

Bell Correlations and Selection Bias

Huw Price

Trinity College, Cambridge

Abstract: Selection artefacts are common in science. A method of selecting samples from a larger population may produce bias, in either direction. It may induce correlations between variables independent in the full population, or mask correlations between variables dependent in the full population. Here we propose a surprising application of these familiar ideas. We argue that they are relevant to puzzling correlations uncovered in quantum theory by John Stewart Bell (Bell 1964). In the light of Bell's work and subsequent experiments it is widely believed that the quantum world is 'nonlocal', in apparent tension with relativity. Many hold that the only alternative is to abandon 'realism', the view that there is an objective world independent of measurement. We propose instead that Bell correlations are selection artefacts, in tension neither with relativity nor realism.

1. Introduction

Selection artefacts are common in science. Methods for selecting samples from larger populations may produce bias in either direction: inducing correlations between variables independent in a full population, or masking correlations between variables dependent in a full population. Missing such bias has practical consequences, where science guides policy.

The case discussed here is a long way from policy, but of unusual theoretical interest. Some of the most puzzling correlations in modern science are those uncovered by the physicist John Stewart Bell in the 1960s (Bell 1964). Bell argued that if the predictions of quantum mechanics (QM) are correct, then there are correlations in the quantum world that cannot be explained by a *local* theory – that is, by a theory in which causal influences propagate at finite speed, as required by special relativity. The relevant predictions have since been confirmed in many experiments. The 2022 Nobel Prize in Physics was awarded to Alain Aspect, John Clauser, and Anton Zeilinger, for work of this kind.

Do these experiments prove that the world is *nonlocal* in Bell's sense? It is often said that the only alternative is to reject *realism* – roughly, the view that the world exists independently of our measurement choices. The experiments in question are described as disproving *local realism* – as showing that either locality or realism must be abandoned.¹

Here we propose an alternative.² We show that the correlations in these experiments may be regarded as selection artefacts. As already noted, selection bias often produces correlations between variables that are actually independent. It turns out that Bell correlations fit this model.

¹ (Shalm et al 2015) claims 'a loophole-free violation of local realism', for example.

² More accurately, a *new* alternative. For discussion of other proposals see (Myrvold et al 2024).




Figure 1: Survivorship bias.³

As we will explain, there are two ways to describe this conclusion. One is to say that the proposal *avoids* nonlocality, preserving locality without sacrificing realism. The other is to say that Bell nonlocality *itself* turns out to be a selection artefact. This is a terminological issue, depending on what we choose to mean by ‘nonlocality’ when the new option is in view. We will return to this question below (§4.1), referring to Bell’s own discussion of the term.

Either way, the upshot is that Bell correlations are much less puzzling than they have seemed, in tension neither with relativity nor realism. We stress that ‘less puzzling’ does not imply ‘not puzzling at all’. The proposal identifies an escape route, but not the underlying mechanisms that enable QM to take advantage of it. Again, more on this below (§4.3).

We begin with a discussion of selection bias in general. Much of this material will be familiar to many readers, but it will be helpful to explain how we use various terms. We will also introduce new terminology, at a couple of points.

2. What is selection bias?

Readers may recognise Figure 1. It represents the distribution of damage found on bombers, returning from missions in WWII. (Think of the red dots as bullet holes.) Clearly, there is a strong correlation between location on the aircraft and the probability of damage at that location. In some parts of the aircraft the probability is high, in other places close to zero.

We can imagine ways in which this correlation might be explained via causal dependencies. There are three main possibilities.

1. The geometry of the aircraft might be influencing the trajectory of the projectiles causing the damage. Perhaps the enemy gunners reliably aim for those places, or perhaps the aircraft have force fields, steering the bullets to those points and not others.

³ Image by Granjean and McGeddon from https://en.wikipedia.org/wiki/Survivorship_bias

2. The projectiles might be influencing the location of the aircraft, the latter shifting itself to ensure that the damage is confined to these areas. (Imagine Wonder Woman dodging bullets.)
3. Some third entity – a *common cause* – might be influencing the location of both the aircraft and the projectiles, to ensure that the impacts have this distribution. (Imagine an embroiderer, controlling a needle with one hand and the position of the piece of work with the other.)

As we know, none of these explanations is correct. The right explanation, as the Columbia statistician Abraham Wald apparently pointed out,⁴ is that the correlation resulted from the fact that we are looking at a biased subset of the data: namely, those aircraft that made it back to base. Aircraft with damage in other locations did not survive, and that's what's skewing the data, producing the correlation.

This is an example of selection bias. This particular variety is called *survivorship bias*, for the obvious reason: the selection variable is 'survival'. Survivorship bias is a subspecies of *collider bias*. In the terminology of causal models, a collider is a variable influenced by two or more causes – a node where two or more arrows 'collide', in the graphical notation of direct acyclic graphs (DAGS; see Figure 2). Conditioning on a collider – that is, selecting cases in which the collider variable takes a particular value – typically produces a correlation between the contributing causes, even if they are in fact independent. (Again, think of the bomber example.)

We propose that Bell correlations are similar to collider bias – indeed, they, or some of them, may actually *be* collider bias. But it will be helpful to start further back, stressing first that not all selection bias is collider bias.

2.1 Selection bias without colliders

Why is selection bias not always collider bias? For two reasons, important in different ways. First, there are non-physical cases of selection bias, e.g., in mathematics. Imagine a study of the correlation between the parity (oddness or evenness) of successive digits of the decimal expansions of two irrational numbers, say π and e . Let π_n and e_n be the n th digits in these expansions, respectively. For each n , there are four possibilities: the two digits π_n and e_n may be even–even, even–odd, odd–even, or odd–odd. Let's assume that these cases are equally likely. This means that there's nothing that the parity of π_n tells us about that of e_n and vice versa. (If that's not true about π and e , we can choose different numbers.) Now let \mathbf{S} be the set of natural numbers such that π_n and e_n are not both odd. Within \mathbf{S} , π_n and e_n are highly correlated: if one is odd then the other is even. \mathbf{S} behaves like a collider variable, but it is a matter of logic, with no causation in the picture.

⁴ The story may be apocryphal: <https://www.ams.org/publicoutreach/feature-column/fc-2016-06>.




Figure 2: A collider




Figure 3: A common cause

The second reason does apply in the familiar, causal world. A well-known property of common causes (see Figure 3) is that their effect on correlations is the opposite of a collider. The correlations they produce *disappear* if we condition on the common cause – if we select cases in which the common cause variable takes one particular value. The explanation is simple. If K co-varies with L because they both co-vary with a common cause J, then if we remove the variation of J, holding it fixed, we remove the co-variation between K and L.

Here's an example. Suppose that K and L are rare diseases of mice. Both diseases are genetic disorders, caused by different abnormalities of the genes J that code for white hair colour. Few white mice get these rare diseases, but no grey mice do, because they lack the relevant genes.

Imagine we don't know these facts, and we're interested in whether diseases K and L are correlated. We examine large numbers of mice, looking for one or both conditions. If our lab breeds only white mice we will miss the correlation, because we are holding fixed the common cause J. This is selection bias, but here it is masking a correlation, not producing one.

This variety of selection bias is called *range restriction* (Dahlke & Wiernik 2019). Unlike collider bias, it comes in two forms, depending on when and how the range is restricted. As described, we could call it *preselection bias*. The sample population of mice is produced with the restriction already built in. We could imagine the same error as *postselection bias*. Suppose now that the lab breeds mice of all the available colours, but that our assistant, Charlie, favours

the white ones.⁵ Believing that we are kinder to mice than are other scientists in the lab, Charlie always selects white mice for our study. This masks the correlation between K and L in just the same way.

2.2 The Minimal Selection Bias Criterion (MSBC)

With these examples to guide us, let's formulate a test for selection bias. What does such bias necessarily involve? At a minimum, a distinction between a larger population of cases and a smaller sub-population, the latter picked out of the former by some sort of selection process. Let's call these populations the *super-ensemble* and the *subensemble*, respectively. We get selection bias if there are correlations that *differ* between the super-ensemble and the subensemble. Call this the *Minimal Selection Bias Criterion* (MSBC).

We have seen that the difference in correlations required by MSBC may go in either direction. Selection may *remove* as well as *produce* correlations. We have also seen that a super-ensemble need not be an ensemble of *actual* cases. When preselection is involved, the selection may be one that prevents some of the super-ensemble from coming into existence in the first place. In such cases, much of the super-ensemble consists of possible but unrealised cases – it is partly a *virtual* super-ensemble, as we might say. This will be important in the QM case: ordinary preparation of the initial state of a Bell experiment is a preselection of this kind.

2.3 Correlating and decorrelating forks

We noted an important difference between colliders and common causes. When we fix a common cause (e.g., the variable J in Figure 3), we *remove* or '*screen off*' the correlation between its two effects (K and L). Colliders work the other way. Fixing a collider variable (F in Figure 2) typically *produces* correlations between the other two variables (G and H).

It will be helpful to have descriptive labels for these two kinds of cases. The philosopher Hans Reichenbach called the common cause case (Figure 3) a *conjunctive fork* (Reichenbach 1956). We will call it a *Decorrelating Fork*, to make its effect on correlations explicit. We will call the collider case a *Correlating Fork*, and refer to the variable at the vertex (F in Figure 2) as a *Correlator*. The immediate advantage of this notation is that it is general enough to apply to non-causal cases, such as our mathematical example. The set **S** acts as a Correlator with respect to the parity of π_n and e_n .

⁵ As Charlie puts it: 'I dint know that mice were so smart. Maybe thats because Algernon is a white mouse. Maybe white mice are smarter than other mice.' Daniel Keyes, *Flowers for Algernon* (Keyes 1959), p. 7.




Figure 4: A two-particle Bell experiment

In the familiar macroscopic world, there is a striking time-asymmetry between Correlating Forks and Decorrelating Forks. Compare Figures 2 and 3. Reflecting the time-asymmetry of causation itself, Correlating Forks have their central vertex in the future, while Decorrelating Forks have it in the past. As we shall see, the present proposal rests on the fact that this time-asymmetry no longer holds in the quantum realm. There, Correlating Forks with the opposite temporal orientation are also common, and are the basis of Bell correlations. The core of the proposal is that the entangled states prepared at the beginning of ordinary two-particle Bell experiments are themselves Correlators. Bell correlations are the result of holding such a variable fixed, as we do when we make repeated runs of an experiment with the same initial state. We now turn to the details.

3. Bell experiments

3.1 Introducing Factorizability

Figure 4 depicts a typical two-particle Bell experiment. Pairs of entangled particles are prepared at **C** in an initial state C . In one common version of the experiment, the particles are photons, prepared in the so-called *singlet state*. The particles are directed to experimenters at **A** and **B**, who each choose a setting for their measurement device, and record an outcome. In the photon version, the measurement devices are polarizers.

The settings and outcomes both take two possible values, which for convenience we represent as 0 and 1. We use a and b as variables for the settings at **A** and **B**, respectively, and A and B for the outcomes. QM allows us to calculate the joint probabilities of the four possible combinations of outcomes, given the initial state C and the four possible combinations of settings. We write these probabilities as $P(A,B|a,b,C)$.

The intuitive idea of Bell's condition of locality is that when the experiments at **A** and **B** are widely-separated – in particular, if they are spacelike separated in the sense of special relativity – then the probabilities at **A** and **B** should be independent, for a given *C*. The probability for an outcome at **A** should not depend on the setting or outcome at **B**, and vice versa. If this is true, then the probabilities are said to be *Factorizable*:

(Factorizability)
$$P(A,B|a,b,C) = P(A|a,C)P(B|b,C)$$

Bell showed that Factorizability implies a numerical constraint on the expected values of such experiments, a kind of constraint now known generically as a *Bell inequality*. Bell realized that QM predicts that such constraints are violated in some cases, predictions now confirmed in many experiments. In most cases, the actual Bell inequality used is the so-called Clauser-Horne-Shimony-Holt (CHSH) Inequality, derived by these authors in a refinement of Bell's argument (Clauser et al 1969).

Experimental violations of Bell inequalities are the basis for the claim that Factorizability fails in QM, and hence that the quantum world is nonlocal. However, we will now argue that failure of Factorizability is a selection artefact, requiring no spacelike dependency of any other sort. We will approach this conclusion in several steps.

3.2 Bell correlations by postselection (I): classical toy models

Starting gently, let's observe that postselection makes it easy to produce Bell correlations in a *classical* toy model. Suppose that two experimenters Alice and Bob each choose a 'setting' bit, 0 or 1, and toss a coin to generate an 'outcome' bit. They pass the resulting ordered pairs to Charlie, giving him an ordered binary 4-tuple (a,A,b,B) . Charlie discards some of these 4-tuples and retains others, using a probabilistic algorithm to make the cut. The probability that his algorithm will retain a particular result (a,A,b,B) is set to $P(A,B|a,b)$, where the latter is the probability predicted by QM in a standard Bell test as in Figure 4.

This procedure ensures that in the long run, the correlations in Charlie's ensemble of retained results match those predicted in the real Bell experiment. But they are merely collider bias. The probability that a result will be retained ('survive') depends both on Alice's two bits and Bob's two bits, so we have a collider – and in our new terminology, a Correlator. The correlations emerge when we hold fixed this Correlator, ignoring the discarded cases.

It is easy to see that Charlie could produce much stronger correlations, e.g., by discarding everything except perfect matches. It would still be selection bias, requiring no causal influence between Alice and Bob. Note also that Alice and Bob can be spacelike separated. As selection artefacts, the correlations between them pose no challenge to special relativity.




Figure 5: Bell experiment with four possible initial Bell states.

Let's make this classical toy model a little more sophisticated. The real Bell experiment we had in mind can be performed with any of the four so-called Bell states as the initial preparation. Let's write these four states as C_0, C_1, C_2, C_3 , where C_0 is the singlet state. (See Figure 5.) For each of these states, there is a probability distribution $P(A,B|a,b,C_i)$ giving the predicted joint probability of outcomes (A,B) given settings (a,b) and initial state C_i .

Charlie can use these probabilities to sort results from Alice and Bob into four hoppers H_0, H_1, H_2, H_3 , such that (i) no results are discarded and (ii) the frequencies in each hopper match the predicted results of a Bell test with the corresponding initial state. One way to see that this is possible is to imagine the real Bell experiment performed with the initial state C_i chosen at random, with equal probability for the four options. The inverse probabilities $P(C_i|a,b,A,B)$ are the probabilities needed by Charlie's algorithm, for sorting a result into hopper H_i . The results within each hopper exhibit Bell correlations, as before, but again, these are selection artefacts.

3.3 Bell correlations by preselection: a quantum toy model

Leaving the classical model behind, consider the QM model just described – a Bell experiment with the initial Bell state chosen at random on each run, with equal probability for each of the four possibilities. This device generates an ensemble of results in which (i) there are no Bell correlations in the ensemble as a whole, but (ii) there are such correlations in each of the four sub-ensembles defined by the initial state C_i . These facts are what guaranteed that the same would be true in the four-hopper version of our classical toy model.

On the face of it, then, we have met the MSBC test for treating *some* Bell correlations as selection artefacts, in a real QM case. In a case in which selection produces correlations, MSBC calls for two things: an uncorrelated super-ensemble, and a selection procedure that picks out correlated subensembles of this super-ensemble. In the case just described, this selection is performed by the random device that controls the initial state C_i . Bell correlations

emerge in the subensembles defined by a fixed value of C_i . The initial state C_i is therefore a Correlator, in the sense defined above.

Thus we have a toy model using real QM components, in which it is easy to see how Bell correlations can be represented as selection artefacts. If we could argue that the toy model is actually *general* – a good representation of *arbitrary* Bell tests with these components – then we would have something interesting. We'd have an argument that such Bell correlations are selection artefacts.

For the moment, though, we just have the toy model. In this model, the uncorrelated super-ensemble is constructed artificially, by an imposed random choice of the initial Bell state. Why not regard this as gerrymandering, rather than a construction with application in the real world?

3.4 Experimental control and the difference between past and future

We are proposing that in the kind of experiment depicted in Figure 5, Bell correlations arise because we fix a Correlator at \mathbf{C} when we prepare the initial state. Putting the choice of the initial state in the hands of a random device reduces our degree of control, but it doesn't remove it altogether. The probabilities of the four possible initial states are still controlled by us – we built the randomiser, after all, and we could have dialed in any distribution we chose of the probabilities of the four possible initial states.

What would the probabilities be if there were *no experimental control whatsoever*? If the four initial states were equally likely in that case, we could take the case with the randomiser to be a model of Nature's own probabilities.

Readers may find this question perplexing. Of course we control the initial states of our experiments, at least in principle! How could it make sense to throw that away? But consider this. Whatever this initial control of experiments amounts to, it is time-asymmetric. We don't control the final states of experiments in the same way. This should give us pause, if we are inclined to insist that initial control is something fundamental, something whose absence doesn't make sense. If that were true, it would amount to a fundamental time-asymmetry.

Yet it is widely believed that fundamental physics, including quantum theory, is time-symmetric.⁶ Observed time-asymmetries are standardly attributed to a boundary condition in the early universe, imposing very low entropy at that point. If this story is correct, then our own existence, and our time-asymmetric ability to control our environments, is also attributable to this low entropy past (no doubt somewhat remotely).

This standard picture provides the resources, in principle, to consider the case without initial control. Whether in classical or quantum versions, the picture involves a time-symmetric

⁶ Or CPT-symmetric, if we want to be careful.

measure of possible histories for the universe, from which our observed history emerges by imposing the low entropy past. All manner of unlikely possibilities may be considered in an imaginary regime lacking the low entropy constraint. Random pieces of quantum optics will be an easy trick, compared to Boltzmann Brains (Carroll 2017).

For present purposes, the only point we need is that this ‘unconstrained’ case, in which the initial state of the Bell experiment is not under experimental control, is not illicit in some deep theoretical sense. The question about the natural measure on the space of unconstrained possibilities at **C** is a respectable one, albeit an unusual one.

From this point, there are two paths our argument might take. First, we could simply *postulate* the measure we want. In other words, we could postulate that the example with a randomiser is an adequate model of the unconstrained case. This postulate comes at very modest cost – why should we assume that the measure is anything else? – and it has a generous payoff. It ensures that the Bell correlations in experiments with the form of Figure 5 are entirely explicable as selection artefacts, emerging from an uncorrelated super-ensemble when we fix the value of the Correlator at **C**. The proposal could then be regarded as an hypothesis, relying on this assumption about the measure.⁷

The second strategy makes a virtue of the ordinary asymmetry between past and future. As we just observed, we normally have control of the initial but not the final conditions of our experiments. We have been interested in the case in which this initial control is imagined to be absent. Could we find a model for that case in a version of our experiment in which the relevant variable sits in the future, rather than the past? In the future it would be uncontrolled anyway, with no act of imagination needed. This thought leads us to a different category of Bell experiments, which also happen to be interesting for other reasons.

3.5 Bell correlations by postselection (II): real-world quantum cases

Figure 6 depicts an important experimental protocol, used for some of the best tests of Bell correlations. We can call it the *W-shaped protocol*, referring to the shape of its arrangement in spacetime. It relies on so-called *entanglement swapping*. Two pairs of entangled particles are produced independently, one pair at the source **S₁** and the other at the source **S₂**. One particle from each pair is sent to a measurement device at **M**, which performs a Bell state measurement on the two particles – a measurement with four possible outcomes, shown in Figure 6 as M_0 , M_1 , M_2 and M_3 .

⁷ Readers may object that the postulate is unnecessary. If the Universe favoured some other measure in the unconstrained case, then the proposal could proceed as before, except that some of the relevant preselection would be being done by the Universe itself. We set this suggestion aside here, but see (Price & Wharton 2025, §6), for relevant discussion.

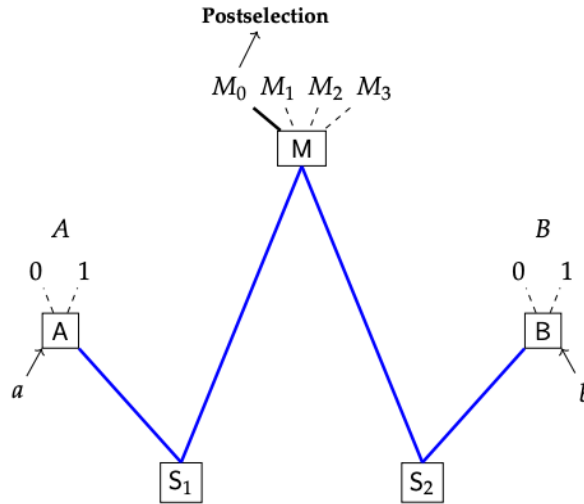


Figure 6: A W-shaped Bell experiment with entanglement swapping

The remaining two particles go to measuring devices at **A** and **B**, respectively. The experimenters each choose a setting and record an outcome, in the usual way. The expected results are as follows. There are no **A–B** Bell correlations in the ensemble of results as a whole, but such correlations emerge if we postselect on any of the four outcomes at **M** – e.g., on M_0 , as shown in Figure 6. As noted, some of the best real Bell experiments use this protocol.

It is easy to see that the case meets our MSBC test for the presence of selection artefacts. The super-ensemble is the set of results as a whole, and the sub-ensembles are the four sets of results we obtain by postselection on a single value of M at **M**. In the sense implied by MSBC, Bell correlations in such experiments simply *are* selection artefacts – there’s no question about the matter. Because the selection is inducing rather than masking correlations, the output variable M is a Correlator, in our new terminology.

Two clarifications. First, the diagnosis just given applies to the Bell correlations ‘across the W’ between **A** and **B**. It says nothing, as it stands, about the entanglement in the two wings of the experiment, needed to produce the results at **M** that in turn yield postselected correlations between **A** and **B**. We set this point aside here, with the following note. If we have a diagnosis that works in ordinary two-particle (‘V-shaped’) Bell experiments, then we could expect to apply it here, to deal with the wings as well. The W experiments would then be hybrid cases, involving selection at all three vertices: preselection at **S**₁ and **S**₂, and postselection at **M**.

Second, our diagnosis has said nothing about a significant feature of the geometry of these W-shaped experiments, a feature that makes a big difference to how they are represented in orthodox QM. The measurement at **M** can have three possible locations with respect to those at **A** and **B**.⁸ It may be in their absolute future – so-called *delayed choice* entanglement swapping (DCES) – in their absolute past, or somewhere in between. In the orthodox QM

⁸ We are ignoring some hybrid possibilities, e.g., cases in which **A** and **B** are not spacelike separated from one another.

picture with collapse, this makes a big difference to the description of the case. In the DCES case the measurements at **A** and **B** collapse the state of the second particle in each pair, influencing the states reaching **M**. In the absolute past case, the influence goes the other way: the measurement at **M** affects the states of the particles received at **A** and **B**.

In the DCES case, as several authors have noted, this means that we have something like a causal story, for how information at **A** and **B** can make a difference to the outcome at **M**. It becomes plausible to regard **M** as a collider, in other words, and hence to treat the resulting **A–B** correlations as a case of collider bias.⁹ But this story won't work, apparently, in the other possible *W* geometries.

However, we stressed earlier that not all Correlators are colliders. If we were trying to force Bell correlations into a box labeled *Collider Bias* then the geometry of the *W* cases would certainly be a concern. But MSBC doesn't mention causal structure, merely statistical structure: namely, a difference in the correlations in the super-ensemble and selected subensembles. By this purely statistical criterion, it seems indisputable that the Bell correlations are selection artefacts, associated with a Correlator at **M**, in all the *W* geometries. We need to postselect on the result at **M**, wherever **M** stands in spacetime with respect to **A** and **B**.

In sum, our model applies easily to the *W* cases – including, as we said, to some of the best recent Bell tests. To take us back to *V* cases, note that QM itself provides a flat distribution over the four possible outcomes at **M** – all four possibilities have equal probability. This seems a promising analogue for the corresponding probabilities at **C**, in the two-particle case we were discussing in §3.4 – the imaginary case in which no distribution of probabilities for inputs at **C** is supplied by the experimenter.

In principle, these experiments could be conducted so that Alice and Bob, confined in their laboratories, do not know whether the experiment as a whole has the *V*-shaped geometry of Figure 5 or the *W*-shaped geometry of Figure 6. The two experiments can be configured so that the joint probability of outcomes (A,B) given settings (a,b) and outcome M_i at **M** in the *W* case, is identical to the joint probability of outcomes (A,B) given settings (a,b) and initial state C_i at **C** in the *V* case – so that $P(A,B|a,b,M_i) = P(A,B|a,b,C_i)$, in other words. In this case, it is as though the vertex at **C** has simply been 'flipped' into the future, keeping the probabilities the same, but changing inputs to outputs.

3.6 Summary

We have proposed that Bell correlations be regarded as selection artefacts, the relevant selection being the choice of the initial state at **C** in the *V*-shaped case of Figure 5, or the

⁹ See (Gaasbeek 2010), (Egg 2013), (Fankhauser 2019), (Bacciagaluppi & Hermens 2021), (Price & Wharton 2021), and (Mjelva 2024).

output state at **M** in the W-shaped case in Figure 6. In either case, this meets the MSBC requirements, an uncorrelated super-ensemble being readily at hand in both cases.¹⁰

In the V-shaped case the super-ensemble may be largely a *virtual* ensemble, in the sense that many of its members are never brought into existence. However, we noted that this is common in cases of range restriction where the source of selection bias is a common cause. (Remember our white mice.) Once we have Correlating Forks with the temporal orientation of Figures 4 and 5, preselection from virtual ensembles should also be expected for Correlators. Just as with the white mice, the artefacts produced by selection in Figure 5 arise equally from preselection or postselection. We can choose to run the experiment so that it is locked to a single initial state at **C**, or we can take what comes, and throw away the cases we don't want. Either way, the resulting Bell correlations meet the MSBC condition to count as selection artefacts.

4. Discussion

4.1 Nonlocality – diagnosis or cure?

Now to a question deferred at the beginning. Should we regard this proposal as a way of *explaining* Bell nonlocality, or as a way of *avoiding it altogether*? This is a subtle matter, for it depends on what we mean by 'locality'. We can explain this point with reference to Bell's own discussion of the term. Bell's most extended treatment of locality is in his late piece 'La nouvelle cuisine' (Bell 1990), written a few months before his untimely death. He begins there with what he calls the *Principle of Local Causality*.

The direct causes (and effects) of events are near by, and even the indirect causes (and effects) are no further away than permitted by the velocity of light. (Bell 1990, 239)

He then says that this 'principle of local causality is not yet sufficiently sharp and clean for mathematics', and follows this immediately with a warning:

Now it is precisely in cleaning up intuitive ideas for mathematics that one is likely to throw out the baby with the bathwater. So the next step should be viewed with the utmost suspicion. (1990, 239)

As we'll see, Bell was right to be cautious, for he himself misses an important baby – it is selection bias.

After this warning, Bell formulates a second principle of local causality, motivated by simple considerations of light cone structure, and the like. It is very close to Factorizability itself, which Bell takes to be an immediate consequence. We don't need the details of the second principle

¹⁰ Does such a super-ensemble need to be unique, where it is partly virtual? Our mice suggest not. The correlation due to the common cause will be visible if we add to our sample grey mice, black mice, or both, for example.

here,¹¹ but it is worth noting that Bell makes a virtue of the shift to talk of probability, rather than cause and effect:

Note, by the way, that our definition of locally causal theories, although motivated by talk of “cause” and “effect”, does not in the end explicitly involve these rather vague notions. (1990, 240)

We’ll return to Bell’s cautiousness about causation below (§4.3).

When Bell introduces Factorizability explicitly, he says that it is ‘consequent on local causality’ (1990, 244), meaning his second version of the latter. For our purposes, it will be sufficient to restate Factorizability as we had it earlier, in a slightly simpler form than Bell’s.

(Factorizability) $P(A,B|a,b,C) = P(A|a,C)P(B|b,C)$ ¹²

About this, Bell says the following (we have changed the notation slightly to match our own):

Now this formulation has a very simple interpretation. It exhibits A and B as having no dependence on one another, nor on the settings of the remote polarizers (b and a respectively), but only on the local polarizers (a and b respectively) and on the past causes, $[C]$... We can clearly refer to correlations which permit such factorization as ‘locally explicable’. Very often such factorizability is taken as the starting point of the analysis. Here we have preferred to see it not as the *formulation* of ‘local causality’, but as a consequence thereof. (Bell 1990, 243)

It should now be clear that we have two options. The first is to take the usual path, *defining* locality as Factorizability. In this case, the present proposal implies that locality fails, because Factorizability fails. But the failure of Factorizability is a selection artefact, rather than an indication of any direct causal influence from **A** to **B**. It is a benign form of nonlocality, in other words, in no tension with relativity or realism.

Once we see this, we may prefer the second option. It is to follow Bell, distinguishing Local Causality and Factorizability. The present proposal then shows that Bell was wrong to conclude that the latter is an automatic consequence of the former. Factorizability fails, though Local Causality does not. Bell was right to be suspicious of the step from one to the other. The baby discarded with the bathwater is the possibility of selection artefacts.

As we said, the choice between these two options is a terminological matter. The need to make this choice should not obscure what the two options have in common. In either version, the

¹¹ For the details, see the discussion in (Myrvold et al 2024), to which we are indebted here.

¹² The main modification we have made to Bell’s formulation is to remove possible ‘hidden variables’ λ from the conditions of all these conditional probabilities. Our present framework is purely operational, and this simplification makes no difference to the conclusions drawn below.

present proposal avoids the seeming tension between Bell correlations, on the one hand, and the conjunction of relativity and realism, on the other.

4.2 Correlators and entanglement

Next, a comment on the relation of the present proposal to the understanding of entanglement. Quantum entanglement was named by Erwin Schrödinger in 1935. Schrödinger was responding to a now-famous paper by Einstein, Podolsky and Rosen (EPR 1935). He notes that in the two-part quantum systems that EPR discuss, two components that have just interacted cannot be described independently in the way that classical physics would have allowed. As he says:

When two separated bodies that each are maximally known come to interact, and then separate again, then such an *entanglement of knowledge* often happens. (Schrödinger 1935a, §10, emphasis added)

Schrödinger stresses the centrality of this point to the new quantum theory.

I would not call that *one* but rather *the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought. (Schrödinger 1935b, 555)

Our proposal offers a diagnosis of this ‘entanglement of knowledge’, at least in the cases of entanglement relevant to Bell experiments. The defining characteristic of collider bias (and more generally, in our new terminology, of Correlator bias) is that it yields correlations between variables that are actually independent, in the relevant super-ensemble. The phrase ‘entanglement of knowledge’ seems a good fit for this more general phenomenon. When a Correlator is known to be fixed, it does induce an epistemic dependence between one arm of the Correlating Fork and the other.

Does this mean that our term Correlator is just a new word for entanglement? That would be an uninteresting proposal, but no, for two reasons. First, Correlator is a more general category, with many other examples both in the physical world (ordinary colliders) and elsewhere (e.g., in mathematics). So it is not mere relabelling to propose that quantum entanglement, or some cases of it, *also* fall into this category.

Second, and more importantly, the diagnosis we offer of Schrödinger’s ‘entanglement of knowledge’ relies on more than just the presence of a variable behaving as a Correlator. It also requires a selection process, such as preparation of an initial state, to pick out a particular value of that variable. This means that we are not simply giving entanglement a new name. We are proposing an explanation *of how it arises*, at least in the cases relevant to Bell correlations. It is a selection artefact, the result of holding fixed a Correlator.

4.3 Whence these quantum Correlators?

Finally, we noted at the beginning that although the proposal offers an escape route from the tension between realism and relativity, it doesn't explain how the quantum world manages to avail itself of this loophole. One way to emphasise this point is to stress that the proposal is purely *operational*. It is based simply on the correlations between inputs and outputs of Bell experiments, and the operational notions of preparation and postselection. It offers a diagnosis of some structure present at this operational level – the Bell correlations themselves – but tells us nothing about how that structure arises. *Why* do the initial states of Bell experiments behave as Correlators? The proposal is silent on this matter. It is also silent on the issue of the *strength* of the correlations. It doesn't tell us why Bell correlations do not permit signalling, for example.

As in many medical cases, silence on these aetiological questions does not call into question the clinical diagnosis, or (in the present case) its relevance for issues of locality. Like the diagnosis itself, the relevance to locality rests on MSBC, whose application has turned out to be straightforward. Varying the metaphor, this means that we can see that the proposal flies, even if we don't know the mechanisms that keep it in the air.

Not everyone in quantum foundations is attracted to the search for mechanisms. For those who are, a natural approach is to invoke the framework of causal models, thought of as an abstract first step on an intended path to something more concrete. We close with some recommendations for such an approach, drawing on the present discussion and some of Bell's own remarks about causation in physics.

Our main proposal suggests a question for any application of causal models to the cases we have discussed. In the familiar, classical, physical world, Correlators are also colliders, with the temporal orientation of Figure 2. We noted that in the DCES version of the W case in Figure 6, the central measurement result at **M** also looks like a collider, in orthodox QM – we have what looks like a causal story, via collapse.¹³ Hence the question: should the other W cases and the V case *also* be modelled as involving colliders, if we are interested in underpinning the operational proposal with a causal model?

If we answer 'Yes' we will be envisaging *retrocausality* in some cases. This idea has often been suggested in QM,¹⁴ and it may be playing a role here. This point connects with Bell's own caution about the use of the term causation in physics, mentioned above. In an early piece he expresses one of his reservations, and his view of its resolution, as follows:

In this matter of causality it is a great inconvenience that the real world is given to us once only. We cannot know what would have happened if something had been different. We cannot repeat an experiment changing just one variable; the hands of the

¹³ In DCES cases the most straightforward causal story treats outcomes as well as settings at **A** and **B** as contributing causes at **M**. We set this aside here, but see (Price 2026, §11) for discussion.

¹⁴ For an overview, see (Friedrich & Evans 2023), (Wharton & Argaman 2020).

clock will have moved, and the moons of Jupiter. Physical theories are more amenable in this respect. We can *calculate* the consequences of changing **free** elements in a theory, **be they only initial conditions**, and so can explore the causal structure of the theory. (Bell 1977, 101; emphasis added)

Bell puts a time-asymmetry in by hand here, with the phrase in bold. He himself doesn't seem to have been inclined to investigate the source of this time-asymmetry. As noted above (§3.4), a standard view associates it with the presence of a low entropy boundary condition, in the early history of the universe.

If we wish to discuss causality in an imaginary regime without such a boundary constraint, we can remove the time-asymmetry that Bell puts in by hand, by deleting the emphasised phrase. We can calculate the consequences of changing other elements, and be liberal about what we count as a free or exogenous element. With these tweaks to Bell's own conception of causality in physics, we get a time-symmetric picture.¹⁵ In this picture, we get colliders where we want them, in all the experimental geometries.

Some readers may feel that the resulting time-symmetric conception of causality is too thin to deserve the name. Why not simply speak of statistical dependence? This is largely a terminological preference, though some accounts of causation are committed to the conclusion that there is no true causality in microphysics.¹⁶ Again, an advantage of our shift from 'collider' to 'Correlator' is that it frees us from taking a stand on these issues, terminological or not.

These points do not give us answers to the questions posed at the beginning of this subsection. They don't give us a *mechanism* for the operational correlations on which our main proposal relies. But they do suggest some guidelines for the search for such mechanisms. First, we should be cautious, as Bell was, about the use of causal language, making sure we understand what we mean by it. Second, we should be *more cautious* than Bell was about relying on intuitive time-asymmetries, no matter how straightforward they may seem. Again, it is widely accepted that the source of observed time-asymmetries lies not in fundamental theory but in a distant low entropy boundary condition. Unless we wish to challenge this orthodoxy – or to accept that QM is not a fundamental theory – we should take care to respect it, in discussing the ontological foundations of quantum phenomena.¹⁷

¹⁵ The Future-Input-Dependent models of (Wharton & Argaman 2020) are a step in this direction, but do not envisage treating outcomes as contributing causes, in the sense mentioned in fn. 13.

¹⁶ David Papineau is a recent example: '[The] affinity with statistical mechanical accounts of thermodynamic behaviour argues that causation is an essentially macroscopic phenomenon, depending on structures which emerge only with respect to coarse-grained aspects of physical systems' (Papineau 2025, 14).

¹⁷ This piece is greatly indebted to Ken Wharton. It builds on (Price & Wharton 2025) and joint work over many years. I am also grateful to Carlo Rovelli, David Papineau, Richard Healey, Michael Hall, Arthur Fine, Simon Friedrich, Jason Grossman, Gerard Ciepielewski, George Musser, Zack Savitsky, and Travis Norsen. I first heard the suggestion that preparation is preselection from Gerard Milburn. It took me some time to understand what he meant, and discussions with Heinrich Päs and his group in Dortmund played a crucial role.

References

Bacciagaluppi, G. & Hermens, R. 2021. Bell inequality violation and relativity of pre- and postselection. arXiv:2002.03935

Bell, J. S. 1977. Free variables and local causality. *Epistemological Letters* 15, 79–84. Reprinted in (Bell 2004), 100–104. Page references are to the latter version.

Bell, J. S. 1990. La nouvelle cuisine. In *Between Science and Technology*, A. Sarlemijn and P. Kroes (eds.), 97–115. Amsterdam: Elsevier. Reprinted in (Bell 2004), 232–248. Page references are to the latter version.

Bell, J. S. 2004. *Speakable and unspeakable in quantum mechanics*, 2nd edn. Cambridge: Cambridge University Press.

Carroll, Sean. 2017. Why Boltzmann Brains are Bad. <https://arxiv.org/abs/1702.00850>

Clauser, J. F., Horne, M. A., Shimony, A., and Holt, R. A. 1969. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23: 880–884.

Dahlke, J. A. and Wiernik, B. M. 2019. Not Restricted to Selection Research: Accounting for Indirect Range Restriction in Organizational Research. *Organizational Research Methods*, 23(4): 717-749. <https://doi.org/10.1177/1094428119859398>

Egg, M. 2013. Delayed-choice experiments and the metaphysics of entanglement. *Foundations of Physics*, 43: 1124–1135.

Einstein, A., Podolsky, B. and Rosen, N. 1935. Can quantum-mechanical description of physical reality be considered complete? *Physical Review* 47: 777–780.

Fankhauser, J. 2019. Taming the delayed choice quantum eraser. *Quanta* 8, 44–56. arXiv:1707.07884

Friederich, Simon and Peter W. Evans. 2023. Retrocausality in Quantum Mechanics. *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/win2023/entries/qm-retrocausality/>>.

Gaasbeek, B. 2010. Demystifying the delayed choice experiments. arXiv:1007.3977

Keyes, Daniel. 1959. Flowers for Algernon. *Fantasy and Science Fiction* 16(4): 5–30.

Mjelva, Jørn. 2024. Delayed-choice entanglement swapping experiments: no evidence for timelike entanglement. *Studies in History and Philosophy of Science* 105: 138–148.

Myrvold, Wayne, Genovese, Marco and Shimony, Abner. 2024. Bell's Theorem. *The Stanford Encyclopedia of Philosophy (Summer 2024 Edition)*, Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/sum2024/entries/bell-theorem/>>.

Papineau, D. 2025. Causal inference and the metaphysics of causation. *Synthese* 206, 204.

Price, Huw. 2026. Retrieving the Baby: Reichenbach's Principle, Bell Locality, and Selection Bias. <https://arxiv.org/abs/2602.16985>

Price, Huw & Wharton, Ken. 2021. Entanglement swapping and action at a distance. *Foundations of Physics* 51: 105. doi.org/10.1007/s10701-021-00511-3

Price, Huw & Wharton, Ken. 2025. Taming entanglement. <https://arxiv.org/abs/2507.15128>.

Schrödinger, E. 1935a. Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften*, 23, 807–812. English translation in (Trimmer 1980).

Schrödinger, E. 1935b. Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society* 31: 555–563.

Shalm, L. K., Meyer-Scott, E., Christensen, B. G., Bierhorst, P., Wayne, M. A., Stevens, M. J., et al. Strong loophole-free test of local realism. *Physical Review Letters* 115(25): 250402

Trimmer, J. D. 1980. The present situation in quantum mechanics: A translation of Schrödinger's 'cat paradox' paper. *Proceedings of the American Philosophical Society*, 124: 3230–3338.

Wharton, K. & Argaman, N. 2020. Bell's Theorem and locally-mediated reformulations of quantum mechanics. *Reviews of Modern Physics* 92, 21002. arXiv:1906.04313