Is there a perfect system for winning at poker?

Robert Northcott
University of Missouri-St Louis

I walked into the room with a plan. After many Friday-night games with friends this was my first tournament in a real casino, I had read the books, and now I had formulated the perfect Hold 'Em betting strategy. While the suckers followed their hunches, instincts and anecdotes, I would be cutting through all that loser-talk. I would be the robot, the machine, remorselessly accumulating. (That was just as well since I was also playing with money I did not strictly have, it having been borrowed out of next month's salary.)

Every year brings scores of books from poker experts. But are there really any can't-lose systems? Maybe just playing the mathematical odds is enough? Or can only formal game theory give you the ultimate best strategy? Once you start to explore these questions more deeply, it turns out you soon find yourself in the territory of modern philosophy of science. Let's begin with the very notion of poker odds in the first place.

Good luck, bad luck
The crucial hand this night turned out to be one where my opponent was staying in seemingly against all poker sense, chasing a very unlikely inside straight. I was set to win big, become the richest at my table, one of the leaders overall … when, by fluke, on the last card he actually made it. This was the definition of a bad beat – I had played correctly, the odds were greatly in my favor, but by pure luck the cards came through for the other guy anyways.

A first question is: what does 'luck' really mean here? *Ought* my opponent to have bet on completing his straight? The odds said that it was crazy unlikely. Most of the time he would *not* have completed his straight, therefore in this situation it was bad strategy to bet so much on such a low probability, and so the fact he *did* then complete it and win big was due to luck, not skill. The skill comes in playing the percentages correctly; the luck in how, after doing that, the cards then fall. No serious player can hope to compete at poker without a working knowledge of card odds. Should I bet 20 chips to win a pot of 200? If it's a 50-50 shot, of course a 20-for-200 gamble is good business, but a 1-in-a-100 shot makes it the play of a fool. In the long run, luck evens out, so a night spent pursuing only gambles with favorable odds will usually prove very profitable – and 'usually' is the way to bet if you want to win. Leave the long shots for the losers in the bar.

At first all this sounds very reasonable. But look more closely, and it seems there is also a sense in which these supposedly objective probabilities are actually *subjective*, reflecting merely our ignorance. In this particular case, presumably in reality the card my opponent was about to be dealt was 'decided' already, sitting waiting on top of the deck. That is, it was *already* certain that he would hit his straight, it's just that neither of us knew that. So in a funny way, therefore his bet *was* the correct move after all.

All the same, still there does remain something objectively compelling about the relevant probability calculations. When the first card is dealt from a full and shuffled pack, for example, an ace of spades really *is* a 1 in 52 chance of coming up, we feel, not 1 in 3, and the rational player needs to know that. What gives? A common move is to reason as follows: given the four face-up cards and my opponents' two hole cards, he knew the 46 remaining cards that the river card could have been. And out of those 46, only 4 (say) would have completed his straight. Therefore his objective odds were 4 in 46. Yet the subjectivist point remains – actually, in this particular case, his true odds were either 0 or 1, depending on whether or not the next card waiting to be dealt actually was one of the four he wanted.

And who's to say that 4 in 46 was the correct objective chance facing my opponent anyways? I don't mean that perhaps he had peeked at the card or otherwise cheated. Rather, imagine a trained magician who noted carefully the cards of the previous hand, and then further noted carefully the croupier's shuffling technique between hands. Close observation would enable the magician still to retain some information about the order of cards in the new hand, and hence objectively to know that the odds in this case were perhaps either slightly above or below 4 in 46. This new probability seems to be *more* objective than the crude old 4 in 46.

Next imagine an even more skilled magician, this time one able to track the shuffled cards even better and perhaps able to fix on a probability now very different from 4 in 46. Perhaps indeed, had she successfully tracked a key card, she might even have noted that that very card was the next one up, and so known that the chance of completing the straight was actually 100 per cent. Similarly, a computer equipped with a camera might also be able to reliably pick out the probability as either 0 or 1. In other words, which number we select as the 'objective' probability in fact seems to depend on very messy *un*-objective personal factors like just how good I am at tracking shuffled cards.

So did I really suffer a bad beat after all? Well, on the assumption that neither I nor my opponent was preternaturally skilled at tracking shuffled cards, yes I did. But if I had been more like the magician, then I might have known that the objective probability was actually rather higher than 4 in 46, and perhaps then I should have seen it coming that my opponent would complete his inside straight.

Imagine an incompetent beginner who forgets his own hole cards when calculating the odds of completing a straight and goes wrong because of that. No one would say he had suffered a 'bad beat'; rather we would just say he had been incompetent. So why should I say that I suffered a bad beat in my case? Was I not just incompetent too? Relative to a magician I was. Generally, all talk of luck and bad beats seems not to be absolutely objective after all. Rather, such concepts are only relative to whatever we happen to deem the normal level of observation skills.

These classic philosophical difficulties in making sense of objective probability tie into an even deeper issue, that of *determinism*. Is any event truly chancy, or is all uncertainty

a result merely of our ignorance? For example, if only we knew the exact micro-composition of a coin, the exact movements of the air molecules and the exact strength with which the coin had been flipped, could we not always calculate with certainty whether it would come up heads or tails? If yes, then saying that a coin flip is '50-50' is merely to express our ignorance of the relevant micro-details rather than to capture any deep physical fact about the world. Under the influence of Newtonian physics, for centuries many scientists and philosophers thought that deep down the universe really is utterly deterministic and predictable in this way. If only we knew every detail, nothing would be uncertain to us.[1] More recently though, quantum mechanics has often been taken to imply that perhaps the universe is fundamentally chancy after all, and uncertainty is *not* a symptom merely of our ignorance. Albert Einstein in turn famously rejected this latter view, quipping that 'God does not play dice'. But many others think he does (God, that is, not Einstein). The controversy continues today.

Here's one final thought for this section. What if we also apply the idea of determinism to the human brain? After all, our brains are presumably part of the physical universe too. But then in principle we should be able to predict other humans' behavior – I could know in advance whether my opponent would fold, and *really* know, with certainty. But then could I not also know my *own* future actions? In fact, what would there be for me to *decide* about at all? After all, all my future decisions, and all my opponents' too, would already be determined and perhaps knowable. In such a situation, poker would become rather boring. So it seems that what philosophers call 'decision theory', one example of which is poker strategy, makes no sense without the assumption of *free will*, that is without the assumption that we are free to decide our actions one way or the other just as we please.[2]

So why not simply assume free will when talking about poker then? After all, once outside the philosophy department that's what we usually do in the rest of everyday life. OK, but then our previous issue arises again – how can we justify declaring one poker strategy *objectively* superior to another? For as we have seen, as soon as we try to do this, we get entangled with messy issues of how much decision-makers know … including ultimately how much they do or do not know about their own and others' brains.

What follows from all this? That any winning poker system making use of 'objective' poker odds, is actually at best only winning relative to a particular level of ignorance or imperfect knowledge. That in turn leaves you, like all ignorant people, vulnerable to rude awakenings. And what we call a 'bad beat' is just our label for such an awakening.

What might have been…

---

[1] A hypothetical omniscient creature to whom nothing would be uncertain is sometimes called 'Laplace's demon', named for the French philosopher and mathematician Pierre Laplace who put forward this metaphor for determinism almost 200 years ago. Not coincidentally, Laplace was also a pioneer of the subjective interpretation of probability.

[2] Nor does this problem necessarily rest on our assuming determinism. For, as the American philosopher Daniel Dennett among others has pointed out, it is not at all clear exactly how a probabilistic quantum mechanics could be used to save free will either.

OK, so perhaps my best strategy was in fact best only relative to my imperfect knowledge, but nevertheless even once taking that point on board a bad beat is still galling. And the hand with the straight was especially galling precisely because it was a big hand, and losing it left me dead in the water and virtually chipless. If the guy had only not made that fluky straight, then I would have been winning…

That now brings up a thorny new issue. Did this one moment really cost me, as I imagined afterwards, the whole tournament? To assess that, we'd need to know what *would* have happened had I instead won that hand. How *would* the rest of the tournament have unfolded? However, this we can never know for sure since of course there is no way of directly measuring or observing things that never actually happened. How then can we say anything sensible about them? Such hypothetical situations are known to philosophers as *counterfactuals*, since they are 'counter' to the facts of what actually did happen. And they are notoriously tricky to handle.

Suppose that the hand in which my opponent made his fluky straight was worth X chips. Then clearly, if I had won it I would have been X chips better off. In order to see what would have happened after that, could we not just 're-run the tape' of what actually did happen, and add X chips to my score? But sooner or later this method would break down. In particular, in the real tournament I was quickly eliminated, after which of course it continued without me. But with an extra X chips I would not yet have been eliminated, so any tape running on without me would then no longer be right.

Perhaps, given the psychological and strategic impact if I had won the hand, subsequent hands would have worked out differently. My spirits would have soared, my opponents' (I hope) sunk, with my extra chip stack I could have commanded the table better, and in the end I would have won the whole tournament. Well, perhaps – and then again, perhaps not. We can guess, but how can we know which guess is the most reasonable?

Here's a further difficulty. In order for my opponent not to have made his fluky straight, the order of the deck would have to have been different to how it actually was. It seems to follow that this would have led to the order being different on subsequent deals too, hence that all the subsequent hands would have featured totally different cards to those that were dealt in reality. Moreover, my different reaction to winning rather than losing the hand with the straight would presumably have had some micro-impact on the dealer's brain, perhaps resulting in her shuffling the cards slightly differently, hence again resulting in totally different subsequent hands. Given these kinds of consideration, any hypothetical extrapolation starts to seem very difficult. Maybe we need to know in exactly what way the order of the deck was different in order to know just what these new hands would have been – for instance, how do we know that if I had won the hand with the straight I might not then have been dealt pocket aces in every other hand subsequently?

These and other complications meant that for many years philosophers were rather skeptical whether any rigorous evaluations could really be made of counterfactual claims at all. It's not just that it is difficult to find out the fact of the matter about them. Rather, it

is the more fundamental point of where there even *is* a fact of the matter – for by definition, of course, with counterfactuals there precisely are no *actual* facts we can appeal to directly.

Perhaps then, given these problems, the solution is just to cleanse ourselves of speaking about these messy counterfactuals? The trouble is though that it seems very difficult to do this, for we all unavoidably talk about counterfactuals all the time. Some examples: he wouldn't have got sick if he'd taken the medicine, if you drop the cup it'll smash, if you don't practice you won't get better, if you'd called and told me I could have done something about it, if the government passed this law then crime would soar … and so on. Generally, it is often impossible without thinking about counterfactuals for us to assign moral praise or blame, to explain some occurrence, or to decide between alternative actions. For instance, it is only right to warn you not to drop the cup if the associated counterfactual claim – that if you were to drop it then it would indeed smash – is also right.

Looking at the above examples more carefully, it seems that there is an intimate connection between counterfactuals and saying that something *caused* something else. (Exactly how intimate is controversial.) For example, claiming that dropping the cup caused it to smash might be read as saying just that *if* you had *not* dropped it then it would not have smashed. Similarly, when I nonchalantly claim that it was only my bad beat with the straight that caused me not to win everything, this can be read as the claim that *if* I had won that hand then I *would* have won the whole tournament. This aspect of causation was first mentioned by the famous eighteenth-century Scottish philosopher David Hume. In recent times, the late American philosopher David Lewis particularly championed the idea that counterfactuals in fact capture causation's very essence.

In support of this program, in the late 1960s Lewis and others proposed a clever and complicated formal system for evaluating just the kind of troublesome counterfactuals we have been looking at. In particular, he proposed an apparatus of 'possible worlds' in addition to our actual world, where the different ways that things *could* have been – but in fact aren't – still do occur, only in these parallel worlds instead. Some possible worlds are 'closer' or 'more similar' to the actual world than are others. A claim about a counterfactual can then be evaluated, roughly speaking, according to whether the possible worlds in which it is true are closer to the actual world than are possible worlds in which alternative counterfactuals are true. What exactly 'similarity' amounts to, and to what extent talk of possible worlds is merely metaphorical, are themselves topics of lively debate. Overall, this and other developments have succeeded in making discussion of counterfactuals philosophically respectable again. But, as is the way of philosophy, plenty of tricky counterexamples have been found both to Lewis's particular approach and to its successors, and it may be that other proposed methods for tackling the issue will eventually prove more fruitful. In any case, so far no single approach has achieved any degree of consensus support.

Therefore, regarding counterfactuals, as yet philosophers are only able to say: can't live with 'em, can't live without 'em. That is, even though we as yet have no agreed way for

ever actually deciding whether one is right or not, still we can't help invoking them all the time.

Among other things, I'm afraid it follows that all those endless arguments in sports will … remain endless, which may be good news for TV sportscasters but bad news for everyone else. Thus: he'd have won if he'd made that shot, he'd have won if he was mentally stronger, what a great career he would have had if he hadn't got injured, they'd never have recovered from 3-0 down, etc etc. Well, no one can ever know! And no one can ever know either whether I really would have won that tournament if it weren't for the bad beat – which is convenient, because that means I can carry on forever explaining to you all why indeed I would have…

Finally, all this has implications not just for evaluating whiny hard-luck stories but also for any winning system at poker. Put most simply, why should I play your system? Your answer must be, at some level, that by doing so I would stand a better chance of winning than otherwise. Or, to phrase it differently, perhaps playing your system will *cause* me to have a better chance. Either way, your pitch makes no sense without some appeal to counterfactuals.

Rationality and psychology
In poker, it can seem that rationality is everything. As we noted, a good working knowledge of the basic odds is essential for any serious player, and certainly for any winning strategy. But is just playing the odds enough? For poker lore is full of the importance of psychology. Well, maybe poker lore is just full of it. On the other hand, maybe psychology really does matter? It turns out that, in practice at least, indeed it does, and in a way that is personal to each particular player. This is therefore another reason to be suspicious of any one-size-fits-all perfect system.

The root of the issue is this: poker's not just a matter of playing the odds, it's also a matter of having the self-discipline to *stick* to playing the odds. Of course, even then you may still lose, perhaps through bad luck, perhaps through an opponents' superior skill. But you can be a lot more certain of losing if you abandon discipline and go on tilt.

One way that happens is through tiredness. Another is through being drunk or stoned – it has been well observed that when drunk it is not so much that you do not still want to win, more that you no longer feel so bad about losing… and that is all it takes to forget about boring discipline and to start trying on fancy plays and spectacular bluffs instead. Even if they don't come off, hey, the champagne's still nice.

Similarly, there are occasions when the problems are inattention and distraction. I once made the mistake of joining a game where I was strongly attracted to one of the other players. If you're concentrating on how to impress her, you're not thinking about the odds and you're not tracking the others' bets. Maybe that's OK, as frankly you care more about impressing her than about the poker. Fair enough, but that does mean you'll fail at the cards. And as it turned out, I failed at impressing her too.

The classic way though to go on tilt is to suffer a bad beat. Maybe some of you never play drunk or horny, but certainly all of you from time to time will suffer bad beats. Feeling angry about your tough break is just another way in which you're not concentrating on the cards. It has been shown that often players lose far more while on tilt afterwards than they did in the initial bad beat. For this reason a key skill for professional poker players in Vegas casinos, for instance, is the ability to recognize when a punter, preferably a rich one, has gone on tilt. Then typically lots of them suddenly materialize around the sucker's table, sharks smelling easy money.

An even more important skill is to be able to recognize when you *yourself* are on tilt. All of us are human, all of us feel angry after a bad beat, therefore all of us go on tilt at least for that moment. The trick is to be able to recognize this, and to follow a strategy to recover. Many expert players, for instance, know that they must leave the table in such circumstances. That is, they know from experience that this is the best way to manage their own temperament. If it is impossible to leave the table, say because you are in a tournament, then alternative tricks of emotion management need to be learnt. It has been shown that usually the very worst response is to think you can play through such episodes by sheer force of will, denying to yourself all the while that you are still on tilt – even though of course you are. After all, no one *not* on tilt ever feels the need to deny it.

In other words, an important component of rational play turns out to be a kind of 'meta-rationality' – that is, being rational about your own (hopefully temporary) *ir*rationality. This is something that formal decision theory finds difficult to incorporate. For on the one hand you need to manage your anger rationally, but on the other hand we are thereby simultaneously assuming that part of you – the angry bit – is incapable of acting rationally in the first place. Technically, decision theory should just be advising you not to be angry and simply to keep on playing the odds. 'Rationality' can thus mean two different things: first, playing the odds soundly; or second, managing your psychology optimally given that you're temporarily incapable of playing the odds soundly, i.e. temporarily incapable of being 'rational' in the first sense.

One other famous example of relevant psychology is the ability to detect 'tells', i.e. unconscious pieces of body language or behavior that give away a player's mood and hence indirectly their hand. When playing with strangers, often it may take a while to learn to read these. In philosophers' jargon, this is a typical problem of causal inference. You might be learning, say, whether an opponent's blink was a symptom of their disappointment, and it takes time to eliminate other possible hypotheses – for instance, that in fact it was just due to a piece of dust and actually signified nothing important at all.

And what if you yourself have a bad poker face? Well, you can always play on the internet.

Summing up, it seems that a winning strategy is not just a narrow matter of playing the odds correctly, or at least it isn't if it's supposed to actually help you win. Rather, we

need to understand 'rationality' in a broader sense that also includes aspects of psychology too. And when it comes to *your* psychology, you know things better than I do – and better than any pricey poker manual does either.

Game theory

Maybe there is one branch of decision theory though that could overturn all the preceding emphasis on psychology. For in principle it can yield a complete optimal strategy for poker that we could all thereafter just follow mindlessly. Poker would become a matter merely of mechanically following this single set of golden rules, and psychology would be rendered as irrelevant as if we were playing tic-tac-toe. This is the promise of *game theory*.[3]

Notwithstanding its name, game theory is not just about frivolous parlor games. Rather, it offers a rigorous and general mathematical treatment of situations of strategic interaction, in other words of situations where your best move depends in part on what *other* people do. Of course, in turn what they do depends in part on what they think *you* will do, so what you will do depends in part on what you think they think you will do, so what they will do depends in part… etc. It is one of the achievements of game theory to cut through this seemingly endless chain of second- and third-guessing to arrive at concrete optimal strategies: all things considered, what you *should* do in this situation is such-and-such. Poker is a particularly good example of strategic interaction, for your optimal betting strategy depends in part on how the other players will bet, which in turn depends in part on how they think you will bet… etc. Indeed it is said that the inspiration for the development of game theory came to its founder, the Hungarian mathematician John von Neumann, precisely while playing in his poker school. (Von Neumann, incidentally, was later the alleged model for the movie character Dr Strangelove.) And many of the pioneers of game theory in the 1940s and 50s, such as John Nash of *Beautiful Mind* fame, were also inspired by it.

The dream is that one day someone will work out for poker a complete game theoretical analysis. Unfortunately the mathematical complications multiply horribly swiftly, and so far no one has got anywhere near achieving such a thing. For now then, you can safely play against a game theorist without fear. Some progress was made in the early days but only with respect to greatly simplified versions, typically with only a single bet each for instance and only one level of permitted raise.[4] Nowadays, even these efforts have rather petered out.

But suppose one day someone did hit the jackpot and a complete analysis of poker appeared, that is of one of the versions played by real people and not just mathematical models. Would that then spell the death of the game? Would we no longer see ESPN specials on the world series of poker, just as today we don't see ESPN specials on the

---

[3] Note though that even this would only address worries concerning the influence of psychology, not the earlier ones concerning counterfactuals and objective probabilities. Note also that many game theorists are of course aware of the further difficulties about to be pointed out in this section.

[4] Interestingly, already in these games it proved optimal to bluff some of the time.

world series of tic-tac-toe? Or maybe the clever theorist would keep his or her knowledge to themselves like a secret golden code, all the better to astonish us by mysteriously sweeping all before them at every table they ever visited? But one recent branch of philosophy of science suggests that all this is overstated. In particular, once we look in detail at how abstract scientific models like those of game theory are applied to real-world phenomena like actual poker games with actual human players, we see that a 'complete' analysis, even were it ever to be discovered, would in reality not be so complete after all.

Begin by understanding that in game theory an 'optimal strategy' is just the best possible strategy given that everyone else is playing their optimal strategies too. Thus in principle the only stable outcome, i.e. the only situation in which no one could improve their chances by changing strategy, is when *everyone* is simultaneously playing their optimal strategies. (This is known as a 'Nash equilibrium'.) Say, for the sake of argument, that in the case of poker this optimal strategy turned out to be some complicated function of your own cards, other players' bets, and the overall chip situation. Call all the players following this strategy Theorists. But let's suppose now that one obstinate player, let's call him Stupidus, is *not* playing the prescribed optimal strategy. Suppose for instance that instead Stupidus insists on always initially raising the maximum, regardless of his cards, but then folding if he is ever called. I think we have all played against characters like Stupidus at one time or another, including on occasion my own opponents...

Notice two important points. First, it is unlikely that the optimal strategy played by Theorists is any longer optimal once pitted against Stupidus. A much better play against Stupidus would be always to call him, thereby winning his maximum raises every time without a fight. Second, not only would the Theorists not be exploiting Stupidus as much as they could. It's even worse than that – they might (depending on the exact parameters) actually be *losing* to him. For, most of the time, Theorists would fold meekly against Stupidus's initial aggression, and when with good cards they finally did make a game of it then Stupidus would just drop out immediately rather than incur big losses. An optimal strategy is only 'optimal' against other optimal strategies. In general, it cannot even guarantee winning, let alone winning optimally. So who is stupidus now?

The only way out would be if the optimal strategy were able to recognize players such as Stupidus and adjust its play accordingly. But given that there are literally an infinite number of possible strategies to sift through, this complicates the mathematical task enormously – a task, remember, that is *already* intractable. There would also now be the problem of just how you would recognize exactly which strategy an opponent is playing. For instance, it would take many rounds to get a sense of what Stupidus was up to. Technically, it seems unlikely the task could be completed for sure in a finite time. And this doesn't even take into account the possibility of Stupidus *varying* his strategy through the game.

So even granting the fantasy of an optimal strategy being discovered for real-life versions of poker, still that would leave us a long way from knowing how to win in actual games against actual players. It seems therefore that in the end game theory will offer no real

replacement for those grizzled rules of thumb and salty poker wisdoms after all. We'll all just have to carry on thinking about the usual suspects: seating order relative to active and passive players, tells, poker probabilities, avoiding going on tilt, etc. Ultimately, for better for worse, and indeed for richer for poorer, not even the appliance of game theory science can furnish us with that can't-lose strategy.

What has philosophy ever done for us?
Therefore while poker may equip you nicely for introductory courses in philosophy of science or in metaphysics, unfortunately it seems philosophy cannot really return the favor. In particular, it cannot tell you how to win – I'm sorry about that. Maybe it can tell you though why to beware hucksters pitching miracle perfect systems. There is no such thing as a single objectively perfect strategy that will win for everyone against anyone at any time. Maybe some plans may at some moments be more useful than others, but always you should make room too for individual variation and special circumstance. And don't forget, in any case there is no way anybody can guarantee the truth of the counterfactual that *if* you were to follow their perfect plan then you really would win. Perhaps in a way this is all good news really, for it suggests that poker will never be reduced to a game for robots.

Nevertheless, in poker clever schemes spring eternal, as do hucksters. For my part, this year I have my own new plan again, and I'm looking forward to trying it out against opponents. I hope one day one of those will be you.