

Coincidences and the Grain of Explanation*

HARJIT BHOGAL

University of Maryland, College Park

Many events in every day life are *coincidences*. When I went to Denver last year I ended up sitting next to the same person on both my flight there and my flight back. That was a coincidence. Owens (1992, pp.8-9) considers the coincidence of him ending up on a cruise with his old enemy. Lando (2017, p. 133) describes the event of ‘the only two people in a supermarket aisle. . . wearing the same pair of 1980’s red Nike Air Ship high-tops’ as a coincidence.

So it is common to take certain everyday events to be coincidences. But as well as being an everyday notion, the notion of coincidence is theoretically important. In particular, there is pressure for scientific or philosophical theories to avoid coincidences. This pressure is evident in influential arguments in both philosophy and science.

For example, Kitcher (2001, p. 71) discusses the discovery of eighteenth century doctor John Arbuthnot who found that more males had been born than females in London in each of the last 82 years. There is a possible microphysical explanation of this fact, one which lays out separately all of the physical details which led to each male or female birth. But, Kitcher argues, such an explanation would make the regularity seem like a ‘gigantic coincidence’. However, there is an alternative evolutionary explanation of the regularity. There is an evolutionary tendency to produce a 1:1 sex ratio at the time of sexual maturity but it is fact about humans that infant mortality is higher in males than females. So, there is an evolutionary pressure for more males to be born than females — this explains the regularity. Kitcher influentially argues that there is value in the higher-level explanation of Arbuthnot’s regularity because it avoids making the regularity seem like a coincidence.

Or consider a scientific example: the cosmic coincidence problem is, very roughly speaking, that it appears that the proportion of energy of the universe coming from dark energy is very similar to the proportion coming from dark matter. However, our current best theories say at most other times of the universe there is not such an approximate equality — either dark energy dominates or dark matter does. This appears to be a coincidence. Again, there is theoretical pressure against such a coincidence. As Arkani-Hamed et al. (2000) note ‘If we believe the data, there appears to be no choice but to believe that we live in a special time in the history of the universe. This possibility is thought by many to be sufficiently

* Thanks to Sharon Berry, Nina Emery, Elizabeth Miller, Peter Tan, Jennifer Wang, an anonymous referee and audiences at the 2018 Vancouver Summer Philosophy Conference and the 2018 CHPS conference at CU Boulder.

distasteful to warrant disbelieving the data.’ Other physicists react differently, by building models that attempt to explain the purported coincidence (e.g. Zlatev et al., 1999). Either way, it is clear there is pressure against accepting the coincidence.

Even though the notion of coincidence is common in everyday life, and plays an important theoretical role, it’s somewhat unclear what makes for a coincidence. The aim of this paper is to give an account of coincidences. In particular, I’m going to be considering what makes a contingent event a coincidence. (So, for example, I’m not going to be considering *mathematical coincidences*, even though they might exist. (Lange, 2010))

I’m going to start by considering what Lando (2017) calls the *Traditional View* of coincidences, before considering her counterexample to that view. The success of this counterexample leaves us in need of new account. I’ll argue, however, that Lando’s suggested new account faces some significant problems.

I’ll then give my account of coincidence — one which draws parallels between the way we think about coincidences and the questions about the correct ‘grain’ or ‘level’ of explanation. In order to do this I’m going to outline an account of explanatory goodness and show how we can use this to give an account of coincidence.

1. Preliminaries

But, before we consider the traditional view, there is some preliminary work to do.

There are (at least) two major parts to our notion of a coincidence. Firstly, coincidences involve there being some kind of striking or compelling ‘match’ between distinct components. If I’m picking numbered balls from three urns and I pick the number 11 in all three urns that is a striking match, and (possibly) a coincidence. If I pick 78, then 9, then 82 there is no such match, and so there is no coincidence. The coincidences we mentioned above similarly involve such a match between components of the coincidence — a match between my flights to and from Denver with respect to the person sitting next to me; a match between the two people in the supermarket aisle with respect to what shoes they have on; and so on.

Secondly, coincidences have a distinct causal or explanatory structure. It is a coincidence if I pick 11, 11, 11 from the three urns at random. It is not a coincidence if the urns have been rigged, say by heating up the balls numbered 11 so I know which ones to pick. Similarly, it was a coincidence that I was sitting next to the same stranger on my flights to and from Denver. It would not be a coincidence if I were sitting next to my partner who was coming on the trip with me.

Like Lando and Owens I’m going to focus on the second part — looking at what the explanatory and causal structure of coincidences is — and I’ll black-box the notion of a ‘match’ between the parts of the coincidence, leaving further investigation for other work. (Though as Lando notes (p. 133), it seems plausible that the relevant notion of matching is best analyzed in psychological terms.)

Given this understanding of coincidences it will be useful to introduce some terminology. There are two ways of describing the situation in a putative coincidence. As we noted, coincidences involve a match between components or parts of the coincidence. That is to say, coincidences involve distinct events sharing a relevant property or feature. What this feature or property is will be different in different cases. So, we can describe a coincidence by noting the feature or property that is shared between events. So, for example, in the urn case we can describe the situation by saying that *the same number*

was picked from all three urns. Call the proposition that describes a coincidence in this way the *matching proposition*.

Similarly, in the case of my flights to and from Denver, the matching proposition is that *the same person was sitting next to me on both flights*. And in the case of Arbuthnot's regularity, the matching proposition is that *in each of the last 82 years, more males were born than females*. (Of course, Arbuthnot's regularity ultimately isn't a coincidence, but it does involve a relevant kind of match between component events, so we can describe it in the same way as actual coincidences.)

But we can also describe the coincidence in a different way too. We can describe the relevant properties of the separate component events in virtue of which there is a match. For example, in the urn case, we could describe the situation by saying that *ball 11 was picked from the first urn, ball 11 was picked from the second urn, and ball 11 was picked from the third urn*. This implies the matching proposition — that the same number was picked from all three urns. Call this proposition that describes the particulars of the relevant events the *particular proposition*.¹

In the Denver case the particular proposition is that I was on United 405 on 8/3/2017 in seat 37E and on United 596 on 8/6/2017 in seat 22A and the stranger was on United 405 on 8/3/2017 in seat 37D and on United 596 on 8/6/2017 in seat 22B. This implies that the same person was sitting next to me on both flights. In the case of Arbuthnot's regularity the particular proposition describes the numbers of male and female births each year — there was 10,341 male births last year, and 10,113 female; there was 10,238 males births two years ago and 9,987 female and so on. This implies that in each of the last 82 years, more males were born than females.

Importantly, the matching proposition is less specific than the particular proposition — the particular proposition entails the matching proposition but not vice versa.

2. The Traditional View

We can now consider what Lando calls the 'Traditional View' of coincidence. The traditional view is that the components of the coincidence — the two, or more, events that exhibit a striking match — are *causally independent*.²

On this view it is a coincidence that I picked 11, 11, 11 from the three urns because the causes of my picking 11 from the first urn were independent of the causes of picking 11 from the second and the third. And it's a coincidence that I sat next to the same person on my flight to and from Denver because the causes of my booking that flight and being assigned those seats are independent of the causes of his booking that flight and being assigned adjacent seats.

Notice that this formulation of the view isn't quite right because the parts of a coincidence will generally have some common causal ancestor if we go back far enough. Presumably the Big Bang is a common causal ancestor of the two parts of most coincidences, for example. The traditionalist could tweak their view in a couple of ways

¹ Of course, there isn't one unique particular proposition for each coincidence — we can describe with more or less specificity the properties of the component events in virtue of which there is a match. It will turn out that this isn't particularly important — the things I say will apply to all the possible particular propositions. So, for simplicity, I'm going to carry on talking as if there is just one.

² Lando notes that Owens (1992), Sober (1984) and Horwich (1982) amongst others hold versions of this view.

— Lando (p. 134) suggests distinguishing between salient causes and background conditions, allowing that coincidences can have common background conditions but not common salient causes. Owens (pp. 8-9) suggests that we can say that events can still be *partial coincidences* if the parts have some common ancestor as long as the important causes (in some sense) are not in common. There's a lot more to say about these ideas, but we don't need to go into detail here because, as Lando convincingly argues, there are counterexamples to the traditional view. And these defenses do not help with those counterexamples.

2.1. Lando's Counterexamples

Lando gives a counterexample to the traditional view — a case which is intuitively a coincidence but where the components of the coincidence are part of the same causal chain. In such a case the components are certainly not causally independent, or anywhere close to causally independent. Here is the main case that Lando discusses:

Pianos A boy is playing with a ball in the courtyard of an apartment complex. He throws the ball too high, and it bounces off of the balcony of one apartment, sails through the air, bounces onto the balcony of another apartment, and finally falls to the ground. On each of the two balconies sits a grand piano. As the ball lands on the first balcony it strikes a note on the first piano, and as the ball lands on the second balcony, it strikes a note on the second piano. On each of the two pianos, the note struck is the high A. (p. 135)

Clearly, it is a coincidence that the same note was struck on both pianos, but it is not the case that the two strikings are causally independent of each other. In fact, the second striking very much depended upon the first. The two components of the coincidence have most of their causal ancestors in common. This looks like a counterexample to the traditional view.

Further, neither of the previously discussed defenses of the traditional view seem to help. The throw of the ball is a salient cause, and that is a common cause of the two parts of the coincidence, so we can't say that the two parts of the coincidence only have non-salient causes in common. Neither can we say that the event is still a partial coincidence because most of the important causes are not shared between the parts — in **Pianos** most of the causal chain is in common between the two parts of the coincidence. Lando discusses, and criticizes, at length other possible responses to the counterexample that the traditionalist could give. She convincingly argues that those responses are unattractive.

But without going into the back and forth of possible defenses of the traditional view, seeing a case like **Pianos**, and seeing how we can pretty easily develop similar cases, should convince us that the traditional account of coincidences is on the wrong track. Coincidences just aren't to do with causal independence. We don't need small tweaks to the traditional view — we need a completely different approach.

3. Causal and Explanatory Relationships

So Lando has convinced us that we are in need of a new account of coincidence. How does she think we should proceed? Well, she doesn't give a full account of coincidence,

though she does raise some relevant considerations and suggest that these could point the way towards a full account. (Her main focus is slightly different — on the implications of her counterexamples for the connection between causation and causal explanation.)

Take, again, the distinction between the matching proposition and the particular proposition in *Pianos*. The matching proposition tells us that the same note was struck on both pianos; the particular proposition tells us that the high A was struck on the first piano and the high A was struck on the second piano. Lando claims that the matching proposition is a *relational fact* and suggests that relational facts must be explained differently to non-relational facts.

Even if we could show, separately, why the note struck on the first piano was an A and why the note struck on the second piano was an A, it seems to me we would not yet have explained why it is that the same note was struck on the two pianos. Why not? This last fact is a relational fact; the former are not. (p. 144)

What does Lando think that we need to do to explain a relational fact then? We need to ‘exhibit some relationship’ between the components of the coincidence. For example in the urn case, in order to explain why the same number was picked from the urns ‘we would have to exhibit some relationship between the numbers on the . . . balls’. (p. 144)

It seems that Lando thinks that in coincidences the relevant relational fact — the matching proposition — is unexplained because we cannot exhibit a relationship between the components of the coincidence.

However, there are some issues with this view. The first issue is less important, so I’ll just mention it quickly. The view puts a lot of weight on the distinction between relational and non-relational facts. Notice that this distinction goes much further than our distinguishing between the matching and the particular proposition in putative coincidences. The latter distinction just involves identifying two different ways that we could describe a certain situation — we could describe the match between components of the coincidence, or we could describe the features of the components in virtue of which there is a match.

But Lando uses the distinction between relational and non-relational facts in a way which suggests that it has much more substance — in particular, the distinction plays an important role in our theory of explanation. Relational facts must be explained differently from non-relational facts — to explain a relational fact we need to exhibit the relationship between the parts, this is not the case for non-relational facts.

On the face of it, then, a kind of disjunctive theory of explanation is required — facts are split into two different classes and we have to give a different story about explanation for each of the classes. Such a disunified theory of explanation is fairly unattractive on its face. But also, it puts a lot of theoretical weight on the distinction between relational and non-relational facts, perhaps more than the distinction can bear (especially given that the distinction between relational and non-relational facts is not especially well theorized in the metaphysics literature).

There are plenty of things Lando could say here, perhaps she could point towards an account of the distinction between relational and non-relational facts, or perhaps she could argue that, contrary to first appearances, she isn’t committing to a disjunctive theory of explanation. But let’s put this aside, because the second issue with Lando’s approach is more important.

The second issue is that is that in **Pianos**, and in similar cases, it's not clear that we can't exhibit a relationship between the two parts of the coincidence.

To see this, consider how we would explain the particular proposition — a non-relational fact, according to Lando — in **Pianos**. That is, the fact that the high A was struck on the first piano and the high A was struck on the second piano. We would note where the pianos were placed and then note how the ball was thrown — how hard and at what angle, with what spin and so on. And we would show how this leads it to hit the high A on the first piano, and how that leads to the ball bouncing to hit the high A on the second piano. Call this the *Trajectory explanation* since it cites the exact trajectory of the ball.

On the face of it, this explanation does exhibit a relationship between the two parts of the coincidence. It shows how the ball hitting the relevant note on the second piano is caused by it bouncing off the note on the first piano in the right way. It does not just explain why the high A was hit on the first piano, and then, totally independently, explain why it was hit on the second piano. It doesn't just conjoin two unrelated explanations. Rather, the explanation tells a single story about both notes being hit.

So, it seems like we could cite all the same factors to give the Trajectory explanation of the matching proposition — the fact that the same note was hit on both pianos — which does exhibit a connection between the parts of the coincidence. And of course, this would be in tension with the view that in coincidences we cannot explain the relevant relational fact — the matching proposition — because we cannot exhibit a relationship between the components of the coincidence.

Lando though, has a response available. She distinguishes *causation* from *causal explanation*. In particular, she claims that causation is 'an extensional relation between coarse-grained events' while '[c]ausal explanation, on the other hand, is an intentional relation between facts of a certain kind, or events under a description.' (p. 143)

This separation opens the possibility for her to accept that there is a causal connection between the two parts of the coincidence in **Pianos**, but *not an explanatory connection*.

By separating causation from causal explanation, however, we can still maintain that there is no explanation for the fact that the same note was struck. This even though the striking of the same note just was, as a coarse-grained event, the striking of an A here and an A there, and those two events have a common cause in the throw. (p. 143)

Even though there is a common causal chain between the two parts of the coincidence, Lando can deny that this is enough for explanation — although there is a causal relationship between the two parts of the coincidence we still cannot exhibit an explanatory relationship.

Although the distinction between causation and causal explanation opens up this logical possibility it's still unclear why we should think, in the particular case, that the causal connection is not enough for explanation. What component needs to be added to the causal connection? We need some positive story about why, in **Pianos**, the Trajectory explanation of the matching proposition fails.

What Lando says at this point is that a legitimate explanation of the matching proposition — one which properly exhibits the relationship between the parts — involves making it such that the parts of the coincidence are *not probabilistically independent* (p. 144). In particular, she claims that 'A on the first piano does not guarantee or make more

probable an A on the second', and this probabilistic independence precludes there being a relevant explanatory relationship between the two parts of the coincidence (even though, again, there is a causal relationship).

However, in **Pianos**, the coincidence is that the ball hit the high A on the first piano and the high A on the second piano, and it seems clear that, at least in some versions of the case, these events are not probabilistically independent. After all, the high A is all the way on one end of the piano, if the ball had struck the other end of the first piano, say the low G, it way well have been impossible, or close to impossible, for the ball to have struck the high A of the second piano. If this is the case then it is more likely that the high A is struck on the second piano given that it is struck on the first piano.

Consequently there would be probabilistic dependence between the two parts of the coincidence. We can't then claim that we have not properly exhibited a relationship between the parts of the coincidence on the grounds of probabilistic independence. (Perhaps there are some versions of the **Pianos** case where there is probabilistic independence between the parts. But the fact that there are versions of the case, which are clearly still coincidences, where there is no probabilistic independence means that Lando cannot appeal to such independence to defend general claims about coincidence.)

Of course, there's more that could be said here. Perhaps Lando could tweak her probabilistic criterion for when we have a genuine explanatory connection in a way which fits with our intuitive judgements about coincidences, or perhaps she could stress the idea that an explanation needs to 'exhibit' the lack of probabilistic independence between the parts of the coincidence, and give an account of what information explanations 'exhibit' which rules out the relevant explanations in **Pianos** and other coincidences.

I'm not going to delve into those options in detail, though both seem difficult to implement. But as it stands, I don't think Lando's probabilistic independence story backs her claim that in cases like **Pianos** we don't have a relevant explanatory relationship between parts of the coincidence even though we do have a causal relationship. Consequently, the view that in coincidences we can't explain the relational fact because we can't exhibit an (explanatory) relationship between the parts of the coincidence seems unsatisfactory.

Perhaps Lando could respond that it's just intuitively obvious that that we fail to exhibit an explanatory relationship in cases like **Pianos**, even if we haven't, so far, got a full account of why that's the case.

Nothing I've said so far rules out taking this kind of attitude. But I'm not sure I share this intuition. And, judging by discussions I've had with a variety of philosophers, people's intuitions seem to vary on this point. Some are committed to the idea that there has to be a kind of explanatory independence — a lack of explanatory relationship — between the parts of the coincidence in **Pianos**; some think that's not at all clear in light of the discussion of this section.

In any case, since it's not wholly clear if, or why, we fail to exhibit an explanatory relationship in cases like **Pianos** I think we have enough motivation to look elsewhere for an account of coincidence. I will start to do that in the next section.

Though, as it happens, one way of developing the approach that I will describe might, in fact, help us to validate the thought that there isn't a (substantial) explanatory relationship between the two parts of a coincidence, even when there is a causal relationship.

4. Coincidence and the ‘Grain’ of Explanation

Here’s where we are so far. Lando’s counterexamples to the traditional view convinced us we need a new account of coincidence. But her discussion of how we might develop this account doesn’t, I think, point in the right direction. In this section and the next I’m going to develop my account. In this section I’ll motivate the approach, in the next I’ll give the account in more detail.

To motivate the approach let’s compare two cases. The first is **Pianos**. This is a coincidence. The second is Arbuthnot’s regularity. Even though this case involves a striking match between component events it is not a coincidence. So, comparing these cases should help reveal what makes for a coincidence.

Let’s start by considering the explanatory structure of **Pianos**. How would we explain the particular proposition in this case — the proposition that the high A was hit on piano one and the high A was hit on piano two? As we noted in the previous section, we would give the *trajectory explanation*. That is, we would describe how the ball was thrown, and how that leads to the ball taking a particular trajectory to hit the high A on the first piano, and how that causes it to bounce onto the high A of the second piano.

How would we explain the matching proposition in **Pianos** — that is, the proposition that the same note was hit on both pianos? It looks like the only thing we could do is repeat the trajectory explanation. There’s no further story about the matching proposition once we have the trajectory explanation.³

Notice, though, that things are different in the case of Arbuthnot’s regularity. How would we explain the particular proposition — the proposition that there were 10,341 male births last year, and 10,113 female; there was 10,238 males births two years ago and 9,987 female and so on? We would have to go into great detail about what caused the exact numbers of births that there were, and of each birth we would go into detail about why it was a male birth or a female one. (Clearly this is possible in principle, even if we don’t have the knowledge to give this explanation in practice.) But in order to explain the matching proposition — the fact that more males were born than females in each of the last 82 years — we don’t need to just repeat this story. We can give the evolutionary story described earlier in the paper — the story that cites the pressure for a one-one sex ratio at sexual maturity and the greater rate of infant mortality among males.

This is a striking difference between the two cases — in **Pianos** we are inclined to explain the matching proposition and the particular proposition in the same way, for Arbuthnot’s regularity we explain the propositions very differently. And this difference seems relevant for the way in which one case is a coincidence and the other is not — as Kitcher (2001, p. 71) notes, it is precisely the existence of this distinct evolutionary explanation of Arbuthnot’s regularity that makes us recognize that the regularity is not a coincidence. The explanation where we separately explain why each birth was a male or female birth makes the regularity appear like a coincidence.

³ Someone might object that this just doesn’t count as an explanation of the matching proposition. I don’t see a particularly compelling reason to deny the trajectory explanation of the matching proposition. But, in any case, I’m going to end up saying that the trajectory explanation is, in a well-defined way, a *bad* explanation, even though it strictly is an explanation. In fact, the difference between saying it is a bad explanation versus saying it is not an explanation at all may well be close to verbal. The key point right now is just the best we can do when trying to give an explanation of the matching proposition in **Pianos** is just to appeal to the trajectory explanation again.

But what exactly is the upshot of our inclination to explain the matching proposition in the same way as the particular proposition in cases like **Pianos**? Put simply, the upshot is that in cases like **Pianos** the matching proposition cannot be explained as well as the particular proposition.

Why is this? The key is the point that we noted in section 1, that the matching proposition is less specific than the particular proposition — the matching proposition holds in more possible worlds than the particular proposition. The matching proposition entails the particular proposition but not vice versa. Consequently, the trajectory explanation is a good explanation of the more specific particular proposition, but a much worse explanation of the less specific matching proposition.

Intuitively, the trajectory explanation is a bad explanation of the matching proposition in **Pianos** because it is too tied to the specific way in which the same note was played in the actual world. Explaining why the same note was played on both pianos by appealing to the exact details of the throw ignores all the non-actual ways in which the same note could have been played.

The trajectory explanation tells us about the precise details of the way that the ball landed on the high A of the first piano, and how it bounced onto the high A of the second piano, but there are many other ways in which the same note could have been played on both pianos — that is, there are many other ways in which the matching proposition could be true. For example, the same note could have been struck on both pianos in virtue of the high-A being struck on both, as in the actual situation, but also in virtue of any other note being struck on both pianos. The trajectory explanation fails to apply in any of these non-actual situations where the matching proposition is true. Consequently, the trajectory explanation of the matching proposition feels too specific and so sub-optimal.

These considerations are very similar to those that come up in the literature on the correct ‘level’ or ‘grain’ of explanation. Consider Putnam’s (1975) classic example of a square peg and a round hole where the hole diameter is the same as length of the side of the peg. I fail to put the square peg in the round hole. What explains this? One possible explanation cites the exact microphysical details of my attempts to put the peg in the hole and shows how the basic physical forces lead to the peg not going in the hole. This explanation seems far too specific and too tied to the details of the actual world, ignoring all the other possible ways I could have failed to put the peg in the hole. A much better explanation just notes the geometric facts about the peg and the hole, and their solidity, to show how the peg can’t fit in the hole.

Another example: we can give an economic explanation of why the dollar declined in value by noting that the central bank cut interest rates thus reducing the incentive for investors to put their money in US accounts and so reducing demand for the dollar. But we could also, at least in theory, give an microphysical explanation that talks about the movement of all the particles that make up the economic system. However, the economic explanation here seems appropriate in a way that the microphysical explanation does not — the microphysical explanation is too specific and too tied to the actual physical details.

More generally, many philosophers agree that often we should favor ‘higher-level’ or ‘less specific’ or ‘more abstract’ explanations over the alternative explanations of the same fact. These philosophers differ in the exact reasons we should sometimes favor less specific explanations. But many philosophers share the basic thought, that we expressed

with respect to the peg-hole and dollar-decline cases, that explanations should be in a certain way ‘robust’ — they should apply to a wide range of possible situations and not only the very specific situation that actually occurred (see, for example, Weslake (2010); Garfinkel (1981); Wilson (1994)).

Alternatively, some philosophers make the point in a different, but closely related way. Often it seems that we want a kind of *proportionality* between explanans and explanandum (Yablo, 1992; Woodward, 2010). That is, we want the specificity of the explanans to ‘match’ the specificity of the explanandum. To adapt very slightly a classic example from Yablo, imagine that a pigeon is trained to peck at a target if and only if the target is red. Then consider one specific case of the pigeon pecking at a target. We can explain this by noting that the target was scarlet. But this seems overly specific because the pigeon would have pecked at other shades of red. The explanans here is more specific — it asserts that the target is scarlet — while the explanandum is more general — it holds in cases where the target is any shade of red. So, there is a lack of proportionality between explanans and explanandum. A better, more proportional, explanation would explain the pecking by noting that the target was red.

In a similar way the trajectory explanation of the matching proposition in **Pianos** is flawed because it fails to be robust and the explanans isn’t proportional to the explanandum.⁴

However, the trajectory explanation of the *particular proposition* — the fact that the high A was played on both pianos — is not flawed in this way, or at least not as flawed. This is because the particular proposition is more specific than the matching proposition, so it more reasonable that the explanation of that proposition be specific too. The trajectory explanation doesn’t fail to apply to lots of other possible ways in which the high A could have been played on both pianos, because there aren’t very many other ways in which the high A could have been played on both pianos. Or, to put it another way, the explanans and the explanandum are somewhat proportional — both the explanans and the explanandum are rare, specific, events.

In summary, then, in **Pianos** the matching proposition is explained in the same way as the particular proposition. This has the consequence that the explanation of the matching proposition is flawed — it is too specific for the explanandum.

Things are very different in the case of Arbuthnot’s regularity. As we noted, in that case we explain the matching proposition differently from the particular proposition — we give the evolutionary explanation of the matching proposition. The evolutionary explanation doesn’t face analogous problems with proportionality and robustness — it isn’t too specific for the explanandum. And again, it is the existence of this good evolutionary explanation that convinces us that Arbuthnot’s regularity is not a coincidence.

Comparing these cases suggests a way to think about coincidences: Coincidences arise when the only way we can explain the matching proposition is to appeal to an explanans

⁴ There are other ways of making the case that we should often favor less specific explanations. In particular, some philosophers argue that we should often reject more specific explanations because they cite things which are not *difference-makers* for the event in question (Strevens (2008) develops this thought in the most detail). Although I won’t focus on this thought in what’s to come I’m happy to accept this diagnosis for the why the trajectory explanation of the matching proposition is flawed — it cites non-difference-makers. And, I think, the account of explanatory goodness I sketch in the next section, along with the account of coincidences built upon that, is very congenial to this thought about difference-making. Though that’s a much longer story than I’m able to go into here.

that gives us a flawed, overly specific, explanation. However this explanans is much more appropriate for explaining the particular proposition.

This is the core idea of the account. In the next section, I'm going to develop a more detailed account along these lines.

5. Explanatory Goodness and Coincidence

The way I'm going to develop the ideas of the previous section is by outlining an account of *explanatory goodness* — that is, an account of when explanations are better or worse. (This is in contrast with an account of *explanatory correctness*, which is an account of when something genuinely is an explanation and when it is not.) The key idea is that in coincidences the explanation of the matching proposition is flawed because it is too specific while the explanation of the particular proposition is less flawed.

Once we have an account of explanatory goodness we can develop this key idea as follows: in a coincidence the matching proposition is explained substantially worse than the particular proposition and this is in virtue of the way that the the explanation of the matching proposition is too specific. (We will have to see the account of explanatory goodness before we can make this idea fully precise.)

Although I am appealing to my preferred account of explanatory goodness I don't mean to claim that the basic idea — that in coincidences the explanation of the matching proposition is flawed because it is too specific — can only be developed using this account. Other accounts of explanation or explanatory goodness may be able to do the job too. But I don't have the space to explore these possibilities here.

Similarly, I don't have space to mount a complete defense of the account of explanatory goodness — though I will note why it is plausible. The point is simply to exhibit one plausible way that we can develop the basic idea that in coincidences the explanation of the matching proposition is flawed because it is too specific, rather than to commit to the exact details of the account of explanatory goodness.

So I'm going to start by outlining the account of explanatory goodness. Then I'll give an account of coincidence in those terms.⁵

The account of explanatory goodness starts by assuming the notion of explanatory correctness, and it builds the account of explanatory goodness from this. The account of explanatory goodness doesn't presuppose any particular account of explanatory correctness, just the generic notion of explanatory correctness is required. But for definiteness, and ease of presentation, I'm going to fix on one very simple account of explanatory correctness. The account is in the spirit of Lewis (1986). It says that A explains B when A gives information about the nexus of causal relations that led up to B.

Given this, we can outline two dimensions of explanatory goodness. The first dimension of explanatory goodness, *PRECISION*, says that explanations are better if in more of (that is, a higher proportion of) the physically possible worlds where the explanans is true, the explanans explains the explanandum. For example, explaining the ice cube melting by the fact it was dropped in warm water scores very highly, but not maximally, on *PRECISION* because nearly all the physically possible microstates that could realize the system are such that dropping the ice cube in warm water does lead to the melting. The

⁵ I discuss this account of explanatory goodness and its connection to questions about the grain of explanation in much more detail in other work.

explanation of the glass of water spilling because I knocked the table on which it was sitting with my knee scores worse on PRECISION, because there are many physically possible worlds where I knock the table with my knee but that does not lead to any spilling (for example, those where I don't knock the table hard enough).

The intuitive force behind PRECISION is that it captures the classic idea that explanations should make the explanandum *expectable* (Hempel, 1965; Salmon, 1989). If A explains B and that scores highly on PRECISION that means that in most of the physically possible worlds where A holds, A explains B. And that implies that in most of the physically possible worlds where A holds, B holds. Consequently, A makes B, in a sense, likely — A should lead us to expect B.

The second dimension, ROBUSTNESS, says that explanations are better if in more of (that is, a higher proportion of) the physically possible worlds where the *explanandum* is true, the explanans explains the explanandum. An explanation of the extinction of a species which cites the introduction of a predator will score higher on ROBUSTNESS than the explanation of the same fact which cites the deterministic fundamental laws and the fact *f* which describes the total state of the world at a particular time *t*. That is because more of the physically possible worlds in which the species goes extinct are such that the extinction is explained by the predator's introduction than are such that the extinction explained by *f*. *f* cannot explain the extinction in many worlds because it holds in only very few physically possible worlds.

The intuitive force of ROBUSTNESS is that it captures the sense in which the explanans should be *required* for the explanandum. Take an explanation that scores badly on ROBUSTNESS and so the explanans holds in few of the worlds where the explanandum holds. In this case the explanans must cite things that are not required for the explanandum, since there are so many worlds where explanandum holds but the explanans does not. In general, explanations which score poorly on ROBUSTNESS are going to cite things that are not close to required for the explanandum. This, I think, is the best way to understand and motivate robustness: It says that we should only cite facts in explanations that are required, or close to required, for the truth of the explanandum.

In general, PRECISION will favor lower-level, more specific explanations - adding more (relevant) detail to an explanation will tend to increase PRECISION by ruling out worlds where the explanans holds but the explanandum does not. ROBUSTNESS, on the other hand, will tend to disfavor very specific explanations, like the explanation of the extinction mentioned above that cites the total state of the world at a particular time.

PRECISION and ROBUSTNESS seem especially plausible as dimensions of explanatory goodness when we see how they interact. ROBUSTNESS tells us that that the explanans shouldn't hold in few possible worlds when the explanandum holds in many. And PRECISION tells us that the explanans shouldn't hold in many worlds when the explanandum holds in few. Together they tell us that there should be *proportionality* between explanans and explanandum. Figure 1 illustrates this.

In the diagrams the boxes represent the space of physically possible worlds and the circles represent the range of worlds where the propositions A and B hold. In all three cases we are considering an explanation of B from A. In 1(a) A holds in a wider range of worlds than B, so there are many worlds where A holds where A does not explain B. Consequently, this explanation will score badly on PRECISION. 1(b) is the opposite case — there are many worlds where B holds, where A does not explain B. So this explanation




Figure 1. PRECISION and ROBUSTNESS together imply proportionality

will score badly on ROBUSTNESS. When PRECISION and ROBUSTNESS are both maximal A and B perfectly overlap. That is, A is perfectly proportional to B.

So this is the account of explanatory goodness. It is one that allows us to identify the right level of specificity of explanation. For example, in Yablo’s pigeon case, we can see that the explanation of the pigeon pecking that says that the target was red is better than the one that says that the target was scarlet because the former explanation scores better on ROBUSTNESS and no worse on PRECISION. Of course, there is a lot more to say about this account of explanatory goodness and how it answers questions about levels of explanation.⁶ But that is not our focus here — our focus is coincidence.

And we can use this account of explanatory goodness to give an account of coincidence. In particular, let’s look at how this account of explanatory goodness applies to the cases of **Pianos** and Arbutnot’s regularity.

As we noted in the previous section, in **Pianos** we give the same explanation of both the particular proposition and the matching proposition. We explain them both with the trajectory explanation — the explanation that cites the exact way the ball was thrown (the angle of release, the speed, the spin on the ball etc.) and then describes how this leads to the ball landing on the first high A, and how it then bounces onto the second high A.

However, the trajectory explanation is a bad explanation of the matching proposition. Intuitively, it is too specific; too tied to the actual ways things went and not the possible ways things could have gone. The account of explanatory goodness fills out this thought. The trajectory explanation of the matching proposition scores well on PRECISION — most (or maybe all) of the ways that the explanans holds are such that the explanans explains the explanandum. That is, most of the possible ways that the ball could be thrown with exactly that trajectory would lead to the ball hitting the high A on both pianos and the matching proposition being true.

But the explanation scores badly on ROBUSTNESS because most of, in fact nearly all, of the ways that the matching proposition could hold are such that it is not explained by the explanans. That is, in most of the physically possible worlds where the same note is struck on both pianos it is not explained by the facts cited in the trajectory explanation. For example, the matching proposition could hold in virtue of the ball hitting the low G on both pianos. But in this case the matching proposition would not be explained by the trajectory explanation — the particular facts about the actual trajectory that are appealed to in the trajectory explanation would not explain the matching proposition in this non-

⁶ To just quickly mention two issues, we will need to rule out certain gerrymandered, unnatural, properties which are specifically designed to form an explanans that is proportional to the explanandum. And we will need to say something about how to define the measure over physically possible worlds.

actual case because they would be false. And of course the same is true for all the other ways in which the matching proposition could hold.

This is illustrated in figure 2. The trajectory explanation of the matching proposition is depicted in 2(a) — the explanans, A, holds in too few worlds in comparison to the explanandum, B, — it fails to be proportional to the explanandum — and so it scores poorly on ROBUSTNESS.


However, the trajectory explanation of the particular proposition doesn't face this problem — or at least not to such a large extent. There are many fewer ways in which particular proposition could hold and so many fewer cases where explanandum holds but the explanans does not. This is shown in figure 2(b). The explanans, A, is the same as in figure 2(a), but since the explanandum holds in fewer worlds the overall explanation scores better on ROBUSTNESS and is more proportional.

In summary then, in **Pianos** the matching proposition is explained badly — worse than the particular proposition — because the explanation scores poorly on ROBUSTNESS.

The same thought applies to the other coincidences we considered. In the urn case, the matching proposition — that the same number was picked from all three urns — is explained worse than the particular proposition — that the number 11 was picked from all three. The only way we can explain why the same number was picked in all three cases is to give the same explanation as the particular proposition — going into detail about why the number 11 was picked in each case. But, there are many other ways in which the same number could have been picked where this explanation does not hold. Consequently, the situation is the same as the one depicted in figure 2 — the explanation of the matching proposition scores poorly on ROBUSTNESS.

Similarly in the Denver case the explanation of the matching proposition — that the same person was sitting next to me on both flights — scores much worse on ROBUSTNESS than the explanation of the particular proposition — that I was on United 405 on 8/3/2017 in seat 37E and the stranger was in seat 37D and so on.

Things are different in the case of Arbuthnot's regularity. Consider the matching proposition in Arbuthnot's regularity — the proposition that more males were born than females in London in each of the 82 years before John Arbuthnot's discovery. Imagine that we explained this in the same way as we would explain the particular proposition — by describing separately the actual details of why each of the births was a male or female birth. This explanation clearly scores poorly on ROBUSTNESS — most of the ways that the regularity could hold would not be explained by these details. (To be clear, this is not to



(a) Trajectory Explanation of the Matching Proposition (b) Trajectory Explanation of the Particular Proposition

Figure 2. The Trajectory Explanation of the Matching and Particular propositions

say that we couldn't explain non-actual instances of the regularity by giving an analogous explanation which points to details about why each of the births was a male or female birth in *that* world, but this explanation would be different from the actual world explanation.)

But again, we don't need to give this explanation — we can give an evolutionary explanation of the matching proposition. And this explanation does much better on ROBUSTNESS. That explanation doesn't cite lots of details of the actual world, rather it just cites general evolutionary pressure and increased rate of infant mortality amongst males. In lots of the other worlds where the regularity holds these factors do explain the regularity, hence the evolutionary explanation scores much better on ROBUSTNESS. (Though obviously, even with this explanation there are some worlds where the regularity holds but this explanation does not apply — so it doesn't score maximally on ROBUSTNESS.)

We can see how the account of explanatory goodness allows us to draw the distinction between coincidences like **Pianos** and non-coincidences like Arbuthnot's regularity. So, we are now in a position to give an account of coincidence. On my account an event is a coincidence if and only if (i) it involves a striking 'match' between two parts of the coincidence and (ii) the matching proposition is explained significantly worse than the particular proposition because the matching proposition scores badly on ROBUSTNESS.

This account, as we have started to see, seems to classify particular cases of coincidence correctly — it gets the right result about Arbuthnot's regularity, **Pianos**, the Denver case and the urn case. And what's more, it seem to get the results for the right reasons — it does seem to be the case that the existence of a more robust evolutionary explanation is what makes Arbuthnot's regularity non-coincidental.

In fact, when Lando (section 7) gives examples of explanations that would dispel a purported coincidence they are examples of exactly this kind — they are cases where we can appeal to a higher-level explanation of the matching proposition that does much better on ROBUSTNESS than explaining the parts of the coincidences separately.

For example, she considers a case where 'a card dealer asks Augie to select a single card from a full, covered deck. The card he pulls at random is the Queen of diamonds. Several hours later Dylan visits the same table. The card dealer asks Dylan to select a card, and the card he pulls is also, as it happens, the Queen of diamonds.' The coincidence is dispelled, though, by noting that 'because of the particular way Augie handles his cards, he makes a mark on the Queen of diamonds as he pulls it out of the deck, causing it to protrude slightly from the rest of the cards. When it comes time for Dylan to choose, his eye is drawn inadvertently to the protruding card.' I agree that this dispels the coincidence. I claim this is because we have an reasonably robust explanation of the matching proposition — the same card was pulled because Augie makes a mark on the card that Dylan is attracted to. This explanation is more robust than the explanation of why the same card was pulled that gives the details of the Queen of Diamonds being pulled in both cases.

I take it that my account fits with the intuitions about such cases. More generally, my account tells a natural story about the explanatory problems with coincidences. In particular, it makes sense of the intuitive idea that in coincidences we can only explain the matching between parts by going looking into the details of each part, rather than telling a more abstract story about the connection between the parts. And it tells us what is explanatorily problematic about this.

I think these are good reasons to look favorably on this account — it seems to get the right results in a variety of cases, and the reasons it gets those results seem natural and intuitive.

6. Comparison to Previous Views

Just before we finish it will be useful to briefly compare my view back to the traditional view and Lando's view.

Take Lando's account. Again, the basic idea was that the matching proposition is a *relational fact* and in coincidences we can't explain the relevant relational fact because we can't exhibit an explanatory relationship between the two parts of the coincidence. For example, in **Pianos** the fact that the same note was struck on both pianos is a relational fact, and we can't explain it because we can't exhibit an explanatory relationship between the two parts of the coincidence — that is, the two strikings.

One criticism I made of Lando's approach is that it puts a lot of weight on the distinction between relational and non-relational facts, seemingly suggesting a disunified account of explanation — that explanation works differently for these two classes of facts. We can now see that my account avoids doing this. I do distinguish between the matching and the particular proposition in cases of coincidence. But there is no sharp metaphysical distinction here — they are just two different ways of describing the same situation. And the distinction has no weight in the theory of explanation — the same account of explanation, and the same measure of explanatory goodness apply to both.

In that respect I think that my account has an advantage over Lando's. Here is another apparent difference between the views, though this one ends up being pretty superficial. My account says that in coincidences the matching proposition is explained badly — it's scores poorly on explanatory goodness (on ROBUSTNESS in particular. Lando says that in such cases the matching proposition, what she says is a relational fact, is not explained at all. The difference between being explained badly or not explained at all is not significant — in fact, it might be close to verbal. And further, it is somewhat easy to adapt my account so that putative explanations that score sufficiently badly on explanatory goodness are, in fact, not explanations at all.

The more substantial disagreement in this region is in what makes for this explanatory deficiency in coincidences — whether we describe that deficiency as a certain proposition being explained badly or not explained at all. Lando's claim is that the deficiency arises from us not being able to exhibit a relevant connection between the parts of the coincidence, and she appeals to a story based on probabilistic independence to say how this can be so when there is a causal connection. I criticized this approach at length in section 3. I, on the other hand, argued that the explanatory deficiency is of the same type as the deficiency that arises when we try to give a microphysical explanation of an economic phenomenon — it's a failure of ROBUSTNESS.

However, we could attempt to minimize (though not remove) this disagreement by fitting my approach into the framework given by Lando. In particular, we might take the goodness of the explanation of the matching proposition to be a kind of measure of whether there is a genuine explanatory relationship between the parts of the coincidence. In **Pianos** the best explanation of the fact that the same note was hit on both pianos scores badly on explanatory goodness because it scores badly on ROBUSTNESS. We might take this to imply that, in the relevant sense, there is only a thin explanatory relationship between the two parts.

Understood in this way the account recovers something like the thought that in coincidences we cannot exhibit a substantial explanatory relationship between the two parts of the coincidence — the idea that there is a kind of *explanatory independence* between the

parts of the coincidence, or at least, something close to explanatory independence (though this thought would be understood in a different way from Lando).

This idea that the explanatory goodness of the best explanation of the matching proposition measures the degree of the explanatory relationship needs more discussion to make it fully plausible. But it looks like a promising option for developing my approach, one that would be amenable to those that who are attracted to Lando's view.

Now consider the traditional approach again. That accounts says that something is a coincidence when two, or more, events exhibit a striking match and those events are *causally independent*.

We have seen counterexamples to the traditional approach — cases where there are coincidences even though there is no causal independence. But our discussion suggests the possibility of counterexamples in the other direction too — cases which are not coincidences but where there is causal independence.

Sometimes we can give good, robust, explanations of the matching proposition by appealing to a common mechanism or principle, rather than a common cause. In fact, the evolutionary explanation of Arbuthnot's regularity might be of this form. When we explain why more males were born than females in each year we don't obviously cite a common cause, rather we show that the same mechanism is operable in each case. An even clearer example might be the fact that in every collision we have observed momentum is conserved. The way we dispel the appearance of coincidence here is not to cite a common cause, but to cite a common law or principle.

It appears that these cases are not coincidences, but that's not because of any causal relation between the separate events. (Of course, a defender of the traditional view might claim that there actually is causal dependence in such cases — any two collisions have some causal ancestors, perhaps the Big Bang — so this isn't a counterexample to their view. But this isn't a dialectically effective response — as we noted in section 2 saying that any two events are not causally independent because they have the Big Bang as a common ancestor would prevent the view from being able to classify any genuine coincidences as such.)

But there is a class of cases where there is something like causal independence between two matching events (again, ignoring worries about how everything has certain shared causal ancestors), and where there is no alternative explanation that appeals to a shared principle or mechanism. In those cases we will have to explain the matching proposition by conjoining disconnected explanations of the two parts of the coincidence. And, for the reasons discussed above, such explanations will be better explanations of the particular proposition than the matching proposition. It is this class of cases that the traditional view gets right.

7. Conclusion

I think the account developed here is an attractive account of coincidence. In particular, it seems the right results in a variety of cases and, as we noted in section 5 it seems to get the results for the right reasons — the reason that Arbuthnot's regularity, and other similar cases, are not coincidental is because there is this robust higher-level explanation of the matching proposition.

But I'm going to finish by mentioning a couple of other positive features of this account of coincidence, though I won't have time to discuss them in detail. Firstly, my account suggests a plausible story about the theoretical norm of avoiding coincidences.

Coincidences, in my view, involve bad explanations. So this points to the possibility of assimilating the norm of avoiding coincidences to principles of inference to the best explanation.

Secondly, as we mentioned in passing in section 2, we might naturally think that things can be more or less coincidental. For example, the fact that I had the same person sitting next to me on my flight to and from Denver becomes seems less coincidental when we point out the attractive features of the flights that might have encouraged both of us to book those particular flights. My account has a way of making sense of this, and even measuring the degree of coincidence. My account says that in coincidences the matching proposition is explained significantly worse than the particular proposition because the matching proposition scores badly on ROBUSTNESS. So, this suggests that an event is more coincidental if this difference in ROBUSTNESS between the matching proposition and the particular proposition is bigger.

So, there is more to be said in order to give a full evaluation of the account, but I hope to have shown that the view of coincidence developed here is a promising one.

References

- Arkani-Hamed, N., L. J. Hall, C. Kolda, and H. Murayama (2000). New perspective on cosmic coincidence problems. *Phys. Rev. Lett.* 85, 4434–4437.
- Garfinkel, A. (1981). *Forms of explanation*. Yale University Press New Haven.
- Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. The Free Press.
- Horwich, P. (1982). *Probability and Evidence*. Cambridge University Press.
- Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford University Press.
- Lando, T. (2017). Coincidence and common cause. *Noûs* 51(1), 132–151.
- Lange, M. (2010). What are mathematical coincidences? *Mind* 119(474), 307–340.
- Lewis, D. (1986). Causal explanation. In D. Lewis (Ed.), *Philosophical Papers Vol. II*, pp. 214–240. Oxford University Press.
- Owens, D. (1992). *Causes and Coincidences*. Cambridge University Press.
- Putnam, H. (1975). Philosophy and Our Mental Life. In *Mind, Language and Reality: Philosophical Papers, volume 2*. Cambridge: Cambridge University Press.
- Salmon, W. C. (1989). 4 decades of scientific explanation. *Minnesota Studies in the Philosophy of Science* 13, 3–219.
- Sober, E. (1984). Common cause explanation. *Philosophy of Science* 51(2), 212–241.
- Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press.
- Weslake, B. (2010). Explanatory depth. *Philosophy of Science* 77(2), 273–294.
- Wilson, R. A. (1994). Causal Depth, Theoretical Appropriateness, and Individualism in Psychology. *Philosophy of Science* 61(1), 55–75.
- Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology and Philosophy* 25(3), 287–318.
- Yablo, S. (1992). Mental Causation. *The Philosophical Review* 101(2), 245–280.
- Zlatev, I., L. Wang, and P. J. Steinhardt (1999). Quintessence, cosmic coincidence, and the cosmological constant. *Phys. Rev. Lett.* 82, 896–899.