Synthese (2011) 179:351–360
DOI 10.1007/s11229-009-9656-z

Quasi-miracles, typicality, and counterfactuals

Dylan Dodd

Received: 1 June 2009 / Accepted: 2 September 2009 / Published online: 16 September 2009

Abstract If one flips an unbiased coin a million times, there are $2^{1,000,000}$ series of possible heads/tails sequences, any one of which might be the sequence that obtains, and each of which is equally likely to obtain. So it seems (1) 'If I had tossed a fair coin one million times, it might have landed heads every time' is true. But as several authors have pointed out, (2) 'If I had tossed a fair coin a million times, it wouldn't have come up heads every time' will be counted as true in everyday contexts. And according to David Lewis' influential semantics for counterfactuals, (1) and (2) are contradictories. We have a puzzle. We must either (A) deny that (2) is true, (B) deny that (1) is true, or (C) deny that (1) and (2) are contradictories, thus rejecting Lewis' semantics. In this paper I discuss and criticize the proposals of David Lewis and more recently J. Robert G. Williams which solve the puzzle by taking option (B). I argue that we should opt for either (A) or (C).

Keywords Counterfactuals · Counterfactual scepticism · Quasi-miracles · Atypical events · David Lewis

1 The puzzle: counterfactual conditionals and weird possibilities

If one flips an unbiased coin a million times, there are 2^{10^6} series of possible heads/tails sequences, and any one of these might be the sequence that obtains. Each sequence is equally likely to obtain. One of these is the series with all heads. So it seems:

1. If I had tossed a fair coin 10^6 times, it might have landed heads every time.

D. Dodd (✉)
Arché—Philosophical Research Centre for Logic, Language, Metaphysics and Epistemology, Department of Logic and Metaphysics,
University of St. Andrews, St. Andrews, Fife KY16 9AJ, Scotland
e-mail: dylan.dodd@gmail.com

However, as Lewis (1986), Hawthorne (2005) and Williams (2008) all note, the following is counted true in everyday contexts:

2. If I had tossed a fair coin 10^6 times, it wouldn't have landed heads every time.

But intuitively, and according to Lewis (1973)'s influential semantics for counterfactuals:

3. 'If p had been the case, q would have been'[1] and 'If p had been the case, q might not have been' are contradictories.

1–**3** form an inconsistent triad.

Furthermore, it seems quantum mechanics tells us that the behavior of subatomic particles is such that, strictly speaking,

4. If I had dropped my pencil, it might have flown off sideways.

But according to common sense:

5. If I had dropped my pencil, it would have fallen to the ground.

3–**5** also form an inconsistent triad.

What should we make of all this? Here are the options:

A. Common sense judgements that **2** and/or **5** are true are mistaken.
B. **1** and/or **4** are false.
C. **3**, and therefore Lewis (1973)'s semantics for counterfactuals, is false.

One of **A**–**C** must be true. But a good case can be made against all of them.[2] In this paper I will criticize solutions to the puzzle that accept **B**, allowing us to deny **A** and **C**. I will focus on the proposals of Lewis (1986) and Williams (2008), but my argument against their proposals will generalize to any solution of this type.[3,4]

[1] Please read my quotation marks as corner quotes where appropriate.

[2] I take the implausibility of **A** to be obvious. To some extent I've already motivated **B**'s being false, and in the rest of the paper I further motivate its rejection. As far as **C** goes, one point against it is that it's inconsistent with Lewis (1973)'s elegant semantics for counterfactuals. Another reason to reject it is provided by DeRose (1999): 'If p were the case, q would have been; and if p were the case, q might not have been' sounds awful, suggesting an inconsistency. On the other hand, Stalnaker (1981) argues for **C**, on grounds independent of this puzzle.

[3] Insofar as I am attacking responses to the puzzle that accept **B**, I will be providing motivation for responses that accept either **A** or **C**. Hájek (MS) argues for **A**, and Eagle (MS) proposes a solution to the puzzle that accepts **C**.

[4] If, with Hájek (MS), we accept **A**, there is an important epistemological consequence, namely we will also have to accept that many if not most of the *a posteriori*, contingent propositions we think we know are *unsafe*, where the safety of a belief is understood in the way described by Williamson (2000) or Pritchard (2005, pp. 162–163). Williamson and Pritchard both say that one only knows that p if one's belief that p is safe. If we accept **A**, we are stuck with the following dilemma: we either have to reject such 'safety theories' of knowledge, or succumb to skepticism—i.e., say that we know much less than we think. I argue *for* this dilemma in my (MS).

2 Quasi-miracles and atypicality

In what follows, I will symbolize '**If it were [had been] the case that p, then it might be [might have been] the case that q**' as '$p \diamond\!\!\rightarrow q$'; '**If it were [had been] the case that p, then it would be [would have been] that q**' will be '$p \,\Box\!\!\rightarrow q$'. According to Lewis (1973)'s influential semantics for counterfactuals:

6. '$p \diamond\!\!\rightarrow q$' is true iff there is some closest world where p obtains/obtained [some closest p-world] which is a world where q obtains/obtained [is a q-world].
7. '$p \,\Box\!\!\rightarrow q$' is true iff all the closest p-worlds are q-worlds.

We have already seen why **A** *seems* to be consequence of **6** and **7**. If one flips an unbiased coin several times, any series of heads/tails has an equal probability of obtaining, and, it seems, the world in which any given sequence obtains resembles our world just as much as a world where any other given sequence obtains. Thus if **6** and **7** are true, it seems that for any sequence S of 10^6 flips of an unbiased coin, 'The coin was flipped 10^6 times $\diamond\!\!\rightarrow$ the flips landed in S' is true. Letting S be the series with all heads, **1** ['*I tossed a fair coin* 10^6 *times* $\diamond\!\!\rightarrow$ *it landed heads every time*'] is true. Assuming **C** ['$p \,\Box\!\!\rightarrow q$' *and* '$p \diamond\!\!\rightarrow \neg q$' *are contradictories*], which is entailed by **6** and **7**, **A** [*Common sense judgements that* '*I tossed a fair coin* 10^6 *times* $\Box\!\!\rightarrow \neg$(*it landed heads every time*)' *and/or* '*I dropped my pencil* $\Box\!\!\rightarrow$ *it fell to the ground*' *are true are mistaken*] is true.

But Lewis (1986) thought **2** ['*I tossed a fair coin* 10^6 *times* $\Box\!\!\rightarrow \neg$(*it landed heads every time*)'] was true and **1** false. He thought the weird possibilities puzzle should be solved by accepting **B** [*1 and/or 4 are false*]. The coin's landing heads each of the 10^6 times would be what he calls a *quasi-miracle*, which he characterizes as follows:

> The quasi-miracle would be such a remarkable coincidence that it would be unlike the goings-on we take to be typical of our world. Like a big genuine miracle, it makes a tremendous difference from our world. Therefore it is not something that happens in the closest worlds . . . My point is not that quasi-miracles detract from similarity because they are so very improbable. They are; but ever so many unremarkable things that actually happen, and ever so many other things that might happen under various counterfactual suppositions, are likewise very improbable. What makes a quasi-miracle is not improbability per se, but rather the remarkable way in which the chance outcomes seem to conspire to produce a pattern.

His example: If a monkey typing randomly produces 950 pages, any 950 page sequence of letters is equally probable to appear. But if the monkey just happens to type a 950 page dissertation on anti-realism, that would be a quasi-miracle, while his typing 950 pages of jumbled letters wouldn't be (1986, pp. 60–61). In our example, a sequence of 10^6 heads is a quasi-miracle, but not a sequence of heads and tails jumbled haphazardly.

Hawthorne (2005) subjected Lewis' quasi-miracles response to the weird possibilities puzzle to heavy criticism. Williams (2008) has put forward a more rigorous response to the puzzle that avoids Hawthorne's objections. Let's now look at Williams' proposal.

Like Lewis, Williams grants **6** and **7**. Also like Lewis, he denies **A**, and puts forward a solution that accepts **B**. He reconciles **B** with **6** and **7** as follows. He divides events into ones that are *typical* and *atypical*, and thinks that if an event is *atypical*, that makes the possibility that it happen less close to actuality. To determine whether an event is typical, one looks at the probability *of a certain set of properties it instantiates* rather than at the probability *of the particular outcome*. For instance, that a series of 10^6 coin flips will be all heads is very improbable, and that it will have roughly the same amount of heads and tails is highly probable. Thus insofar as a series of flips comes out all heads it is atypical, and insofar as it has roughly the same amount of heads as tails it is typical. It doesn't matter that whatever series is instantiated, it was very improbable that series obtained. To focus on *that* fact would be to focus on *a particular outcome*, which doesn't matter for whether an event is typical or not. Rather, an event is typical insofar as it instantiates the right high-probability properties (2008, pp. 407–411).[5]

3 Critique of Lewis and Williams

Let n be an even number such that it would be quasi-miraculous or atypical for a coin that flipped n times to be all heads, but not by much. If n were much smaller, its landing all heads wouldn't be quasi-miraculous or atypical. There's nothing quasi-miraculous or atypical, then, about the coin coming up heads the first $\frac{n}{2}$ flips. So according to both Williams and Lewis, **8** is true but **9** and **10** are both false:

8. I flipped the coin n times $\Box\!\!\rightarrow$ $\neg$(it comes up heads each time).
9. I flipped the coin n times $\Box\!\!\rightarrow$ $\neg$(it comes up heads the first $\frac{n}{2}$ times).
10. I flipped the coin n times $\Box\!\!\rightarrow$ $\neg$(it comes up heads the last $\frac{n}{2}$ times).

Let 'H_i' stand for the proposition that the coin landed heads on the ith flip. Then **8** is equivalent to:

8′. I flipped the coin n times $\Box\!\!\rightarrow \neg[(H_1 \,\&\, \ldots \,\&\, H_{\frac{n}{2}}) \,\&\, (H_{\frac{n}{2}+1} \,\&\, \ldots \,\&\, H_n)]$.

Since they're committed to **8′**, Lewis and Williams must deny one of the following three claims:

11. I flipped the coin n times $\& H_1 \,\&\, \ldots \,\&\, H_{\frac{n}{2}} \Diamond\!\!\rightarrow H_{\frac{n}{2}+1} \,\&\, \ldots \,\&\, H_n$.
12. The closest **worlds where I flipped the coin** n **times** $\& \ (H_1 \,\&\, \ldots \,\&\, H_{\frac{n}{2}})$ are all amongst the closest **worlds where I flipped the coin** n **times**.
13. $8', 12 \vDash \neg 11$.

I shall now argue, however, that **11**, **12**, and **13** are all true.

13 is true because it's an instance of this more general principle:

13′ $p \Box\!\!\rightarrow \neg(q \,\&\, r)$, The closest $p \,\&\, q$-worlds are all amongst the closest p-worlds $\vDash \neg(p \,\&\, q \Diamond\!\!\rightarrow r)$.

[5] I am hoping the reader gets the intuitive idea of how Williams demarcates the distinction between typical and atypical events. For more details, see his paper. The distinction is based on the work of Elga (2004) and Gaifman and Snir (1982).

That **13**′ is true can be shown as follows. Assuming **7** [*‘p $\Box\!\!\rightarrow$ q’ is true iff all the closest p-worlds are q-worlds*], **13**′’s premises imply that all the closest $p\,\&\,q$-worlds are $\neg(q\,\&\,r)$-worlds; equivalently, they’re all $(\neg q \vee \neg r)$-worlds. Since $(p\,\&\,q)\,\&\,(\neg q \vee \neg r)$ entails $(p\,\&\,q)\,\&\,\neg r$, it follows that all the closest $p\,\&\,q$-worlds are $\neg r$-worlds. Therefore, $p\,\&\,q$ $\Box\!\!\rightarrow$ $\neg r$. We’re assuming this to be the contradictory of $p\,\&\,q$ $\Diamond\!\!\rightarrow$ r, since Lewis and Williams are pursuing a solution to the weird possibilities puzzle that accepts **B** rather than **C**. Therefore, $\neg(p\,\&\,q$ $\Diamond\!\!\rightarrow$ $r)$.

Now I’ll defend **12**. What are the closest worlds where *I flipped the coin n times* & H_1 & … & $H_{\frac{n}{2}}$? They’re worlds where I flip the coin n times, and it just so happens, as a chance coincidence, that it lands heads the first $\frac{n}{2}$ times, and where nothing else bizarre, quasi-miraculous, atypical, etc. happens—where nothing else happens that would make these worlds from being as close as they can be to actuality. Are these worlds all amongst the closest worlds where I flip the coin n times? Why wouldn’t they be? Given what’s been stipulated about the value of n, while it may be surprising that the coin landed heads on the first $\frac{n}{2}$ flips, the mere fact that this happens in a world doesn’t prevent that world from being one of the closest worlds where the coin is flipped n times. For a sequence of heads-in-a-row to count against a world being close, it has to be a longer sequence, close in length to a sequence of n flips. A sequence of $\frac{n}{2}$ flips isn’t big enough, *per hypothesi*. And as we’ve seen, the closest worlds where *I flipped the coin n times* & H_1 & … & $H_{\frac{n}{2}}$ are all worlds where nothing else happens in them that prevents them from being as close as they can be to actuality. So **12** is true.

Finally, let’s look at **11**. A prima facie case for **11** can be made as follows. That $H_{\frac{n}{2}+1}$ & … & H_n would obtain may be surprising, but it’s not quasi-miraculous or atypical. Therefore, there’s nothing about $H_{\frac{n}{2}+1}$ & … & H_n being true at a world that prevents the world from being amongst the closest worlds to actuality. There just doesn’t seem to be any basis for denying that **11** is true—Lewis’ and Williams’ appeals to quasi-miracles and atypicality don’t provide any resources for explaining it to be false. I anticipate the following response. Because $(H_1$ & … & $H_{\frac{n}{2}})$ & $(H_{\frac{n}{2}+1}$ & … & $H_n)$ is a quasi-miraculous/atypical event, if **11**’s antecedent holds, then if its consequent were to hold too, a quasi-miraculous/atypical event would occur, making the world further away. Therefore, although **11**’s consequent doesn’t count against a world’s being close per se, it *does* count against a world *where it and **11**’s antecedent holds* from being amongst the closest worlds where **11**’s antecedent holds. And this means that **11** is false.

The key claim is that the fact that **11**’s consequent holds at a world counts against that world’s being amongst the closest worlds where **11**’s antecedent holds (making **11** false), but not against worlds in general from being as close as they can be to actuality. In arguing against this claim, I will adopt the method of Lewis (1973, 1986) of using our intuitions about what counterfactuals are true to guide our intuitions about what worlds are close. An event’s being quasi-miraculous or atypical is supposed to make a world not be amongst the closest worlds. Given our Lewisian methodology about closeness intuitions, we think worlds where the coin lands heads on each flip are not amongst the closest worlds where it’s flipped 10^6 times *because we judge **2*** [*‘I tossed a fair coin* 10^6 *times* $\Box\!\!\rightarrow$ $\neg$(*it landed heads every time*)’] *to be true*. The quasi-miracles and atypicality views are supposed to give us a way of reconciling our intuitions about the truth values of counterfactuals, and thus of closeness, with a line of thought against those intuitions presented briefly in Sect. 1.

It was stipulated that n is an even number that is big enough so that we would have the intuition that the counterfactual **8** [*'I flipped the coin n times* $\square\!\!\rightarrow$ $\neg$(*it comes up heads each time*)*'*] is true, but just barely big enough. If n were a bit smaller, we'd judge **8** false. Given these stipulations, we should choose a value for n so that we have the intuition that **9** [*'I flipped the coin n times* $\square\!\!\rightarrow$ $\neg$(*it comes up heads the first* $\frac{n}{2}$ *times*)*'*] and **10** [*'I flipped the coin n times* $\square\!\!\rightarrow$ $\neg$(*it comes up heads the last* $\frac{n}{2}$ *times*)*'*] are false. It seems that, given that we have the intuition that **9** and **10** are false, we will also have the intuition that **11** [*'I flipped the coin n times* $\& H_1 \& \ldots \& H_{\frac{n}{2}} \Diamond\!\!\rightarrow H_{\frac{n}{2}+1} \& \ldots \& H_n$*'*] is true. I ask the reader to pick a value for n such that she judges **8** to be true and **9** and **10** to be false. Since we're rejecting **C** [i.e., we're accepting that '$p \square\!\!\rightarrow q$' and '$p \Diamond\!\!\rightarrow \neg q$' *are* contradictories], given that we judge **9** and **10** to be false, we'll judge the following to be true:

9′. I flipped the coin n times $\Diamond\!\!\rightarrow$ it comes up heads the first $\frac{n}{2}$ times.
10′. I flipped the coin n times $\Diamond\!\!\rightarrow$ it comes up heads the last $\frac{n}{2}$ times.

Now ask yourself what your intuition is about this sentence: 'If I flipped the [unbiased] coin n times, and it landed heads the first $\frac{n}{2}$ times, it might have also landed heads the last $\frac{n}{2}$ times' (i.e., **11**). I know what intuition I have. When I assign a value to n so that I have the intuition that **9′** and **10′** are both true, I then have the intuition that **11** is true too. It strikes me as *very* counterintuitive to say that this sentence is false, while claiming that **9′** and **10′** are true. So if, following Lewis, we let our intuitions about the truth values of counterfactuals guide our intuitions about closeness, we won't make the objection to **11** just imagined. That is, we won't claim that the fact that **11**'s consequent holds at a world *does count* against that world's being amongst the closest worlds where **11**'s antecedent holds, even though in general the fact that **11**'s consequent holds at a world *does not count* against that world's being as close as can be to actuality. Insofar as we follow this Lewisian methodology, we would only say *that* insofar as we had the intuition that **11** was false but **10′** was true. But in fact when we have the intuition that **10′** is true, we have the intuition that **11** is too.

We've just seen that a problem with Lewis' and Williams' proposals is that they must claim that when we pick a value for n such that **8**, **9′** and **10′** are all true, **11** will be false. I will now strengthen the case against this consequence of their proposals.

I think we can use the following principle as a guide to determining which counterfactuals we judge to be true:

(∗) If I believe a counterfactual of the form '$p \square\!\!\rightarrow q$' to be true, then if I come to believe p, I will be willing to infer q, *assuming that I retain my belief that '$p \square\!\!\rightarrow q$' is true when I come to believe that p.*

I'll start by noting the importance of the final *italicized* clause in (∗). Say I believe that *if Oswald hadn't shot Kennedy, no one else would have* (as many people do). I then come to believe that there was a conspiracy, and that Oswald didn't shoot Kennedy (as some people already believe). It is easy to fill in the background to this case so that I don't infer that no one shot Kennedy, but instead infer that someone else did.[6]

[6] I owe this point to an anonymous referee at *Synthèse*.

So we're imagining a case where first I believe that *if Oswald hadn't shot Kennedy, no one else would have*, I then come to believe its antecedent is true, but I don't infer its consequent—in fact I infer something that contradicts it. However, in this case, at the point I come to believe the counterfactual's antecedent (that Oswald didn't shoot Kennedy) I will immediately give up my belief in the counterfactual. Of course, for how could I believe that *if Oswald hadn't shot Kennedy, no one else would have*, if I believed that *someone else besides Oswald shot Kennedy*?

So, to repeat, I think that (∗) furnishes us with a good guide for determining which counterfactuals we think are true. We can use it as a guide as follows. We ask ourselves if we would be willing to infer the consequent of '$p \mathrel{\Box\!\!\rightarrow} q$' upon learning that p. If the answer is 'yes', that is at least consistent with our thinking the counterfactual is true, and is evidence that we do take it to be true. On the other hand, if the answer is 'no', we then need to ask whether that is because learning that the antecedent is true would mean changing our mind about the truth of the counterfactual – in fact, giving up our belief in it. If we would not be willing to infer q upon learning that p, while at the same time it doesn't seem that learning that p would make us change our minds about the truth of '$p \mathrel{\Box\!\!\rightarrow} q$' (as it would, for instance, with 'if Oswald hadn't shot Kennedy, no one else would have'), then we should think that we don't take the counterfactual to be true.

Following (∗) as a guide in the way just envisioned would lead us to endorse **2** [*'I tossed a fair coin* 10^6 *times* $\mathrel{\Box\!\!\rightarrow} \neg$(*it landed heads every time*)*'*] and **5** [*'I had dropped my pencil* $\mathrel{\Box\!\!\rightarrow}$ *it fell to the ground'*], and this is at the heart of why it would be a cost to have to claim that either **2** or **5** was false. This is why the project of saving their truth, when otherwise plausible theories of counterfactuals seem to have the consequence that they're false, is a worthwhile project. If I learned that someone dropped their pencil, and that there was nothing to obstruct its fall to the ground, I would subsequently judge that the pencil fell to the ground. Thus the counterfactual **5** seems true to me. Similarly, if we learned that someone had flipped an unbiased coin 10^6 times, we would think that some of these flips landed heads, others tails. Thus we think **2** is true. Thus it seems that (∗) is true, and using (∗) as a guide to determine which counterfactuals we think are true gives us reason to think that we take **2** and **5** both to be true.

On the other hand, if we learned that someone flipped a coin three times, we wouldn't believe it didn't land heads each time. Rather, we would think that it just might have landed heads each time, even if it probably didn't. And that's not because learning that the coin was flipped three times would make us change our minds about the truth of 'someone flipped the coin three times $\mathrel{\Box\!\!\rightarrow} \neg$(it landed heads each time)'. In general, it's hard to imagine how *just* learning that a coin was flipped x number of times, and nothing else, would make me change my mind about the truth value of 'the coin was flipped x times $\mathrel{\Box\!\!\rightarrow} \neg$(it landed heads each time)'. Since we think **8** [*'I flipped the coin n times* $\mathrel{\Box\!\!\rightarrow} \neg$(*it comes up heads each time*)*'*] is true, we would think the coin didn't land heads each time upon learning that it was flipped n times. However, we wouldn't think it failed to land heads the last $\frac{n}{2}$ times. Instead, we'd maintain it *might* have landed heads each of the $\frac{n}{2}$ times. Now say we learn that not only is it true that someone flipped it n times, but that it landed heads each of the first $\frac{n}{2}$ times. We aren't told anything about what happened on the last $\frac{n}{2}$ flips. Upon learning this new information, it would be bizarre to reason as follows: “Well, I used to maintain that it might have landed heads on the last $\frac{n}{2}$ flips. But now, as a result of learning that it landed heads on each of **the first** $\frac{n}{2}$ flips, I now think it did *not* land heads on **the last** $\frac{n}{2}$ flips.” Rather, our inclination (or lack thereof) to judge that an unbiased coin did not land heads on each of the last $\frac{n}{2}$ flips will be indifferent to whether or not it landed heads on each of the first $\frac{n}{2}$ flips. Of course! The coin is unbiased after all. That means what happens in the second half of the series of flips is independent of what happens in the first half.

The upshot is that our intuitions about the truth values of counterfactuals, and thus of closeness, are against the philosopher who rejects **11**. First, we have the intuition that **11** is true, not false. Second, we don’t have the intuition that while the fact that $H_{\frac{n}{2}+1}$ & … & H_n is true at a world doesn’t count against that world’s being close per se, we do think it counts against that world being close to a world where H_1 & … & $H_{\frac{n}{2}}$ is true. In fact, our intuitions about *whether* $H_{\frac{n}{2}+1}$ & … & H_n *is true at a world counts against that world being close* don’t discriminate between worlds where I flip the coin n times and H_1 & … & $H_{\frac{n}{2}}$ and the other worlds where I flip it n times. If we let our intuitions about closeness be guided by our intuitions about the truth values of counterfactuals as Lewis taught us, we will judge **11** to be true, and reject the argument for its falsity I imagined previously.

We’ve seen that if we’re going to reject solutions to the weird possibilities puzzle that accept either **C** or **A**, and use Lewis’ or Williams’ proposals to explain why **B** is true instead, we need to deny either **11**, **12**, or **13**. But I have argued that **11**, **12** and **13** are all true. I conclude that we should reject Lewis’ and Williams’ proposals.

The discussion has focused on whether Lewis and Williams can save the truth of **2** [‘I tossed a fair coin 10^6 times $\Box\!\!\rightarrow$ $\neg(H_1$ & … & $H_{10^6})$’]. What about **5** [‘I dropped my pencil $\Box\!\!\rightarrow$ it fell to the ground’]? The way Lewis and Williams would explain the truth of **5** is analogous to their explanation for **2**, and *the problem* with their explanation of **5** is also the same. It seems that quantum mechanics tells us that the pencil might not fall if dropped. Like the coin-flipping case, for the pencil not to fall, a number of individually unexceptional events would have to conspire together to produce the weird possibility. The events would involve the movement of very tiny particles. Let the set of events, all of which would have to obtain for the pencil to fly off sideways, be $\{E_1, \ldots, E_n\}$. In order to maintain that **5** is true, we must claim that ‘I dropped my pencil $\Box\!\!\rightarrow$ $\neg(E_1$ & … & $E_n)$’ is true. According to Lewis and Williams, the reason it’s true is that the event of E_1 & … & E_n occurring is quasi-miraculous/atypical. Now we offer an objection analogous to the objection we gave about coin flipping. There will be an even number n', such that E_1 & … & $E_{n'}$ is quasi-miraculous/atypical, but not by much. For a number n'' slightly smaller than n', E_1 & … & $E_{n''}$ is *not* quasi-miraculous/atypical. Thus while Lewis and Williams will claim that it’s true that

8″ I dropped my pencil $\Box\!\!\rightarrow \neg[(E_1 \,\&\, \ldots \,\&\, E_{\frac{n'}{2}}) \,\&\, (E_{\frac{n'}{2}+1} \,\&\, \ldots \,\&\, E_{n'})]$.

I could give an argument analogous to the one above for each of the following:

11′ I dropped my pencil $\&\, E_1 \,\&\, \ldots \,\&\, E_{\frac{n'}{2}} \Diamond\!\!\rightarrow (E_{\frac{n'}{2}+1} \,\&\, \ldots \,\&\, E_{n'})$.

12′ The closest worlds where I dropped my pencil $\& E_1 \& \ldots \& E_{\frac{n'}{2}}$ are all amongst the closest worlds where I dropped my pencil.

13″ 8″, 12′ $\models \neg$11′.

And as before, **8″** is incompatible with the conjunction of **11′**, **12′** and **13″**.

4 Conclusion

We saw in Sect. 1 that the weird possibilities puzzle shows that we must accept one of the following, in spite of the fact that there's a good case to be made against all of them:

A. Common sense judgements that **2** [*'I tossed a fair coin* 10^6 *times* $\Box\!\!\rightarrow \neg$(*it landed heads every time*)*'*] and/or **5** [*'I dropped my pencil* $\Box\!\!\rightarrow$ *it fell to the ground'*] are true are mistaken.

B. **1** [*'I tossed a fair coin* 10^6 *times* $\Diamond\!\!\rightarrow$ *it landed heads every time'*] and/or **4** [*'I dropped my pencil* $\Diamond\!\!\rightarrow$ *it flew off sideways'*] are false.

C. **3** [*'p* $\Box\!\!\rightarrow$ *q' and 'p* $\Diamond\!\!\rightarrow \neg$*q' are contradictories*], and therefore Lewis (1973)'s semantics for counterfactuals, is false.

My criticism of Lewis and Williams extends to any proposal that accepts **B** and denies **A** and **C**. Any such proposal will have to accept that there are values of *n* such that **8** [*'I flipped the coin n times* $\Box\!\!\rightarrow \neg$(*it comes up heads each time*)*'*] is true but **9** [*'I flipped the coin n times* $\Box\!\!\rightarrow \neg$(*it comes up heads the first* $\frac{n}{2}$ *times*)*'*] and **10** [*'I flipped the coin n times* $\Box\!\!\rightarrow \neg$(*it comes up heads the last* $\frac{n}{2}$ *times*)*'*] are false. But then **11**, **12**, or **13** cannot all be true. But I argued they all *are* true.[7]

I conclude that we should accept **A** or **C** rather than **B**.

Acknowledgements I presented an earlier and cruder version of the argument of this paper at the Basic Knowledge Seminar at the Arché Research Centre at the University of St. Andrews. I would like to thank all the participants there for their feedback. I am also grateful to Lars Bo Gundersen and Michael De for pointing out a problem in an earlier argument I had for the thesis I defend in this paper, to Robbie Williams for discussion, and to Jason Stanley for pointing me to Williams' paper. Finally, I would like to thank my two referees from *Synthèse* for their useful feedback and suggestions.

References

DeRose, K. (1999). Can it be that it would have been even though it might not have been? *Philosophical Perspectives, 13*, 385–413.
Dodd, D. (MS). Safety, skepticism, and lotteries.
Eagle, A. (MS). 'Might' counterfactuals.
Elga, A. (2004). Infinitesimal chances and the laws of nature. *Australasian Journal of Philosophy, 82*, 67–76.
Gaifman, H., & Snir, M. (1982). Probabilities over rich languages. *Journal of Symbolic Logic, 47*, 495–548.
Hájek, A. (MS). Most counterfactuals are false.

[7] These arguments did assume **C** to be false. So they don't pose a problem for someone who combined an acceptance of **B** with an acceptance of **C**. But it seems to me that if we accept **C**, we don't need to also accept **B** to solve the weird possibilities puzzle, and I don't see any other reason for accepting **B**.

Hawthorne, J. (2005). Chance and counterfactuals. *Philosophy and Phenomenological Research, 70*, 396–405.

Lewis, D. (1973). *Counterfactuals*. Oxford: Oxford University Press.

Lewis, D. (1986). Counterfactual dependence and time's arrow, and Postscripts to "Counterfactual dependence and time's arrow". In D Lewis (Ed.), *Philosophical papers*. Oxford: Oxford University Press.

Pritchard, D. (2005). *Epistemic luck*. Oxford: Clarendon Press.

Stalnaker, R. C. (1981). A defense of conditional excluded middle. In G. Pearce, R. Stalnaker, & W. L. Harper (Eds.), *Ifs*. Dordrecht: Reidel.

Williams, J. R. G. (2008). Chances, counterfactuals, and similarity. *Philosophy and Phenomenological Research, 77*, 385–420.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.