

## Reply to Stich and Nichols

ROBERT M. GORDON

In their vigorous, clear, and well-informed critique (ch. 5), Stich and Nichols tell us they are intrigued by the Simulation Theory, but skeptical. My aim in this reply is to answer their skeptical arguments and also, along the way, to show that the theory is even more intriguing than they suppose. In doing this I sharpen, strengthen, and build on what I wrote in the papers they cite.

### 1 Differences of Interpretation

Stich and Nichols have clearly tried to give a fair and accurate account of what they call 'the Gordon/Goldman proposal'. Still, there are some points at which my own view, at least, differs from the one they criticize.

They speak of the 'off-line simulation theory'. This packs into the very name of the theory what I regard as an ancillary hypothesis: that when we simulate – that is, use our imagination to identify with others in order to explain or predict their behavior – many of the same cognitive systems that normally control our own behavior continue to run as if they were controlling our behavior, only they run off-line, with their normal output systems disconnected and their normal input systems usually modified to some degree. This hypothesis is very plausible, I believe, and I do take it for granted in some of my arguments, here and elsewhere. But I think the imaginative identification theory remains attractive even without the off-line hypothesis.

A second point is that I wouldn't want to characterize simulation as using oneself as a model. People gravitate toward this conception because extrapolation from a model is a familiar form of inference whose logical structure is fairly well understood. But to 'use oneself as a model', as I understand it, would be first to simulate *oneself* under counterfactual conditions – to ask what *I* in particular would do or feel – and then to extrapolate to someone else.

This is quite different, I have argued elsewhere, from simulating the other. Goldman and I evidently differ on this point. A related point – on which we also differ – is that simulation as I understand it doesn't demand that one already possess intentional concepts and be capable of applying them in one's own case. I have also advanced the far stronger suggestion – inspired in part by Stich's account of belief attribution (Stich, 1983) – that to ascribe to another individual  $x$  a belief that  $p$  is to assert that  $p$  within the context of a simulation of  $x$  (this volume, ch. 2).

Stich and Nichols offer two general objections to the simulation theory itself (ch. 5, section 5, 'In Defense of the Theory-Theory') and several objections to specific arguments for the theory (ch. 5, sec. 4, 'Arguments in Support of Simulation-Based Accounts'). I'll respond first to one of the general objections, the 'cognitive penetrability' argument. The other general objection I leave to the end, because it concerns developmental findings that remain much in dispute. As to the objections to specific arguments, I shall, given limitations of space, concentrate on those that appear to pose the most serious challenges to the simulation theory.

### 2 The 'Cognitive Penetrability' Argument

In their argument from 'cognitive penetrability', Stich and Nichols challenge us to predict the responses of subjects in certain experiments by imagining ourselves in the role of subject. The experiments uncovered little-known features of human psychology, reflected in generalizations such as the following: people tend to prefer the rightmost item in a display of merchandise, other things being equal. People who know nothing about the experiments and are ignorant of these generalizations tend to mispredict what they would do in such an experiment. This is just as one would expect if their predictions actually depended on their faulty knowledge of human psychology. If instead they arrived at their predictions by truly running their own behavior control systems off-line, then their ignorance should make no difference, Stich and Nichols argue. For, given that their own system is subject to the same bias toward the rightmost item – the 'position effect' – their simulation should *correctly* predict the findings.

Or should it? Consider an analogous use of imagination:

I ask you to visualize two straight lines of equal length, one just above the other. One of the lines has a regular arrowhead at each end, the other has an inward-pointing arrowhead at each end. Now, *which line is longer?*

I have just tried to replicate the Müller-Lyer illusion, using only your visual imagination. There are two problems. One is *methodological*: It's silly to ask, 'Which line is longer?' because I stipulated that the two lines are of *equal length*. Suppose we ask, 'Which line *looks* longer?' There is still no reason to think

we'll get the illusion going. That's because of the second problem, which may be called the problem of *imaginative impenetrability*: Even if visual imagery is a product of the visual-perception system running off-line, it is surely going to bypass some *stages* of visual processing, especially early stages – and the illusion may originate at one of these stages. In that case we will have failed to recreate the antecedent conditions for the illusion.

It's the same with trying to replicate the 'position effect' experiment. First, there is the methodological problem: unlike the subjects in the original experiment, the subject in the imagination experiment must be told that the items on display are identical (and thus of equal quality). Second, there is imaginative impenetrability. When we are actually viewing a display of identical items, position may generate illusory differences. But how can we generate illusory differences without actually viewing a display? As with visualizing the Müller-Lyer figure, visualizing a display evidently fails to 'penetrate' visual processing at the right place. That would explain why, as Stich and Nichols say (p. 151), 'most of us think we would report that the garments looked to be very similar, and then choose one randomly'.

The general problem is that for some types of psychological or psychological phenomena the difference between reality and imagined or simulated reality is crucial. To *imagine* drinking six martinis doesn't make one drunk, not even off-line. Here simulation is inherently inadequate, and an independent source of knowledge is indispensable. That's a shortcoming of *simulation*, but it isn't a shortcoming of the simulation *theory*. It is a *virtue* of the theory that simulation fails to predict phenomena that our common-sense belief-desire psychology fails to predict.

For Stich and Nichols's second example, predicting the results of the lottery ticket experiment, methodological inadequacy is the major problem. In the original experiment, subjects in one group were given a choice of tickets and those in the other group were sold tickets without a choice. Days later, on the morning of the lottery, they were asked to sell back their tickets, and the choice group demanded a much higher price than the no-choice group. Suppose a psychologist wanted to replicate the experiment. But suppose he is short of subjects. So he uses the same subjects over again, first in the no-choice condition and then in the choice condition. Any methodologist would tell us this is a bad experimental design. For one thing, a subject might be influenced on the second run by what he did on the first run. So we shouldn't be surprised if the subjects assigned the *same* price to the two tickets. But now suppose that the psychologist is also short of *time*. So he sells the two tickets only seconds apart; then, a second or two later, he proceeds to buy them back, instead of waiting several days, as in the original experiment. This of course makes the bad methodology far worse.

Finally, suppose that the psychologist is not only short of subjects and short of time, but also short of lottery tickets – or, more simply, suppose he is a philosopher. So he asks his subjects to *imagine* being sold two lottery tickets and then to *imagine* being asked to sell them back. I quote from Stich and Nichols (this volume, p. 151).

Now imagine yourself in both roles – first as a person who had been handed the ticket, second as a person who had been given a choice. What price would you ask in each case? Would there be any difference between the two cases?

What we now have is *simulation* harnessed to very bad methodology.

The third example offered by Stich and Nichols – the belief perseverance experiment – suffers from *both* a methodological problem and a problem of 'imaginative impenetrability'. One can't expect to replicate a belief perseverance experiment if subjects are *immediately* shown that they are getting false 'information'. And it is far from obvious that the mechanism responsible for the perseverance of *beliefs*, whatever it may be, must also be triggered by merely *feigned* beliefs.

Summing up so far: Stich and Nichols are wrong in saying that if we were simulating, then we would have correctly predicted the findings in these experiments. Contrary to what they suggest, simulation has its shortcomings, some inherent and some only when it's harnessed to bad methodology. And that is just as it should be, if the simulation theory is right.

### 3 The Comparison to Grammar

Stich and Nichols claim that two of the arguments Goldman and I advance against the theory-theory of folk psychology would apply to theory-theories in other domains where a theory-theory really is the only plausible account. Their chief example of such a domain is grammar. Replying to argument 1, that no one has been able to state the principles of common-sense psychology, they write (p. 134), 'It has proven enormously difficult to state the principles underlying a speaker's capacity to judge the grammaticality of sentences.' And in answer to the anomalous precocity argument (argument 3), they say that according to generative grammar (p. 136), 'the knowledge structures that underlie a child's linguistic ability are enormously complex'.

But there is a difference. Theory theorists generally assume that the laws of folk psychology can be formulated in terms of the common-sense mental vocabulary. They *must* assume this if they hold, as most seem to, that the common-sense theory implicitly *defines* these terms. Generative grammarians, on the other hand, have no such compunctions about introducing a technical vocabulary. They want to explain common-sense intuitions of grammaticality, but they are not constrained to do this by way of a *common-sense grammatical vocabulary*. The reason is simple. They don't claim that what explains our grammatical competence is *folk grammar* – the principles that underlie common-sense notions of grammar and common-sense explanations why one word string is grammatical and another is not. But theory theorists *do* generally make the corresponding claim: that what explains our competence in predicting behavior is *folk psychology* – the principles that underlie our mental attributions and explanations of behavior.

Why is this difference relevant? Well, for one thing, where there is no constraint on concepts – no requirement that the theory must be framed in terms of *the child's own concepts* – it isn't clear how a *precocity* issue could arise. Analogy: It's all right to say that to catch a ball children have to solve differential equations, as long as this doesn't entail that they have to *understand* differential equations.

What of the argument that no one has been able to state the principles? Actually it is Stephen Schiffer who argues most forcefully:

*But can anyone state so much as a single generalization that fills [the] bill?* (Schiffer, 1987, p. 29; Schiffer's emphasis)

Schiffer's negative answer, argued at length in *Remnants of Meaning* (1987), makes him skeptical of the view that there is a

system of law-like generalizations using the notions belief and desire that is known, or 'used', by plain folk who possess these concepts. (Schiffer, 1987, p. 29)

Yes, there remains serious disagreement about the computational mechanism underlying our grammatical competence: but *of course* that shouldn't lead us to think there is *no* computational mechanism that underlies our grammatical competence. These seem to be vastly different issues. But I won't argue the point. I have never made the complaint that no one has been able to state the principles. My charge has been that whatever common-sense laws and generalizations we may find, they are at best useful heuristics. We are much better predictors than our generalizations are, because *our own* decision-making system warns us when a generalization yields a crazy prediction: that is how we fill in the implicit *ceteris paribus* clause.

#### 4 Simplicity

I turn now to Stich and Nichols's reply to argument 4 (pp. 137–8), that the simulation theory is simpler than a view that posits a special stock of laws corresponding to rules of logic and reason. Stich and Nichols ask: What about the investment in a special *control mechanism* for taking one's decision-making system off-line, feeding it pretend inputs, and interpreting its outputs as predictions of another's behavior? My answer is that it isn't special: we have to have it *anyway* if we're to do hypothetical and counterfactual reasoning. To judge whether a conditional statement is true or false would require (according to both the classical Ramsey account and Stalnaker's modified account) a belief generator that adds the antecedent to one's existing stock of beliefs. It would also require the capacity to use the modified stock of beliefs as a basis for deciding – *off-line*, of course – whether or not the *consequent* is true. Finally, it would require the capacity to interpret this output as a decision whether or

not the *conditional* is true (Stalnaker, 1968). So, contrary to Stich and Nichols's claim, simulation gets its control system as well as its data base pretty much 'for free'.

But I agree with Stich and Nichols that greater precision is needed if the simplicity issue is to be given much weight – precision not only in formulating the theories to be compared but also in stating what the issue is. The general notion of theoretical simplicity isn't a clear one. Moreover, it isn't clear why nature should *answer* to it, that is, prefer the simpler theory. So I propose to put the notion aside altogether and talk instead of *code compression*. It's clear why nature would prefer a coding scheme that offers greater compression, other things being equal: Less code, therefore fewer synapses, therefore smaller brain, smaller energy demand, and so forth. Other things being equal, the more compression a scheme allows, the more probable it is that it's the scheme we've got.

I approach the topic by first considering an experienced Shakespearean actress, who has mastered a variety of roles: she can 'become' characters as diverse as Ophelia, Titania, and Lady Macbeth. How are these distinct characters represented in her brain? A very plausible answer, I think, is that for each character what is stored is not a set of *facts* about the character's mental life, such as an inventory of Lady Macbeth's mental states, processes, and tendencies; but rather a set of *operations*, namely the set of changes or adjustments the actress makes as she mentally prepares each evening to 'become' Lady Macbeth. Information about each character is stored as a procedure rather than in declarative form; and specifically, as a transformational procedure. Here is a partial analogy: video signals may be digitally transmitted or stored in compressed form by coding only the differences between each frame and its predecessor, rather than coding each frame in its entirety. Thus a frame is represented as a set of changes from the preceding frame – except for the initial frame, which would be coded in its entirety. One of the *disanalogies* is that in the case of the actress, the initial 'frame' – information about the actress herself – isn't coded at all. She needs no inventory of *her own* mental states, processes, and tendencies.

The simulation theory – together with the ancillary off-line hypothesis – tells a similar story about our common-sense knowledge of the mental lives of particular others. Such knowledge, insofar as we have it, is not a kind of speculative knowledge but rather a kind of know-how: it is knowing how to transform ourselves into the other. The various people we know are represented in their mental aspects as sets of transformational operations. But unlike the actress we aren't called upon to transform ourselves audibly or visibly; we keep it to ourselves. And whereas the actress transforms herself into characters in *fictional* worlds, we typically match these transformational operations with particular bodies in the actual world.

A theory, it is true, could gain some of the advantages of such a system by assigning *default values* to all parameters and then coding only *deviations* from those values. But the default values would have to come from somewhere. Even if they were just based on one's own case, the theory would have to be

fed information about one's own case. In other words, the initial 'frame' would have to be coded. Simulation avoids altogether this initial investment in information acquisition and storage.

### 5 The Developmental Arguments

I have left the developmental arguments and counter-arguments to the end, because the findings on which they are based remain in dispute and there is a considerable amount of work in progress.

Stich and Nichols try to neutralize argument 6 (the 'Maxi' argument, pp. 141-4) by showing on the one hand that only one of several possible versions of the simulation theory predicts the experimental results, and on the other hand that a version of the theory-theory would make the same prediction. They also offer a developmental argument of their own, based on experimental results (Wimmer, Hogrefe and Sodian, 1988) which, they say, are readily explained by the theory-theory but not by the simulation theory. I shall argue that there is good reason not to trust the results Stich and Nichols cite in favor of the theory-theory, and further, that they are mistaken about the developmental implications of the two theories.

According to Wimmer et al. (1988), young children are much stingier in attributing knowledge to others than they are in attributing it to themselves. This seems at odds with the simulation theory, for 'the answer [the child] comes up with [as to whether the other knows] is *not* the one that she herself would come up with, were she in the subject's place' (Stich and Nichols, this volume, p. 150). But the result can be explained by a version of the theory-theory, according to Stich and Nichols.

Unlike the findings of Wimmer and Perner (1983), which have been replicated several times, this alleged disparity in first- and third-person attributions is by no means well established. (Indeed, it is in conflict with other recent findings, as Goldman shows in detail in this volume, ch. 9.) And there are independent reasons to distrust the finding. Whereas Wimmer and Perner (1983) were chiefly concerned with children's predictions of Maxi's behavior, Wimmer et al. (1988) ask children to say whether a particular individual 'knows' the answer to a certain question, namely, 'What is in the box?' And there are good reasons to be cautious in drawing inferences from children's applications of the predicate 'know' (or *wissen*, the term Wimmer et al., were working with). The verbal evidence appears to be too liberal a measure of attributions of knowledge to oneself and too conservative a measure of attributions of knowledge to others.

It is too liberal a measure of attributions to oneself because there is a cheap and easy procedure for giving generally plausible answers to questions about what one knows, namely, the 'answer check procedure' described by Wimmer et al.: to say whether one knows the answer to a question, simply check to see whether one has an answer to the question.<sup>1</sup> Virtually any speaker can be trained in this procedure, without any understanding of the sentence

form 'x knows that p'. Likewise, a speaker may learn that it is optional to preface any assertion, such as, 'There's a pencil in the box', with the formula, 'I know that' or the formula, 'I believe that'. But from this it should not be inferred that the speaker is capable of meta-representation, that is, of believing that he knows or believes something.

On the other hand, the verbal evidence appears to be too conservative a measure of attributions of knowledge to others. Children about 5 years of age reportedly deny that another 'knows what is in the box', even after observing the other *looking* into the box or *being verbally informed* of its contents. Stich and Nichols infer that 'these children have not learned that people will come to know *that* p by seeing or being told *that* p'. But a more plausible account would be that the children haven't yet mastered the word 'know'. For there is a good deal of behavioral evidence that even four-year-olds are well aware that if one reports to X the fact that p, or if one makes the fact visually evident to X, then the fact that p becomes *accessible to X as a possible basis for action or emotion*. And if the fact that p is accessible to X as a possible basis for action or emotion, *then X knows that p*. Note, for example, that four-year-olds are reasonably skillful at hiding facts from others. If they don't want others to get angry or to tease *them* about something they did, they don't *tell* them they did it — much less *show* them, 'Look what I did!' And they ask others not to tell: 'Don't tell Bea I had an "accident"!' my son at age 3½ implored his mother, 'she'll tell everybody.'

I conclude that the alleged disparity between self-attributions and other-attributions is probably only a verbal disparity. Without supplementation by further research, the results reported in Wimmer et al. (1988) support *neither* the hypothesis that the subjects do recognize that *they themselves* can learn by seeing or being told nor the hypothesis that they fail to recognize that *others* can learn by seeing or being told. In short, there is reason to doubt that the developmental facts are as Stich and Nichols portray them.

Now consider what Stich and Nichols have to say about the developmental implications of the theory-theory and the simulation theory, respectively. They are right about two things. There is a possible version of the theory-theory that would explain the results of the Maxi experiment. And there is a possible version of the theory-theory that would explain the results of the ways-of-knowing experiments. But Stich and Nichols overlook a rather important point: these are *different* versions of the theory-theory; indeed, they are, or at least they appear to be, *incompatible* versions. I'll explain.

In chapter 2 of this volume, I argued that if it were by acquiring a *theory* that children develop the capacity to predict the behavior of others, the capacity to predict actions on the basis of the other's beliefs would develop, so to speak, as a single package: it should not matter whether a belief was true (in agreement with the child's own belief) or false (contrary to the child's own belief). And the Maxi experiment seems to show that in fact it does *not* develop as a single package. But as Stich and Nichols rightly point out, the theory-theory can accommodate the Maxi experiment by stipulating that (this volume, p.143):

[A]t a given stage of development, children have mastered the part of the theory that specifies how beliefs and desires lead to behavior, though they have not mastered the entire story about how beliefs are caused. At this stage, they might simply assume that beliefs are caused by the way the world is; they might attribute to everyone the same beliefs that they have. A child who has acquired this much of folk psychology would (incorrectly) attribute to Maxi the belief that the chocolate is in the cupboard.

Stich and Nichols argue that the theory-theory can also accommodate the ways-of-knowing experiments of Wimmer et al. (1988). It does this by stipulating that children acquire the fragment of folk psychology that associates knowledge or belief that *p* with its verbal expression, *saying* that *p*, while still in the dark to what causes knowledge or belief that *p*. Only later do they acquire the generalization that people will come to know or believe that *p* by seeing or being told that *p*, and later still the generalization that people acquire knowledge or belief by *inference* from what is perceived.

Both versions of the theory-theory, the one that accommodates the Maxi experiment and the one that accommodates the ways-of-knowing experiments, hold children to be initially naïve about the causes of knowledge or belief. *But they hold them to be naïve in exactly opposite ways!* According to the Maxi experiment version, the three-year-old folk psychologist is ultra-liberal about knowledge, ascribing it to others without asking whether the other has access to the facts, that is, a *way* of knowing what is the case. But to accommodate the ways-of-knowing experiments, the three- and four-year-old must be portrayed as ultra-conservative, ascribing knowledge only if there is a recognized way of knowing that would give the other access to the facts in question. The theory-theory can't have it both ways.

Moreover, each of these versions of the theory-theory is unattractive in itself. Especially bizarre is the *ultra-conservative* version, which implies that until age five children regard others, including their own care-givers, as presumptively windowless creatures, cut off from the world and consequently unresponsive to the environment. The problem facing the *ultra-liberal* version is this: If the initial supposition is that beliefs are never at variance with the facts, if they introduce no dimension of variability, then *why would children posit beliefs at all?* Is it that three-year-olds can't tolerate action at a distance and so must posit a proximate cause inside the head? Not likely.

It isn't surprising that for each of these experimental findings there is a possible version of the theory-theory that explains it. That is so because, as Stich and Nichols grant, the theory-theory 'is compatible with (but does not entail) lots of developmental patterns' (p. 144). Taken merely as the general claim that people use a theory to explain and predict behavior, the theory-theory is noncommittal as to whether people acquire the theory all at once or piecemeal, and if piecemeal, which pieces come first. But the simulation theory is not similarly neutral on these matters, contrary to what Stich and Nichols suggest. Try to devise a version of the simulation theory that would conflict with the finding that children predict behavior simply on the basis of

the actual facts (the *actual* current position of the chocolate) before they develop the ability to consider the other's *beliefs* (for example, Maxi's belief that the chocolate is where he had left it). This is not just a matter of devising scenarios in which the capacity to generate 'pretend'-beliefs is developed before or simultaneously with the capacity for off-line decision-making. What is proposed is that from the start – from the very moment that the child begins to use off-line decision-making to predict the behavior of others – rather than naïvely projecting his own belief base to the other, he starts *making adjustments to it*. But if he doesn't at first naïvely project his own belief base and learn that this sometimes *leads to error*, what *reason* would he have to make adjustments? And how would he learn *when* to make adjustments? And how would he know *what* adjustments to make? Unable to learn any of this from his own experience, the child must either be innately equipped with the information and motivation or somehow acquire it from others. That is, children can skip the initial naïve stage and move directly into the sophisticated stage of the four-year-old, but only by using resources not derived from simulation. Left to simulation alone, they will do just as we actually find children doing in the Maxi experiment. The simulation theory is not neutral.

### Conclusion

I have replied to what appear to be the most serious challenges posed by Stich and Nichols, and I believe I have successfully answered each challenge. But much has been gained from the dialectic. The 'cognitive impenetrability' argument has been rebutted, but we learned something about the *limits* of simulation. The simplicity argument has been made more precise with my sketch of a new account of the representation of other minds. The appeal to the Maxi experiment still stands, but it is now strengthened by a more careful contrast between the developmental implications of the simulation theory and those of the theory-theory. If Stich and Nichols's aim was to stimulate the development of the simulation theory – and I think it was, in part – then they have been successful, too.<sup>2</sup>

### Notes

- 1 Actually, this is more appropriate to belief than to knowledge, as I argue in my 'Reply to Perner and Hower', this volume, ch. 8.
- 2 My thanks to John Barker and Russell Treholme for helpful comments on an earlier draft of this reply.

### References

- Goldman, A. I. 1989: Interpretation psychologized. *Mind and Language*, 4, 161–85. Reprinted as ch. 3 in this volume.

- Schiffer, S. 1987: *Remnants of Meaning*. Cambridge, MA: MIT Press.
- Stalnaker, R. C. 1968: A theory of conditionals. In N. Rescher (ed.), *Studies in Logical Theory*. Oxford: Blackwell, APQ Monograph no. 2.
- Stich, S. 1983: *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.
- Wimmer, H., Hogrefe, J. and Sodian, B. 1988: A second stage in children's conception of mental life: Understanding informational access as origins of knowledge and belief. In J. Astington, P. Harris and D. Olson (eds), *Developing Theories of Mind*. Cambridge: Cambridge University Press.
- Wimmer, H. and Perner, J. 1983: Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13, 103–28.

## 8

## Reply to Perner and Howes

ROBERT M. GORDON

The simulation/theory debate has become an area of fruitful exchange between experimentalists and philosophers. Needless to say, there are misunderstandings. One of these is apparent in Perner and Howes's application of the label 'Cartesian.'

Perner and Howes correctly attribute to me the strong claim 'that simulation is *sufficient* for providing children with intentional concepts like belief, knowledge, desire, etc.' But then they go on to say that 'This strong claim is often referred to as "Cartesian"' (p. 161). I am not sure who characterizes it that way, but I want to make it very clear that it reflects a serious misunderstanding, at least of my position. According to the Cartesian doctrine, as they note, 'the mind knows what state it is in (i.e. must have the concept of what kind of state it is)'. But if mental concepts can – and, as I believe, do in fact – arise out of the procedure of *simulating others*, then it is *not* true that we already, independently, possess the concepts, much less that we already are capable of applying them in our own case, as the Cartesian doctrine would maintain. If we get our intentional concepts from simulating others, or if at least we *can* acquire them in that way, then we do not, or at least we *need* not, contrary, incompatible positions. It is true that some proponents of a simulation theory (Paul Harris, 1989, p. 57, for one), do seem to endorse some sort of Cartesianism. But in that case, they could not also hold, as I do, that mental concepts arise out of the procedure of simulating others.

The assumption that the Simulation Theory is Cartesian plays a large role in Perner and Howes's argument against the theory. For the relevance of their experimental data depends on whether the Simulation Theory would hold, as they suppose, that it is by *introspecting their own states* that their subjects decide, for example, whether John thinks he knows the answer to a question. But the focus of my discussion will be on a different assumption they make: