

The Seduction of Elegance

Shared Mathematics and the Parts of the World

Luke F. Walton

Independent Researcher · lukefwalton.com

ORCID: 0009-0005-9263-1954

Preprint · v1.0 · October 2026 · Not yet peer-reviewed

Abstract. ML and AI systems are often large and hard to interpret, yet useful. It is easy — perhaps even seductive — to confuse structure for the parts, especially when structure is found through an opaque training process. A black hole, defined by its boundary, invites a similar seduction. Its entropy is proportional to that surface’s area, an exact relation without an inventory of what fell in. Physics cuts space into subsystems. Biology too names its parts. Many fields score what the parts contribute with one mathematics of mereology. Despite these futuristic topics, the study of parts and wholes is as old as philosophy itself. Plato’s Statesman considers the name “barbarian,” and from it the expectation of one kind. But sharing scores is not equivalent to finding the same parts. Even independent methods can produce an epistemically empty agreement when the representation used to compare them erases the differences at issue. The mathematics agrees because each field handed it an order. A structure or mathematics can be useful, or match a division the system has, or be the inventory. One does not automatically describe another. Wimsatt counted it false or misleading to take the entities a powerful decomposition yields as all that matters. Indeed, usefulness does not hand us correspondence. Correspondence does not hand us the inventory. Agreeing on a quantity of the whole is agreeing about the whole. This paper argues against seduction in mereology, which has legitimate promise in many fields. One must not mistake local success for an ontological one.

Keywords: mereology; Möbius inversion; quantum mereology; robustness; representation; emergence

1. Greek or Barbarian?

The Stranger of Plato’s Statesman describes how the Ancient Greeks divided humankind. They separated the Greeks from everyone else, giving all the other peoples, countless and unrelated in blood or language, a single name barbarian. “Because of this single name they expect it to be a single kind” (διὰ ταύτην τὴν μίαν κλῆσιν καὶ γένος ἐν αὐτὸ εἶναι προσδοκῶσιν; Politicus 262d; Burnet

1900–1907). Plato corrects this a page later: every kind (εἶδος) of a thing is a part (μέρος) of it, yet a part is not thereby a kind (263b). What would tell a part from a kind he does not say. In this paper, let’s take up the name. It can be doing the work an inquiry has not done.

Let’s consider three examples. For physics: a Hilbert space, the space a quantum system is written in, does not come factored into subsystems. For biology: a gene-expression state does not come labeled with the cell types it will settle into. For technology: a vector of activations does not come with the concepts it encodes written on its coordinates. The parts had to be found, by a procedure with a criterion inside it, but the procedures use one seductive word, part.

Several sciences that consider parts share a mathematics for what these parts contribute. Jansma (2025a) showed that their measures of contribution, from epistasis to the Shapley value, are one operation — Möbius inversion — run on a partial order of said parts. Where two procedures return different kinds of object (such as a tensor factor and a class of histories), a common space can be built in which both are read as sets of observables. These shared words, theorem, and comparisons teach a shared expectation: one kind of thing, perhaps the inventory of the world. They may be. But sharing the score is not finding the same parts. The exactness lives in the handed order (§3), and two selectors can agree in a representation that cannot show the difference between their parts (§6).

Jansma does not claim this, but it is the role of the philosopher to consider the ontological implications of technological change, and demand robustness in its account. Three things said of a decomposition can come apart: useful for a purpose, prediction first among them; a match to a division the system has; or the inventory, the decomposition rather than one of several. The leap is treating the mathematics as what links them. Usefulness does not give correspondence, and correspondence does not yield the inventory.

The inference fails in two directions. Where two disciplines are taken to have found the same parts because their accounting agrees, a difference between the parts can go uninvestigated or assumed, while the difference between these parts is where a discovery would be. Where the mathematics is taken to have shown that the world has one decomposition, or, from the other side, that it has none because every decomposition is conditional on a procedure, a conclusion has been drawn that the mathematics did not pay for. The first loss is to science. The second is a monism or a nihilism arrived at through the seduction of formalism. Local success is not ontology.

Wittgenstein’s question is what the word “part” is doing in these practices, and whether it does one thing across all of them. A family of uses need share no common essence (Wittgenstein 2009, §§65–67).

The question can be made exact in nature’s most perplexing object: a black hole.

2. Before the Whole, Consider the Hole

A black hole is defined by its boundary. Bekenstein (1973) proposed that it carries an entropy proportional to the area of its horizon, and Hawking (1975) showed that quantum effects make it radiate as a hot body at a temperature set by its surface gravity. Therefore, if it isn't absorbing more energy than it radiates, the black hole loses mass over time.

Page (1993) then asked what happens to the information if the black hole disappears altogether. If its formation and evaporation can be described by an S matrix, the information about what fell in would be expected to come out in the radiation. His estimate was that it may come out so slowly, or so spread out, that an analysis perturbative in the Planck mass over the hole's mass would never see it. Now consider our question. The relation of entropy to the area is the one the theory uses, and the area law admits quantum corrections. What fell in was never an inventory of parts. An account that pictured the whole's information as contents in a container is insufficient, and what the area law says about where information is stored is a contested inference drawn from these observations and mathematics.

So, the question needs a partition. Inside and outside is a cut into two subsystems, and the entropy one side assigns to the other is defined across that cut. In the holographic models, the geometry has been argued to come from the entanglement across that cut (Ryu and Takayanagi 2006; Van Raamsdonk 2010), but the entanglement is defined only once the cut is chosen. Those results hold inside a stated holographic framework, and in that framework the models are handed the cut. The selection of parts becomes a question with physical consequences. The hole sharpens the question about the whole.

Let's zoom in from black holes to the quantum case. Quantum systems are written in what's known as Hilbert space, and such a space does not come factored. A composite system is written as a tensor product of subsystem spaces. Entanglement and locality are defined relative to that product, and the product is not unique: a state entangled under one factorization can be a product state under another. Zanardi, Lidar, and Lloyd (2004) made that point operational. The structure is induced by the interactions and measurements one has access to, so being a subsystem is relative to what one can do with the system. Decoherence uses the same product. Zurek's account of how an environment selects a subsystem's robust pointer states begins after system and environment have been distinguished, an assumption he lists as an unstated axiom (Zurek 2003).

Levins and Wimsatt asked what agreement among independent procedures establishes about a thing itself. Robustness is their test for whether usefulness has reached correspondence: multiple means of determining the thing. When Wimsatt compares two decompositions, he states a narrower test. It asks whether the parts' boundaries coincide, and these parts must be regions of points. Two parts coincide when any two points in one of them fall inside a single part of the other, and conversely.

A factorization of a Hilbert space supplies no such points. As he states it, the test has no points to apply to. Carrying it across takes a translation: something in the factorization has to stand in for the points.

Before we can ask how the whole fits together, we must ask what constitutes its parts. This is the problem of quantum mereology. Three questions concern us: Which parts behave classically? Which parts make physical interactions local? Which parts preserve their informational identity?

First, we need an operator for a given quantum space’s energy to track how it changes in time. This is known as a Hamiltonian. Carroll and Singh (2021) ask which factorizations of a finite-dimensional Hilbert space, given its Hamiltonian, give a system approximately classical behavior.

A caveat: they require pointer states that stay robust and entanglement with the environment that grows slowly for initially unentangled states. They run over two-way splits whose dimensions are fixed in advance, they say the algorithm’s feasibility in practice is unclear, and they say that generic Hamiltonians admit no such decomposition at all.

Cotler, Penington, and Ranard (2019) argue that when a Hamiltonian is local in some factorization the spectrum alone almost always fixes that factorization up to local unitaries, and give what they call an effective proof of that uniqueness; analytic results conditional on one example together with a numerically found example for translation-invariant nearest-neighbor Hamiltonians on a small number of qubits.

Zanardi et al. (2024) propose that the natural subsystems are the ones that scramble information least, and set themselves apart from Carroll and Singh in their own words: the framework is purely quantum and unconcerned with emergent classicality. It selects only among a family of subdivisions that operational constraints have already fixed, since “such a division into simpler components is by no means unique as it depends on the questions one is trying to address, the experimental limitations of the observer, and the physical regime one is operating in” (Zanardi et al. 2024, 1).

Adil et al. (2026), starting from decoherence, find for a fixed Hamiltonian several factorizations whose pointer states are robust, so that pointer robustness by itself need not select one, and find that the decohering subsystems “do not necessarily align with our classical notion of locality.”

Each of these programs buys a determinate answer by paying a purpose. Carroll and Singh fix the dimensions of the split and ask for classicality. Zanardi et al. drop classicality and keep a family someone has handed them. Cotler, Penington, and Ranard assume that a local Hamiltonian exists and ask what the spectrum then fixes. Zanardi et al. are right that the division depends on the question one is asking, and the paper takes that as the premise of the inquiry. A physicist, asked what a part is, answers with a criterion and a purpose, and a second physicist answers with another.

Where any two of them agree, the agreement would be a fact about the Hamiltonian that made them agree, or about an assumption they share.

Naturally, something can be useful without being complete. Newtonian celestial mechanics takes a planet as a point at its center of mass. The point is an idealization the calculation uses; it is accurate to the precision the calculation needs, so the theory remains in use. Two senses of a good description come apart in the case: good because the calculation it supports is good, and good because it says what the system consists of. Quantum mereology is where the formalism writes that gap down, so that nobody can mistake it for a failure of nerve. The same gap is written into the mathematics that scores the parts once they are chosen, and there it is easier to miss, because that mathematics is exact.

3. When the Mathematics Agrees, What Agrees?

So let's look at the score. Williams and Beer (2010) asked how the information two sources carry about a target is distributed between them. Their answer was a lattice: redundant, unique, and synergistic shares. The lattice is exact once the sources are given. It still doesn't fix the redundancy measure.

Jansma (2025a) showed how much of this kind of accounting is one piece of mathematics. Möbius inversion is the combinatorial operation that recovers a function on a partially ordered set from its sums over lower elements (Rota 1964; the arithmetic case is Möbius 1832). This inversion is what partial information decomposition does when it extracts its atoms from a redundancy measure; what epistasis analysis does when it extracts gene interactions from fitness measurements; what interaction information does; and what the Shapley value does when it attributes a cooperative payoff to players.


These fields share a structure: a quantity measured on wholes, a partial order of parts, and the question of what each part added. The unification is a mathematical result that deserves credit as such. Jansma is also explicit that the result begins after a choice. “A Möbius inversion only solves the inverse problem relative to a chosen mereology,” and “while there are no universal criteria for a ‘good’ decomposition, the mereology should be motivated by knowledge of the system’s structure and the macroscopic property of interest” (Jansma 2025a, in the discussion). So, a choice comes first.

Where does the unification live? In the order. Inclusion–exclusion holds on every powerset, whatever the elements are, and the incidence algebra of a locally finite partial order is a fact about the order. What the fields were shown to share is the form of a question: given a quantity on wholes and an order of parts, what did each part add? It is a structural fact about genes and players both that

each supports a quantity on wholes which an order of parts decomposes. The operation would run unchanged on any order whatever was ordered. The fields were asking the same form of question. The exactness is what makes it feel like the same parts. That is the Stranger’s inference with a theorem in place of a name.

Jansma’s informal account of the same mathematics opens on three rings, the Borromean rings (Jansma 2025b). Any two of them can be separated. The three cannot, and a description that stops at the pairs misses the link. What the picture records is how the rings are arranged. The rings are already the parts.

(a)



(b)




Figure 1: Borromean rings, after Jansma’s informal account (2025b): any two come apart, and the three do not. Beside them, the powerset of three elements ordered by inclusion. The sign on each node is (-1) raised to three minus the node’s cardinality, the coefficient of that node in the contribution of the whole set (Jansma 2025a). The alternating sum takes the three elements as given.

He marks two limits himself. The inversion needs the macroscopic quantity defined on every part, so temperature and flocking, the textbook emergents, are out of its reach. And in the field theory he treats, the inversion returned infinite sums over Feynman diagrams rather than the contribution of an individual diagram, so that “the microscopic interactions therefore inherit their justification from the mereology, and are not guaranteed to be ‘real’ or useful” (Jansma 2025a, in the discussion).

The inversion tells you what the parts contributed. It does not tell you which parts to inquire of.

He also relates the accounting across two decompositions when one is a coarse-graining of the other, as a Galois connection, and gets the renormalized couplings of a one-dimensional Ising model out

of it. That is a relation between two decompositions that have already been specified. Which one to privilege is still outside the mathematics.

A restriction the accounting imposes is not a finding about the system. The inversion needs a locally finite partial order with the whole at its top, and a quantity of the whole to invert over the parts below it. Jansma keeps the first condition, so continuous mereologies are out. Two orders are in play, and they need separate names. Call the first the lattice of parts: the parts of one factorization of a finite-dimensional space, a finite lattice, over which the inversion runs without difficulty. Call the second the space of factorizations, ordered by refinement: a continuum with no canonical meet, the one Carroll and Singh parametrize by a continuous family of unitaries and Zanardi et al. minimize a rate over. The space of factorizations is locally finite too, since every interval in it is finite, but the inversion still has nothing to do there.

Infinitely many factorizations lie below the undivided whole, so the sum the inversion would undo is not a finite sum. And no quantity of the whole has been defined on the factorizations, which is what the inversion would have to be given before it could recover a contribution. Nothing in the accounting ranges over the second order. Read the exactness of the lattice of parts as a verdict on the space of factorizations, and a limit of the math has turned into a fact about the system: finitely many parts, parts that cannot be split, or a decomposition that survives a finer cut.

The hotel now named for Hilbert is the old reminder that an intuition about parts need not survive a change in the structure that represents them. An infinite hotel with infinite guests can take in a new guest by moving each guest over one room; a collection already infinite can take in a new member and not grow. Somebody chose these parts first.

4. Parts Found for Something

Two fields have a word for what survives a coarse-graining, and in one case the two words have been proven to name one thing. The renormalization group was finding parts before anyone called them that, and its word is relevance: an operator is relevant when its scaling dimension is low enough to govern behavior at a larger scale. Information theory has the same word with a different definition. Under an information bottleneck, what is relevant is the information about a target that survives compression (Tishby, Pereira, and Bialek 1999).

Koch-Janusz and Ringel (2018) turned that search into an algorithm: a learner that picks the coarse variable of a lattice block keeping the most information about the block's distant environment, with no physics told to it. Gordon et al. (2021) proved what the algorithm was doing. For statistical systems a field theory describes, with short-range interactions and a large buffer between the block and its environment, the degrees of freedom the bottleneck keeps are the operators of lowest scaling

dimension. These are two definitions of a part, shown to be one inside a scoped theorem. It needed a target. Relevance means relevance to something, and the something was a region of the lattice past a buffer, in a block the user fixes.

Computational mechanics did the same thing for time series. Shalizi and Crutchfield (2001) define a process’s causal states as the classes of histories that predict the same future, and prove that the machine built from those states is the unique minimal maximally predictive representation of the process. The parts of a process, on that account, are whatever distinctions prediction requires, and nothing else. Still et al. (2010) tie that criterion to the bottleneck’s: when the relevance variable is the process’s own future and the distribution is known, the bottleneck recovers the causal-state partition in the limit where the penalty on model complexity vanishes. That is the second theorem, and both are anchored in the bottleneck; neither reaches the quantum criteria. What an earned identity looks like is again a theorem with a target and a scope. Where Still et al.’s theorem holds, the bottleneck and the causal states coincide, but two procedures proven to coincide don’t corroborate each other.

Millions of parts, found by a procedure with a criterion inside it: the phenomenon is older than the machines, and the machines industrialize it. Bengio, Courville, and Vincent (2013) describe what a trained network holds as structure useful for the objective it was trained on, and the work since reads it as a direction, a subspace, or a cluster (and, on one published hypothesis, a direction in a language model’s representation space; Park, Choe, and Veitch 2024). Locatello et al. (2019) are the brake. Unsupervised learning of disentangled representations is impossible without inductive biases on both the model and the data, and across large experiments the disentanglement a method reports depends on hyperparameters and the random seed more than on which objective was used. The data alone does not select a decomposition; something has to be brought. Every procedure above that found parts brought it: a distant region, a future to predict, a labeled concept, a Hamiltonian and a purpose.

Biology is where you want one principle for the whole thing. A cell type, on this reading, is a part of an organism’s state space that the dynamics makes stable. Huang et al. (2005) drove one cell line along two biochemically distinct paths in gene-expression space that diverged and then converged on a common state, evidence for an attractor reached from different directions. A subsystem that stays robust under its Hamiltonian, a cell type that stays stable under its regulatory dynamics, a representation that stays useful under its objective. The resemblance is real, and it is the moment a reader reaches for one principle organizing stars, cells, and minds. The objects differ, a region of a state space and a factor of a Hilbert space, and the resemblance is not a theorem.

The next question is whether a boundary found in a model is a boundary in the world. Bruineberg et al. (2022) distinguish a Pearl blanket, which lives in a model, from a Friston blanket, posited

as a boundary in the world, and argue that the mathematics does not pay for the second. A decomposition found in a representation is a decomposition of the representation until something further is shown.

Let's ask what the word part is doing in each of these practices. It names a factor whose pointer states survive a dynamics; an operator of low scaling dimension; a class of histories that predicts; a direction in an activation space; an attractor of a regulatory network. Each is a part for something, and the something is in the criterion. The uses resemble one another the way games do. That is a reason to look for a theorem, and Gordon et al. and Still et al. each found one, inside a scope. A theorem for two of the uses does not give the family an essence. That's the list. What would it mean for two of its entries to be the same?

5. Things Come Apart

The rings hold together as three, but three claims made about any decomposition do not.

A decomposition is useful when the calculations and predictions it supports are good for the purpose they were built for. It corresponds when the division it makes is the system's and not only the method's. It constitutes when it is the inventory, so that any other division of the same whole is a redescription of it and not a rival.

That a division survives changes of procedure and purpose which have nothing to do with it is the evidence that it corresponds, not the definition. Constitution in that sense is stronger than being a constituent. What a discovery of parts claims is correspondence. Constitution is a further result, the rival divisions excluded, and the exactness of a description does not do that excluding. That further result is the claim a powerful decomposition is seduced into making.




Figure 2: Three claims, the separation the rings refuse. Usefulness for prediction does not entail correspondence with something in the world, and correspondence does not entail that the decomposition constitutes the world. The dashed arrow is evidence from usefulness toward correspondence, and it establishes neither. No arrow is drawn from correspondence to constitution: the correspondence of one decomposition is evidence about its parts and not about the inventory until the rival decompositions have been excluded.

Aristotle marks a whole off from a total by whether position makes a difference: where the arrangement of a quantity’s parts changes nothing he says a total ($\pi\alpha\nu$), and where it does, a whole ($\delta\lambda\omicron\nu$) (Metaphysics $\Delta.26$, 1024a1–10; Ross 1924). Wimsatt’s conditions for a system property to be a mere aggregate of the properties of its parts are that distinction made operational, and they are relative to the decomposition: intersubstitution of parts, size scaling, decomposition and reaggregation, and the absence of cooperative or inhibitory interactions (Wimsatt 2007, 280–281). The four are almost never all met, and their failure is his emergence. A decomposition that comes closer to the conditions is easier to calculate, and the move he names a functional localization fallacy follows: “a move from the claim that a decomposition is particularly powerful or revealing to the claim that the entities and forces it yields are all that matters,” a “nothing but” claim he counts false or misleading if it says the preferred entities are the only real ones (Wimsatt 2007, 305). That is the seduction of elegance in the calculation, ease of calculation taken for discovery, named by a philosopher who likes the heuristics.

Usefulness does not entail correspondence, and the cases in which the two come apart are the ones in which a decomposition that works does not survive a change that should not have mattered. The epicycles are the familiar case: a decomposition of a planet’s motion into circles, exact to the precision of the tables for fourteen centuries, which did not survive a second, independent procedure — dynamics found the motion contained no circles — and whose circles nobody now counts among the parts of the motion. Locatello et al.’s experiments are the contemporary case. The disentangled factors a method reports vary with the random seed more than with the choice of objective, and a factorization that moves with the seed may be a division of the data or an artifact of the run;

what the representation is good for does not say which. Dennett (1991) counts a pattern real when some description shorter than the data's bit map captures it. Wallace (2003) takes the reality of macroscopic structure to depend on the usefulness of the theories that include it. By their measure a factorization that compresses the data as well on one seed as on another is as real on either, which is permissive by design and is not where the parting is. Here's what their criterion doesn't ask: whether two useful factorizations are one division or two. Survival under changes that have nothing to do with the purpose is Wimsatt's robustness, and it is evidence that one decomposition corresponds. Whether two that survive are the same parts waits on the comparison. Usefulness is evidence for correspondence, and how strong the evidence is depends on what the usefulness survived.

Correspondence does not entail constitution: cells and molecules are both real parts of an organism, and neither division is thereby the organism's inventory. Adil et al. find several factorizations of one Hamiltonian in which pointer states stay robust, each with features of its own; if each survives what it should, that survival is evidence that each corresponds, and the plurality, if the evidence holds, tells against any one of those factorizations being the sole division of that Hamiltonian. It is not a failing in any one of them. Grant Dennett the word for the patterns that survive. The same data can bear more than one such pattern, according to the noise the description tolerates, and what his criterion then returns is several patterns per data set, cells and tissues and gene-expression attractors among them; he does not say it singles out one inventory.

The philosophical arguments that run from entanglement to a thesis about composition all begin after the parts have been chosen. Schaffer's priority monism keeps the inventory of parts fixed while it reassigns priority among them: "The world has parts, but the parts are dependent fragments of an integrated whole" (Schaffer 2010, 33), and the entanglement in that argument is defined for a system and "its n components" (51), a definition that does not choose the components. Ismael and Schaffer (2020, 4151 n. 23) grant, in the course of defending a fundamental ontology, that the partial trace which assigns a part its state needs a decomposition into parts already in hand. Calosi and Tarozi (2014) ask what mereology the parts satisfy once the factorization is given, and formal mereology, from Simons (1987) onward (Varzi 2026), is handed its inventory the same way.¹

6. The Type of a Comparison

Agreement can look earned in three places. First, the word was the Stranger's. Second, the mathematics was Jansma's: the inversion agrees across fields because every field handed it an order, and the agreement is about the form of the question. Third, the comparison. Two selectors can be said

¹Van Inwagen's (1990) view that almost nothing composes is a different claim, as is the claim that the parts are fictions, and the finiteness the accounting imposed is not a count of the parts.

to agree only once both outputs are read into one representation, the type of the comparison, and the condition on a type is plain. The type has to keep the respect in which the parts would differ. A type that identifies every decomposition of the whole shows nothing. A type that counts one factorization written in two local bases as two invents a difference.

The adequate type shows a difference exactly where the parts differ, so that the transformations it ignores are the ones that preserve the parts, local unitaries and relabelings among them. The condition has a name in the lineage Wimsatt's robustness descends from. Campbell and Fiske asked of a battery of tests not only that different methods converge on one trait but that they tell different traits apart, discriminant validity beside convergent validity (Campbell and Fiske 1959; Thompson 2026), and a type on which every decomposition agrees fails the discriminant half. What fixes the respect is the question the comparison serves, and where two selectors return objects of one kind the formalism fixes the type as well.

Zanardi, Lidar, and Lloyd made a subsystem the thing a set of observables induces: a tensor-product structure is a family of commuting subalgebras of operators whose join is the full algebra, and two factorizations are the same when the families are.

Let's draw it on two qubits. Conjugating the two one-qubit algebras by a controlled-NOT, with the first qubit as control, gives a second pair, the algebra generated by $\sigma_z \otimes 1$ and $\sigma_x \otimes \sigma_x$ and the algebra generated by $\sigma_z \otimes \sigma_z$ and $1 \otimes \sigma_x$, neither of which lies inside either of the first. The Bell state $(|00\rangle + |11\rangle)/\sqrt{2}$, entangled under the first factorization, is a product state under the second. A two-dimensional factor has no proper refinement, so the two factorizations have no common refinement, an instance of the incompatibility Pasqualini and Fortin describe. The coupling $\sigma_z \otimes \sigma_z$ is an interaction between the parts of the first factorization and a field on one part of the second, which is Zanardi, Lidar, and Lloyd's own example with a different unitary, and which is Wimsatt's interactional complexity in quantum dress: a system is interactionally complex to the extent that its subsystems under one decomposition cross the boundaries another decomposition draws (Wimsatt 2007, 181). Two qubits is the smallest space in which this happens, the gate is the canonical entangler, and the coupling is the canonical interaction. That difference is the one that matters. The Hamiltonian is part of what every selector is given, and three types fail to show it. I know of no published comparison of selectors that uses any of them.

First, the algebra all the parts generate is the full algebra for every complete factorization. Second, the spectrum of the Hamiltonian is unchanged by any change of factorization, and here $\sigma_z \otimes \sigma_z$ and $1 \otimes \sigma_z$ have the same spectrum, so a spectral type cannot tell the coupling from the field. Third, a type that counts two factorizations the same when a unitary carries one to the other, which is not a quantity of the whole and is the natural type when only the Hilbert space is given, identifies every two-by-two factorization of two qubits and so identifies an interaction with a field.

The case also lies outside Cotler, Penington, and Ranard’s result for a plain reason, that on two qubits every Hamiltonian is two-body in every factorization, so locality constrains nothing and the spectrum has nothing to pick out; the example shows the freedom their result closes and does not test it. Where their result holds, with enough factors that locality is a constraint, a quantity of the whole does pick out a factorization, and that is the sense in which a frame-invariant fact can pay for parts.

Nothing turns on the gate. Any unitary that is not local, a product of one-qubit unitaries up to a swap of the qubits, carries the first factorization to a second. The two agree on every quantity of the whole, everything no change of factorization alters, because the whole has not changed. They differ in which operators are local. The Hamiltonian is one of the operators whose locality changes. The difference between the factorizations shows up in the Hamiltonian itself, which every selector is handed.

Type and respect can come apart: a field on the first qubit, $\sigma_z \otimes 1$, is a field on one part under both factorizations. The parts differ. The Hamiltonian reads the same. What you have then is a degeneracy, several factorizations one Hamiltonian cannot tell apart. The type governs whether the parts are the same; the respect governs whether the difference is one the selectors could have seen. An interaction is defined on the order that produced it, as Jansma’s coefficients are. What each selector returns is fixed by the Hamiltonian together with its criterion, so the criteria are where they differ.²

Where two selectors return objects of different kinds, the formalism fixes no type and the choice is real. A causal state is a partition of the histories of a measured record, and it enters a space of observables only through the observable whose record it compresses. Two observables give two partitions and two embeddings. So the reading that puts a class of histories beside a tensor factor is itself a selection, and a comparison that reports agreement after it has to say which observable was chosen and why the choice did not make the agreement. The shape of Locatello et al.’s result recurs here as an analogy. Something has to be brought, and what is brought decides what can be found.

7. What an Agreement Returns

Robustness, in the technical sense, is what would join usefulness to correspondence. Levins called a result several independent models recover a robust theorem, “our truth is the intersection of independent lies” (1966, 423), and Wimsatt made robustness the working criterion for “the reality or trustworthiness of the thing which is said to be robust” (2012, 74), with two caveats that bind here. The robustness is illusory when the means of detection supposed to be independent are not

²A script in the supplement checks these identities computationally.

(64, 84). And a system can have robust boundaries while “the different criteria used to decompose it into parts produce radically different boundaries” (76); Wimsatt points that case at the holist, for whom the disagreement of the criteria is a war cry, and it points as well at anyone who would take one criterion’s boundary for the system’s. For selectors the proviso asks what they share: a target, a handed family, a vocabulary, and whether the sharing drives the agreement. Sharing is not copying. A comparison has to show that.

The type is a different matter from independence: two selectors with nothing in common agree on the generated algebra for nothing, so that agreement is not the illusion a shared assumption makes, and it is empty all the same. The two theorems above are where the sharing is proved rather than suspected. Two procedures agree inside a scope, and the common structure that produces the agreement is identified, which Weisberg counts as a robustness analysis carried through (Weisberg 2006). Wimsatt’s corroboration is agreement by means that can fail one at a time.³ A theorem that two criteria coincide on a structure makes their agreement a fact about the criteria, so the means cannot fail separately where the theorem holds. I want robustness in his sense. For the selectors that have not been tied by a theorem, the common structure Weisberg’s analysis looks for is what has not been found, and finding it would be a theorem of the kind Gordon et al. found.

Here’s where I need more from Wimsatt than he claims. His test for two decompositions asks that the parts coincide both ways (§2; Wimsatt 2007, 182). Without the converse, a part contained in a larger part would count as the same part, and a refinement would pass as sameness. What the two-qubit case applied was a translation, inclusion among the parts’ algebras, and the translation returned his verdict, parts that cross one another’s boundaries; the translation is not in his account, and it is the type the formalism fixes where the procedures return objects of one kind. Where they return objects of different kinds the formalism supplies no carrier for a translation, and the type is a choice. That’s the gap. Sameness of parts is typed, and where no type has been fixed two selectors can each recover a robust decomposition and whether they recovered the same parts stays open until one is chosen and defended. An account of what it is for two selectors of different kinds, not yet tied by a theorem, to have found the same parts, rather than for each to have found reliable ones, is what robustness leaves to be written. I can state the condition that account has to meet. I haven’t written the account.

An agreement, when it comes, admits three readings, and the readings are finer than a yes or no. First, the selectors can return the same parts. Second, they can return compatible descriptions at different resolutions, which is the case Rosas et al. (2024) order in a lattice of closed coarse-grainings and the case Jansma’s Galois connection relates on the accounting side. Third, they can return incompatible decompositions, which Pasqualini and Fortin supply, and which may still share

³The debate over what independence requires is live (Thompson 2026).

some of their parts' observables, as the two-qubit pair does. A divergence may be a difference of parts, a difference of scale or timescale, a degeneracy among several factorizations, or a loss in the translation into one family of observables, and it is not, on its own, a finding that a part is a part for a purpose. An agreement under independence returns three things. First, evidence for a correspondence it does not establish. Second, a conclusion about the methods. Third, a conclusion about the assumptions that made the agreement possible. Bruineberg et al. add the test that runs before any of the three readings, whether what each procedure recovers has been shown to be in the world rather than in the model. Incompatibility settles that these divisions are not one division. That is a result about the decompositions compared. It is not, by itself, a finding about what constitutes the world. Sameness is evidence about the parts recovered, and the inventory claim waits on the rivals the comparison did not include.

Neither a universal mereology nor its denial follows. The monism infers from a shared accounting, or a shared stability, to one inventory. The nihilism infers from the conditionality of every selector to the conclusion that there are no parts, or no trustworthy account of them. Neither has a holder among the authors cited. The nearest to the first is the program that seeks a privileged quasi-classical factorization, whose authors say generic Hamiltonians admit none; the nearest to the second is Pasqualini and Fortin's relativity of parthood, which stops short of denying parts. The two inferences run in opposite directions from the same place, a representation, on the same unpaid premise, that a fact about the representation is a fact about the world. The premise has two steps, from the frame to the model and from the model to the world. A frame-invariant fact fixing a factorization pays the first, and the second is paid, where it is paid, by what the decomposition survives; neither is paid by the agreement of the accounting or by the conditionality of the selectors. I decline both. Neither has been earned. A comparison run under Wimsatt's proviso, in a space that keeps the distinctions it tests, is how the possibility of genuinely common parts would be examined. To need the inventory before that comparison, or to take the lack of one as the discovery that there are no parts, calls the result early.

8. Part-icular

A universal theory of parts would be a discovery. These sciences have more instruments for it than any earlier age had. What they have not yet given it is a warrant. The counterfeit has three places: the word, the handed order, and the comparison.

The objection is plain. Of course a representation has to be chosen. When independent methods find the same structure, why should the choice trouble anyone? It shouldn't, once the respect has been named, the type has been shown to keep it, and the methods have been shown independent.

Then the agreement is the beginning of a discovery. I haven't forbidden that inference. What it settles is the parts those methods found.

The black hole shows the size of it. The physics asked what it preserves before it had an inventory of what fell in, and the answer it has is a proportion of entropy to the area of a boundary. Exact, and what fell in stays uninventoried. The same mathematics is how composition gets found, when it is found. The mistake is to take a relation between descriptions for a relation between things. The descriptions now organize what is retrieved, what is counted as similar, and which distinctions survive. The mistake is cheap.

The Stranger's division was a name, and the expectation was a kind. Ours is a name, a theorem, and a space of comparison, and the expectation is one inventory of the world. The Stranger said, show the part to be a kind, then never finished. Where neither a theorem nor a laboratory has settled whether another procedure would find the same parts, the comparison is still open. The procedures are exact enough that the question has a form. That exactness is the seduction; it makes one name feel like one kind.

Disclosures

Competing interests. The author is the founder of Surmado, Inc., which builds AI systems. The paper is written in a personal capacity, without employer involvement and without funding.

Generative AI. Generative AI systems from OpenAI, Anthropic, and xAI were used for research assistance, source location and verification, partial drafting assistance, and editorial review under the author's direction. The framing, the argument, and the final text are the author's, and the author takes full responsibility for the work.

Prior version. A longer synthetic working paper, *Mereology and the Möbius Inversion: What Does It Mean to Find the Same Parts?*, is deposited on Zenodo (concept DOI 10.5281/zenodo.23463198). The present article is a reframing of its argument.

References

- Adil, Arsalan, Manuel S. Rudolph, Andrew Arrasmith, Zoë Holmes, Andreas Albrecht, and Andrew Sornborger. 2026. "Search for Classical Subsystems in Quantum Worlds." *Physical Review D* 113 (10): 103535. <https://doi.org/10.1103/vj4x-96fk>
- Bekenstein, Jacob D. 1973. "Black Holes and Entropy." *Physical Review D* 7 (8): 2333–2346. <https://doi.org/10.1103/PhysRevD.7.2333>

- Bengio, Yoshua, Aaron Courville, and Pascal Vincent. 2013. “Representation Learning: A Review and New Perspectives.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8): 1798–1828. <https://doi.org/10.1109/TPAMI.2013.50>
- Bruineberg, Jelle, Krzysztof Dołęga, Joe Dewhurst, and Manuel Baltieri. 2022. “The Emperor’s New Markov Blankets.” *Behavioral and Brain Sciences* 45: e183. <https://doi.org/10.1017/S0140525X21002351>
- Burnet, John, ed. 1900–1907. *Platonis Opera*. 5 vols. Oxford: Clarendon Press.
- Campbell, Donald T., and Donald W. Fiske. 1959. “Convergent and Discriminant Validation by the Multitrait-Multimethod Matrix.” *Psychological Bulletin* 56 (2): 81–105. <https://doi.org/10.1037/h0046016>
- Calosi, Claudio, and Gino Tarozzi. 2014. “Parthood and Composition in Quantum Mechanics.” In *Mereology and the Sciences: Parts and Wholes in the Contemporary Scientific Context*, edited by Claudio Calosi and Pierluigi Graziani, 53–84. Synthese Library 371. Cham: Springer. https://doi.org/10.1007/978-3-319-05356-1_3
- Carroll, Sean M., and Ashmeet Singh. 2021. “Quantum Mereology: Factorizing Hilbert Space into Subsystems with Quasiclassical Dynamics.” *Physical Review A* 103 (2): 022213. <https://doi.org/10.1103/PhysRevA.103.022213>
- Cotler, Jordan S., Geoffrey R. Penington, and Daniel H. Ranard. 2019. “Locality from the Spectrum.” *Communications in Mathematical Physics* 368 (3): 1267–1296. <https://doi.org/10.1007/s00220-019-03376-w>
- Dennett, Daniel C. 1991. “Real Patterns.” *The Journal of Philosophy* 88 (1): 27–51. <https://doi.org/10.2307/2027085>
- Gordon, Amit, Aditya Banerjee, Maciej Koch-Janusz, and Zohar Ringel. 2021. “Relevance in the Renormalization Group and in Information Theory.” *Physical Review Letters* 126 (24): 240601. <https://doi.org/10.1103/PhysRevLett.126.240601>
- Hawking, Stephen W. 1975. “Particle Creation by Black Holes.” *Communications in Mathematical Physics* 43 (3): 199–220. <https://doi.org/10.1007/BF02345020>
- Huang, Sui, Gabriel Eichler, Yaneer Bar-Yam, and Donald E. Ingber. 2005. “Cell Fates as High-Dimensional Attractor States of a Complex Gene Regulatory Network.” *Physical Review Letters* 94 (12): 128701. <https://doi.org/10.1103/PhysRevLett.94.128701>
- Ismael, Jenann, and Jonathan Schaffer. 2020. “Quantum Holism: Nonseparability as Common Ground.” *Synthese* 197 (10): 4131–4160. <https://doi.org/10.1007/s11229-016-1201-2>

- Jansma, Abel. 2025a. “Mereological Approach to Higher-Order Structure in Complex Systems: From Macro to Micro with Möbius.” *Physical Review Research* 7 (2): 023016. <https://doi.org/10.1103/PhysRevResearch.7.023016>
- Jansma, Abel. 2025b. “Complex Systems and Quantitative Mereology.” Blog post, abeljansma.nl, 28 January 2025. <https://abeljansma.nl/2025/01/28/mereoPhysics.html>
- Koch-Janusz, Maciej, and Zohar Ringel. 2018. “Mutual Information, Neural Networks and the Renormalization Group.” *Nature Physics* 14 (6): 578–582. <https://doi.org/10.1038/s41567-018-0081-4>
- Levins, Richard. 1966. “The Strategy of Model Building in Population Biology.” *American Scientist* 54 (4): 421–431.
- Locatello, Francesco, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. 2019. “Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations.” *Proceedings of the 36th International Conference on Machine Learning*, PMLR 97: 4114–4124.
- Möbius, August Ferdinand. 1832. “Über eine besondere Art von Umkehrung der Reihen.” *Journal für die reine und angewandte Mathematik* 9: 105–123. <https://doi.org/10.1515/crll.1832.9.105>
- Page, Don N. 1993. “Information in Black Hole Radiation.” *Physical Review Letters* 71 (23): 3743–3746. <https://doi.org/10.1103/PhysRevLett.71.3743>
- Park, Kiho, Yo Joong Choe, and Victor Veitch. 2024. “The Linear Representation Hypothesis and the Geometry of Large Language Models.” *Proceedings of the 41st International Conference on Machine Learning*, PMLR 235: 39643–39666.
- Pasqualini, Matías, and Sebastian Fortin. 2026. “Towards a Tensor Product Structure-Grounded Mereology.” *Entropy* 28 (6): 627. <https://doi.org/10.3390/e28060627>
- Rosas, Fernando E., Bernhard C. Geiger, Andrea I. Luppi, Anil K. Seth, Daniel Polani, Michael Gastpar, and Pedro A. M. Mediano. 2024. “Software in the Natural World: A Computational Approach to Hierarchical Emergence.” Preprint, [arXiv:2402.09090](https://arxiv.org/abs/2402.09090).
- Ross, W. D., ed. 1924. *Aristotle’s Metaphysics*. 2 vols. Oxford: Clarendon Press.
- Rota, Gian-Carlo. 1964. “On the Foundations of Combinatorial Theory I. Theory of Möbius Functions.” *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 2 (4): 340–368. <https://doi.org/10.1007/BF00531932>

- Ryu, Shinsei, and Tadashi Takayanagi. 2006. “Holographic Derivation of Entanglement Entropy from the anti-de Sitter Space/Conformal Field Theory Correspondence.” *Physical Review Letters* 96 (18): 181602. <https://doi.org/10.1103/PhysRevLett.96.181602>
- Schaffer, Jonathan. 2010. “Monism: The Priority of the Whole.” *Philosophical Review* 119 (1): 31–76. <https://doi.org/10.1215/00318108-2009-025>
- Shalizi, Cosma Rohilla, and James P. Crutchfield. 2001. “Computational Mechanics: Pattern and Prediction, Structure and Simplicity.” *Journal of Statistical Physics* 104 (3–4): 817–879. <https://doi.org/10.1023/A:1010388907793>
- Simons, Peter. 1987. *Parts: A Study in Ontology*. Oxford: Clarendon Press.
- Still, Susanne, James P. Crutchfield, and Christopher J. Ellison. 2010. “Optimal Causal Inference: Estimating Stored Information and Approximating Causal Architecture.” *Chaos* 20 (3): 037111. <https://doi.org/10.1063/1.3489885>
- Thompson, Morgan. 2026. “William Wimsatt, Donald Campbell, and the Idea of Robustness.” *Acta Biotheoretica* 74: 33. <https://doi.org/10.1007/s10441-026-09543-6>
- Tishby, Naftali, Fernando C. Pereira, and William Bialek. 1999. “The Information Bottleneck Method.” *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, 368–377. [arXiv:physics/0004057](https://arxiv.org/abs/physics/0004057).
- van Inwagen, Peter. 1990. *Material Beings*. Ithaca, NY: Cornell University Press.
- Van Raamsdonk, Mark. 2010. “Building Up Spacetime with Quantum Entanglement.” *General Relativity and Gravitation* 42 (10): 2323–2329. <https://doi.org/10.1007/s10714-010-1034-0>
- Varzi, Achille. 2026. “Mereology.” In *The Stanford Encyclopedia of Philosophy* (Fall 2026 Edition), edited by Edward N. Zalta and Uri Nodelman. <https://plato.stanford.edu/archives/fall2026/entries/mereology/>
- Wallace, David. 2003. “Everett and Structure.” *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 34: 87–105. [https://doi.org/10.1016/S1355-2198\(02\)00085-0](https://doi.org/10.1016/S1355-2198(02)00085-0)
- Weisberg, Michael. 2006. “Robustness Analysis.” *Philosophy of Science* 73 (5): 730–742. <https://doi.org/10.1086/518628>
- Williams, Paul L., and Randall D. Beer. 2010. “Nonnegative Decomposition of Multivariate Information.” Preprint, [arXiv:1004.2515](https://arxiv.org/abs/1004.2515).
- Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.

Wimsatt, William C. 2012 [1981]. “Robustness, Reliability, and Overdetermination.” In *Characterizing the Robustness of Science: After the Practice Turn in Philosophy of Science*, edited by Léna Soler, Emiliano Trizio, Thomas Nickles, and William C. Wimsatt, 61–87. Dordrecht: Springer. First published in *Scientific Inquiry and the Social Sciences*, edited by Marilyn B. Brewer and Barry E. Collins (San Francisco: Jossey-Bass, 1981). https://doi.org/10.1007/978-94-007-2759-5_2

Wittgenstein, Ludwig. 2009. *Philosophical Investigations*. 4th ed. Translated by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Chichester: Wiley-Blackwell.

Zanardi, Paolo, Daniel A. Lidar, and Seth Lloyd. 2004. “Quantum Tensor Product Structures Are Observable Induced.” *Physical Review Letters* 92 (6): 060402. <https://doi.org/10.1103/PhysRevLett.92.060402>

Zanardi, Paolo, Emanuel Dallas, Faidon Andreadakis, and Seth Lloyd. 2024. “Operational Quantum Mereology and Minimal Scrambling.” *Quantum* 8: 1406. <https://doi.org/10.22331/q-2024-07-11-1406>

Zurek, Wojciech Hubert. 2003. “Decoherence, Einselection, and the Quantum Origins of the Classical.” *Reviews of Modern Physics* 75 (3): 715–775. <https://doi.org/10.1103/RevModPhys.75.715>
