

Why Throwing 92 Heads in a Row is not Surprising

Martin Smith, University of Edinburgh

When we first meet the title characters of Tom Stoppard's *Rosencrantz and Guildenstern Are Dead*, they are betting on coin throws. Rosencrantz has a standing bet on heads, and he keeps winning, pocketing coin after coin. We soon learn that this has been going on for some time, and that no less than 76 consecutive heads have been thrown, and counting – a situation which is making Guildenstern increasingly uneasy. The coins don't appear to be double-headed or weighted or anything like that – just ordinary coins – leading Guildenstern to consider several unsettling explanations; that he is subconsciously willing the coins to land heads in order to cleanse for some repressed sin, that they are both trapped reliving the same moment in time over and over again, that the coins are being controlled by some menacing supernatural force. He then proposes a fourth hypothesis, which suggests a change of heart; nothing surprising is happening at all and no special explanation is needed. He says '...each individual coin spun individually is as likely to come down heads as tails and therefore should cause no surprise each time it does.' In the end 92 heads are thrown without a single tail, when the characters are interrupted.

Absurdist plays sometimes feature extraordinary or fantastical events that aren't given any explanation, and serve to create a sense of disorientation and dislocation – like humans transforming into animals in Ionesco's *Rhinoceros* or objects suddenly bursting into flame in Beckett's *Happy Days*. Perhaps Rosencrantz and Guildenstern's run of 92 heads is another example of this. And yet, in one way, this run of heads is precisely *not* like these other events. What is remarkable about the opening scene of *Rosencrantz and Guildenstern* is that, although it does succeed in creating a feeling of unreality, once you drill down into the details, *nothing extraordinary actually happens*. As Guildenstern himself points out, all that we have here is a sequence of 92 perfectly ordinary events, none of which *needs* any explanation.

My aim here is to defend Guildenstern – or his last hypothesis anyway. I will argue that there really is *nothing surprising* about throwing 92 heads in a row and that Guildenstern more or less explains why – though we may want to expand upon his reasoning a bit. I should say right away that I don't think Guildenstern has especially high standards for what should count as surprising, and neither do I. It's surprising if I flick the light switch and the room remains dark. It's surprising if a work colleague tells me she'll be at the meeting at 3 and then doesn't show. It's surprising if I park my car on the street and then return an hour later to find it gone. Those are all surprising things – but if I threw 92 fair, ordinary coins and every one of them came up heads, then that wouldn't be surprising. This, at any rate, is what I'm going to try and argue, using Guildenstern's own reasoning as a starting point.

When I say that there's nothing surprising about a run of 92 heads, I'm not making a prediction about what I, or anyone else, would *feel* if we were actually confronted with such a thing. My claim is that we *shouldn't* feel surprised, that we have *no reason* to feel surprised and, if we do feel surprised, then we're being irrational. If this is right, then it's not just some curiosity about coins – it has a much broader significance. The question of what we should *believe*, given our limited

evidence, is one over which a great deal of ink has been spilled – and it is also one that we *all* face, in some form, each day of our lives. If it's not surprising to throw 92 heads in a row, then one very standard, familiar answer to this question is suddenly thrown into doubt.

But let's stay with coins for the time being. Guildenstern seems to reason like this: Since there is nothing surprising about any one particular coin landing heads, there's no point in the sequence of 92 consecutive heads at which anything surprising actually happens. If it's unsurprising for an event e_1 to happen and it's unsurprising for an event e_2 to happen, does that mean that it's unsurprising for e_1 and e_2 to *both* happen? Not necessarily. It might be unsurprising if I leave for work at 8:30am and unsurprising if I arrive at work at 8:31am – but it might be very surprising if I leave for work at 8:30am *and* arrive at work at 8:31am. Clearly, though, these two events are *connected* – when I arrive at work will depend, in part, on when I leave for work, and that's why it would be surprising for both of these events to occur, even though neither event would be surprising on its own.

But what if we consider two completely unrelated, independent events? If it's unsurprising for me to leave for work at 8:30am and it's unsurprising for my work colleague Anna to arrive at work at 8:31am, and our morning commutes have absolutely nothing to do with one another, then it does seem to follow that it would be unsurprising for me to leave at 8:30 and Anna to arrive at 8:31. Consider the following, which we might call the *conjunction principle*: If it's unsurprising for event e_1 to happen, and it's unsurprising for event e_2 to happen, and these two events are independent of one another, then it's unsurprising for e_1 and e_2 to both happen.

When we flip 92 coins, what we have are 92 completely independent events – how one coin lands has nothing at all to do with how other coins have landed previously or how other coins will land subsequently. Coins can't predict the future and are not aware of the past. If it's unsurprising for the first coin to land heads, and it's unsurprising for the second coin to land heads, and these are independent events, then by the conjunction principle it's unsurprising for the first two coins to land heads. If it's unsurprising for the first two coins to land heads and it's unsurprising for the third coin to land heads and these events are independent, then it's unsurprising for the first three coins to land heads. And so on right up to 92 – or even further if we wish. If it's not surprising for any particular coin to land heads, and we accept the conjunction principle, we end up with the result that 92 heads in a row is not surprising. This, I think, is a fair reconstruction of Guildenstern's reasoning. But is it right?

In the 1950s and 60s, the economist George Shackle developed a precise mathematical theory of surprise, and the conjunction principle is actually very like one of the principles of Shackle's system – his 'axiom 7'. According to Shackle, when two events e_1 and e_2 are independent the surprisingness of e_1 & e_2 (which is measured by a number between 0 and 1) is equal to the surprisingness of e_1 or the surprisingness of e_2 – whichever is higher. As such, if e_1 and e_2 are both completely *unsurprising* (each have surprisingness values of 0) then e_1 & e_2 must be completely unsurprising too. By using this principle over and over again we can prove that, for any series of n events $e_1 \dots e_n$, if they're mutually independent and all unsurprising then e_1 & e_2 & \dots & e_n is unsurprising.

There were a number of problems with Shackle's system – problems which he never quite managed to resolve. Indeed, one very thing that he struggled with was properly fitting the notions of dependence and independence into his theory. A better, more complete, mathematical treatment of surprise is provided by *ranking theory*, first described by Wolfgang Spohn in the 1980s. Ranking theory is a powerful formal framework that has a number of potential applications, surprise being one. On this approach, when two events e_1 and e_2 are independent, the surprisingness of e_1 & e_2 (measured now by a positive integer 0, 1, 2, 3, 4 etc.) is equal to the sum of the surprisingness of e_1 and the surprisingness of e_2 . This is a consequence of what Spohn calls the *law of conjunction (for negative ranks)*, and our conjunction principle is really just a special case of this – if e_1 and e_2 are independent and are both completely unsurprising (both have surprisingness values of 0) then e_1 & e_2 is completely unsurprising too. Once again, by using this principle over and over again we can infer that, for any series of n events $e_1 \dots e_n$, if they're mutually independent and all unsurprising then e_1 & e_2 & ... & e_n is unsurprising.

Put less formally, the idea that seems to be at work in both of these formal treatments of surprise is that the surprisingness of a conjunction e_1 & e_2 must be a function of the surprisingness of e_1 , the surprisingness of e_2 and the connection between them. As such, if e_1 is completely unsurprising and e_2 is completely unsurprising and there's *no* connection between them, then there's nowhere for the surprisingness of e_1 & e_2 to come from. If we are allowed to make use of the conjunction principle then, in a way, we can give a 'proof' that 92 heads in a row is not a surprising event.

You may want to complain at this point that I'm missing something obvious – namely, that it's very *unlikely* for someone to throw 92 heads in a row. And if something very unlikely happens, then that's got to be surprising, doesn't it? Surely any 'proof' that seems to show otherwise is just some sort of trick and no real proof at all. Surely this shows that the conjunction principle has to be wrong. It is indeed very unlikely to throw 92 heads in a row – perhaps even more unlikely than you might guess at first. Assuming that the probability of any one coin landing heads is 0.5, and that the coin throws are mutually independent of one another, the probability of 92 coins landing heads in a row is equal to 0.5^{92} – that is, $0.5 \times 0.5 \times 0.5 \times 0.5 \dots$ 92 times. And that is a *very* small number – approximately 0.000000000000000000000000000002 or 1 in 5000 trillion trillion. This figure is too small to even properly get one's head around. This is (much) less than the chance of two people being asked to randomly choose a single grain of sand from anywhere on the Earth and happening to choose *exactly the same one*. Surely if a 1 in 5000 trillion trillion event were to actually happen, then this would be near miraculous and certainly very surprising. If you have this reaction, then you're in good company.

In the 1760s the polymath Jean le Rond d'Alembert questioned whether it was even *possible* to observe a long run of a single outcome when two equally likely outcomes could result on each trial. In a work published in the 1840s, Antoine-Augustin Cournot, one of the pioneers of the mathematics of probability theory, claimed that it was a 'practical certainty' that an event with a very low probability won't happen. But this kind of idea is perhaps put most starkly by Émile Borel –

another major figure from the history of probability theory. In his *Les probabilités et la vie* (*Probabilities and Life*), first published in 1942, Borel stated that ‘Events with a sufficiently small probability never occur’. Borel referred to this a ‘law of chance’ – indeed he once said it was the *only* law of chance. It’s now sometimes known simply as ‘Borel’s law’. A natural first reaction to Borel’s law is to think that it can’t be exactly right – after all, improbable things do sometimes happen – but that it’s close to being right. Perhaps what we should say is not that improbable events *never* happen, but that it’s very rare for them to happen and surprising when they do – or something like that. This would be enough to give us the result that 92 heads in a row is surprising.

I’ve come to think, though, that Borel’s law is not even close to being right – in fact it’s almost the exact *opposite* of the truth. One very important difference between the mathematics of probability and the mathematics of surprise (on both Shackle treatment and the ranking theoretic treatment) is that we can have a set up in which every possible outcome is highly *improbable*, but we cannot have a situation in which every possible outcome is highly *surprising*. Improbability and surprisingness cannot be the same thing. Come back to the case of the 92 coin throws. If we are going to throw 92 coins in a row, then we know in advance that there is going to *some sequence* of 92 results – if not heads every time (HHHHHHHH...) then it’s going to be some mixture of heads and tails in some sort of order (HTTHHTHT... or THTHTHH... etc.). And here’s the thing – each one of these sequences is *just as unlikely* as 92 heads in a row. In fact, each of these sequences has a probability of 0.5^{92} . If we throw 92 coins in a row then a one in 5000 trillion trillion event *is bound to happen* – and, as such, we shouldn’t be surprised when a one in 5000 trillion trillion event *does happen*.

If I’m surprised by 92 heads in a row, on the grounds that it’s so unlikely, then I’d have to be surprised by *any* sequence that came up – surprised *no matter how the 92 coins land*. This already seems like a bad enough result, but it goes much further than coin throws. Perhaps you just took a breath. Nothing unlikely about that you might think – or is there? If you just took a breath, then it must have had some precise duration, there must have been some precise volume that was inhaled and exhaled – indeed there must have been some precise number of oxygen molecules, water molecules, carbon dioxide molecules etc. that entered and left your lungs and so on. We don’t know what these numbers are of course, but we know this: It is exceedingly unlikely that your breath should have had precisely the duration, precisely the volume, involved precisely the number of oxygen, water, carbon dioxide molecules etc. that it did. In fact, this could be even *more* unlikely than throwing 92 heads in a row.

It’s not only when we repeatedly flip coins that something unlikely is bound to happen – something unlikely is bound to happen with every intake of breath, every heartbeat, every step. If I’m surprised by throwing 92 consecutive heads, based just on its low probability, then I should be in a state of constant amazement. This is why I say that Borel’s law is the opposite of the truth. According to Borel’s law unlikely things never happen – and yet, in a sense, *everything* that happens is an unlikely thing, once it is seen in sufficiently *high definition*. Another example: It’s likely that my phone will ring at some point over the next week. So when my phone does ring, isn’t that a likely event? Well, yes and no. Whenever my phone rings it will have to be at a particular second of a particular minute of a particular hour of a particular day – and it was always extremely unlikely that it would ring at precisely that second. The only reason the claim ‘my phone will ring at some point

over the next week' is likely is because there are so many different ways in which it could be true. But every one of these ways is extremely unlikely. The only reason this claim is likely is because it is so *nonspecific* about what will actually happen – but whatever *does* happen will be an unlikely thing.

To sum up so far: The conjunction principle allows us to prove that throwing 92 heads in a row is not surprising, even though it's extremely unlikely. One reaction to this is to reject the conjunction principle and dismiss Guildenstern's reasoning. But this would be hasty – because the fact that an event is extremely unlikely gives us no reason, in and of itself, to think that the event is surprising.

Another possible reason to think that 92 consecutive heads must be surprising is because, when we flip 92 fair coins, what we should *expect* to happen is to get roughly the same number of heads and tails – and 92 heads would fly in the face of this expectation. In a sense, I think it's right that we should expect to get roughly the same number of heads and tails – but talk about 'expectations' is somewhat ambiguous. Probability theorists define the 'expected value' of a random variable to be the probability weighted average of the possible values that the variable could take. The number of heads in 92 coin throws can be considered a random variable and, if the coins are fair and the throws are independent, then its expected value is indeed 46.

If we plot the probabilities of obtaining n heads in 92 fair, independent coin throws, then this will approximate a normal distribution or 'bell curve' with its peak at 46. We can calculate that the probability of getting exactly 46 heads is around 8.3%, the probability of getting between 40 and 50 heads is around 73.8%, the probability of getting between 30 and 60 heads is around 99.9% and so on. Obviously, the outcome in which we get 92 heads is located right in the extreme tail of the curve (over 9 standard deviations above the mean, if we want to put it in these terms). Does this mean that we should regard 92 heads in a row as a surprising result? Undoubtedly, there are cases in which it's surprising to observe an extreme divergence from an average – but is this one of those cases?

Consider the claim 'there will be between 40 and 50 heads' – the kind of thing that we're meant to 'expect'. Although we can calculate this claim to be approx. 73.8% likely, it's not as though there is some special force compelling the coins to fall in this way. The claim 'there will be between 40 and 50 heads' is a bit like the claim 'my phone will ring at some point over the next week' – the only reason it is likely to be true is that there are so many different ways in which it could be true (each of which is extremely unlikely). In fact, there are about 3700 trillion trillion different sequences of 92 coin throws that feature between 40 and 50 heads in some combination. This is a large set, but there's nothing special about the sequences that make it up – no reason to *prefer* them over the 1300 trillion trillion or so remaining sequences. As we've seen, all of the sequences are equally likely and any one could come about just as easily as any other. In fact, we could pick *any* set of 3700 trillion trillion sequences, on whatever basis we like, and it will be approx. 73.8% probable that the actual sequence will be one of the ones in the set. But this doesn't mean that we should suddenly regard the sequences outside the set as *surprising*. If we took some sequence from inside the set, and some sequence from outside the set, there would be no reason at all to regard the latter as being any more surprising than the former.

So yes, there is one sense in which we should ‘expect’ to get around 46 heads – we should regard this as highly likely, or assign it a high probability. The set of sequences in which we have around 46 heads covers a large proportion of the total set of outcomes. But there’s another sense in which we shouldn’t ‘expect’ to get around 46 heads – we shouldn’t *believe that this is going to happen*. We shouldn’t believe that the sequences outside the set won’t come up, while keeping an open mind about the sequences inside the set. There are no grounds for this – the sequences are all on a par.

This leads right to the final topic that I want to discuss – the relation between surprise and belief. While questions about what is surprising do have some interest in their own right, what makes them really significant is precisely the way in which they seem to be bound up with questions about what we should believe (and this, indeed, is precisely why Shackle and Spohn were interested in surprise). Generally speaking, surprise is what we experience when the world doesn’t match our beliefs – if we believe that something is going to happen and it doesn’t, or we believe that something isn’t going to happen and it does, then that’s surprising for us. *Surprise* is a guide to *belief*.

Furthermore, if we have reason to believe that something isn’t going to happen, then we have reason to be surprised if it does happen. Or, to put it another way, if we have no reason to be surprised if a certain event happens, then we have no reason to believe that it won’t happen – we should keep an open mind about it. *Rational surprise* is a guide to *rational belief*. If it’s right that we have no reason to be surprised by throwing 92 heads in a row, it follows that we shouldn’t believe in advance that this won’t happen.

Questions about when we are justified or rational in believing things have been discussed a great deal by philosophers, scientists, legal theorists and by many others. Many of those who have considered such questions have converged on the view that *probabilities* should be our guide when forming beliefs – that we should believe those things that are likely to be true, disbelieve those things that are likely to be false, and otherwise suspend judgment. Call this the *probability principle*. The thought seems to be that believing something is a bit like taking a gamble on what the world is like – and if the odds are in my favour then the gamble should be a rational one. In many ways this is a very appealing and powerful picture – but, in the end, I don’t think that belief is like this, and I don’t think the probability principle is correct. It’s very likely that the coins won’t land THTTHHTH... and very likely that they won’t land TTHTHTHH... and so on. While it might be perfectly rational for me to *bet* on these things, it’s not rational for me to *believe* them. If the coins did land THTTHHTH or land TTHTHTHH, then it would not be rational to be surprised. In actual fact, exceptions to the probability principle are all around us – for some large number n , it’s very likely that n won’t be the exact number of oxygen molecules that I inhale on my next breath, but it’s not rational for me to believe this.

What should we believe then? If probabilities are not the key to justified, rational belief, then what is? Earlier on I gave some examples of things that I thought would be genuinely surprising – like this: If I park my car on the street and then return an hour later to find that the car is no longer

where I parked it, then that's surprising. While this may well be an unlikely event, what seems more significant is that it's an event that demands *explanation* of some kind. Perhaps someone broke into the car and stole it. Perhaps I parked illegally and the car was then towed. Perhaps I didn't properly apply the handbrake and the car rolled away... Whatever the truth, it can't 'just so happen' that the car is now gone and there's nothing more to the story. This isn't the sort of thing that can 'just so happen' – there *has* to be more to the story. If I park my car on the street, it would be natural to believe, in an hour's time, that it's still there. I think it can be *rational* to believe this too, because there would have to be some explanation if it turned out to be false.

It's obvious that, if the coins land THTTHHTH... or land TTHTHTHH..., there doesn't have to be any special explanation for this – after all, the coins had to land in some way, and it might just as well have been these ways as any other. When it comes to a run of 92 heads, we might be tempted to think that there really does have to be some explanation in this case – thus our temptation to think that this would be a surprising event. Psychologists have found that people are generally unwilling to accept that outcomes exhibiting some striking *pattern* could arise through a purely random process. It's for this reason that people are generally reluctant to pick consecutive numbers like '1234567' in the lottery. People will often try to find some deliberate, intentional process behind a patterned outcome – even preferring something supernatural (as Guildenstern does initially) to admitting that there's no special explanation to be found. But '1234567' is as good a lottery pick as any other – a fair lottery could produce this result just as easily as any other result. And throwing 92 fair coins could produce a sequence of 92 heads just as easily as any other sequence.

What, in any case, is the *point* of surprise? What would we be missing in our lives if we never felt surprised by anything – if we greeted everything with a shrug? I think we would be missing something crucial – for part of the purpose of surprise is to spur us into *action*. If an event surprises us, then that prompts us to investigate *why* and *how* it happened – to try and explain it. There is something agitating about a feeling of surprise, that only abates when a satisfying explanation is found. This is why surprise would be such an inappropriate reaction to a sequence of 92 coin throws, like THTTHHTH... As unlikely as this sequence might be, it can just so happen that this is the sequence that came up, and there's nothing more for an investigation to reveal or unearth. Surprise simply isn't the right response to an event like this.

And the situation is really no different when it comes to a 'patterned' outcome like Rosencrantz and Guildenstern's run of 92 heads. When faced with this result, of course it is sensible to check (as Guildenstern does) whether the coins are double-headed or weighted or anything of that kind. Having observed a run of 92 heads in a row, one should regard it as very *likely* that the coins are double-headed or weighted. But, once these realistic possibilities have been ruled out, and we know they don't obtain, any remaining urge to find *some explanation* (no matter how farfetched) becomes self-defeating. As difficult as it may be to accept, there *doesn't have to be* an explanation for this – and it's not rational to relentlessly search for one.

Roughly speaking, it's rational to be surprised by an event if and only if that event requires investigation and explanation. And, going back to Guildenstern's reasoning, this seems to predict that the conjunction principle is indeed sound. The conjunction principle, remember, states that if

it's unsurprising for event e_1 to happen, and it's unsurprising for event e_2 to happen, and these two events are independent of one another, then it's unsurprising for e_1 and e_2 to both happen. If there's no need to explain the fact that e_1 occurred, and there's no need to explain the fact that e_2 occurred, and the events have nothing whatsoever to do with one another, then there's no need to explain the fact that e_1 and e_2 both occurred.

When I park my car on the street, it's rational to believe that it will still be there an hour later. If it isn't, then it would be rational to be surprised and to look for an explanation. If my work colleague tells me that she will be at the meeting at 3, it's rational to believe that she will. If she isn't, then it would be rational to be surprised and to look for an explanation. It's rational to believe that the lights will come on when I flick the switch. It's rational to feel surprised, and seek an explanation, if I flick the switch and the room remains dark. There are many things that we can rationally believe – but the claim that we won't throw 92 heads in a row is not one of them. I can rationally regard it as extremely likely that I won't throw 92 heads in a row, but I can't rationally believe it.

These ideas about rational belief are, of course, very sketchy and I won't try to pursue them further here. Maybe they aren't even on the right track at all. But what I hope I have shown here, at the very least, is that there is a different way of looking at surprise and belief, and that a 'Guildensternian' theory of surprise can be defended. I mentioned at the outset that it's common for absurdist plays to feature fantastical events that are left unexplained. Another very common trope in absurdist drama is for characters to reason in nonsensical ways and to jump to bizarre conclusions. Guildenstern's first three hypotheses about the coin throwing episode are indeed bizarre. And so too, I suppose, is his fourth hypothesis – but it also, just maybe, happens to be true.

Acknowledgments

This paper was presented at the University of Edinburgh in April 2016. Thanks to all of those who participated on this occasion. Thanks also to Campbell Brown and Stephan Leuenberger for helpful discussions of this material and to Philip Ebert, Meredith Regan, Steph Rennick and Ed Wilson for providing valuable comments on earlier drafts of this paper.

Further Reading

Bar-Hillel, M. and Wagenaar, W. (1991) 'The perception of randomness' *Advances in Applied Mathematics* v12(4) [A survey of some early research on human judgments of what is random and what isn't.]

Borel, E. (1946) *Probabilities and Life*, Baudin, M. trans. (New York: Dover Publications) [The discussion of what has come to be known as 'Borel's law' is on pages 2 and 3].

Shackle, G. (1967) *Decision, Order and Time 2nd ed.* (Cambridge University Press) [Contains Shackle's most detailed presentation of his ideas about surprise. His axioms, and his struggles over axiom 7, can be found in chapter X]

Spohn, W. (2012) *The Laws of Belief* (Oxford University Press) [Spohn's definitive presentation of ranking theory and its various applications. Discusses Shackle and surprise in section 11.1. The law of conjunction for negative ranks is principle 5.16 in chapter 5]

Stoppard, T. (1967) *Rosencrantz and Guildenstern are Dead* (London: Faber) [The coin flipping scene opens the play].