

Forthcoming in *Russell's Republic: The Place of Causation in the Constitution of Reality*. Huw Price and Richard Corry, eds. Oxford University Press.

Isolation and folk physics¹

Adam Elga

There is a huge chasm between the sort of lawful determination that figures in fundamental physics, and the sort of causal determination that figures in the ‘folk physics’ of everyday objects. For example, consider a rock sitting on a desk. In everyday life, we think of the rock as having a fixed stock of dispositions--the disposition to slide on the desk when pushed, to shatter when struck by a sledgehammer, and so on. When a strong interaction comes the rock’s way, the rock’s dispositions determine how it will respond. More generally, we think of the behavior of an ordinary object as being determined by a small set of conditions. The conditions typically specify the object’s dispositions to respond to various sorts of interference, and describe the sorts of interference that the object in fact encounters. Call this the ‘folk model’ (Norton this volume).

In fundamental physics, no small set of conditions suffices to determine an ordinary object’s behavior. Instead, differential equations describe how the exact physical state of the world at one time² lawfully constrains its state at other times (Russell 1913). In the worst case--the case of non-local laws--one would have to specify the entire state of the world at one time, in order to determine the state of even a small region at some future time.³ And even if locality holds in the sense of relativity theory (so that no influences travel faster than light), one would still have to specify the state of a huge region of the world (Field 2003:

439). For example, in order to determine what a rock will do in the next 0.05 seconds, one would have to specify the exact present state of the entire Earth.

What to make of this chasm between the two sorts of determination? One reaction would be to utterly renounce the folk model. The most extreme portions of Russell (1913) can be seen as advocating that reaction. That reaction would be too extreme. The folk model is useful. It seems to capture important features of the world. So there must be *something* right about it. What?

Norton (this volume) answers: the folk model is approximately correct, in certain limited domains. Here is the idea. When an old scientific theory is superseded by a new one, sometimes the new theory allows us to derive that, in a certain limited domain, the old theory is approximately correct. For example, one can derive from General Relativity that the classical theory of planetary orbits is approximately correct, provided that spacetime isn't too curved. Norton argues that the folk model can be recovered as approximately correct, in an analogous way.

On this picture, our fundamental laws have a very special feature. They are such as to make the folk model approximately true in certain domains (including the domain of the mundane comportment of medium-sized dry goods).⁴ Not all laws have that feature. Some fundamental laws, represented by perfectly respectable differential equations, do not make the folk model even approximately true in any domains at all.

So there is a story to tell about how *our* laws yield the approximate truth of the folk model in certain domains. My aim is to tell part of that story.

1. Extreme locality

Why do ordinary objects behave in ways that fit the folk model? The first step in answering is to explain why any objects exist in the first place. Why isn't there just an amorphous soup, for example, or no stable matter at all? Answering such questions is beyond the scope of this paper.⁵ But the questions are still worth posing, in order to emphasize that their answers are not obvious. It is not automatic that a system of fundamental laws should allow for stable objects--huge quantities of particles that tend to move as a unit. Thankfully, our laws do.

Now: given that there are relatively stable middle-sized objects, why does their behavior roughly accord with the folk model? In answering, it is best to start by considering the *extreme locality* of the folk model.

Suppose that our fundamental physical laws are local, in the sense that the speed of light is the maximum speed of signal propagation.⁶ That sort of locality guarantees that the rock on your desk is isolated from very distant goings-on. For example, when it comes to what the rock will do in the next ten seconds, it absolutely does not matter what is going on now at the surface of the sun.⁷ That's because not even light could get from the sun to your rock in ten seconds.

Locality, in the above technical sense, gives a certain guarantee that objects are isolated from distant goings-on. But the guarantee only concerns *very* distant goings-on, since the speed of light is so high. Example: for all the guarantee says, the observable behavior of your rock in the next second depends on whether someone blinks right now at the opposite

end of the Earth. (Remember that light can cover the diameter of the Earth in a twentieth of a second.) Indeed, for all the guarantee says, everything on the Earth could be so massively interconnected that any change anywhere on the Earth would make huge, unpredictable differences everywhere else within a twentieth of a second.

Thankfully, things aren't nearly so interconnected. When it comes to the behavior of the rock sitting on your desk in the next second, you can pretty much ignore what's going on right now at the other end of the Earth. It's not that you have a guarantee that your rock's imminent behavior is *utterly independent* of what's going on at the opposite end of the Earth. For if there were a supernova there right now, that would certainly make a big difference to your rock in less than a second (Field 2003: 439). You don't have a guarantee of *absolute* isolation. Instead, you have an assurance that subject to some very weak background conditions (for example, the condition that there will be no gigantic explosions, and that the mass of the Earth will remain roughly unchanged), the rough behavior of your rock is independent of what is going on at the opposite end of the Earth. Indeed, you have a similar assurance that the rough behavior of your rock is independent of what is going on down the block.⁸ That's what it means for the generalizations that figure in the folk theory to be 'extremely local'.

In short, when it comes to the rough behavior of your rock, you can often treat it as if it were isolated from distant influences. Note the qualifications, though. Your rock isn't *really* isolated from distant influences. For example, the exact microscopic trajectories of the rock's molecules are sensitive to goings on in the next room, due to gravitational effects. But who

cares about the exact trajectories of rock molecules? When it comes to getting around in the world, the rough macroscopic behavior of rocks is much more important.

Now: *why* is it that you can treat the rough macroscopic behavior of your rock as if it were independent of distant influences? There are two factors.⁹

The first factor is that the forces acting on your rock from afar are either negligibly tiny or nearly constant. There are four forces to consider: strong, weak, electromagnetic, and gravitational. The strong and weak nuclear forces fall off in strength dramatically at greater than atomic-scale distances. Electromagnetic forces are stronger and longer-ranged, but--around here anyway--there are few strongly charged macroscopic objects. So when it comes to the strong, weak, and electromagnetic forces, objects are only subject to tiny distant influences.

The one remaining force, gravitation, operates with significant strength even at long distances. But in our neighborhood of the universe, mass distributions do not rapidly and massively fluctuate. So gravitational forces on Earth don't change much over short distance or time scales. As a result, we can treat the gravitational force as a fixed background. Relative to that background, changes in distant matters make only negligible gravitational differences to medium-sized objects on Earth.¹⁰ The bottom line is that differences in distant matters of fact only make for tiny differences in the forces acting on ordinary objects.

Here enters the second factor: statistical-mechanical considerations show that tiny differences in the forces acting on the rock are very unlikely to affect its rough macroscopic behavior. How does that go? On the assumption of determinism, one standard story is that the fundamental laws supply a probability distribution over initial conditions of the universe. That

probability distribution induces a probability distribution over the states of typical rocks sitting on tables--a distribution that counts it as unlikely that small differences in forces would affect the rough behavior of those rocks

Putting the two factors together, we can conclude that differences in distant matters of fact are unlikely to make a difference to the macroscopic behavior of your rock. So when the folk model says that the behavior of your rock depends only on the nature of the rock, and on the strong interactions that come the rock's way (e.g., shaking of the table), it doesn't go too far wrong. Furthermore, what goes for your rock goes for many ordinary objects.

2. Default behavior, and the importance of isolation

The folk model ascribes default--or 'inertial'--behaviors to many systems. And the model says how such systems are disposed to deviate from their default behaviors, if they encounter interference (Maudlin 2004). For example, the default behavior of the rock on your desk is to just sit there and do nothing. And your rock is disposed to slide along the desk if pushed.

Given this framework, one can partially represent the causal structure of a situation with a graph. Each node in the graph represents a system, and arrows represent interactions in which one system perturbs another from its default behavior.¹¹

However, the framework is useless unless systems can to a large degree be treated as isolated from their surroundings. For consider a system whose rough behavior is sensitive to a wide range of variations all over the place. That system won't have a single behavior that is stable enough to be usefully treated as a default. Furthermore, no manageable list of deterministic dispositions will capture interesting regularities about the behavior of the

system. And if one insisted on representing such a system in the sort of causal graph described above, the system would require so many incoming arrows that the graph would be useless.

In contrast, the folk model is useful to us because so many systems can be treated as isolated from so much of their environments. As a result, the generalizations that figure in the folk model are fairly simple, and the associated causal graphs (of ordinary situations) are fairly sparse (Woodward this volume).

So it really is crucial to the applicability and success of the folk model that many systems can be treated as isolated from much of their environments. The same is true of many special sciences. For example, consider the second law of thermodynamics, according to which closed (isolated) systems never decrease in entropy. Strictly speaking, the law *never* applies to reasonable-sized systems, since long-range gravitational effects ensure that such systems are never completely isolated. The law has practical applications only because many systems (e.g., gasses in sealed canisters) can be treated for many purposes as if they were isolated.

3. Sensitive systems

The rock sitting on your desk can be treated as isolated from much of its environment. But not all ordinary systems can be treated as isolated in this way. For example, some rocks may be precariously balanced; others may be a hairsbreadth away from cracking in half. In other words, some systems are *sensitive*: their rough behavior depends sensitively on small

differences in forces--and hence on distant matters, even when the dependence is not mediated by strong nearby interactions. How well does the folk model fit such systems?

In general, of course, the folk model needn't fit such systems very well at all. But here on Earth, many sensitive systems are either **detector-like** or **quasi-chancy**. And the folk model accommodates detector-like and quasi-chancy systems quite well. I will explain each category in turn.

3.1 Detector-like systems

A system is *detector-like* if it is sensitive to distant influences, but only those of a very particular kind (or of a small number of kinds). Examples: devices that measure tiny seismic vibrations, fancy light detectors, cosmic ray detectors, and spy devices that eavesdrop on distant rooms by bouncing lasers off window panes. The rough behavior of such systems is sensitive to distant influences. And such influences needn't be mediated by a strong interaction. For example, a good light detector can register the presence of a single photon. But such systems are not sensitive to just *any* old distant influence. Seismic detectors only are sensitive to vibrations in the ground, light detectors to light (and usually just light coming from a particular direction), and so on.

Detector-like systems can easily be incorporated into the folk model. For though they are sensitive to distant influences, they are sensitive to only very particular distant influences. So they can be treated as isolated, *excepting* the particular influences that they detect. Think of it this way. The folk model is useful because so many objects can be treated as isolated from so much. Insensitive systems (such as rocks on tables) are insensitive to distant

influences, and so are well described by the folk model. But detector-like systems are *almost completely* insensitive to distant influences, since they are sensitive to such a narrow range of distant influences. As a result, they too are well described by the folk model.

So much for detector-like systems. Before turning to quasi-chancy systems, however, we will need some background on the nature of statistical explanation.

3.2 Statistical explanation

Some physical processes are downright ruled out by fundamental dynamical laws. For example, classical mechanics downright rules out a process in which a motionless, isolated particle suddenly accelerates though it is subject to no force.

That is one way for fundamental laws to explain why a particular sort of process does not occur. But it is not the only way. For example, consider a process in which an unsuspected boulder hovers in midair. Such a process is perfectly compatible with the fundamental dynamical laws (all that is needed is a sufficient imbalance between the number of air molecules hitting the bottom of the rock, and the number hitting the top). But nevertheless our laws make such a process exceedingly *unlikely* (Price 1996). Perhaps the laws are indeterministic, and they ascribe a very low chance to such a process. Or perhaps the laws are deterministic, and only a very small range of lawful initial conditions lead to the occurrence of such a process. In the former case, the explanation appeals to the objective chances that figure in the indeterministic laws. In the latter case, the explanation appeals to an objective probability distribution over initial conditions of the universe (see Loewer 2004 and Albert 2001). But either way, the explanation depends on an objective probability distribution

over lawful histories. That distribution determines which sorts of histories the laws count as likely or typical, and which sorts the laws count as unlikely or anomalous.¹² The following discussion will appeal to such objective distributions.

3.3 Extreme quasi-chancy systems

Some systems are sensitive to a great range of distant influences. An extreme example of such a system is the *Brownian amplifier*: a device that includes a tiny speck of dust haphazardly floating in a sealed glass container. The amplifier makes a sound every ten minutes: a whistle if the speck's last fluctuation was to the left, and a beep if it was to the right.¹³ Even under the assumption of determinism, the Brownian amplifier is *incredibly* sensitive to distant influences. For example, consider a tiny change in the amplifier's distant environment: displacing a one-pound moon rock by one foot. Very shortly, the pattern of sounds produced by the amplifier will be completely different in the original and the displaced-rock scenarios.

It seems therefore that the amplifier *cannot* be treated as isolated (or even as almost isolated). For its behavior--even its rough macroscopic behavior--depends on the state of pretty much every chunk of matter for miles around. But there is a trick. Think of the amplifier as having a chancy disposition: the disposition to beep-with-chance-50%-and-whistle-with-chance-50%. That *chancy* disposition is *stable* with respect to distant influences. In other words, if you treat the amplifier as if it were a chancy device, faraway goings-on will not affect the *chances* you should ascribe to it.

The fundamental laws license your treating the amplifier in this way. Here is why. Restrict attention to creatures with powers of observation and control rather like ours. The

fundamental laws make it extremely unlikely that such creatures can do better in predicting or controlling the amplifier's behavior than to treat it as a device that produces a random sequence of beeps and whistles.

To evaluate this claim, consider a gambler who repeatedly places bets (at fair odds) on what sound the amplifier will make next.¹⁴ If the gambler can do better than to treat the amplifier as a 50/50 chance device, then she'll likely win money over a long sequence of bets. She can do that only by following an appropriate rule. For example, suppose that the gambler thinks that the machine is very likely to produce the same sound as it last produced. Then she will follow the rule 'Bet on beep whenever the last sound was a beep'.¹⁵ Or suppose that she thinks that the machine is somehow coupled to the value of the Euro. Then she might follow the rule 'Bet on beep whenever the Euro just increased in value with respect to the Dollar'.

The fundamental physical laws make it extremely unlikely that any such rule would allow the gambler to cash in. That is the sense in which the laws make it unlikely that we can do better than treat the amplifier as a 50/50 chance device, in predicting its behavior.¹⁶

Here is why the laws make it unlikely that the gambler cashes in. For the purposes of this discussion we may assume that the laws are deterministic.¹⁷ Now consider rules of the form 'Bet on beep whenever condition *C* holds', where condition *C* is one that the gambler is capable of detecting before placing her bet. The laws make it very likely that, of the sounds that Brownian amplifiers make after condition *C* holds, about half are beeps. For example, the laws make it likely that, of the sounds that Brownian amplifiers make immediately after the Euro has risen with respect to the dollar, close to half are beeps.

And why is *that*? Why is it that creatures like us are not capable of sensing conditions that distinguish upcoming beeps from whistles? After all, we have assumed determinism, and so we have assumed that such conditions exist. (One such condition is a gigantic particle-by-particle specification of every state of the world that leads to the machine beeping next.) The reason is that creatures like us cannot detect such complicated conditions. We can detect relatively simple macroscopic conditions, such as ‘the rock is on the left side of the desk’. With the help of special apparatuses, we can even detect a certain very limited range of conditions concerning microscopic matters. But remember that the Brownian detector is sensitive to the position of nearly every chunk of matter for miles around. It is also sensitive to the detailed trajectories of the air molecules in the chamber that holds the dust speck. A condition that managed to pick out upcoming beeps would have to put detailed, horrendously complicated constraints on all of these matters, and more. Creatures like us have no hope of detecting such conditions.

There is one loose end. Where did the 50% come from? Why is it that for any condition we can detect, that condition is followed by beeps approximately 50% of the time? Why not 10%? Or no stable percentage at all? Strevens (1998) has offered a beautiful answer in the tradition of Poincaré’s (1905) ‘method of arbitrary functions’.¹⁸ Here is the basic idea.

Suppose that you must choose a color scheme for a black-and-white dart board.¹⁹ Your goal is to guarantee that close to half of the darts thrown at the board land in a black region. The thing for you to do is to choose a scheme that (1) alternates very rapidly between black and white, and (2) is such that in any smallish square region, about half of the region is colored black. One such scheme is an exceedingly fine checkerboard pattern.

Such a scheme is a good idea because the distribution of landing-places on the board is likely to be relatively smooth.²⁰ For example, it is unlikely that any player will be accurate enough to cluster all of her throws in a single square millimeter of the board. As a result, conditions (1) and (2) make it extremely likely that about half of the darts will land on black. In other words, the board's color scheme allows a weak *qualitative condition* on the distribution of landing-places (that the distribution is smooth) to provide a near-guarantee that about half of the darts land on black (Strevens 1998: 240).

The dynamics of the Brownian amplifier accomplish an analogous trick. Any smooth probability distribution over the state of the amplifier's environment makes it very likely that the amplifier will beep about as often as it whistles. Indeed, something stronger is true. Think back to the gambler who suspects that an increase in the strength of the Euro indicates an upcoming beep. Any smooth probability distribution over the state of the amplifier's environment makes the following very likely: of the sounds that the amplifier makes immediately after the Euro has gotten stronger, close to half are beeps. In other words, a gambler who guides her betting by the condition of the Euro is unlikely to cash in. And the same is true for any other condition simple enough for creatures like us to detect. That is why, no matter what (relatively simple) condition a gambler uses to select her bets, the laws count it as very unlikely that she will cash in by betting on the amplifier. In other words, the laws count it as very likely that creatures like us (who are trying to predict the sounds that the amplifier will make) can do no better than to treat the amplifier as if it were a 50/50 chance device.

So the folk model *can* apply to the amplifier. The amplifier is *quasi-chancy*: it has stable dispositions to *act as a certain sort of chance device*. Understood in this way, it is reasonable to treat the amplifier as isolated from distant influences.

The same is true of many other systems that are sensitive to a wide variety of distant influences. By the time a system is sensitive to a wide enough range of distant influences that it no longer counts as Detector-like, it very often ends up *so* sensitive that we can do no better than treat it as if it were genuinely chancy. The folk model can accommodate such quasi-chancy systems by ascribing to them stable dispositions to produce particular chance distributions.

3.4 Less extreme quasi-chancy systems

The Brownian amplifier is the most extreme sort of quasi-chancy system, since creatures like us absolutely cannot do better than treat it as a 50/50 chance device. It is worth considering systems that are less extreme in this respect. Consider, for example, roulette wheels in casinos in the 1970s. One might think that humans cannot do better than treat such wheels as roughly uniform chance devices. But it turns out that measurable conditions (the initial velocities of the wheel and the ball) yield significant information about what quadrant the ball will land in. Indeed, gamblers have attempted to use such information to exploit Las Vegas casinos (see Bass (1985), as cited in Engel (1992: 96)).

So unlike the Brownian amplifier, people *can* do better than to treat 1970s roulette wheels as uniform chance devices. Nevertheless, doing better requires the sort of information and knowledge of detailed dynamics that few people possess. So for most purposes, it is still a

good approximation to think of the wheels as if they were uniform chance devices. The same is true for many ordinary systems. Such systems are not so sensitive as to *utterly rule out* that creatures like us could do better than treat them as chance devices. But they are sensitive enough that they are indistinguishable from chance devices by people with the sort of information ordinarily available. Such systems include precariously balanced rocks, leaves fluttering to the ground, light bulbs that are poised to burn out, crash-prone computers, and perhaps even the ping-pong-ball devices used to choose lottery numbers.

There is a tradeoff between simplicity and generality in whether to treat such systems as chancy. On the one hand, one can treat them as deterministic, in which case they will count as sensitive to a wide range of factors, and as having quite complicated dispositions. The nodes that represent the systems in causal graphs will have many incoming arrows. In representing the systems this way, one gains generality at the cost of complication. On the other hand, treating such systems as if they were chancy will simplify matters greatly, since they will count as isolated from much more of their environments. Such representations gain simplicity at the cost of accuracy and generality.

4 Conclusion

Folk models of everyday situations are enormously useful. What makes them useful is that so many ordinary objects can be treated as isolated from so much of their environments. As a result, we can often ascribe to objects salient default behaviors, from which they may be perturbed by interactions of only very particular kinds.

An important reason we can treat so many systems as isolated in this way is a combination of *insensitivity*, and *supersensitivity*. Some systems are insensitive to small differences in forces and initial conditions. In combination with an appropriate statistical assumption, this licenses us to treat such systems as isolated. Other systems are very sensitive to differences in forces and initial conditions--but many such systems are *so* sensitive that we do better to treat them as chancy devices. By ascribing chancy dispositions to such systems, we can again treat them as mostly isolated. Again, this is licensed by an appropriate statistical assumption.

In attempting to reconcile the folk model of causation with fundamental physical laws, Russell focused on dynamical laws. Little wonder, then, that he thought a reconciliation was impossible. For as we have seen, the dynamical laws do *not* on their own underwrite the usefulness or approximate correctness of the folk model. They do so only in conjunction with statistical assumptions: either probabilistic laws, or laws that supply a probability distribution over initial conditions.²³

References

- Albert, D. Z. (2001). *Time and Chance*. Boston: Harvard University Press.
- Bass, T. A. (1985). *The Eudaemonic Pie*. Boston: Houghton Mifflin.
- Diaconis, P. and Engel, E. (1986). 'Some Statistical Applications of Poisson's work'. *Statistical Science*, 1(2):171–174.
- Engel, E. (1992). *A Road to Randomness in Physical Systems*. Berlin: Springer.
- Field, H. (2003) 'Causation in a physical world', in M. Loux and D. Zimmerman (Ed) *Oxford Handbook of Metaphysics*. Oxford: Oxford University Press.
- Keller, J. B. (1986). 'The Probability of Heads'. *American Mathematical Monthly*, 93(3): 191–97.
- Lieb, E. (1976). 'The Stability of Matter'. *Review of Modern Physics*, 48: 553–569.
- Lieb, E. (1990). 'The stability of matter: From atoms to stars'. *Bulletin of the American Mathematical Society*, 22: 1–49.
- Loewer, B. (2004). 'David Lewis' Humean Theory of Objective Chance'. *Philosophy of Science*, 71 (5): 1115-1125.
- Maudlin, T. W. (2004). 'Causes, Counterfactuals and the Third Factor'. In J. Collins, N. Hall, and L. Paul (Eds). *Causes and Counterfactuals*. Oxford: Oxford University Press.
- Norton, J. D. (this volume). 'Causation as Folk Science'.
- Poincaré, H. (1905). *Science and Hypothesis*. New York: Dover.
- Price, H. (1996). *Time's Arrow and Archimedes' Point*. New York: Oxford University Press.

Russell, B. (1913). 'On the Notion of Cause'. *Proceedings of the Aristotelian Society*, 13: 1–26.

Strevens, M. (1998). 'Inferring probabilities from symmetries'. *Nous*, 32: 231–46.

Strevens, M. (2003). *Understanding Complexity through Probability*. Cambridge: Harvard University Press.

van Lambalgen, M. (1987). 'Von Mises' Definition of Random Sequences Reconsidered'. *Journal of Symbolic Logic*, 52(3): 725–755.

Woodward, J. (this volume). 'Causation with a Human Face'.

Notes

¹For helpful discussion and correspondence, thanks to the Corridor Group, Sheldon Goldstein, Chris Hitchcock, and Jim Woodward.

²Here for convenience I speak as if rates of changes of physical quantities at a time count as part of the state of the world at that time.

³Here I have assumed determinism, but a similar point holds under indeterministic laws.

⁴Compare Norton (this volume: section 4.4): ‘Our deeper sciences must have quite particular properties so that [entities figuring in a superseded theory] are generated in the reduction relations.’

⁵For a particularly accessible introduction to quantum-mechanical derivations of the stability of matter, see Lieb (1990). For more technical treatments, see Lieb (1976) and the references therein.

⁶More carefully: suppose that the physical state at any point of spacetime is nomically determined by the state on time-slice of the back light cone of that point.

⁷Here the relevant ‘now’ is, say, the one determined by your rest frame.

⁸You have that same assurance even if the laws turn out to violate relativistic locality (due to quantum-mechanical entanglement, for example).

⁹I am indebted in the following two paragraphs to helpful correspondence with Jim Woodward.

¹⁰The same cannot be said of larger objects. For example the tides depend on differing gravitational forces (due to the moon) at different places in the Earth. So when modeling the

tides, we are not free to think of the Earth as being in a near-constant gravitational field.

Thanks here to Frank Arntzenius.

¹¹Such graphs are similar to the ‘interaction diagrams’ from Maudlin (2004: 439).

¹²Of course, such explanations are worthwhile only when the histories in question are grouped into relatively natural categories. For example, the laws will count *any* single history—specified in microscopic detail—as extremely unlikely. But that fact doesn’t make every history anomalous. Thanks to Roger White for raising this objection.

¹³The Brownian amplifier is a variant of a device described in Albert (2001).

¹⁴Assume that the gambler must bet several minutes in advance of the sound in question.

¹⁵To ‘bet on beep’ is to bet that the next sound the machine will make will be a beep.

¹⁶The connection between randomness and ‘invariant frequencies under admissible place selections’ is inspired by Von Mises’ definition of an infinite random sequence (see van Lambalgen 1987).

¹⁷The arguments carry over in a straightforward way under indeterminism.

¹⁸See also Keller (1986), Diaconis and Engel (1986), and Engel (1992).

¹⁹This example is adapted from Diaconis and Engel (1986).

²⁰There are three subtleties here. *First*: the notion of smoothness employed here is *not* the technical notion of being continuous and infinitely differentiable. Instead, in the present case it is that the probability density for different dart locations varies slowly on the length scale set by the fineness of the checkerboard pattern. (An appropriately generalized notion of smoothness applies to other cases. Throughout this paper, it is this notion of smoothness that I employ.) *Second*: the smoothness of a distribution depends on how the space of outcomes is

parameterized. In the present case, the relevant parameterization is the natural one, in terms of distances on the board as measured in standard units (Strevens 1998: 241). *Third*: since the conclusion concerns a whole sequence of throws, the relevant distribution is over *sequences* of landing places. That way, the resulting notion of smoothness rules out bizarre dependencies between the throws. For example, it rules out a distribution according to which the first throw is uniformly distributed, but subsequent throws are guaranteed to land in the same spot as the first one.

²³ Compare to Field (2003): ‘the notion of causation, like the notions of temperature and entropy, derives its value from contexts where statistical regularities not necessitated by the underlying [dynamical] physical laws are important.’