

Musings on Information and Knowledge

Robert J. Aumann

Abstract. An invited contribution to a symposium on "Information and Knowledge in Economics," to appear in *Econ Journal Watch*. Topics discussed include the distinction between information and knowledge; awareness; logical omniscience; the cost of calculation; semantic and syntactic models of knowledge, and the equivalence between them; and common knowledge of the model. Finally, some aspects of the symposium contributions of Ken Binmore, Jim Friedman, and Eric Rasmusen are discussed.

1. Some Distinctions

The first Symposium question is, "Is there an important distinction between information and knowledge?" Of course, it depends on what is meant by these terms. An advantage (disadvantage?) of formal reasoning is that there, that kind of question does not arise. You must first define your terms, and then it's usually easy to tell whether there is or is not an important distinction. Some people might say that even informally, the question has no substance until you've said what you mean.

One could interpret the question in terms of common usage. Some of the Symposium materials (specifically, the quotations from Kenneth Boulding's *Beyond Economics* and from Paul Weiss) suggest that information is the raw material from which knowledge is manufactured; that information is purely factual, whereas knowledge is something deeper—information that has been distilled, digested, internalized, processed or somehow transformed into an idea or principle. It's an interesting thought, but doesn't reflect common usage. You can know perfectly mundane things—whether it rained yesterday, or when the train left.

A distinction in common usage is that information can be impersonal, whereas knowledge is personal. It is people who know things; knowledge isn't "out there," has no existence of its own that is independent of people's minds. On the other hand, information exists outside of people's minds; it can be gathered, so to speak. You might say to your travel agent, "Get me information on flying to New Zealand." It's quite possible that nobody in the world knows how much it costs to get from Jerusalem to Christchurch, but the information is "there," the travel agent can gather it.

Another distinction in common usage is that "knowledge" has a more precise connotation than "information." For example, a probability estimate might be considered information, but it isn't knowledge. A detective might say, I have information on who committed the murder, but I don't know.

Possibly neither of these distinctions is very important to economics.

There is an important distinction in scientific usage: Information can be measured, knowledge cannot. For example, the minimal number of bits needed to transmit some piece of information can be taken as a measure of the "amount" of information being transmitted. This is closely related to the idea of entropy as a measure of information. Knowledge, on the other hand, is a 0–1 affair; you either know something, or you don't.

There's also an important distinction between "knowledge theory" and "information theory." The former refers to partition models of knowledge, the syntax of knowledge, common and mutual

knowledge, axiomatics, and so on. Sometimes this is called “formal epistemology,” and when there are several agents, “interactive epistemology.” On the other hand, “information theory” deals with information transmission, noisy channels, entropy, and so on. Though related, the two are really quite different. Both are highly relevant to economics.

Two important issues that may be relevant to a possible distinction are awareness and logical omniscience. These are worth separate treatments.

2. Awareness

In most formal models of knowledge—particularly, those used in economics—agents are aware of the possibilities, but do not know which obtains. But in real life, an important component of ignorance (the lack of knowledge) is ignorance of the possibilities. Columbus, when setting out westwards for India, had no idea that there might be another continent between Europe and Asia; that possibility did not enter his mind. When we do scientific research, we often have no idea what we will find. Indeed, that is the most interesting case; and it is quite different from, say, not knowing the time.

Experimental economists report that subjects in experiments tend to overestimate low probabilities; that is, they react to low probabilities as if they were higher. A possible explanation is in terms of unforeseen circumstances. The subject raises the low “official” probability because he instinctively takes unforeseen circumstances into account. Or perhaps not instinctively, but from experience; he has seen or heard of too many “safe” drugs that caused irreparable harm, “safe” ships that sank, or reputable experimenters who lied—or failed to tell the whole truth—to experimental subjects.

The basic difficulty with building a model that explicitly takes unforeseen circumstances into account is that usually (though not invariably) the agents are taken to know the model, so the unforeseen circumstances become foreseen. Though there has been work that addresses the problem of awareness, I know of none that has “caught on,” that is truly satisfactory.

One could perhaps distinguish between “knowledge” and “information” by associating knowledge with the traditional concept—where agents are aware of all the possibilities, but do not know which obtains—and associating information with the more amorphous situation, where agents are not even aware of the possibilities. That is, one might perhaps say that one seeks “knowledge” about the time or the weather, and “information” about the results of some future exploration or scientific research.

3. Logical Omniscience and the Cost of Calculation

A fundamental axiom of knowledge theory is that if you know p and you know that p implies q , then you know q . That sounds perfectly harmless and even obvious. But it entails logical omniscience: that agents know everything that follows logically from anything that they know. It follows that they know all logical tautologies, and in particular, all theorems of mathematics. For example, that Fermat’s “last theorem,” which was proved only in the 1990’s after remaining open for 350 years, was always known to all mathematicians (indeed to everyone)—a patently absurd proposition.

This matter is of great economic importance. A significant shortcoming of economic theory is that it fails to take the cost of calculation into account. Under logical omniscience, all results of all calculations are already known to all agents, so calculations are in effect cost-free.

In real life, obviously, computations are costly. Computation—in particular, cost calculating—is important to almost any economic activity. But there is a fundamental economic difficulty with the matter of cost calculation. Namely, in order to know how much to calculate—when doing an optimization, say—one must know beforehand how much this will cost. But this in itself involves a calculation, indeed one that is likely to be more complex than the optimization itself. One is thus led to an infinite regress of ever more complex calculations, from which there seems to be no escape. This is a very serious problem, of great economic importance, for which there exists no solution that is even remotely satisfactory.

Again, one could perhaps distinguish between “knowledge” and “information” by associating knowledge with the traditional concept—where logical omniscience does obtain—and information with the more amorphous situation, where logical omniscience is not assumed. That is, one might perhaps say that one seeks “knowledge” about the time or the weather, and “information” about some cost calculation.

Both this and the awareness distinction are in the spirit of what was said in the first section—that information is less precise, more amorphous than knowledge.

4. Common Knowledge of the Model

Having responded to the first Symposium question, we skip to the last: “What other thoughts do you have on knowledge and information in economics?”

Recall that a proposition is commonly known if all concerned (the “players”) know it, all know that all know it, all know that, all know that, and so on ad infinitum. In the “Game Theory” entry in the *New Palgrave Dictionary of Economics* (Macmillan, 1987, Volume 2, p. 473), we wrote as follows: “The common knowledge assumption underlies all of game theory and much of economic theory. Whatever be the model under discussion, whether complete or incomplete information, consistent or inconsistent, repeated or one-shot, cooperative or noncooperative, the model itself must be assumed common knowledge; otherwise the model is insufficiently specified, and the analysis incoherent.”

Today, 18 years later, this formulation no longer seems appropriate; in a sense, it is even incorrect. For example, to get Nash equilibrium in a game G , one need assume common knowledge neither of the game nor of the players’ rationality.¹ It is important that G indeed be the game, and that the players indeed be rational; but common knowledge of these items is not needed.

In another sense, though, it is correct, but unnecessary. What we were saying in 1987 is that in addition to explicit assumptions like zero sum, complete information, one-shot, or whatever, it is also implicitly assumed that the “model” is commonly known by the players; and that this assumption—the “common knowledge assumption”—is outside the model, and cannot be, or is not, stated within the model itself. We now understand that it is not an assumption at all, but a truism, a tautology, a “theorem;” it follows from more fundamental considerations.

To understand why, we must delve a little into interactive epistemology, which provides tools for analyzing what the players know—about the world and about each other’s knowledge. There are two parallel formalisms, the semantic and the syntactic. As a practical matter, the semantic formalism is simpler to use, so has become standard in economic applications—in spite

¹ Aumann, R.J., and A. Brandenburger (1995), “Epistemic Conditions for Nash Equilibrium,” *Econometrica* 63, 1161-1180.

of being conceptually roundabout and a little difficult to fathom at first. It consists of a set of possible states of the world (or simply states), and for each player, a partition of the states into information sets, which tells us what that player knows at each state.² No doubt, some readers are familiar with this formalism; for the others, it is not necessary, nor would it be useful, to describe it further here.

The syntactic formalism, on the other hand, is conceptually entirely straightforward, but in practice a little awkward. It is simply a formal language, which incorporates the basic propositions under discussion (like “Yesterday it snowed in Jerusalem”), logical connectives (“and”, “or”, “not”, “implies”), and a way of saying “Ann knows that ...,” Ann being a generic player. Fewer of our readers will be familiar with this formalism; but again, it is neither necessary nor useful to describe it further here.

On the face of it, there seems to be a fundamental conceptual difference between the two formalisms. The syntactic formalism is simply a language; it has no substantive content, does not say anything about the real world. It does not say that it snowed yesterday in Jerusalem, or that Ann knows that it did. It just provides a way of saying these things, if we want to say them.

But the semantic formalism does seem to have substantive content. Namely, that it is commonly known among the players that the states in the formalism are the only possible ones, and that the knowledge partitions really do describe the players’ knowledge. In particular, each player is assumed to know the knowledge partitions of all the others. At the least, this appears to be substantive; and indeed, it isn’t clear what justifies it.

There is another, related difference between the two formalisms. Given the players and the basic propositions, the syntactic formalism is canonic; there is only one syntax. But the semantic formalism is not; it is possible to construct the states and the information partitions in many essentially different ways.

But in fact, the semantic formalism has no substance, either; it, too, is just a language. Indeed, it is entirely equivalent to the syntactic formalism. To see this, we construct the states explicitly in terms of the syntactic formalism. Conceptually, a “state” is simply a complete specification of all (relevant) aspects of the world. Thus in each given state, every proposition is either definitely true or definitely false. So we may think of a state simply as a list of propositions, which contains, for each proposition, either that proposition or its negation. In addition, we must require that the list be logically consistent, i.e., that each list contain the logical consequences of the propositions in that list. Call such lists complete and coherent.

Thus we may think of a state simply as a complete coherent list of propositions. There are, of course, infinitely many such states; but the syntactic formalism is perfectly explicit and transparent, so we have a good understanding of how a state looks, and indeed of how the whole “universe”—the set of all states—looks.

So much for the states. How about the information partitions? Where do they come from? Why are they known to all the players, indeed commonly known?

Well, it turns out that the information partitions are implicit in—can be read off from—the states. Two states are in different information sets of Ann if and only if she knows, in one state, that the other is not the true state. That is, there must be something that Ann knows in one state, that

² Explicitly, two states are in different information sets of a given player if and only if in one state, the player knows that the true state is not the other.

she does not know in the other; formally, a proposition p such that the proposition “Ann knows p ” is in the list describing one state, and not in the list describing the other.

So the semantic formalism is just a different way of writing the syntactic formalism. The two are entirely equivalent, and neither has any substantive content. Both are just languages; they make no assertions—embody no assumptions—about the real world.³

The semantic formalism just constructed is called the universal formalism. Clearly it’s canonic, since it was constructed in a canonical way from the syntactic formalism. Though there also are other semantic formalisms, they’re all naturally embedded in the universal formalism; the universal semantic formalism encompasses everything of epistemic interest.

So, coming back to the “assumption” of common knowledge of the model: There is nothing substantive there. The “model” may be taken to be simply the universal formalism; as we have seen, this is equivalent to the syntactic formalism, which has no substance—is just a language.

It is of course possible to assume that various substantive items—zero-sum, complete information, the payoffs, the players’ rationality, whatever—are in fact common knowledge. But such assumptions can and should be made explicitly, within the model. Unlike what we thought in 1987, assumptions that are implicit—“outside the model”—are never needed.

5. Comments on the Symposium Materials

5.1. Binmore

In the symposium materials, Ken Binmore writes, “Game theorists usually assume that the rules of the game and the preferences of the players are common knowledge. In analyzing a game, they typically need also to assume that the fact that all the players subscribe to appropriate rationality principles is also common knowledge, although they are seldom explicit on this point.” We respectfully disagree.

(1) Ever since the ground-breaking work of Harsanyi⁴ on games of incomplete information, game theorists have not assumed that “the rules of the game and the preferences of the players are common knowledge.” To be sure, sometimes one does assume this; but for close to forty years, it has not been a part of the game-theoretic canon. Ken surely knows this; it’s possible that the quote is out of context. Anyway, we thought it important to set the record straight.

(2) The second sentence is also inaccurate. As pointed out at the beginning of the previous section, “typically”—for Nash equilibrium, say—one must assume that the players indeed “subscribe to appropriate rationality principles,” but not that this is common knowledge.

5.2. Friedman

Jim Friedman writes, “Usually, ... games of complete information are characterized by each player knowing the entire structure of payoffs of the game, by each player knowing that all players possess this information, and by all players knowing that all players have this information. There is

³ Other than underlying assumptions like awareness (of the basic propositions) and logical omniscience. For a formal treatment of these ideas, see, e.g., Aumann, R. J. (1999), “Interactive Epistemology I: Knowledge,” *International Journal of Game Theory* 28, 263-300.

⁴ “Games of Incomplete Information Played by Bayesian Players,” Parts I-III, *Management Science* 14 (1967-8), 159-182, 320-334, 486-502.

... an important conceptual distinction ... between (a) a complete information game in which complete information is common knowledge and (b) a complete information game in which each player does not actually know whether the other players also have complete information." Again, we respectfully disagree. The first sentence describes second order mutual knowledge of the information in question, whereas what is needed is common knowledge. Therefore, the second sentence is also inaccurate. In a complete information game, complete information is commonly known, so option (b) is impossible; this follows from the general theorem that when something is commonly known, then it is commonly known that it is commonly known.

5.3. Rasmusen

Eric Rasmusen writes: "For clarity, models are set up so that information partitions are common knowledge. Every player knows how precise the other players' information is, however ignorant he himself may be ... Making the information partitions common knowledge is important for clear modeling" Here, we do not disagree, but the thrust seems misplaced. As explained in the previous section, it is tautological that the partitions are common knowledge.