The Selection Gap: What Petri Nets Add to the Freedom Proof

Geir Isene
https://isene.com

2026-07-02

Abstract

Formal systems cannot ground themselves. The Trans-Existential Grounding (TEG) framework argues from this that existence requires a trans-existential ground, and that only pure potential, free will, can serve. This paper adds a second, independent gap, drawn from concurrency theory.

Petri net semantics treats conflict, two lawfully enabled transitions competing for the same token, as a primitive the formalism itself cannot resolve: both continuations are lawful, and the theory permanently declines to select between them. Causality in net theory is a partial order with no global clock, following Carl Adam Petri's explicitly physical postulates. Lawfulness therefore does not entail determinism; the inference from "there are laws" to "the future is fixed" requires an extra assumption. Attempts to internalize a scheduler either reproduce the conflicts or destroy the concurrency, in structural parallel to the Gödelian containment argument. Bell's theorem rules out a hidden schedule at nature's base. The laws cannot ground themselves, and the laws cannot select the history: the grounding gap and the selection gap admit the same answer.

1 An old tool, a new direction

Years ago I explored Petri nets, with their boxes, circles and tokens. A way to draw how processes flow and wait and synchronize. I used them on business processes, hunting bottlenecks and resource constraints. Then I filed them under useful engineering and moved on.

Coming back now, from the free will side, I see something I missed. Petri nets are more than a modeling tool. They are a mathematical theory of lawful systems that refuse to be clockwork. And the refusal is not hand-waving: it is exact, formal and proven. That gives the Freedom Proof[1] a second gap to stand in.

The first gap is Gödelian. Formal systems cannot ground themselves [4], so existence needs trans-existential grounding. A grounding gap. Net theory opens a different one: even granted the laws, the laws do not select the history. A selection gap. This paper is about the second gap, and why one answer fills both.

To my knowledge, and after a deliberate search of the philosophical literature, no published work builds a free-will argument from net-theoretic conflict and concurrency. Technical texts note in passing that the resolution of conflict "is a matter of philosophy" and move on. The nearest neighbor is Pratt's case for partial orders against global interleaving [13], which is not about free will at all. The ground appears to be open.

[1] The TEG framework and its Gödelian argument are developed in Isene [5].

2 Three primitives

A Petri net has places, transitions, and tokens. Places hold tokens. A transition fires when every input place holds a token. Firing consumes one token along each input arc and produces one token along each output arc. Tokens are not conserved: a transition with two inputs and one output merges, and that is how nets model synchronization.

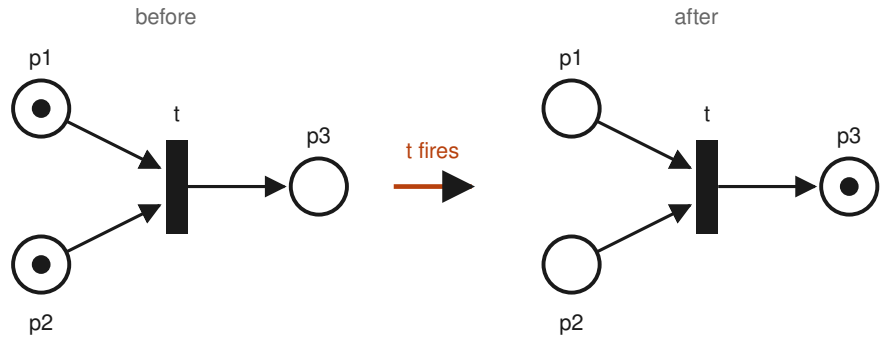

Figure 1: The whole machine. Transition t is enabled because both input places hold a token; firing consumes one token per input arc and produces one per output arc. Two tokens in, one out: tokens are not conserved.

That is all of it (Figure 1). The firing rule is completely formal. No randomness, no hidden parameters, no observer. You can easily simulate one on paper.

But something familiar is absent: a clock. Nothing in the formalism says *when* a transition fires, or which fires first. There is only local enablement. A transition sees its own input places and nothing else.

3 Conflict

Now put one token in a place feeding *two* transitions (Figure 2). Both are enabled. Firing one disables the other.

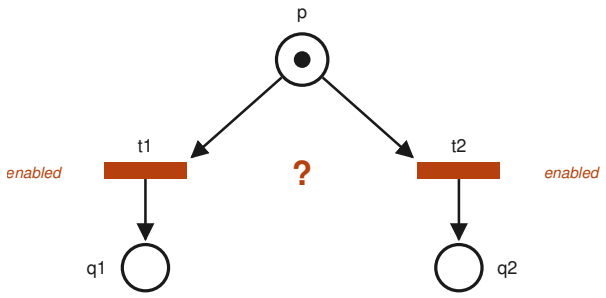

Figure 2: A conflict. One token, two enabled transitions. Both continuations are lawful; the formalism declines to select between them.

Net theory calls this a conflict, and the semantics declines to resolve it. Both continuations are

lawful. The theory does not say which one happens. It says, precisely and permanently, that this is a branch point the system itself cannot close.

You may object and say this is sloppiness, an unfinished corner of the math. It isn't. Conflict is a primitive of the theory, as basic as causation [10, 15]. Concurrency theory even has a technical term, *confusion*, for the tangle where concurrency and conflict meet: fire two independent events in one order and a choice arises, fire them in the other order and that choice never exists (Figure 3). There is a whole hierarchy of net classes sorted by how much of this they allow [9]. So the structure of choosing is mapped with full rigor, while what is chosen stays open. Nothing in the formal machinery can close it.

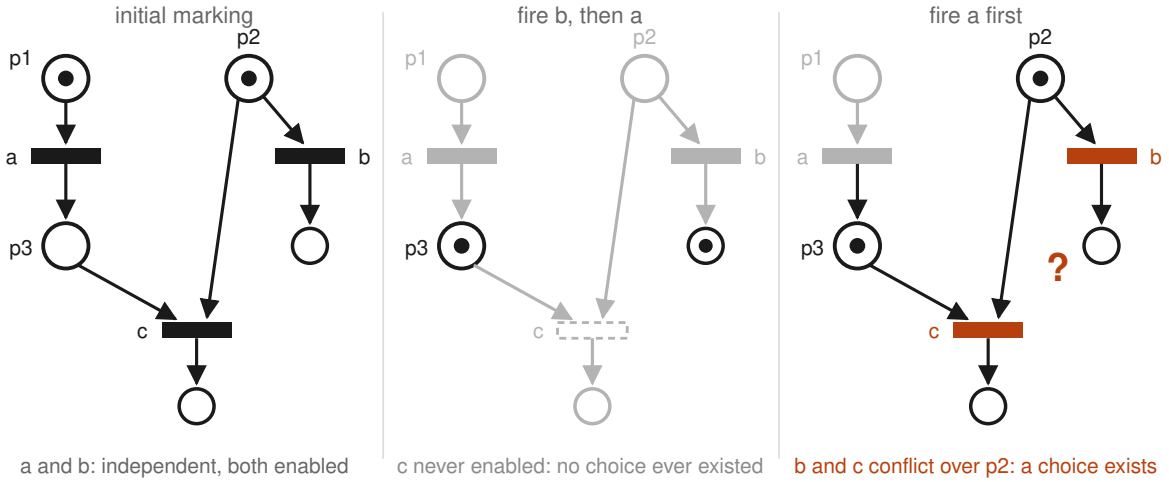

Figure 3: Confusion: whether a choice exists depends on the order of independent events. From the initial marking (left), a and b are concurrent. Fire b before a (middle) and c is never enabled: no conflict ever arises, and b fired without a rival. Fire a first (right) and b and c come into conflict over the token in $p2$: now there is something to choose. The order of two independent firings decides whether a choice ever existed.

A determinist looks at a law-governed system and infers a fixed future. Here is a system that is entirely law-governed, discrete, checkable by hand, and its future is a branching tree. Lawful does *not* mean determined. That inference was always an extra assumption smuggled in. Now we can point at the mathematics that pulls the two apart.

4 No global clock

Carl Adam Petri built the theory this way deliberately. His 1962 dissertation, *Kommunikation mit Automaten* [11], rests on two physical postulates: there is an upper bound on the speed of signals, and an upper bound on how densely information can be stored. Relativity and the uncertainty principle, taken seriously as constraints on computation [14]. From there he rejected centralized control outright. No universal observer, no global now. He wanted a theory of information flow that respects physics instead of quietly assuming a cosmic clock.

So in net theory, causality is a partial order [10]. Event A causes B, or B causes A, or neither (Figure 4). In the third case the events are concurrent, and there is *no* fact of the matter about which came first. Not an unknown fact. No fact.

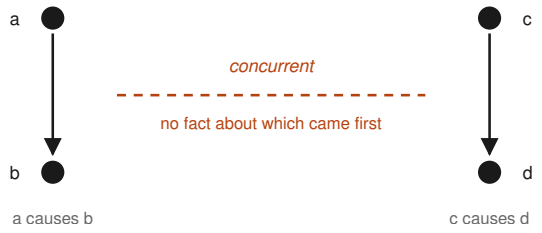

Figure 4: Causality as a partial order. Within each chain the order is fixed; between the chains there is none. Concurrency is not ignorance about the order. There is no order.

Petri himself went further than the textbooks. In his late work, which reads net theory as a kind of discrete physics [12], even the partial order is not assumed. Causality and concurrency are axiomatized as local relations, and ordering is something you derive, not something you start with [6]. The whole program moves in one direction: strip away assumed global structure.

The determinist picture needs the opposite: one global timeline where each total state fixes the next. Net theory shows that picture is optional. Full lawfulness on a partial order works fine, with independence and branching built in from the start. The block universe is one way to read physics. Lawfulness itself does not force it on you.

5 The scheduler lives outside

In engineering we do resolve conflicts. We add a scheduler. The operating system picks which thread runs; the arbiter circuit picks which signal wins.

But the scheduler always sits *outside* the net. Try to pull it inside, encoding it as more places and transitions, and you land in one of two spots. Either the bigger net still contains conflicts, so the problem has only moved. Or you have sequentialized everything, and the concurrency is gone; what remains is a different and poorer system.

Concurrency theory does offer net classes that look like counterexamples. Add priorities and every conflict resolves by rule: the higher priority always wins. Inhibitor arcs go further still. But watch what the priority rule does. It does not resolve an open choice; it deletes it. Where the base net had two lawful continuations, the prioritized net has one. That is the second horn again: no choices left, determinism by construction. Whether the world is wired like that is an empirical bet, and it comes up below.

The shape of this failure is worth staring at. Adding components to a formal system does not let it ground itself, and Lawvere showed that even climbing to category theory keeps the Gödelian limit [7]. In the same way, adding machinery to a net does not let it resolve its own choices. You get a bigger net with the same open points, or no choices left at all. Resolution comes from outside the formalism, or the branching remains.

I want to be careful here. This is a structural rhyme, not the same theorem. But it is a rhyme the argument leans on, and it holds: twice now, formal systems need something they cannot contain. Once for their grounding. Once for their history.

6 Where Bell enters

The obvious objection: the openness is just abstraction. A real CPU resolves conflicts through microphysical detail the net leaves out. Zoom in far enough and the schedule was always there.

That objection assumes determinism at the bottom, and Bell's theorem blocks its straightforward version [1]. Nature, in experiment after experiment, exceeds every correlation limit that *local* hidden pre-set answers allow. There is no local schedule to zoom in on.

I should be precise about what Bell forbids, because two deterministic escapes remain. Bohmian mechanics keeps determinism by adding a pilot wave that coordinates outcomes across any distance, instantly [2]. Superdeterminism keeps it by letting the universe's initial conditions fix every future measurement choice in advance. Both escapes are live. And both are instructive, because look at what they install: a hidden, globally coordinating structure that no local law contains. A cosmic scheduler. The determinist saves his schedule by conceding this paper's shape and relocating the scheduler outside the local laws. The openness gets closed, but only from outside.

Then the pieces fit together. Net theory shows that lawful systems can have genuinely open branch points needing outside resolution. Bell shows there is no local schedule at nature's base, and the surviving deterministic readings must import a global one. TEG reads the result: indeterminate plus not-random equals chosen [5]. Something resolves the branch points, and it is not a mechanism, because mechanisms are exactly what the formalism exhausts.

The stochastic move changes nothing. Engineers sometimes bolt probabilities onto conflicts and call it a stochastic net. But the probabilities are an annotation, added from outside the causal structure; the base theory stays silent about resolution. And randomness as metaphysics faces the same TEG analysis it always did. Pure chance grounds nothing. It renames the gap.

7 What this does not prove

Conflict in nets is underdetermination, not self-reference. It is not a fourth instance of the Cantor-Gödel-Tarski diagonal. The Freedom Proof does not need it to be, and claiming it would invite a correct objection.

Net theory also does not prove that conflict resolution is choice. A determinist can still say the world happens to be a conflict-free net, one where no branch points ever arise. That is a consistent position; the net theorists even have a name for the tame class, free-choice nets, where conflict and concurrency are kept structurally apart [9]. But look at what the position costs. Determinism becomes an empirical bet about the world's specific wiring, no longer a consequence of lawfulness. Bell then makes the bet very expensive. The free ride from "there are laws" to "the future is fixed" is over.

Nor does the argument touch compatibilism. This paper inherits TEG's frame, where free will requires genuinely open alternatives, so its target is the inference from lawfulness to a fixed future. A compatibilist who places free will inside determinism disputes the frame, not this argument. That debate belongs with the framework itself, not with this companion.

That is the real contribution, and it sits in exactly the right place: the selection gap is formally respectable, and closing it from inside has been tried and fails.

8 Two gaps, one answer

Gödel: the laws cannot ground themselves. Petri: the laws cannot select the history. Both times, the formal apparatus points outside itself.

TEG argues the ground must be pure potential, the capacity to choose, because only that demands no further ground [5]. The selection gap fits the same answer with no forcing. Pure potential is precisely what a branch point calls for: something that can resolve without being mechanism, since every mechanism just extends the net.

In the theory of unfoldings, all possible runs of a net form one branching partial order, and any single completed history is a maximal configuration: one full path through the tree, every conflict resolved one way [10, 8, 3]. The mathematics describes the tree completely and says nothing about which configuration gets to be actual (Figure 5). That silence is the selection gap in its sharpest form.

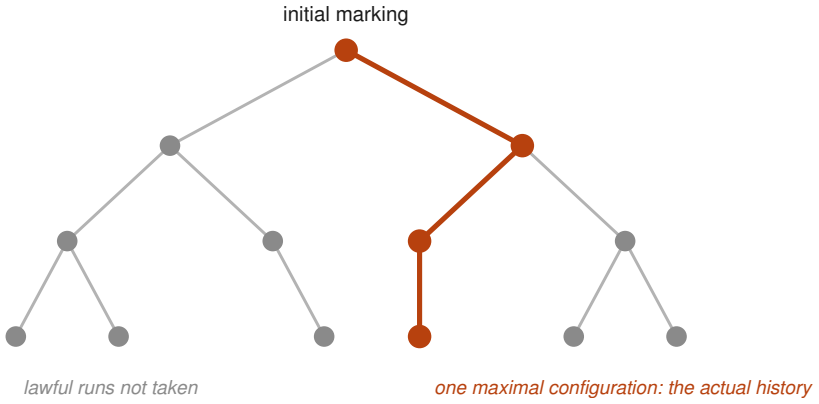

Figure 5: The unfolding: every lawful run of the net in one branching structure. The actual history is one maximal configuration. The mathematics describes the whole tree and is silent about which path gets to be actual.

One ground, two gaps, one answer. Existence gets its being from choice, and its history from choosing. The mathematician who mapped concurrency without a clock left the choosing outside his nets. I think he left it exactly where it lives.

References

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

[2] Bohm, D. (1952). "A Suggested Interpretation of the Quantum Theory in Terms of 'Hidden' Variables, I and II". *Physical Review*, 85(2): 166–193.

[3] Esparza, J., Römer, S., Vogler, W. (2002). "An Improvement of McMillan's Unfolding Algorithm". *Formal Methods in System Design*, 20(3): 285–310.

[4] Gödel, K. (1931). "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I". *Monatshefte für Mathematik*, 38: 173–198.

[5] Isene, G. (2025). "A Mathematical Framework for Free Will: Beyond Determinism, Randomness, and Computational Limits". PhilPapers: `https://philpapers.org/rec/ISEAMF`.

[6] Kummer, O., Stehr, M.-O. (1997). "Petri's Axioms of Concurrency: A Selection of Recent Results". *Application and Theory of Petri Nets 1997*, LNCS 1248, Springer: 195–214.

[7] Lawvere, F.W. (1969). "Diagonal arguments and cartesian closed categories". *Category Theory, Homology Theory and their Applications II*, Lecture Notes in Mathematics 92, Springer: 134–145.

[8] McMillan, K.L. (1995). "A Technique of State Space Search Based on Unfolding". *Formal Methods in System Design*, 6(1): 45–65.

[9] Murata, T. (1989). "Petri Nets: Properties, Analysis and Applications". *Proceedings of the IEEE*, 77(4): 541–580.

[10] Nielsen, M., Plotkin, G., Winskel, G. (1981). "Petri Nets, Event Structures and Domains, Part I". *Theoretical Computer Science*, 13(1): 85–108.

[11] Petri, C.A. (1962). *Kommunikation mit Automaten*. Doctoral thesis, Technische Hochschule Darmstadt. Published in Schriften des Rheinisch-Westfälischen Instituts für Instrumentelle Mathematik an der Universität Bonn, Heft 2, Bonn.

[12] Petri, C.A. (1996). "Nets, Time and Space". *Theoretical Computer Science*, 153(1–2): 3–48. DOI: 10.1016/0304-3975(95)00116-6.

[13] Pratt, V. (1986). "Modeling Concurrency with Partial Orders". *International Journal of Parallel Programming*, 15(1): 33–71.

[14] Smith, E. (2015). *Carl Adam Petri: Life and Science*. Springer.

[15] Winskel, G. (1987). "Event Structures". *Petri Nets: Applications and Relationships to Other Models of Concurrency*, LNCS 255, Springer: 325–392.