

Is it bad luck or the ref's fault?

Robert Northcott

My son and I couldn't believe it. England out on penalties yet again. Months of anticipation ended by heartache and misery, another generation's World Cup dreams dashed forever in the cruelest way. Because a penalty shoot-out is just a lottery, a matter of who gets lucky on the day. We all know that, yes?

But luck is a slippery concept. When we look more closely, it's impossible to make good sense of it without engaging a whole range of philosophical issues surrounding probability and explanation. Is luck really something objective, or just a convenient scapegoat? Let's see.

What is Luck?

I am English (couldn't you tell?). England has a particularly miserable record in penalty shoot-outs. For some time, we seriously argued that there was no point even in practicing them. In the face of that attitude, persistently losing shoot-outs starts to seem less like bad luck and more like simple negligence. Now even England practices them beforehand—but there is preparation, and there is preparation. These days, it is becoming common for goalkeepers to have all previous penalties by the opposition players on film, even viewing them on an iPod during the event (as Ben Foster did to help Manchester United top Tottenham in the 2009 Carling Cup). Researchers have pinpointed the optimal ordering of takers, as well as precisely what kind of penalty is most likely to succeed.¹

If you don't take advantage of these methods, twenty years ago maybe you could shrug your shoulders and blame bad luck but these days that won't cut it. Imagine entering a penalty shoot-out fielding a goalkeeper who was blind. No one would then blame the presumably inevitable defeat on bad luck, as it is obvious that being able to see improves your odds greatly. But there is good evidence that these other methods of preparation also improve your odds. So why should going in without any practice be thought any less culpable than going in with a blind man?

At the root of the notion of luck is a second one, that of *probability*. 'Bad luck' implies that it isn't your fault. The probabilities were in your favor, nothing more could be done, the dice just happened to fall the wrong way. Perhaps you prepared for a shoot-out super-professionally but by chance your keeper happened to dive the wrong way each time and so you lost anyway.

How far can we push this? The key phrase here is 'by chance'. It implies that, assuming your keeper has prepared all he reasonably can, he can't then be blamed if still he guesses the wrong way. But is any event truly chancy, or is all uncertainty merely a result of our ignorance? Is it really true that the opposition's penalties couldn't have been predicted better? Maybe a super-advanced alien scientist, by minutely examining the kicker's brain, could have predicted successfully which side he would shoot. Or, similarly, which way a goalkeeper would dive. In which case, diving or shooting the

¹ Since you ask, apparently a shot into the top metre of the goal scores 99 per cent of the time. This suggests that, if skilful enough, players should be perfecting that. And the optimal order is for your best penalty-taker to take the fifth penalty, the second best the fourth, and so on back to the first.

wrong way is not bad luck rather it is merely incompetence. After all, the alien scientist would never get it wrong.

Presumably, no one considers it reasonable to expect quite *that* degree of preparation. But maybe in the future a keeper who dives the wrong way will seem no more acceptable than one who is blind. In other words, all talk of luck seems not to be absolutely objective after all. Standards of judgment change. What was bad luck twenty years ago may just be negligence today, even though the physical event is exactly the same. Rain delays at the Wimbledon tennis championships used to be blamed on bad luck with the weather. More recently though, the reaction became, “why don’t they build a roof?” And so, they did. The point is that all depends on what we judge a normal or reasonable level of preparation to be.

This doesn’t apply only to penalty shoot-outs, of course. Is hitting the post bad luck, or just inaccurate shooting? Is getting injured bad luck, or just inadequate conditioning or concentration? Again, it all depends on what is considered normal or reasonable.

Fate and Determinism

These difficulties tie into a deeper issue, that of *determinism*. Is it true that, if only we knew the exact micro-composition of a coin, the exact movements of the air molecules around it and the exact strength with which it had been flipped, we would then know whether it was going to come up heads or tails? Similarly, if only we knew the exact state of every neuron in a goalkeeper’s brain, could we not then calculate with certainty whether he would dive right or left? If yes, then saying that a coin flip is 50-50, or that it’s 50-50 which way a keeper will dive, 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 everything in the universe really is determined in this way. If only we knew every detail, nothing would be uncertain to us. More recently though, quantum mechanics has been taken to imply that perhaps the universe is fundamentally chancy after all, and uncertainty is *not* always a symptom of our ignorance. The controversy continues today.

In the case of penalties, we are applying the idea of determinism to the human brain. After all, our brains are presumably part of the physical universe too. But if, alien-style, I was able to predict other humans’ behavior, could I not also predict my *own* future actions? In fact, what would there be for me to *decide* about at all? That’s fatalism. My future decisions, and all my opponents’ too, would already be determined. I could know in advance exactly what would happen and who was going to win. Penalty shoot-outs would become rather boring. So would all of sport for that matter, and plenty else besides.

It is often complained that the World Cup’s knockout format does not reliably identify the best team because it is too chancy. But, in light of determinism, what is the “best” team? One that would have won, say, sixty percent of tournaments but, as it happened, not this one? But perhaps if we knew the state of the world in enough detail we’d also know for sure who was going to win the tournament, no doubts or percentages required. Any sixty percent figure would then be just a reflection of our ignorance, as

would any “best” team claim—that team would merely have *seemed* best to us, not actually *been* the best.

A hypothetical omniscient creature to whom nothing is uncertain is sometimes called “Laplace’s demon”, named for the French philosopher and mathematician Pierre Laplace who put forward this metaphor for determinism almost two hundred years ago. Such a demon would have no truck with sixty percent chances of winning. To its all-knowing eye, everything would be either a one hundred percent chance or zero.

Can Luck Be Objective?

So is there nothing objective to luck or probabilities after all, rather just ignorance? That conclusion would be too hasty; we mustn’t throw out the baby with the bathwater. It is true that all depends on whatever level of knowledge or preparation we consider reasonable. And, in turn, what we consider reasonable may be partly arbitrary—but also partly not. Let’s focus on the glass half-full. Perhaps talk of luck can be rehabilitated as objectively respectable. Sure, it has to be relativized to a particular understanding of reasonableness but if everyone shares that understanding, where’s the problem?

This approach has promise. The details of defining this relativized variety of objective probability turn out to be tricky, but all we need here is the rough idea. When we can agree on what’s reasonable and what isn’t, things start looking good again for objective judgments of luck. And often we *are* able to agree plenty enough. Consider a single penalty. Everyone accepts that, in the absence of bizarre alien technology, it is not reasonable to expect fantastic mind-reading skills. For that reason, assuming no other effective method is available either, all would agree that it is indeed a matter of luck whether the kicker guesses correctly which way the keeper will dive.

Here’s another example: “penalty shoot-outs are more of a lottery than a regular game.” The best way to examine this is to think about the chances of a good team beating a bad one. If that chance is 75-25, say, then outcomes are likely to reflect merit rather than mere luck; if only 52-48, then we are clearly in lottery territory. So the million-dollar question becomes: what are the numbers for a regular game and for a penalty shoot-out? It’s only fair to castigate shoot-outs for being a lottery if it’s their number that’s closer to 50-50.

Such numbers only make sense once we understand what ‘good’ and ‘bad’ mean. In the shoot-out case, perhaps a good team corresponds to one that makes all the specialized preparations mentioned earlier, plus maybe having a goalkeeper especially good at saving penalties. And perhaps even—a factor beloved of amateur psychologists everywhere—the confidence that comes from a good history in such shoot-outs. Define a bad team as one lacking all these attributes. What are the odds then of the good team winning the shoot-out? It’s hard to estimate precisely without further research but it seems plausible that although the good shoot-out team would be favored, a good regular team would be favored in a regular game even more. In that case, we can say—objectively—that shoot-outs are indeed more of a lottery than regular games. (So perhaps England can pin some of the blame on bad luck after all.)

Notice a few further points here: first, this doesn’t mean that shoot-outs are a *total* lottery. They’re not 50-50 shots; just closer to 50-50 than are the regular games. (So it’s probably not *just* bad luck that England have lost so many.) Second, it’s perfectly possible for the same team to be rather bad at the regular game but pretty good at shoot-

outs. They may therefore often win a shoot-out even after having been poor otherwise. But that's not necessarily luck—that's just being good at an especially vital aspect of World Cups. Blame the rules committee, not the winners. Third, maybe you disagree with some of my values for these numbers? OK, but the real point is to nail down exactly which numbers are relevant. How you then fill in their values is a separate matter.

All this also raises a separate question: What figures should we *want*? Presumably 50-50 would make a sport pretty boring because it would be totally fluky, akin to a World Cup of coin-tossing. But equally, in the eyes of many, 100-0 would be pretty dull too—there should always be some room for an upset. Perhaps how best to strike the balance is ultimately just a matter of taste. Different sports offer different answers. The best player wins the tournament frequently in tennis, but only rarely in golf. Baseball is nearer the chancy end, as in a regular game a good team will beat a bad team at best 55 or 60 per cent of the time. Rugby and American football are rather more predictable. Football seems to be somewhere in the middle.

Whose Fault Is It?

So we see that we can't always blame everything on bad luck. What (or who) else can we blame it on, then? It turns out to be tricky to pin everything on anything. Playing the blame game is easy to do but irritatingly difficult to justify.

There are few things more fun than whining about referees, for instance. But does it ever really make sense to pin the blame for defeat on one refereeing decision? After all, any defeat has myriad causes, such as the other goals, missed chances, injuries to star players, and so forth. Aren't these other causes equally as important? Claims of bad luck can be judged objectively if—but only if—we agree on standards of reasonableness. Something similar now turns out to apply when assigning blame too.

A recent strand in the philosophical literature defines a *cause* in terms of deviations from a default state. For example, physically speaking a stationary football's default state is to remain stationary and undisturbed, like a sleeping dog. A cause is then any deviation from this default state, in other words a disturbance that makes the ball move, such as being kicked. That may just sound like a roundabout way of saying that kicking the ball caused it to move, but the real point is how things are relativized to what we take to be the 'natural' default state. After all, strictly the ball's physical robustness was also a cause of its moving—otherwise, upon being kicked the ball would merely have shattered rather than moved. But we take this robustness for granted and thus it is disregarded as a cause. Because virtually all of us share this judgment of what's default and what's not, so we also all share the thought that being kicked—rather than its own robustness—is what caused the ball to move. If, in contrast, we had thought that a ball's robustness naturally fades over time, then the deviation from the default would have been the ball's *sustained* robustness and that *would* now have been fingered as a cause, just as we are happy to give causal credit to a marathon runner's unusual stamina.

So, alas for family values, deviance is where it's at—no deviation, no causation. Let's bring this back to referees. If our default expectation is that a referee gets easy decisions correct, then it's certainly a deviation from that if he refuses our team an obvious last-minute penalty. Such a critical mistake is endorsed as indeed a cause of our defeat. But many other candidates are not. For instance, the forty-yard shot that didn't go into the top corner, or the routine save that wasn't fumbled—in such cases, the actual

outcome was exactly the default expectation, and no deviation means no cause. The same applies to ‘causes’ such as the presence of oxygen in the atmosphere. Sure, there’s a sense in which the oxygen is a cause of our losing the game as without it the game presumably would have been abandoned as impossible. But unless the tournament is being held in outer space, happy levels of oxygen are clearly assumed by all so again there’s no deviation here. Once given our shared default expectations, that is, it is perfectly objective to single out the referee’s bad decision as the cause of defeat and to ignore the oxygen and the missed forty-yarders. In the same way, to explain a shoot-out defeat we usually pick out the missed penalties not the ones that were scored.

Exploring this point more widely, we can now appreciate better all those pundit clichés that concern credit and blame. Take the favorite mantra of the smarter-than-thou, namely that, “defense wins championships.” What, if anything, might this claim amount to? After all, obviously many things win championships, not just defense. Sometimes, no doubt, the comment just is half-baked. But here are a couple of more charitable readings:

1) The difference between a good and bad defense matters more than that between, say, a good and bad attack, or good and bad tactics. (I’ll assume sufficiently shared understandings of ‘good’ and ‘bad’ here.) Imagine someone saying it’s important to know the rules of the game. Well, sure, but even relatively ignorant players know the rules well enough to play, so in practice it’s not a big factor. The claim here is that in a similar way having a good rather than bad attack is not a big factor, or at least not as big as having a good rather than bad defense.

2) Alternatively, perhaps attack and defense matter equally. But whereas anyone can appreciate dazzling forward play, good defense tends to go unnoticed because it is less eye-catching. Emphasizing the latter, therefore, is a useful corrective to this common bias. Parroting “defense wins championships” serves to highlight that defense *also* wins championships, not that it does so *more* than other factors.

My own feeling is that there may be some of both 1 and 2 going on. Ultimately though, claim 2 concerns our psychology, so I won’t pursue it here. But claim 1 is about picking out some causes rather than others, precisely what we’ve been talking about. Because the necessary understandings of ‘good’ and ‘bad’ are widely shared, I think we can adjudicate claim 1 objectively. An initial crude check, for instance, would be to see whether winning a league correlates more with having the best goals-against rather than goals-for record.

Regrettably, pundits are rarely explicit about what they do mean exactly, or even aware that they ought to be. Nevertheless, that doesn’t mean that *we* can’t spell it out for them. Here’s another common pearl of wisdom: “the game will be decided in midfield.” What could be meant by picking out the midfield (or wherever) in particular? One answer is that the difference between a good and bad outcome was especially great there. To put this another way, whereas the match-ups elsewhere were unlikely to be particularly one-sided, in midfield there was a realistic chance of one team dominating the other. Assuming that the default expectation is for rough parity in all areas, the biggest deviation—and so the most important cause—would therefore be in midfield. This claim

is, in principle at least, objective. Moreover it may be true of some matches more than others, in which case it would actually be informative too—assuming it was right.

The same applies to much else emanating from blogs and TV studios: “at the highest level, you need quality,” “England will never win, because of their tactical naivety,” “you need a good squad these days, not just a first eleven,” “forwards are match-winners,” and so on. With effort, it may really be that all of these can be translated from mere hot air into something objectively testable.

Choosing the Scapegoat

Alas, often several different causes all qualify as deviations. As well as the referee’s blunder, our striker also missed an open goal and our best player got injured. The default expectation, of course, is that an open goal will not be missed—that’s why missing one is so memorable—and that, unless he’s made of balsa wood, a player will not be injured. Accordingly, both of these other causes are endorsed by the default-deviation criterion and so there is no objective justification for singling out the referee’s decision alone.

Then we are back to our original question—how can we justify picking out some causes over others? In terms of objective causation, it seems that we can’t. John Stuart Mill, a nineteenth-century English philosopher, termed this, “the problem of causal selection.” But often, as Mill himself noted, it’s really other factors that do the work. Blaming the referee, for instance, is a well known way to divert attention from the equally blameworthy poor tactics or goalkeeping blunder. Indeed, so convenient is this tactic that you often hear it even when the referee hasn’t really made a mistake at all.

Frequently though, it’s *moral* feelings that direct us towards one cause rather than another, as when we blame an accident on a driver’s drunkenness rather than on the bad weather. Likewise in football, if a coach or player is especially irritating, of course we’ll seize gleefully on their mistakes rather than on those of others. And in a mirror image, if a player is one we especially like we tend to credit his contributions in particular. Perhaps it’s because rule-enforcers are inherently irritating anyway that we are especially likely to pick out referees.

Really, much the same was true of luck as well. Does a reasonable level of preparation for penalty shoot-outs include proper practice? Maybe it’s *because* I have long been irritated by the English players’ lump-headedness that I am more inclined to say ‘yes’. So next time you’re cursing bad luck or singling out someone to blame, ask yourself whether you aren’t just dressing up your own prejudice. It’s a very common human tendency.

As often, then, overall the truth is nuanced. Philosophy tells us that sometimes we really are objectively entitled to blame bad luck or the referee (or both). But only sometimes.