforwardly from the fact that the cardinality of any expressible list cannot be as large as the total number of objects about which questions arise. There are nondenumerably many real numbers. About each the question, e.g.,

(16) What is the largest prime natural number smaller than it?

arises. But of course these cannot all be formulated in one language, since at most denumerably many questions can be formulated in a single language.

Finally it follows from the fact that many questions implicate concepts of a very special kind, i.e., determinables[7] of which the following are examples: 'distance in inches', as in

(17) What is the distance in inches between John's two nostrils?

or 'mass', as in

(18) What is the mass of the moon?

or 'D-structure', as in

(19) What is the D-structure of 'Jean est facil à convaincre'?

'Voltage', 'chemical composition', and 'anatomical structure' are further examples. Determinables can express functions that map objects on questions. But to grasp a determinable one has to have a fair amount of theory. Thus, before Newton's gravitational theory became available, no one was in a position to entertain the questions that we are able to express in English by using the expression 'gravitational mass'. Aristotle – rational though he was – not having the concept of mass, would not have been in a position to know that he did not know the answer to

(20) What is the mass of the Parthenon?

or to include it in a specification of his ignorance. It seems unlikely that anyone will ever grasp all the determinables that generate sound questions about actual objects, events, or states of affairs.

To what subset of his ignorance must an ignoramus then apply his policy to be deemed rational? The enumerable subset consisting of all and only the questions that he is in a position to formulate and to entertain (i.e., whose presupposition he knows to be true, and whose determinables he has grasped, etc.)? But the construction of such a list would require an infinite amount of time, since the list would include, e.g., questions about each natural number, i.e., the list would be infinitely long: rationality would defeat itself by preventing itself from getting off the ground!

There may be a way out of this paradox.[8] But I do not know what it is. So let us set it aside for the time being by dropping the notion of rational ignoramus at time t as too strong for present purposes, and let us replace it with a weaker notion: that of a rational ignoramus relative to a set of questions Q at time t:

(21) A person P is *a rational ignoramus at time t with respect to a finite set of questions* Q if and only if, at time t, Q is an *accessible*[9] proper subset of P's ignorance, and P has a rational policy for selecting the next question in Q to be eliminated from P's ignorance.

Let us now think about that rational policy. It is presumably designed to maximize certain values and to minimize certain costs. What value should it seek to maximize?

3. THE VALUE OF QUESTIONS AND THEIR TYPES

3.1. *Intrinsic Value and Added Value*

The correct answer to a question can be of value to us for many different reasons. For present purposes I will sort these under three headings. It may provide us with an *intellectual* benefit, i.e., satisfy our curiosity. There are different degrees of curiosity. And hence, from that point of view, answers to different questions may turn out to have different values. Let us call that the *gosh value* of an answer, a measure of the intellectual pleasure derived from coming to know the answer or maybe the relief of no longer not knowing the answer.

We may also prize the answer to a question because it enables us to construct something, to repair something, to find something, to obtain some *material* benefit. Let us call that the *cash value* of the answer.

I have defined the gosh and cash value of an answer in terms of certain properties of that answer. Of course one may not know the value of the answer before coming in possession of it or while still contemplating the question. Yet estimates of such values should obviously play a role in any policy designed to select questions to be answered. To refer to such estimates, when we talk of *the gosh and cash values of a question* (rather than of an answer), let us mean the estimated value of the gosh and cash value of the answer to that question.

Let us call the sum of the gosh value and cash value of a question the *intrinsic value of the question*. Intrinsic value is obviously a notion that has to be relativized to individual preferences, to time, to background knowledge, and perhaps to other circumstances. So when we talk about the intrinsic value of a question, let that be shorthand for the intrinsic value of that question as estimated by our hypothetical rational ignoramus in the light of his circumstances at a time when he is using his selection policy.

3.2. *Added Value*

Besides having gosh and cash value, an answer may be of further value because it puts one in a position to obtain answers to other questions with values of their own. Thus if one knows the width and the length of a floor, and one knows the formula for computing the area of a rectangle one is thereby put in a position to find the answer to "What is the area of the floor?" If the latter is of value, then surely some of that value should accrue to the answers to "What is the length of the floor?" and to "What is the width of the floor?" Let us call that the *added value of the answer*. And let us use *added value of the question* as shorthand for the estimated value of the answer relativized for preference, time, background knowledge, and other relevant circumstances.[10]

3.3. *Value Adders*

Questions get added value through devices and algorithms, that is, *value adders*, that transform their answers into answers to other questions. Let us look at these value adders for a bit. They come in various shapes and forms.

Consider

(22) All men are mortal

and

(23) Is Fido a man?

If the answer to (23) is "yes" then (22) puts one in the position to compute the answer to

(24) Is Fido mortal?

If the answer is "no" then it does not. That sort of a value adder is a *gappy* value adder, or at least gappy for a certain question. It generates added value to that question only for some answers and not for others. What the correct answer is matters to it. If it does not like the answer, it does not add value.

Contrast that with the pair

(25) The distance between A and B is twice the distance between B and C.

and the question

(26) What is the distance between A and B?

Here (25) puts one in the position to compute the answer to

(27) What is the distance between A and B?

no matter what the answer to (26). It adds value to the question (26), no matter what the answer turns out to be. Let us call that sort of value adder an *erotetic* value adder relative to a certain question.

But now consider the formula for the pendulum

(28) $T = 2\sqrt{1/g}$.

It will add value to the question

(29) What is the length of the pendulum?

for any pendulum under consideration that is influenced only by local gravitational force. But it is much more powerful than that. It will add to the value of any question (and no matter what the answer) constructed on the determinable

(30) What is the length of ... (name a pendulum) ... ?

the value of the corresponding question of the form

(31) What is the period of ... (name same pendulum) ... ?

Let us call that sort of value adder *super-erotetic* value adders.[11] We have here a hierarchy that can clearly be extended further. Thus differential equations of a certain kind will generate super-erotetic value adders, and theories – or at least certain kinds of theories – will in turn generate super-super-erotetic value adders.

3.4. *Science and Value Adders*

Many contributions to science are designed – either explicitly or implicitly – to provide us with value adders. The value adders they provide can vary in attractiveness along a number of dimensions: they can vary in reliability; they can vary in precision; they can vary in subject matter; and of course, they can vary in power along the hierarchy that I have just sketched. They can vary in the amount of gosh and cash value of questions under their jurisdiction. The practical sciences are built around a concern about questions with high gosh and cash value. Theoretical sciences seek powerful value adders. Primitive sciences tend to wallow at the level of gappy value adders of low precision, questionable reliability, remote subject matter, and little cash value, but very high gosh value. Theology, much psychology, a fair amount of what comes under the heading of semantics and pregenerative linguistics belong here. Of course, this does not mean that they should not be pursued: low level sciences have a way of blooming into high level ones. And high level sciences like physics, biology, chemistry, and more and more of linguistics of the right sort, spew forth value adders that shine in all directions: precision, reliability, power, and intrinsic value of many of the questions in their domain.

But value adders vary not only along the valuational dimensions just described. They also vary drastically in ontological character and in the competences that they demand of their users. To see this, let us look briefly at some banal examples.

3.4.1. ***Models***. What is a model? Think of small models of a big airplane, or of plastic models of DNA molecules, or of maps. What makes one of these objects a model of another? The fact that by finding out certain things about the model one can then find out certain things about the thing modeled. By counting the stacks on the plastic model of the Queen Mary one can find out the number of smokestacks on the real Queen Mary. By counting the number of black spheres on a wire and wood model of methane, one can find out the number of carbon atoms in a real methane molecule. Artificial models are designed with just that sort of relationship in mind. To fix ideas, think of maps and of the territories which they represent. By measuring distances on a map one can find the answer to the following sorts of questions: "What is the distance between such and such a point on the map, and such and such another point on the map?" Anyone versed in map reading can translate the answer to such questions about the map into questions about the territory mapped, e.g., "What is the distance between such and such a place in the territory and such and such another place in the territory?"[12]

Thus, as a first approximation, we might define a model as follows:

(32) M is a model of O relative to a set of triples $\langle Q_m Q_o A \rangle$ if and only if in each triple Q_m is a set of questions about M, Q_o is a set of questions about O, and A is an algorithm that translates any answer to a member of Q_m into an answer to a member of Q_o, and *correct* answers to the former into *correct* answers to the latter.

In the case of maps as in the case of other models, the set Q_m and the set Q_o are sets whose members are built around common determinables (e.g., 'distance'). But that is not an essential feature. Population graphs, for instance, do not have it.

Maps are clear instances of super-erotetic value adders, and so are other models that conform to our definition of a model. This definition is almost identical to that of a super-erotetic value adder.

Maps and other models are physical super-erotetic value adders[13] but there are other kinds of physical super-erotetic value adders, such as gauges. A pressure gauge, for instance, is a device that enables one to obtain the correct answer to instances of "What is the pressure of ... (name a sample of gas to which gauge is connected) ... at ... (name a time) ... ?" from instances of "What is the position of the needle at ... (name the same time) ... ?".

But there are also ways of conceptualizing things that yield super-erotetic value adders. I'll briefly mention two.

3.4.2. *Natural Kinds*. A class of objects constitutes a natural kind only when each member can serve as a model for every other member. In other words, associated to every natural kind, there is (or is presumed to be) a set of *projectible questions*, i.e., when the answer to one of these questions is found for one member of the kind, one obtains the answer to the same question for every member of the kind, and the answer is the same.[14] For instance, by finding the answer to

(33) What is its boiling point under standard conditions?

for one sample of water, one *ipso facto* gets the answer for every sample of water. And the same goes for

(34) What is its density?
(35) What is its freezing point?

3.4.3. *Categories*. Natural kinds arrange themselves into categories. For example, tigers are animals, a *category* that includes cats, dogs, and earthworms, as well as tigers. Samples of water are samples of substances, and substances include gold, alcohol, iron, cellulose, etc. Tokens of the word 'Apotheosis' are tokens of an English word. English words include 'dog', 'cat', 'apocalypse', 'Ticonderoga'. Members of such categories share projectible questions. Thus every animal kind shares with every other animal kind the projectible question 'What is its anatomical structure?' though, of course, different kinds get different answers. Every substance kind shares with every other kind the projectible question 'What is the boiling point under standard conditions?' (but gets a different answer).

It is at the level of categories that formulae come into play. Consider the formula governing simple pendula. Simple pendula that have the same length and period can be viewed as forming a natural kind, and all simple pendula together, collected regardless of their length, can be viewed as forming a category; the formula is then a device that enables us to compute the answer to one projectible question for any pendulum whatever its kind, from the answer to other projectible questions about the same pendulum. The formula covers the entire category. Of course, we do not usually think of simple pendula as forming natural kinds: nature has created very few natural simple pendula. But formulae that enable us to compute the boiling point for any substance from the structure of its molecular bond, or to compute across species some traits of plants from other traits, function essentially in the same way as the formula of the pendulum.

4. TYPES OF COSTS

Let us return to our ignoramus. One might assume that his rational policy would select questions to be answered according to their total (gosh plus cash plus added) value, i.e., go for the ones with greatest total value. But that is too simple a view. It overlooks at least two likely situations: (a) the answers to the questions with highest total value may be beyond reach; (b) even if they are within reach, the cost attached to finding them may be too high.

The first situation will obtain whenever the answers cannot be had "directly" by observation, or through intuition, or by asking someone, and when no appropriate (i.e., sufficiently reliable) value adders are available that will produce them from answers that can be had directly, or that are otherwise forthcoming. A rational policy will obviously pass over such questions, regardless of their total value.

This fact, by the way, puts quite a burden on our rational ignoramus; it requires him to find out whether the answer to the most highly valued question is obtainable, and if it is not, to go on searching until he comes to a question whose answer is obtainable. How is he to determine the likelihood of getting an answer? And how is he to estimate the value of these *subsidiary* questions about questions (and their cost) raised by his very attempt to use his policy? And how are *we* to analyze the place of such higher-order questions in a rational scheme?

The second situation is created by the existence of costs. That is a very messy topic. There are many different kinds of costs: cost in time required for retrieval of information, for computation, for mastering value adders; financial resources needed for instrumentation, for assistance; and emotional costs such as boredom, anxiety, or frustration. How these costs are to be measured and compared is far from obvious. Their bearing on the formulation of a rational policy, on the other hand, is all too obvious.

5. M.I.T.

A rational ignoramus, as described so far, is someone engaged in making decisions based on assessments of risks, assessments of gains, and assessments of costs, in short, someone engaged in a game-theoretical situation. Simulating his situation as an actual game may therefore help us see a little more clearly the strategic demands on a rational policy. The name of the game is clear: it must be either *Minimizing Investigative Travails* or *Maximizing Investigative Trade-offs* but in any case, be a name that abbreviates as *M.I.T.* Other aspects of the simulation are less obvious. Here is one way of picturing them.

Think of the game as played by two agents, *A* and *B*. *A* is pitched against *A* and against luck. *B* plays a subsidiary role.

M.I.T. is won or lost in sets.

- A set consists of *n* games.
- At the beginning of each game, *A* announces a question.
- During the game *A* can do a number of things: (1) He can put questions to *B* – these may not include the one he announced at the outset of the game, or of any previous game, or of any future game, and must pertain to specific individuals, or events, or state of affairs. (If *B* answers at all, *B* answers truly. And *B* answers whenever he can.) Each question costs *A* 10 points, whether *B* answers it or not. (2) He can buy value adders. Gappy value adders cost 100 points each; erotetic value adders cost 1000 points each; super-erotetic value adders cost 10,000 points; theories cost 100,000 points. (3) He can compute; cost of computation is 1 point per minute.
- A game ends as soon as *A* has figured out the answer to the question he announced, or after one year, whichever comes first.
- The score for the game is computed as follows: if *A* gets the answer to his question, he scores the intrinsic value (i.e., gosh plus cash value) of the question minus the costs he has incurred. If he does not get the answer to his question, he scores a negative value amounting to the costs he has incurred during the game.
- The games are not completely independent. When *A* starts a new game, he can come equipped with all the answers and all the value adders he has accumulated in previous games.
- If at the end of a set, *A* has scored 100,000 points, (i.e., the equivalent of a theory) he is declared a winner and gets a prize called ***tenure***, and playing comes to an end or goes on, but for higher stakes (e.g., an endowed chair, a Nobel prize). If he gets less than 100,000 points he must go on playing another game, only this time he has to be the research assistant, and *B* gets to ask the questions.

M.I.T. is too simple and unrealistic to represent adequately the strategic decisions faced by a real rational ignoramus even at real M.I.T. But it is complex enough to indicate the network of problems that his policy must handle, the different roles that knowledge will play in its execution, and to bring out that the best strategy in most situations will be far from obvious.

6. WHY BOTHER WITH ALL THIS?

The time has come to stop wondering about rational ignoramuses and to start acting like one. And so let us ask: of what value are the questions that we have been asking? What could their gosh or cash value possibly be?

As far as I can see, it is not very great. Nor do I think that these questions are likely to accrue added value from elsewhere. Their value is of a different nature. The value of their answers lies in their potential as question *generators*! They will generate at least two sets of questions with high gosh value.

First, they generate a set of questions in philosophic epistemology. Does our rational ignoramus embody plausible normative aspirations, i.e., does he represent cognitive traits and intellectual attitudes that we aspire to embody, and does his policy exhibit standards by which we measure aspects of our worth and that of other human beings? Or does he represent an unattainable, illusory, and uninteresting ideal of no relevance to a responsible epistemic conscience?

Second, they generate a set of cognitive-psychological questions. Are human beings rational ignoramuses of a sort at points of their development? More specifically, are infants rational ignoramuses? To conjecture that they are is to conjecture that at various stages of their development they are endowed with representations of part of what they do not know, with weighings attached to some of these representations, with concepts of determinables, with question raising principles, and with value adders; it is to conjecture that they may come innately equipped with such things, and that they may guide their attention by such things. It is to invite ourselves to find out specific features of that endowment and of its evolution. Thus the answer to our questions about rational ignoramuses should raise new issues about the nature of innate and acquired cognitive capacities.

But note that this is a dimension of value which has been given no weight so far in the policy of our rational ignoramus: it is not gosh or cash or added value, but the value that a question derives from the fact that its answer may open up new fields of ignorance, a value even less predictable than the others. Let us call it *golly* value, and stop.

NOTES

[1] In what follows, I will indulge in a number of idealizations and simplifications that a lengthier discussion would avoid. For instance, I will pretend that all questions have at most one correct answer. That is blatantly false, but should not affect the main thrust of the discussion. I will also disregard the fact that one may be said not to know the answer to a question when one mistakenly takes a false answer to be a correct answer. In fact, I will use the word 'know' rather loosely.

[2] Most of these topics will be discussed in a larger work still in preparation.

[3] I use 'interrogative sentence' to refer to linguistic entities and 'question' to refer to what these linguistic entities express. Some complex issues related to the existence of propositions, senses, Fregean thoughts, etc., are masked by this terminology. However nothing in what follows hinges on the stand one takes on any of these issues.

[4] 'Possible answer' is a semantic notion that covers correct, incorrect, plausible, implausible, etc., answers and excludes only things that are not answers at all.

[5] Of course, someone who thinks that there is a King of France might think that he does not know. But he would be mistaken. Admittedly, such a person will gain new information by learning that there is no King of France. But to learn that is not to learn the age of the King of France. Similarly, there was a time when people thought that they did not know the specific gravity of phlogiston, the ratio of fire to earth in wood smoke, the distance from the earth to the sky. Later discoveries did not cure that ignorance; they dissolved it.

[6] That "some thing of a certain character" need not be a physical object but can be a state of affairs or fact or event. Thus "Why does copper turn green when exposed to air?" presupposes that copper turns green when exposed to air.

[7] No "theory" of determinables, i.e., explicating their semantics, is currently available, and hence the class is difficult to demarcate, though its most typical members are easily recognized.

[8] There may be a rational policy whose application does not require that questions be selected or eliminated or ranked one by one, but that treats some of them *en bloc*. In all the cases that I have examined difficulties similar to the above reemerge.

[9] That is, P can express each member of Q, P knows – or has at least good reasons to believe and no good reason not to believe – that its presupposition is true, P knows that it arises, P is in a position to know that he does not know its answer.

[10] These "other circumstances" include many sorts of contingencies that cannot be discussed here. For instance, added value depends not only on the availability of value adders – see below – but also often on the availability of answers to other questions required by the value adder.

[11] Super-erotetic value adders turn into erotetic value adders when provided with specific objects as arguments (in the case at hand when '... (a pendulum) ...' is replaced by an expression referring to a specific pendulum).

[12] Of course, many questions that can be answered by looking at the map do not translate into questions about the territory, e.g., "What is the name of the publisher printed on the map?" And the converse is also true, i.e., many questions about the territory can usually not be answered through questions about the map.
[13] At least in part: the algorithms are obviously not physical.
[14] This feature, though essential, is not sufficient to characterize natural kinds.

Department of Linguistics and Philosophy
Massachusetts Institute of Technology
Cambridge, MA 02139
U.S.A.