

Thinking about Thinking about Thinking about Thinking (about Poker)*

Jonathan Ellis**

When I am getting ready to reason with a man I spend one-third of my time thinking about myself and what I am going to say, and two-thirds thinking about him and what he is going to say.

– Abraham Lincoln

Remember that childhood game “Odds or Evens” you used to play in order to settle important disputes such as who gets the last slice of pizza? There was only one element of skill to that game: trying to figure out what the other person would throw. But that wasn’t easy. If your opponent was savvy, that meant trying to figure out what he thought you were going to throw. And that sometimes meant figuring out what he thought *you* thought *he* was going to throw.

Thinking about what other people think is something we do all the time, not just when playing games: You told your friend to meet you in the restaurant at 7:00, but when it’s 7:15 and she’s not there, you realize she must think you’re meeting in *front* of the restaurant. The way we anticipate and explain other people’s behavior is by anticipating and understanding what they *think*.

Philosophers call this “thought attribution,” and top poker players are remarkably good at it. The ability to attribute thoughts to other people is especially important in the No-Limit Texas Hold’em tournaments that have, through television and the Internet, swept the globe in recent years. What is required to succeed in this game is not merely attributing “first-order” thoughts to other players, but attributing to them thoughts about your thoughts about their thoughts. In fact, that’s the minimum. The kind of thinking done by the very best players (T. J. Cloutier, Howard Lederer, and Daniel Negreanu are especially good at it) is much more complex than that.

Here’s an example. It’s the third day of the World Series of Poker, and you peer down at your cards to see handsome “Big Slick” (Ace-King) looking up at you. Your heart pounding, you raise three times the pot and get called by one player. When the flop comes 2-6-9 (or “rags”—no high cards), you face a decision. Should you bet in the hopes that your opponent will fold? Or will he reason that it’s likely that you raised before the flop with a high Ace and are now bluffing? As all good players know, some players are more apt than others to bet with Ace-King (A-K) in this situation. Now, you’re a good player, and your opponent knows this, so he knows you’re aware that this is the situation. And you know he

* Published in *Poker and Philosophy*, ed. E. Bronson (Chicago, IL: Open Court, 2006).

** Department of Philosophy, University of California, Santa Cruz. Cowell Faculty Services, University of California, Santa Cruz, CA 95064. E-mail: jellis@ucsc.edu.

¹ T. J. Cloutier and Tom McEvoy, *Championship Hold’em* (Las Vegas, NV: Cardsmith Publishing, 2000), p. 202.

knows you're aware of this. In fact, he's already seen you twice in the tournament *not* bet in a similar situation. So perhaps he thinks you'd do that again. If so, then if you bet, he'd probably fold, thinking that you must have a pair. But, taking it one more level, perhaps he thinks you would think that he would think this, in which case your plan would backfire entirely. What to do!

First, you scan your mind for anything your philosophical heroes might have said about this predicament. Didn't Socrates say something about this in one of Plato's dialogues? You seem to recall:

Socrates: Good morning, Gyges, what brings you this early to the card room?

Gyges: Hello, Socrates. I have been so wound up thinking about the argument you made yesterday that I could not sleep and needed to clear my head.

Socrates: What troubles you, Gyges? Do you not agree that it is better to have more chips than to have fewer?

Gyges: It is certain.

Socrates: And do you not agree that if you've raised before the flop with Big Slick, and then the flop comes rags, then a bet could be seen as a bluff and would thus be imprudent?

Gyges: By Zeus, you are correct. I must then check, Socrates.

Socrates: Very well, I must be off to join Cephalus and Glaucon at the 2-4 table.

While in truth it is difficult to find ancient philosophical passages contemplating the sublime joy of hitting your set on the flop, philosophers have nonetheless had some important insights about thought attribution that should prove illuminating for poker players and fans.

Thinking about Rabbits

Consider the following scenario first introduced by Willard V. O. Quine (1908-2000), and later discussed by Donald Davidson (1917-2003), two eminent American philosophers. Suppose we are in a foreign land and come upon an unknown tribe of human beings. We'd like to communicate with this tribe and so embark upon the project of learning their language. How would we go about ascertaining what they mean by their words? We would start, these philosophers say, by observing their verbal utterances and studying how these utterances correlate with what is going on around them. For instance, if we noticed that they utter "Gavagai" every time a rabbit passes by, we would form the hypothesis that what they mean by "Gavagai" is "Rabbit," or "There's a rabbit," or something like that. Future observation could reveal that this is not what they mean, but starting with hypotheses like these is the only way to *begin* to determine what these people mean by their words.

Now, notice that in forming the hypothesis that the natives mean "Rabbit" by "Gavagai," we make a substantive assumption. In forming this hypothesis, we assume that

when they uttered “Gavagai,” they *believed* (or thought) that there was a rabbit passing by. Of course, this would be a reasonable assumption to make: the natives appear to have good eyesight, and there was indeed a rabbit passing by in every case. However, if we made no supposition at all as to what the natives believe, we would have no reason to think that what they mean by “Gavagai” is “Rabbit.” The words that a speaker chooses to utter are a function not only of what she means by them but also of what she believes. Just as the classical physicist describes a moving object’s direction along a plane as a “vector” of its horizontal and vertical speeds, Davidson characterizes what a person *says* as a vector of her meanings and beliefs. One utterance could be the result of countless combinations of meaning and belief.

Moreover, in forming our assumption, our *own* beliefs played a significant role. We didn’t just make any assumption about what the natives believed, we assumed that what they believed about their environment was precisely what *we* believed about their environment: that there was a rabbit passing by. Davidson pays special attention to this. In order to determine what a person means, we need to start by assuming that she and we hold many of the same beliefs. And if we do not hold many of the same beliefs, Davidson claims, we simply will be unable to determine what she means. This is true not only for ascertaining what someone means, but also for ascertaining what someone *believes* and *thinks*. Even to attribute to someone a thought that you think is absolutely false, you and she must hold many of the same beliefs. Take my friend’s conviction that the reason she’s having a bad day is that she’s a “Scorpio,” and Scorpios suffer when Mercury’s in retrograde. In order for me to attribute this belief to her (which I believe is patently false), she and I need to share many *other* beliefs, such as that people are born in particular months, that Mercury’s a planet, that people have bad days, and so on. Widespread agreement, according to Davidson, is the only possible background against which disagreement can be identified. (Scorpios suffer when *Jupiter’s* in retrograde, of course, not Mercury.)

This thesis has been applied to a wide variety of philosophical issues—in ethics, in metaphysics, in epistemology—but it also has significance for poker. To see why, let’s return to that dilemma you face with Big Slick. Once again, the flop comes rags and you think to yourself:

Maybe I can still win this hand if I get him to fold. This guy’s seen me twice in this tournament not bet when I had A-K and the flop came rags, so if I bet, he’ll probably think I raised before the flop with a high pocket pair.

This initial line of thinking is already a complex one involving many different thoughts, including one 2nd-order thought: *so if I bet, he’ll probably think I raised with a high pocket pair*. This is a thought of the 2nd order because it is a thought about a thought. Notice that this 2nd-order thought is not about what your opponent *thinks* but about what he *would* think if you bet. In poker, we are often trying to influence our opponent’s behavior (to get him to call, fold, and so on). And a primary way to influence someone’s behavior is to influence what he thinks. A central part of poker is thus *figuring out what you can do to get your opponent to have whatever thought you think would lead him to do what you want him to do*. Really, poker’s about trying to place false thoughts in other people’s minds—when you’ve got a good hand, you want them to think falsely that you don’t, so they’ll call. (For the ethics of this, see Barnhill and Solomon’s chapter in this volume.) Still, to figure out

what someone *would* think, you have to know a good deal about what he *does* think, as we shall see.

Sharing Thoughts at the Poker Table

Let's return to your predicament with Big Slick. Before you throw some chips in trying to move your opponent off his hand, you remember that this is T. J. Cloutier you're up against, who writes, in a moment of quiet self-examination, that he has "got more moves than a mongoose."¹ So, careful not to breathe or fumble your chips in such a way as to betray the weakness of your hand, you retreat once again into your own mind:

But T.J. probably knows that I know that he remembers that I didn't bet with A-K twice before. In which case, if I bet, he might think I'm just betting because I think he'll think I must not have A-K. In which case, he might call or even raise.

This is another complicated line of thinking, now involving a thought of the 4th order: you have a thought in which you attribute to T. J. a 3rd-order thought, namely, that you're just betting because you think he'll think you must not have A-K. This thought that you attribute to T.J. is a 3rd-order thought, because it's a thought in which he attributes to you a 2nd-order thought, namely, that he'll think you must not have A-K. Confusing, I know. Adding subscripts to keep the orders straight, your 4th-order thought again is:

Thought₄: If I bet, he might think₃ I'm just betting because I think₂ he'll think₁ I must not have A-K.

We can already see the Davidsonian thesis at work. In order to have good reason to attribute to T. J. the 3rd-order thought you do, you have to assume that you and he agree on many things. Some of the beliefs you must assume you and he share are:

- that A-K is a hand people like to raise with;
- that one strategy players sometimes employ when the flop comes rags after they've raised before the flop with A-K is to bet, hoping their opponent will fold;
- that twice before, you didn't bet with A-K in this situation;
- that T. J. *remembers* that twice before you didn't bet with A-K;
- that *you* remember that twice before you didn't bet with A-K.

And so on. If there were any *one* of these beliefs that you didn't have or didn't think T. J. had, you wouldn't be justified in forming the 4th-order thought.

The idea that successful thought attribution depends on shared belief explains the surprising fact that an experienced player will sometimes fare quite poorly against a novice. The pro cannot assume, for instance, that the novice knows that some players are more apt to bet with A-K when the flop comes rags than others—or that it's significant that the flop has come with all low cards. It can thus be difficult for pros to deduce what a novice would think if the pro were to act in such-and-such way. And if the pro did assume his opponent shared some of his own sophisticated thoughts about A-K (which he might well do, before he realizes his opponent is a novice), he might undermine his own cleverness. (In "Odds or

Evens,” the kiss of death is being *two* steps (not one) ahead of your opponent.²⁾ But perhaps I’m just trying to convince myself that I’d have a chance against a pro. After all, the pros would still be able to capitalize on all my physical tells. And unfortunately, I tend to giggle when I have a good hand. (Or do I?)

Back to your mental sparring with T. J., which might not stop at the 4th order. You might well take it one step further. After thinking₄ *If I bet, he might think₃ I’m just betting because I think₂ he’ll think₁ I must not have Ace-King*, you might then think₆:

Maybe he’d think₅ I’d think₄ this, in which case I should bet.

That’s a thought of the 6th order. Now what about *your* thoughts—that is, my patient reader’s—in reading this essay? The thoughts involved in your comprehending the last few sentences were of the 7th order. Those sentences were about a thought of the 6th order, so your *comprehending* them was of the 7th. I could have titled this essay “Thinking about Thinking about Thinking about Thinking about Thinking about Thinking (about Poker)”. (And if you understand why, then I could have added another “Thinking”.)

You might wonder whether the pros really do think these long, intricate thoughts at the tables. Spelled out explicitly, the 6th-order thought is:

Maybe he’d think₅ I’d think₄ that if I bet, he might think₃ I’m just betting because I think₂ he’ll think₁ I must not have Ace-King, in which case I should bet.

Do the pros really say all this to themselves at the table? Well, that depends on whether *saying* all that to themselves is required for their having that thought. Precisely what is involved when someone thinks a thought? Sentences in the mind? Words? Mental images? What *is* thinking, for that matter? Or, as Oxford philosopher Gilbert Ryle (1900-1976) is said to have articulated the question, “What *is it* that the man in Rodin’s sculpture is *doing*?” Thinking about Big Slick of course, but Ryle wasn’t asking what the man in the sculpture is thinking *about*, but rather what his thinking—or any thinking, really—*is*, or *involves*. According to some players, feeling and the “subconscious” play a significant role in their thoughts at the table. Crowd favorite Doyle “Texas Dolly” Brunson writes about a particular situation, “. . . even though I might not *consciously* do it . . . I recall that this same play came up (or something close to it) and this is what he did or what somebody else did. So I get a feeling that he’s bluffing or that I can make a play here and get the pot. But, actually my subconscious mind is reasoning it all out.”³ But these matters bring us quickly to the heart of some central debates in contemporary philosophy of mind and psychology.

Thinking about Daniel Negreanu

By emphasizing the number or height of the levels of thought employed in expert poker games, I do not mean to nourish the perhaps specious assumption that those who excel at

² Since in this game there are only two possible moves (throwing one finger or two), being two levels of thought higher than your opponent will lead you right back to throwing what he wants you to throw.

³ Doyle Brunson, *Super System* (New York: Cardoza Publishing, 1978), p. 430.

poker think at extremely high levels—as if T. J. Cloutier is so successful because he can go to the 16th, while most other pros go only to the 14th or 15th. My sense, actually, is that the levels are not so high. Not that the thoughts are not complex or impressive, but that more crucial to expert play is making *sets* of thought attributions, each of which is typically of the 4th order or below, and then inferring what those attributions together imply.

Consider a hand that professional poker player Dan Harrington deems the most ingenious in recent history.⁴ It was down to the last two players at the 2004 Plaza tournament, Daniel Negreanu and Freddy Deeb. Negreanu, with \$337,700 and in the little blind, has A♥-7♦ and calls the \$1,600 big blind. Deeb, with \$342,300 and in the big blind, has A♠-K♦ and raises \$7,000. Negreanu calls. The flop comes K♠-6♥-2♥. Deeb bets \$16,000, and Negreanu calls. The turn is the 4♣. Deeb checks, Negreanu bets \$30,000, and Deeb calls. The river is the 4♥. Deeb bets \$65,000, Negreanu raises \$100,000, and Deeb folds.

Negreanu had the weaker hand but still managed to win the pot against an opponent who held A-K and flopped top pair. How did he do this? You may think Negreanu simply got lucky, that he bluffed on the river as anyone can, and his opponent happened to fold. But this is far from true. In order for a bluff to be successful against an experienced player who holds a decent hand, many things need to be in place. Two crucial ones are:

- (1) that your opponent believes that, given what you've done on each round of betting so far, the only hands you could reasonably have would all beat him.
- (2) that your opponent is capable of folding when he believes this.

The great players are those who can recognize the rare situations in which they know (2) is true and can *make* (1) true. As Harrington sees it, Negreanu believed that if he raised \$100,000 on the river, Deeb would review what Negreanu had done on each round of betting and would conclude that the only cards he could reasonably have, consistent with how Negreanu acted on each round, would make Negreanu either a flush or full house.⁵ But why raise only \$100,000? \$100,000 is a small raise given the amount of money already in the pot. The answer is that Negreanu knew that Deeb is particularly leery of “lure” bets (relatively small bets made to lure a call). He also believed Deeb knew that he (Negreanu) is capable of making such a bet when he has a good hand. So Negreanu figured that if he made a small raise, Deeb might think he was trying to lure a call, which would also be consistent with his having a flush or full house. Talk about more moves than a mongoose!

The reason I've traversed this hand so thoroughly is to show that Negreanu's thinking appears not to involve any thoughts higher than the 4th order. His central thoughts were probably:

- Thought₄: Deeb knows₃ I'm capable of thinking₂ that a smaller bet is more likely to make my opponent think₁ it's worth it for him to call.

⁴ Dan Harrington and Bill Robertie, *Harrington on Hold'em*, Vol. 2 (Las Vegas, NV: Two Plus Two Publishing, 2005), pp. 435-42. I am indebted to Harrington's work for the following analysis.

⁵ For the complex line of reasoning Harrington attributes to Negreanu for this conclusion, see Harrington, pp. 437-41.

- Thought₂: Deeb's particularly likely—more so than other players—to think₁ that a small bet is a lure bet and to fold his hand in response.
- Thought₁: Deeb can't beat a flush.
- Thought₂: If I raise, Deeb will go₁ through everything I've done so far in the hand and conclude₁ that my most likely hands are a flush and a full house.⁶

From this set of thoughts Negreanu reasons that if he bets \$100,000, Deeb will fold.

We can thus distinguish between two skills the expert player must have: the ability to make individual thought attributions, and the ability to deduce from all the attributions he's made what play would be the most effective. The second is no less important than the first. Yet another skill is the ability to *anticipate* when those rare sets that allow for ingenious plays like Negreanu's might materialize later in the hand. Negreanu had little reason to call on the flop unless he anticipated the possibility of successfully bluffing later on. Indeed, he knew which cards would afford him precisely that opportunity.

What sets of thought attributions allow for such clever plays? It might be fruitful to lay out a variety of these plays and scrutinize the relations among the sets involved. For instance, Negreanu remarked about a play Erick Lindgren once made against him, "I didn't think he could be so stupid. But it wasn't stupid. It was like a step above. He knows that I know that he wouldn't do something so stupid, so by doing something so quote-unquote stupid it actually became a great play."⁷ How does the set of attributions Lindgren made and capitalized on compare to the one Negreanu employed against Deeb? Do they share a similar structure? If valuable sets fall into different groups, would a taxonomy of them help players identify and anticipate them more easily at the table? Are there promising sets that remain undiscovered? This may be fertile territory for the serious poker player.

My sharp reader may protest: How do you know Negreanu's thoughts about Deeb didn't go higher than the 4th order? For instance, maybe Negreanu thought₆ about the possibility that Deeb would reason₅ that Negreanu had these thoughts₄ but concluded₆ that Deeb wouldn't reason₅ this way. Or maybe Negreanu thought₈ that Deeb would think₇ Negreanu would figure₆ Deeb *would* reason₅ this way and so would fold to a small bet. Perhaps Negreanu did think all this—it's consistent with all the data. Only the pros can tell us whether their thoughts go this high. Harrington never intimates that they do. But then again maybe that's all just to keep us from knowing what these guys *really* think. You never know in poker—which is why it's all so interesting to an epistemologist.⁸

⁶ Some of these thoughts are based on yet other attributions, again of the 4th order or below. For instance, Negreanu makes this last attribution on the basis of thought attributions such as *Deeb couldn't think I have K-6 off-suit, because he has to assume I wouldn't have called his pre-flop raise with that, or He couldn't think I have pocket 5's, because he's got to figure I would have folded on the flop.*

⁷ Kevin Conley, "The Players," *The New Yorker*, July 11 & 18, 2005.

⁸ Epistemology is the branch of philosophy that studies, among other things, what is required for having justification to believe something.

It is true that no matter what level your present thought about your opponent is at, you can *always* wonder: “But then again, he might *think* I think this.” My point is that taking it one step higher than your opponent in this way may not be the central element of poker, as it perhaps is in other games, like Odds or Evens, or even Rock, Paper, Scissors. (Incidentally, it’s becoming potentially lucrative to excel at Rock, Paper, Scissors too. In the spring of 2005, the president of a Japanese electronics company organized a game of Rock, Paper, Scissors between Christie’s and Sotheby’s, in order to settle which auction house would sell his company’s \$20 million art collection.⁹ Christie’s won. Christie’s’ people had researched the thought attribution that is typically done in Rock, Paper, Scissors—they talked to 11-year olds (really)—whereas Sotheby’s had assumed the game was simply “a game of chance.”)

Still, if your opponent anticipates what you’re thinking, you’re in trouble. Here are two things to keep in mind when wrestling with how and when to stop those “But maybe he thinks I think this” iterations. First, if you’re playing someone who can out-think you, one tactic is to choose “randomly.” You might use the second-hand on your wristwatch. In “Odds or Evens,” you could throw two fingers when the hand is past the Six, and one when it’s not. (My sister caught on to this one sooner than I’d have liked.) Choosing randomly disarms your opponent’s skill entirely. Likewise, if you wish to bet your A-K only 25% of the time, bet it only when the second-hand’s past the Nine. Second, the more complex your own reasoning, the less confident your opponent can be that your reasoning is that complex, even if the possibility that it is crosses his mind, and so the less prudent it would be for him to risk his chips on his suspicion that it is. Even if it *had* occurred to Deeb that Negreanu may be making the ingenious move he made, it would have been difficult for Deeb to risk the rest of his chips on this thought. The possibility that Negreanu was thinking everything that he was had to be, from Deeb’s vantage point, more remote than the possibility that Negreanu had a flush or full house. Yet another skill of the great poker player, then, is the ability to detect situations in which, even if your opponent *does* intuit what you are up to, he cannot prudently act on that intuition.¹⁰ Notice, finally, the depth of agreement that must again exist for a play like Negreanu’s to be successful: *making such a move requires putting the other person in a position where you both agree on what he has to do.*

How Does Negreanu Do It???

All of this leads at last to what is the most impressive and, to me, bewildering element of good poker. How do you know what your opponent is thinking? You may think I’ve already addressed this epistemological question, but I haven’t really. I discussed the nature of some of the thoughts you might attribute to your opponent when you have A-K and the flop comes rags, but I didn’t say which of them you *should* attribute. I explained that expert plays typically involve drawing inferences from a set of thought attributions of the 4th order and below, but I didn’t indicate how and when one is justified in making the individual

⁹ Carol Vogel, “Rock, Paper, Payoff: Child’s Play Wins Auction House an Art Sale,” *New York Times*, April 29, 2005.

¹⁰ It should be clear by now that, contrary to a common supposition, what makes a good poker player is not solely or even primarily an aptitude for “math” or probability. Au contraire! It is necessary to be proficient with probability, but it is certainly not sufficient.

thought attributions in the first place. Or to put it a different way, how on Earth did Negreanu know all that about Deeb?! Even if Davidson is right that you need to begin by assuming a great deal of agreement, how much and about what?

In everyday life, we learn a lot about what people think from what they say. Even in the case of the rabbits, we had the natives' words to go on. But in poker, no one's going to tell you anything. And unless you've got a friend who already knows your opponent and can tell you how he plays—that he's loose-aggressive, say, or a calling station, a flounder, or a rock (or that your entire table is a rock garden, for that matter)—the only way to determine how your opponent thinks is to infer it from what he *does*.

But inferring how your opponent thinks from what he does is extremely difficult. And the reason goes back to those rabbits. Remember I said that what words you utter on a particular occasion is a vector of what you believe and what you mean? Similarly, what you *do* is a vector of what you believe and what you *want*. You might wait outside the restaurant because you believe your friend is going to be outside and you want to see him, or you might do so because you believe he's going to be inside and you *don't* want to see him. One and the same action could be the product of countless sets of beliefs and wants.

Suppose the guy two seats to your left raises before the flop. You don't know anything about him, except that he's wearing mirrored sunglasses, a green visor, and a t-shirt that says "Lucky Larry," so you pay close attention to see what you can learn. Before the flop, Larry gets called by one player, and then bets when the flop comes 2-6-9 (just as you were considering doing earlier against T. J. with your A-K). His opponent folds, and Larry wins the pot. What can you conclude about how Larry plays? Not much. His two actions (his raise and subsequent bet) could have been the product of many different sets of beliefs and wants. Larry might have raised with A-K, believed he was weak on the flop and wanted his opponent out of the hand. Or maybe he raised with the beautissimos pocket A-A and wanted a call. Or he had a middle pair (you know, like Snowmen or the Speed Limit), thought it was weak and wanted a fold. He might even have been trying to steal the pot before the flop with nothing but Woolworths, gotten lucky when a 5 and 10 came on the flop and then wanted a call. These are just some of the many combinations of belief and want that could have led Larry to act as he did. And that's all assuming he's a decent player, that he's rational, and so on.

Now, it is true that in this last example you did not get to see Larry's cards, and sometimes you do. And then it's easier. But not much. For instance, suppose Larry did get called on the flop. He and his caller then both checked the turn and river, and Larry showed A-K. Can you conclude now that he's the kind of player that typically bets A-K when the flop comes rags? Not really. After all, Larry could be betting here to make it *appear* that he's that kind of player. Or perhaps Larry has a policy of betting only once in twelve times with A-K and a ragged flop, and this was that once (his second-hand had passed the Eleven). Or maybe the reason he bet is that he had a read on what his caller was thinking. Sometimes you need to think about what your opponent thinks about what *another* opponent thinks.¹¹

Of course, the more you see Larry show Big Slick in this kind of situation, the more confident you can be that he's the sort of player that typically bets A-K when the flop comes

¹¹ This is different from the kind of thought attribution I've been discussing so far. Here you'd be attributing to someone thoughts not about *your* thoughts but about someone else's thoughts. This kind of attribution involves three people, not two. These are just two of the many forms of thought attribution human beings undertake. Self-attribution (attributing thoughts to yourself) is another.

rag. You must steer clear here, though, of that vexatious “Fallacy of Unrepresentative Sample”: whenever Larry decides *not* to bet in such a situation, he will often subsequently fold, and you will not take note of his A-K.¹² But more importantly, this precise situation just doesn’t happen often enough. Consider all the things that must be in place: Larry has to be dealt A-K, he has to raise before the flop, one other player (and only one) has to call him, the flop has to come rag, Larry has to bet on the flop, his opponent has to call, and finally, Larry has to show his cards at the end of the hand. Entire tournaments could go by without that happening even once to Larry.

So how, in just a couple of hours, does anyone ever learn how someone else thinks at the poker table? The answer, I presume, has a lot to do with patterns that players have observed in past experience—patterns like “People who play fewer hands are less likely to bet their A-K when the flop comes rag.”¹³ How many hands Larry plays *is* something you can observe in a brief period of time. Equipped with a battery of such heuristics (and no doubt they’re often more sophisticated), you can begin almost immediately to sketch schemas of particular players on the basis of them. Until new information about a player presents itself, at which point you’ll modify your schema in light of it, you attribute and act on the basis of the schema you have. This is another, quite different way in which shared thoughts are involved and assumed in thought attribution at the poker table. If such patterns of thought and play did not exist among poker players, this strategy would not be available.

Now, when you apply these patterns to new players, you’ll be taking a significant risk. There’s no way around it. There will be many exceptions to the patterns, and your application will sometimes back-fire. The exceedingly cautious might choose not to make such applications, but this is a losing strategy: these players won’t be using information that their opponents will. In a way, to be risk-averse in poker is to take the biggest risk of all.

The same is true when constructing a schema for a player *without* employing such patterns, as I was trying to do with Larry. With Larry, I was in each case hesitant to conclude that Larry often bets A-K when the flop comes rag, because there were always other possible explanations of his behavior. Two very important points: First, if I had known that those other possibilities occurred relatively infrequently in poker (if people hardly ever play Woolworths that way), I would have been more justified to assume that Larry often bets A-K. Knowing how common particular thoughts and plays are among poker players is of decided significance, as it can allow you to ignore some otherwise worrisome possibilities. This too is available from experience. Second, since there will *always* be other possibilities (no matter how much evidence you’ve amassed), every adjustment you make to a schema will require *assumptions*, which could be false. Again you must embrace risk. Risk is required even to *develop* a schema. This might be especially difficult for the cautious epistemologist.

But here’s the really tough but critical question. How confident must you be about a particular assumption for it not to be *too* dangerous for you to revise a schema in light of it? Intuiting this threshold is, I suspect, one of the great, unexplored skills of poker.

¹² You can’t judge a bag of apples from the apples at the top, for those at the top are the last to rot.

¹³ Another part of the answer involves physical tells (Larry may sniffle when he has a good hand). But many experts claim this is less significant an aspect of the game than is often believed.

There is a happy note to all this. And it has to do with the shared belief and agreement that is guaranteed at the advanced poker table. I mentioned at the outset that poker has swept the globe in the last few years. Almost six thousand people traveled to Las Vegas in July 2005 to play in the World Series of Poker. And that was just for the Main Event. There were 44 other events, including some for women only and some for seniors. Sitting at the tables were players from Vietnam, France, Lebanon, Canada, Brazil, South Africa, Iran, Sweden, Costa Rica, India, Australia, Japan, Russia, and many, many other countries. And for the most part, everyone at these kinds of events is decent to one another, often friendly, chatting about past hands, discussing strategy, ribbing one another. There is a palpable camaraderie. I hardly place the solution to the world's problems in the hands of poker, or poker players (though some are amazingly charitable with their winnings). However, it is at least worth noting that in our day it is a rare institution that places so many people from such disparate backgrounds in such amicable proximity. And an even rarer one that continually requires such deep agreement among them. It is nice to see. Or at least, I think I think this.¹⁴

¹⁴ Thanks to Darse Billings, Eric Bronson, Claire Ellis, Jennifer Ellis, James Engel, Preston Greene, Jed Grodin, Matt Manfredi, Tom McEvoy, and Alan Schoonmaker for their help on this essay.