The Ground Must Choose

Free Will as Trans-Existential Ground

Geir Isene
isene.com

2026-08-24

Abstract

Elsewhere I have argued that existence cannot ground itself, and that any adequate ground must lie outside existence. It must carry no determinate properties, be able to actualise what has properties, and sustain the resulting order. I call that requirement trans-existential grounding. The argument for it stops at a set of conditions. This paper asks what satisfies them, and the answer completes what I call the Trans-Existential Grounding framework, or TEG.

The answer turns on a single step. A ground that actualises must settle which possibilities become actual. Deterministic settling requires fixed rules, and rules are properties, which puts the ground back inside what it was meant to support. Random settling cannot maintain the persistent structure we observe. What remains between mechanism and chance is selection: the ground chooses. I set this out formally. The openness it requires is supported by the failure of determinism at classical scales and corroborated by Bell's theorem. Consequences follow for computation, for responsibility, and for why a chooser ends up constrained by its own past choices.

The identification is an argument from elimination, not a proof, and I mark the difference clearly. What can be proved is that a ground is required. What can be argued is what it must be like.

1 Where this starts

I take as given the result of a companion paper [5]: existence, understood as everything with properties, exhibits formal structure that cannot account for its own consistency, and so requires grounding from outside itself. That requirement is trans-existential grounding, and the framework completed here, which names what does the grounding, is the Trans-Existential Grounding framework. I abbreviate it TEG and use the abbreviation only for the framework, never for the requirement, since the requirement stands whether or not the framework built on it succeeds.

Four conditions fall out of the earlier argument. Three of them arrive there by separate roads, the Gödelian route giving the first, sufficient reason and the property regress giving the third, and the composite is set out in that paper. An adequate ground must be

1. outside existence, since anything inside is part of what needs grounding;
2. able to actualise, since a ground that produces nothing grounds nothing;
3. free of determinate properties, since properties invite the question why these rather than others and place their bearer back inside existence;
4. able to sustain order, since the structure of existence persists rather than merely beginning.

Readers who reject that argument will find little here to persuade them, and I do not re-argue it. My question is narrower and follows only if the first is settled. Given those four conditions, what meets them?

I argue that only unrestricted decision does, and that the reasoning to this conclusion is philosophical rather than mathematical. It is predicted by the earlier argument, since a formal argument conducted inside existence can establish that a ground is needed without being able to derive, from inside, what lies outside. The identification has to be made by elimination and by considering what the work requires. That is what I do here.

2 The bridge

Condition (2) is where the argument turns. A ground that actualises faces a question at every moment of actualisation: which of the available possibilities becomes actual?

Something settles this. Call the settling *selection*, without yet assuming anything about its character. There are three candidate characters, and two of them fail.

2.1 Mechanism fails on properties

Suppose selection proceeds by rule. Given the prior state, a fixed procedure determines the outcome.

A fixed procedure is a determinate property of whatever runs it. It could have been another procedure, and the question why this one arises exactly as it did for the fine structure constant. By condition (3), a ground with determinate rules is a ground with properties, which places it inside existence, and the regress that motivated the whole inquiry resumes.

There is a second problem, independent of the first. A ground operating by fixed rules is a formal system, and the companion argument shows formal systems cannot ground themselves. Grounding existence with a mechanism grounds one formal system in another, which postpones the difficulty by one step and inherits it whole.

2.2 Chance fails on order

Suppose instead that selection is random: the ground actualises without any principle, and outcomes fall where they fall.

The immediate difficulty has nothing to do with randomness being unattractive. It cannot do the work condition (4) requires. Structure in existence persists. The fine structure constant holds its value across billions of years and across every region we can observe. Chemistry works today as it worked yesterday. Mathematical truths do not fluctuate.

Random actualisation gives no reason for any of this. A ground selecting without principle would produce drift, since nothing would hold the constants. It would produce no stable regularities for physics to describe, and it would in particular fail to explain why the order it produced by chance in one instant is the order it produces in the next. One may reply that randomness within constraints yields stability. That reply concedes the point, since the constraints are then doing the work, and constraints are rules, which returns us to mechanism and its failure on properties.

There is a subtler version worth addressing. Perhaps selection is random and the appearance of order is a selection effect, since only ordered regions support observers. The reply is the one that meets the multiverse in the companion paper. A selection effect explains why we find ourselves somewhere ordered. It does not explain why anything ordered is available to be found in. Selection presupposes something to select from, and an ensemble with a measure over it and a rule fixing

what gets realised has properties, which places it inside existence and leaves the grounding question exactly where it was.

2.3 What remains

Between rule-governed determination and unprincipled chance there is one further possibility. The ground settles which possibility becomes actual, does so without a fixed procedure, and does so in a way that maintains coherence across successive actualisations.

That is choice. The word carries associations that do not apply here: deliberation, weighing of reasons, conscious experience. Those are structures inside existence and none of them is attributed to the ground. What is left when they are stripped away is choice in the minimal sense, selection that is neither forced nor arbitrary.

Theorem 1 (The selection requirement)**.** *Given conditions (1) to (4), the ground's mode of actualisation is neither deterministic nor random.*

Proof. Suppose actualisation is deterministic. Then it proceeds by fixed rules. Fixed rules are determinate properties, contradicting (3). Suppose actualisation is random. Then successive actualisations are unconstrained by prior ones, and the persistence of structure in existence has no account, contradicting (4). The two suppositions are exclusive and exhaust rule-governed and non-rule-governed settling. So actualisation is neither. □

The theorem is negative, and its force depends entirely on the conditions it inherits. What it establishes is that the familiar dichotomy of determinism and randomness does not exhaust the options for a ground of this kind. This is the same gap that appears in the free will literature under a different description, where libertarians are challenged to say how an undetermined choice differs from a chance event. The answer presented here is that the dichotomy was never exhaustive to begin with, and that a ground outside existence is where the third option lives.

3 The framework

Let $\mathcal{E}$ be existence as before. Let $\mathcal{P}$ be the domain of pure potential, understood as what has not been actualised and possesses no properties. Let $\mathcal{W} \subseteq \mathcal{P}$ be the capacity for unrestricted decision.

Definition 1 (Disjointness)**.** $\mathcal{W} \cap \mathcal{E} = \emptyset$.

$\mathcal{E}$ was defined as what has properties and $\mathcal{P}$ as what has none, so the intersection is empty by construction. The substantive claim is this: that something can in some sense be real without having any properties at all.[1] Condition (3) forces that claim, and the rest of the framework rests on it.

Definition 2 (Actualisation)**.** Actualisation is a map $\alpha : \mathcal{W} \times T \to \mathcal{E}$, taking a decision and a moment to an actualised state of affairs.

[1] One might object that *real* is itself a property, so that the claim places $\mathcal{W}$ back inside $\mathcal{E}$. The objection proves too much. On the same reading, *has no properties* would be a property, and the phrase would contradict itself whatever it was applied to. What $\mathcal{E}$ collects are determinate characters, features that could have been otherwise and so invite the question why these rather than others. Words like *real*, *property-free* and *outside existence* mark the role $\mathcal{W}$ plays rather than a character it bears, in the spirit of Kant's observation that existing is not one more feature of a thing. The companion paper puts the same point grammatically: a ground with no properties is not a thing among things, and calling it a thing may already be a mistake of grammar [5].

Two features of α deserve comment. It is indexed by T, so actualisation is continuous rather than a single originating act. Existence is not started and then left to run; it is sustained, and on this account the sustaining and the originating are the same operation applied repeatedly. Second, α is not computable, since a computation is a rule-governed procedure and Theorem 1 rules that out. This is a consequence of the framework rather than an assumption within it, and I return to it in Section 6.

A third feature matters more than either, and the language of Section 2 may have obscured it. Throughout I have written of the ground settling which possibilities become actual, which can be read as choosing among options laid out in a world that is already there. That reading is wrong, and the framework is stronger than it.

$\mathcal{E}$ is the image of α. Nothing in existence lies outside what is actualised, so there is no arena within which $\mathcal{W}$ operates and no furniture it rearranges. The possibilities among which the ground settles are in $\mathcal{P}$, which is not part of existence at all. What α yields is the whole of $\mathcal{E}$: the objects, the laws they obey, the regularities, the structure that Section 2 relied on. All of it, and at every moment, since T indexes the map rather than merely starting it.

Two consequences worth stating plainly. First, existence emerges from free will; free will does not emerge from existence, and is not an inhabitant of it. The direction of emergence in this framework is the reverse of the usual one, where agency arises from physical organisation. Second, the grounding is total rather than partial. It is not that existence stands on its own and requires supplementing at the edges. It is that there is no moment at which existence obtains independently of being actualised, and no property within it that arrives from anywhere else.

This is what makes the account a candidate for grounding at all. A ground that pushed things about inside an existing world would leave the world unexplained, which was the failure that Section 1 inherited from the companion paper [5].

Remark 1 (On the empty set). The natural formal image of a property-free ground is $\{\}$. The empty set has no members and no internal structure, yet the entire cumulative hierarchy is built from it by iteration. That a domain with nothing in it can generate everything is a fact about set theory rather than an argument about metaphysics, and I offer it as an illustration of the shape of the claim, not as support for it.

4 Indeterminacy is real

The framework requires that nature is genuinely open, not merely unpredictable by us. If every event were fixed in advance by prior physical states, there would be nothing for α to settle, and the account would be idle.

The strongest support comes from the classical scale rather than the quantum one, and it is not mine. Potter, Ellis and Mitchell argue that physical determinism fails at every level, and their classical case runs as follows [9]. No physical state is specified to infinite precision, so the properties of a classical system are to some degree indefinite as a matter of ontology rather than ignorance. Chaotic and turbulent systems are sensitive to exactly those indefinite details, which means their macroscopic evolution is under-determined by their current state. Matter in the wet and warm makes this vivid: a large molecule in solution undergoes on the order of 10^{14} collisions per second, and the uncertainty relation makes those interactions fundamentally rather than merely epistemically indeterminate.

Their argument continues into neural systems, since their aim is to show that an organism's decision-making is not deterministic. I take the general result and leave that extension where they made it. Nothing here requires it. The claim this paper needs is that the physical order is open at

classical scales, which holds of a turbulent fluid and a collapsing star as much as of any nervous system. Where in a body the interface between chooser and world happens to sit is a question this framework does not answer and does not depend on, and I would resist any reading on which the brain is where the choosing gets done. A switchboard is not a cause.

They call the resulting picture pervasive indefiniteness, and the reason I lead with it is that it needs no amplification story. An argument from quantum measurement always invites the reply that indeterminacy stays confined to the microscopic and averages away before anything macroscopic happens. On this account it never had to travel, because classical systems are themselves indefinite by default. Openness is the baseline condition rather than an intrusion into an otherwise settled order.

Bell's theorem corroborates from the quantum side [2, 1, 6]. Bell showed that any theory assigning definite pre-existing values to measurement outcomes, with no faster-than-light influence, obeys an inequality bounding the correlations between distant measurements. Quantum mechanics predicts violations, experiments find them, and the loophole-free versions of the last decade have closed the escapes that once made the result contestable. Local hidden variables are excluded, so nature does not carry the answers around in advance, and something is settled at measurement that was not settled before.

What Bell does not exclude deserves equal care, since two deterministic interpretations still hold up. Bohmian mechanics is non-local and fully deterministic: the wave function together with the initial particle positions fixes every outcome, and no experiment yet distinguishes it from standard quantum mechanics. Many-Worlds is deterministic at the level of the universal wave function, which never collapses at all. An argument for openness resting on Bell alone answers none of these.

The classical argument answers the first. Bohmian trajectories are fixed only if the initial positions are specified to infinite precision, and infinite precision is what pervasive indefiniteness denies to any physical state. A Bohmian world whose positions carry finite information is open for the same reason any chaotic classical system is. Determinism is an artefact of idealised initial data rather than a feature of the world.

Many-Worlds fails differently. It does not settle outcomes by another mechanism; it denies that anything is settled, since every outcome occurs on some branch. As a ground it inherits the multiverse objection of the companion paper. An ensemble of branches simply adds more existence rather than an account of existence [5]. It also leaves unexplained the one datum every observer possesses, that experience runs down exactly one branch with definite frequencies.

Two cautions. Neither argument establishes that free will performs the settling, and I am not claiming that either does. What they establish is that the settling is real, which removes the objection that determinism leaves no room for an account like this one. The inference from openness to choice is the argument of Section 2, and it rests on the failure of randomness to sustain order, not on any experimental result. Section 8 sets out where I part company with Potter, Ellis and Mitchell on what the openness is for.

The second caution concerns direction. Bell's theorem in fact assumes a form of experimenter freedom, since the derivation requires that measurement settings be chosen independently of the hidden variables. Superdeterminism denies this and thereby evades the theorem's conclusion. I take superdeterminism to be a coherent position with an extraordinary cost, namely that the correlations between our choices and the systems we measure are arranged in advance, which makes experiment unintelligible as a source of knowledge. I do not think a view that undermines the practice by which it would be tested can be assessed on its evidence.

Superdeterminism is also the brute fact move wearing local dress. Asked why our measurement settings correlate with the hidden variables, it answers that they simply do, arranged that way from the beginning, with nothing further to say. The companion paper gives the general reason that

answer never informs [5]. Here it costs more than usual, since the correlation it declines to explain is precisely the one that would otherwise be evidence.

5 Constraint from one's own choices

An account on which a chooser grounds existence faces an obvious complaint from experience. We do not feel like the ground of anything. We feel constrained by circumstance, by history, by the bodies we have and the situations we did not select.

The framework has a specific answer, and I think it is the most interesting consequence it yields.

Actualisation is not free of cost to the actualiser. Each act of α converts potential into something definite, and the definite persists. What persists must be maintained if the structure it belongs to is to continue. So the actualiser accumulates commitments, and each commitment is a constraint on further actualisation.

Two quantities move in opposite directions here, and distinguishing them clears up a good deal of confusion. Physical entropy measures disorder within what has been actualised, and it rises on its own. Call the second quantity volitional entropy: the accumulated commitment of the actualiser to what it has already actualised. It rises with every act of α that is sustained, and it is felt as constraint.

The structure of a game makes this concrete. Before play, one may do anything. Upon sitting down to play chess, the board is eight by eight, bishops move diagonally, and the alternatives are gone. Nobody imposed those rules from outside; they arrived with the decision to play, and they bind for as long as the decision stands. On this account the laws of physics have the same character with respect to the ground: not external constraints on a will, but the standing commitments of one.

This gives the felt experience of constraint a place in the framework without treating it as evidence against the framework. A being that is free in the relevant sense and has been actualising for a long time should feel exactly as constrained as we do.

6 Computation cannot choose

If α is not rule-governed, then no rule-governed system implements it. This has a consequence for artificial intelligence which I state carefully, since the topic invites overreach in both directions.

A computation is the execution of a procedure. Given its inputs and its state, a deterministic machine produces one output; the mapping is fixed in advance by the program, whoever wrote it and however complex it is. Adding a random source does not change the character of the result, it changes the distribution. Neither operation is selection in the sense of Theorem 1, since one is mechanism and the other is chance, and the theorem excludes both.

So on this framework no computational system chooses, at any scale, with any architecture.

What follows is narrower than it may appear. Nothing here says a machine cannot equal or exceed human performance at any task we can specify, and I expect machines to do so at most of them. Nothing here says a machine cannot produce outputs indistinguishable from those of a chooser, including outputs we find surprising, creative or moving. Behaviour is a poor guide to the presence of $\mathcal{W}$, and the framework predicts that it should be. The claim is about what a system is doing when it produces an output, not about how good the output is.

I take this to be a real difference from Penrose [8], who locates non-computability inside physics, in quantum gravitational processes. Any such process is in $\mathcal{E}$ and governed by law. So, on the argument of the companion paper it is subject to the same grounding requirement as anything else. Relocating the mystery to an exotic corner of physics does not remove it from physics.

7 Responsibility

If choice is real and choices actualise, then their consequences are owned by the chooser in a strong sense.

The compatibilist tradition has developed sophisticated accounts of responsibility that do not require indeterminism, and I want to be clear that this framework does not compete with them on their own ground [3, 4]. Compatibilism asks what makes it appropriate to hold someone responsible, and answers in terms of the right kind of causal history, responsiveness to reasons, and identification with one's motives. These are good answers to that question.

My interest is a different question. If $\mathcal{W}$ is what actualises, then responsibility is not a practice we adopt because it produces good outcomes. It is instead a structural feature of what a chooser is. A being whose choices bring states of affairs into existence stands to those states in a relation no account of causal history can dissolve.

There is an observation from outside philosophy that bears on this, and I offer it as observation rather than argument. In some decades of coaching people through difficult decisions, I have noticed that the desire to be excused and the attraction to determinism arrive together with some regularity. The relief of "I could not have done otherwise" is genuine. It is also, as far as I can tell, the beginning of a diminished capacity to choose, since a faculty one declines to exercise is one that grows harder to reach. This proves nothing about whether determinism is true. It suggests something about the cost of believing it, which is a separate matter but a real one.

8 Related work: indeterminism without a ground

Potter, Ellis and Mitchell have recently argued that the free will debate is framed around a premise that modern physics does not support [9]. Determinism, they contend, fails at quantum scales and at classical ones, the latter because no physical state is specified to infinite precision and because chaotic and turbulent systems amplify that imprecision into macroscopic openness. Their paper is the closest recent neighbour to this one and it deserves a careful placing, since we agree about a great deal and diverge at exactly one point.

Three agreements first, and they are substantial.

We agree that the future is genuinely open, and Section 4 leans on their argument for it rather than on my own. Their case needs no amplification story, since on their account classical systems are indefinite by default, and anyone persuaded by them has the premise this paper requires without having to accept anything about quantum measurement. That is a stronger footing than Bell alone provides, and I would rather borrow it and say so than reach for the weaker support because it is mine.

We agree in rejecting what they call determinism-plus-randomness, the picture on which the world runs deterministically and randomness occasionally intrudes to open a fork. They reject it on empirical grounds, observing that neural systems are built to be insensitive to single quantum events. I reject it on structural grounds in Section 2, since mechanism fails the property condition and chance fails the order condition. Different arguments, the same target, and it is the picture that most libertarian accounts and most objections to them presuppose.

We agree that randomness is not what confers freedom. Their formulation is better than mine and I adopt it: randomness is the noisy backdrop against which the organism acts, not a gift that hands it liberty.

The divergence is about what the account is for. Their positive proposal is that the low-level laws are not causally comprehensive, which leaves causal slack, and that higher-order constraints

exploit the slack so that organisation acquires causal power. Agency emerges from the way systems are organised at macroscopic scales.

That is an account of how agents exert control within a lawful order. It is not an account of the lawful order, and it does not claim to be. Their question is where control comes from, given openness. Mine is what produces the openness, the laws, and the ordered arena in which constraints could operate at all. On the framework of Section 3 there is no independent arena for higher-order constraints to exploit, since the arena is itself what α yields at every moment. Their emergence runs from physical organisation to agency. Mine runs from will to existence, and existence includes the organisation.

So the two accounts are not rivals over the same ground, and I want to resist a reading on which they are. A referee might reasonably ask why trans-existential grounding is needed when Potter, Ellis and Mitchell obtain libertarian free will from physics and organisation alone. The answer is that they obtain agency, which is a different and more tractable thing, and that they obtain it by helping themselves to precisely what the companion paper argues cannot be helped oneself to: a consistent lawful structure whose own grounding is left open. Their argument begins where mine ends.

Two further points of contact deserve mention. Their notion of pervasive indefiniteness is structurally close to $\mathcal{P}$: both are the not-yet-definite from which definite things resolve. The difference in placement is the whole disagreement. They locate indefiniteness inside the physical world as a feature of its states, and hold that it resolves apparently at random. I locate it outside existence, and hold that its resolution is what existence is made of.

Their treatment of the luck objection is also sharper than the one I give in Section 9, and I recommend it alongside mine. Where I deny the dichotomy the objection depends on, they reverse its burden. In a world open by default, the puzzle is not how an undetermined choice escapes being merely lucky. It is how an organism exerts sufficient control to make anything happen reliably at all.

9 Objections

The luck objection. If a choice is not determined by prior states, what makes it the agent's choice rather than a lucky accident? This is the standard objection to libertarian accounts and it is a good one [10, 7].

The objection has force against accounts that locate free will in undetermined events *inside* physics, since an undetermined neural event is exactly a chance event, and chance is not agency. The framework here does not locate choice inside physics. $\mathcal{W}$ is not an indeterministic process in the brain; it is the ground of the process. The dichotomy the objection relies on, determined or lucky, is the dichotomy Theorem 1 denies. Whether the denial succeeds is the right place to press.

You have renamed the mystery. One might grant everything and observe that "a property-free ground that selects without rules" explains nothing and merely labels the thing we could not explain.

I accept a version of this. The framework does not tell you how selection works, and it cannot, since a mechanism of selection would be a rule and rules were excluded. What it does is different: it derives constraints on the ground from the structure of the problem, and shows that the familiar candidates violate them. That is less than an explanation and more than a name. Whether it is enough is a judgment I cannot make for the reader.

Causal closure. Physics is causally closed, so nothing outside it makes a difference to what happens inside it.

Causal closure is a methodological commitment of enormous productive value. It is not an

empirical finding. On this framework α does not inject energy or violate a conservation law, since it settles which of the physically available possibilities is actualised rather than adding a force. Where the possibilities are genuinely open, as Section 4 argued they are, selection among them intervenes on nothing. It determines which physics we get.

Mental causation. How does something with no properties make a difference to something with properties?

The question assumes a world view the framework denies. It assumes two items, a world already standing there and a ground outside it that reaches in and nudges. Interaction of that kind is what substance dualism has never managed to explain, and if $\mathcal{W}$ stood to $\mathcal{E}$ as one thing to another, the objection would be fatal.

Section 3 says the relation is not of that kind. $\mathcal{E}$ is the image of α, so there is no world standing independently that $\mathcal{W}$ impinges upon. Everything in $\mathcal{E}$ is actualised out of $\mathcal{P}$, at every moment. T indexes the map rather than merely starting it. Nothing in existence persists on its own between actualisations. Actualisation is continuous creation rather than interaction.

It converts a problem about interaction into a problem about creation: not how something outside reaches in, but how anything property-free creates anything at all. $\mathcal{P}$ is pure potential rather than nothing, and $\mathcal{W}$ within it actualises. The how is not a question that can take an answer, since any answer would be a rule, and rules belong to $\mathcal{E}$. Something with no properties at the very least stops regression or answers by brute fact.

10 What this does not establish

The companion paper proves a requirement. This paper does not prove what satisfies it.

The argument of Section 2 is an elimination, and eliminations are only as good as their enumeration. I have argued that mechanism fails on properties and chance fails on order, and that selection is what remains. Someone who exhibits a fourth mode of actualisation, or who shows that one of the two failures is not a failure, refutes the identification while leaving the requirement intact.

Nor does anything here establish that human beings possess $\mathcal{W}$, as opposed to being structures within $\mathcal{E}$ that the ground actualises. That further step needs an argument I have not given, and I flag its absence because the framework is often read as though the step were made. What I have argued is about the ground of existence. What we are in relation to that ground is a question I have views about but not yet definite proof of.

11 Conclusion

If existence requires a ground outside itself, and if that ground must actualise without possessing properties while sustaining the order it produces, then it settles what becomes actual by a means that is neither mechanical nor random. Mechanism reintroduces the properties the ground was required not to have. Randomness cannot account for the persistence of the structure we observe. Selection is what is left, and selection under no rule is what the word choice means when everything else is stripped from it.

The openness the account requires is supported from both sides: classical states carry finite information, and Bell violations show that outcomes are not fixed in advance. The framework then predicts that a long-actualising chooser experiences its own past commitments as constraint, that no computational system chooses however capable it becomes, and that responsibility is structural rather than conventional.

That completes TEG as a framework: trans-existential grounding as the requirement, and unrestricted decision as what meets it.

The two halves are not held with equal confidence and should not be. The requirement is proved. The identification is argued, by elimination, and an elimination is only as good as its enumeration. I have tried throughout to keep the two apart, since running them together is the failure mode of every argument in this territory including several of my own. A reader who takes the requirement and rejects the identification has taken the part that was established and declined the part that was proposed, which is the correct way to read this pair of papers if the second does not convince.

References

[1] Aspect, A., Dalibard, J. and Roger, G. (1982). "Experimental Test of Bell's Inequalities Using Time-Varying Analyzers". *Physical Review Letters*, 49(25): 1804–1807.

[2] Bell, J.S. (1964). "On the Einstein Podolsky Rosen Paradox". *Physics*, 1(3): 195–200.

[3] Dennett, D.C. (2003). *Freedom Evolves*. Viking Press.

[4] Frankfurt, H.G. (1969). "Alternate Possibilities and Moral Responsibility". *Journal of Philosophy*, 66(23): 829–839.

[5] Isene, G. (2026). "Why Existence Cannot Ground Itself: Gödelian Incompleteness and the Structure of Physical Law". Companion paper.

[6] Hensen, B. et al. (2015). "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres". *Nature*, 526: 682–686.

[7] Kane, R. (1996). *The Significance of Free Will*. Oxford University Press.

[8] Penrose, R. (1994). *Shadows of the Mind*. Oxford University Press.

[9] Potter, H.D., Ellis, G.F.R. and Mitchell, K.J. (2026). "Reframing the free will debate: the universe is not deterministic". *Synthese*. Preprint: arXiv:2503.19672.

[10] van Inwagen, P. (1983). *An Essay on Free Will*. Oxford University Press.