From Truth to Trust: Negotiating Objectivity through Calibrated Instruments, Media, and Institutions

Nagarjuna G.
nagarjuna@iiserpune.ac.in
Visiting Faculty, IISER Pune
Former Professor, TIFR, Mumbai

April 4, 2026

Abstract

This paper argues that truth-conditional analysis, while adequate for predicative assertions (class membership, taxonomy), does not reach the *functional relations*—proportionalities, invariances, lawful dependencies between variables—that constitute the distinctive form of modern scientific knowledge. Following Cassirer's insight that modern science shifted from substance-concepts to function-concepts (Cassirer, 1923), I argue that the epistemology appropriate to functional relations is an account of *trust conditions*: the depth and calibration of the mediating chain through which a claimed functional relation is grounded in the actual world. An instrument is possible when a functional relation between variables exists; calibration checks that the relation holds within tolerances. Mathematics is the formal language of *possible* functional relations; science maps these to the actual world through calibrated instruments.

I develop this through a layered account of epistemic mediation, from the agent's biologically calibrated sensory apparatus (Layer 0, predominantly predicative) through physically fabricated instruments (Layer 1, where functional relations enter) to cascaded multi-instrument practice (Layer N). Trust depth is generated through *procedural calibration*—disciplined, inspectable, reproducible procedures that constrain interpretive variance—of which *physical calibration* is the special case that most powerfully eliminates reflexivity. I call the resulting position *calibrational realism* and situate it in relation to entity realism (Hacking, 1983), constructive empiricism (van Fraassen, 1980), operationalism (Bridgman, 1927), and recent work in measurement theory (Tal, 2017; Chang, 2004).

Keywords: trust conditions, functional relations, objectivity, instruments, calibration, procedural calibration, scientific realism, measurement, philosophy of science

AI assistance declaration. This paper was developed through iterative dialogue between the author and Claude Opus 4.6 (Anthropic), used as a writing and research assistant. All philosophical arguments, conceptual innovations, and editorial decisions are the author's. In addition, Claude assisted in type setting the paper using LaTeX. The author takes full responsibility for the content.

Contents

1 Introduction: The Problem of Objectivity

Two chemists in different laboratories independently measure the melting point of aspirin and report 135 °C. Their agreement is unremarkable to the working scientist but deeply puzzling to the philosopher. What explains it? The standard answers in philosophy of science divide along familiar lines. The scientific realist says the chemists have accessed a mind-independent fact: aspirin really does melt at 135 °C, and their instruments tracked that fact (Psillos, 1999). The constructive empiricist says we should believe only that the instruments performed reliably within the domain of observable phenomena, remaining agnostic about any deeper metaphysical commitment (van Fraassen, 1980). The social constructivist says the agreement is an achievement of shared training, conventions, and institutional norms—that "135 °C" is a social accomplishment, not a discovery (Latour, 1987). The conventionalist says the agreement depends on arbitrary but entrenched choices—the Celsius scale, the definition of "melting point," the protocol for heating rate (Duhem, 1954).

Each answer captures something real. Yet the debate among them has persisted for over a century without resolution, and for a revealing reason: all parties tend to privilege the *proposition*—"the melting point of aspirin is 135 °C"—as the primary unit of epistemic evaluation. They ask: is this proposition true (realism)? empirically adequate (van Fraassen)? socially stabilized (constructivism)? conventionally determined (Duhem)? In each case, the proposition floats free of the material process that produced it.

This paper proposes a different starting point. I argue that truth-conditional analysis, while adequate for predicative assertions ("snow is white," "aspirin is an analgesic"), does not reach the *functional relations* that constitute the distinctive form of scientific knowledge. The proposition "the melting point of aspirin is 135 °C" is not primarily a class-membership assertion; it is a claim grounded in a functional relation between thermal input and instrument response—the physical proportionality of the thermometer, the calibration protocol, the institutional standards that govern laboratory practice. I call the conditions under which such functional relations are epistemically warranted and publicly stable the *trust conditions* of the measurement.

The re-centering from truth conditions to trust conditions is not merely terminological. It relocates the unit of epistemic evaluation from the proposition to the *trust chain*: the sequence of physically and conventionally calibrated mediators through which a phenomenon is translated into a publicly available token. It places instruments—usually treated as peripheral aids to observation—at the center of epistemology, as the arbitrators of objectivity rather than its servants. And it yields a new position in the realism debate, which I call *calibrational realism*: the view that scientific progress consists in the progressive externalization of trust conditions from subjective phenomenal experience to proportional invariant bases of external material.

But the thermometer that measures aspirin's melting point is only one point on a vast spectrum. Consider four cases that span the range of trust-producing micro-worlds:

- A **geometry box**—compass, protractor, ruler—creates a micro-world for spatial relations:

angles, lengths, and proportions that can be measured, reproduced, and communicated independently of the user's estimation.

- An **MRI scanner** creates a far more complex micro-world: multiple transducers capture nuclear magnetic resonance signals, and computational layers perform Fourier transforms, spatial encoding, and image reconstruction before a clinician sees a diagnostic image. The final output is not a single proportional reading but a synthesized representation produced by cascading dozens of functional relations through computational mediation.
- A **randomized controlled trial** testing a new drug creates a procedural micro-world: randomization, blinding, and pre-registration constrain the behavior of human agents so that the treatment-outcome relation can manifest with reduced confounding. The "instrument" here is the trial design itself—a procedurally calibrated structure that manages reflexivity among human subjects.
- A **court of law** applying rules of evidence creates an institutional micro-world: cross-examination, appellate review, and adversarial process constrain interpretive variance among judges and jurors so that the evidence-judgment relation approximates procedural fairness.

Each of these is a *constructed micro-world*: a bounded, calibrated space in which specific functional relations can manifest with reduced interference. The history of human knowledge-making is, in large part, a history of constructing such micro-worlds along two convergent paths. The first path runs through *physical instruments*—from the geometry of Euclid and the simple machines of Archimedes, through Galileo's pendulum and the lenses of the early microscopists, through the thermometer, the transducer, and the electronic sensor, to the computational measurement systems of contemporary science. Each step creates a new micro-world in which a specific functional relation is isolated, packaged, and made publicly available. The second path runs through *social institutions*—from simple collaborating groups and guilds, through structured procedures for dispute resolution and collective decision-making, to the laboratories, journals, peer review systems, and metrological standards that characterize modern science (Shapin and Schaffer, 1985; Knorr-Cetina, 1999). Each step creates a procedurally calibrated action space in which the behavior of agents is constrained to reduce interpretive variance.

The two paths converge in the modern laboratory, which is simultaneously a site of physical instrumentation and a social institution—an *epistemic culture*, in Knorr-Cetina's term (Knorr-Cetina, 1999), organized to produce trustworthy knowledge through the coordinated deployment of calibrated instruments, disciplined procedures, and institutional oversight. Edwin Hutchins' studies of navigation as distributed cognition across persons and instruments (Hutchins, 1995) illustrate this convergence concretely: the ship's course is determined not by any single observer but by a system of calibrated instruments, procedural routines, and coordinated human actions—a trust chain distributed across a sociotechnical micro-world.

The trust-conditions framework provides the vocabulary for describing this entire landscape. From the geometry box to the MRI scanner, from the village council to the Supreme Court, from the

astronomer's notebook to the AI-augmented data pipeline—all are constructed micro-worlds in which functional relations are isolated, calibrated, and validated. They expand the affordances of the world (Gibson, 1979), transforming given environments into navigable, calibrated action spaces—what niche construction theory calls organism-modified environments (Odling-Smee et al., 2003). We have, over millennia, progressively transformed habitat into what might be called *calibrated habitat*: constructed spaces in which trust chains are materially embedded. The epistemology this paper develops is situated in that landscape. It tells the story of how intuitive, phenomenal truth—the kind available at Layer 0—was progressively supplemented by calibrated trust: the kind produced through physically and procedurally calibrated micro-worlds of ever-increasing depth and sophistication.

The argument proceeds as follows. Section 2 develops the core conceptual argument: that truth conditions evaluate predicative relations (class membership), while modern science operates through functional relations (proportionalities between variables), and that trust conditions are the epistemology appropriate to functional relations. Section 3 presents the layered account of epistemic mediation. Section 4 argues that instruments are not peripheral tools but the constitutive locus of objectivity. Section 5 shows that the trust-conditions framework dissolves the discovery/invention dichotomy, since instrumented observation is a form of engineered intervention through which phenomena manifest. Section 6 extends the argument to simulations, showing that they function as virtual minilabs where trust depth can be tested before physical experimentation. Section 7 states and defends calibrational realism in relation to the main alternatives in the realism debate. Section 8 addresses objections. Section 9 draws out implications for the history and demarcation of science. Section 10 concludes.

2 From Truth Conditions to Trust Conditions

2.1 Truth conditions and the logic of predication

The standard picture in analytic philosophy takes meaning to be given by truth conditions. Following Tarski, the meaning of "snow is white" is given by the conditions under which that sentence is true—namely, when snow is white (Tarski, 1944). This framework, extended to scientific statements, underwrites the correspondence theory of truth: a scientific proposition is true when it corresponds to the way the world is, independently of any observer.

Observe the logical form of the paradigmatic example. "Snow is white" is a *predicative* assertion: it assigns an object (snow) to a class (white things). The Tarskian T-schema—" 'Snow is white' is true iff snow is white"—evaluates a *class-membership relation*: is this particular a member of that kind? This is the logic of *predication*, and it is the logic that truth-conditional semantics was designed to evaluate. It is entirely adequate for assertions of this form, and much of common-sense language consists of just such assertions: "That is a cup," "The sky is blue," "Gold is a metal." These are claims about what things *are*—about kind-membership, genus-species relations, taxonomic classification.

But is this the form that scientific knowledge characteristically takes? When a scientist reads

"37.2 °C" from a clinical thermometer, the truth-conditions framework tells us that this statement is true if and only if the patient's temperature is 37.2 °C. It evaluates whether the reading falls in the right class. But it tells us nothing about *how* the mercury column's height *varies as a function of* the patient's thermal state, what physical proportionality the instrument exploits, or what would need to go wrong in that functional relation for the assertion to fail. The proposition is evaluated in isolation from the functional relation that produced it.

This isolation is not merely an omission; it reflects a deeper limitation. Truth-conditional analysis was built for predicative relations—class membership, subset inclusion, kind assignment. But the distinctive form of modern scientific knowledge is not predicative. It is *functional*.

2.2 From substance to function: the objects of scientific knowledge

Ernst Cassirer, in *Substance and Function* (Cassirer, 1923), identified a transformation in the form of scientific knowledge that occurred between Aristotelian natural philosophy and modern science. Aristotelian science sought knowledge of *substances*—essences, universals, genus-species hierarchies. Its logic was predicative: to know something was to assign it to a kind, to locate it in a taxonomy, to identify its essential properties. The method appropriate to this form of knowledge was classificatory: observe particulars, abstract common properties, arrange into genera and species. Induction, as codified by Bacon, was the method of discovering such taxonomic regularities.

Modern science, beginning with Archimedes and fully realized in the work of Galileo, Kepler, and Newton, seeks knowledge of a fundamentally different kind: *functional relations*—proportionalities, invariances, and lawful dependencies between measurable variables. The law of the lever does not predicate a property of a substance; it states that the product of weight and distance from the fulcrum is invariant across the two arms. Newton's law of gravitation does not assert that bodies "possess" a force; it states that the attractive interaction between any two bodies is proportional to the product of their masses and inversely proportional to the square of their distance. As Cassirer observed, "the force of gravitation is not 'inherent' in any one thing, but is essentially a relation between things... There is left only mathematical proportionality" (cf. Cassirer, 1923).

This shift from substance-concepts to function-concepts has profound epistemological consequences. Consider the difference between two claims about snow:

- "Snow is white." This is a predicative assertion—snow belongs to the class of white things. It is a sub-type relation between two categories. Truth conditions evaluate it directly: the statement is true iff snow is white.
- "Snow reflects most of the visible spectrum." This is a functional assertion—it describes a relation between the material disposition of snow (its microcrystalline structure) and the spectral distribution of reflected light. It can in principle be expressed as a function: $\text{reflectance}(snow, \lambda) > threshold$ for $\lambda \in$ visible range.

The second claim is not merely more precise than the first; it is a different *kind* of knowledge. It does not assign an object to a class but characterizes a functional relation between physical variables.

And this is the form that scientific knowledge characteristically takes: Boyle's law ($PV = k$), Ohm's law ($V = IR$), the Stefan-Boltzmann law ($j = \sigma T^4$), the dose-response curves of pharmacology, the reaction-rate equations of chemistry. Even qualitative scientific claims—"increasing CO_2 concentration raises global temperature"—assert relations between potential variables, not class memberships.

Cassirer was right about the shift from substance to function, and his insight has been independently confirmed by subsequent work in the philosophy of science: the non-statement or structuralist view of theories (Suppes, 1962; Stegmüller, 1976; Suppe, 1977), which holds that scientific models are mathematical structures specifying relations among quantities rather than sets of propositions, makes the same point from a different angle. But Cassirer drew neo-Kantian conclusions—the function-concept as an a priori form of thought—rather than epistemological-instrumentalist ones. He did not ask: *how are functional relations grounded in the actual world?* That is the question this paper answers. The answer is: through calibrated instruments that physically instantiate functional relations, and through the trust chains that validate them. An independent line of argument developed in Nagarjuna (1994) arrives at the same distinction from the side of logic and methodology. That work distinguishes *structure-dependent* knowledge—knowledge whose meaning arises only within a relational structure, as "loss" and "profit" make sense only within the structure of a transaction—from *structure-independent* knowledge, such as counting objects at a bank counter, which requires no relational context. All scientific knowledge, on that account, is structure-dependent: the meaning of "frequency," "mass," or "temperature" presupposes a structure of functionally related parameters. The work further argues that inversion—the discovery and coupling of inverse functional relations—is the generative logical operation through which such structures are constructed: science advances not merely by classifying objects into kinds but by constructing and coupling functional relations (transforms, proportionalities, reciprocals) into coherent systems.

2.3 Trust conditions as the epistemology of functional relations

If truth is a predicate of *statements*—assertions of class membership, evaluated by the Tarskian apparatus—then a different epistemic category is needed for *functional relations*. That category is *trust*. Truth asks: does this object belong to this class? Trust asks: does this functional relation hold under specified conditions, with what fidelity, and through what mediating chain?

An instrument becomes possible when a functional relation between variables exists. The mercury thermometer is possible because temperature and mercury-column height stand in a functional relation: $T = f(h)$. If mercury expansion were random with respect to temperature—if no function related the two—no thermometer could be built. The instrument *physically instantiates* a functional relation, and calibration checks that the instantiation holds within tolerances. The trust we place in the instrument is trust in the functional relation it embodies.

This reframing sharpens the distinction between *meaning* and *grounded meaning*. Any token—a number, a reading, a data point—acquires *meaning* through its relational context: the network of concepts, associations, and inferential connections within which an interpreting agent situates

it (Nagarjuna, 2026). The string "CO_2 concentration: 421 ppm" is meaningful to anyone who understands the relevant concepts, regardless of how the string was produced. But meaning and *grounded* meaning are not the same. The string is *grounded* only when the functional relation between atmospheric CO_2 and sensor response has been physically instantiated in a calibrated instrument and validated through a trust chain—calibrated sensor, validated data pipeline, peer-reviewed protocol. Without such a chain, the string describes a *possible* functional relation (a mathematical possibility); with it, the string describes an *actual* one (a physically validated claim about the world).

Mathematics, on this account, is the formal language of *possible* functional relations. Mathematicians have been building this language for centuries—algebra, calculus, differential equations, group theory, category theory—and what they have constructed is a vast space of possible functional structures: possible worlds in which variables stand in specified relations to one another. This space has aesthetic as well as epistemic value—the beauty of a mathematical structure is the beauty of a possible functional world, with its internal coherence, symmetry, and balance. Science is the enterprise of mapping these possible functional relations to the *actual* world, through calibrated instruments that physically instantiate the relations and trust chains that validate the mapping. On this view, the formal semantics of trust conditions is nothing other than the formal semantics of mathematics: the vocabulary of possible functional relations is exactly what mathematicians have been developing, and trust conditions specify when and how those possible relations are grounded in actuality.

What entitles us to *truth-talk*—to treating a token as a candidate for truth rather than mere mathematical possibility—is the depth and calibration of its grounding chain. A claim whose trust chain is deep enough and well-calibrated enough is one we are warranted in calling true; a claim that lacks such grounding remains meaningful (it describes a possible functional relation) but epistemically unsecured (we do not know whether it holds in the actual world). The traditional question "under what conditions is this statement true?" is thus supplemented by a prior question: "under what conditions is this functional relation *grounded*?" The trust-conditions framework provides the answer.

The *trust conditions* of a token are the conditions under which the mediating chain that produced it is epistemically adequate—the conditions under which the chain can be relied upon to translate between the phenomenon and the terminal token with sufficient fidelity. These conditions are not subjective preferences. They can be specified, evaluated, and compared along at least four dimensions:

1. **Physical calibration.** Is each mediating layer fabricated on the principle of proportional material response? Are the instruments calibrated against known standards?

2. **Transparency.** Can the mediating layers be inspected, tested, and independently verified? Or is the chain opaque—a black box whose internal workings are inaccessible?

3. **Reproducibility.** Can different agents, using the same chain, arrive at the same terminal

token? Reproducibility is the social test of a trust chain's reliability.

4. **Institutional stabilization.** Is the chain maintained by institutions that enforce standards, train practitioners, and correct errors? Peer review, metrology, accreditation—all are institutional mechanisms for stabilizing trust chains.

These dimensions are not binary but graded. A kitchen thermometer has shallower physical calibration than a platinum resistance thermometer in a national standards laboratory, but both are physically calibrated in a way that a verbal report of "it feels warm" is not. The trust-conditions framework replaces the binary objective/subjective distinction with a graduated assessment of trust depth.

2.4 Not relativism, not reliabilism

Two objections must be addressed immediately.

First, the shift from truth to trust might seem to invite relativism: if grounding depends on what agents trust, and agents can trust different things, then grounded knowledge is relative to agents. But this conflates *dependence on agents* with *arbitrariness*. Trust chains are not arbitrary; they can be evaluated along the four dimensions just specified. A trust chain grounded in physically calibrated instruments, transparent procedures, reproducible results, and institutional oversight is *objectively better*—more reliable, more resistant to error—than one grounded in hearsay and phenomenal impression. The framework is agent-centred but not relativist: it provides normative criteria for evaluating trust chains that are independent of any individual agent's preferences.

Second, trust conditions might seem to reduce to process reliabilism (Goldman, 1979). Goldman argues that a belief is justified when it is produced by a reliable cognitive process. Trust conditions share the emphasis on the process of production, but differ in two critical respects. First, reliabilism specifies reliability abstractly—as the tendency of a process type to produce true beliefs—whereas trust conditions are grounded in the *specific physical calibration* of material devices. The unit of evaluation is not an abstract process type but a concrete chain of physically fabricated mediators. Second, reliabilism inherits the truth-conditions framework: a process is reliable insofar as it tends to produce *true* beliefs, where truth is evaluated predicatively (does the belief correctly assign an object to a class?). The trust-conditions framework operates at a different level: a chain is trustworthy insofar as the *functional relations* it embodies—the proportionalities between physical variables—hold within known tolerances. Truth-talk is what we are *entitled to* when trust chains are deep enough. Truth is not denied but relocated—it is downstream of grounding rather than the independent standard against which reliability is measured.

Helen Longino's social account of objectivity—that objectivity is achieved through transformative critical interaction within a scientific community (Longino, 1990)—is an important precursor. The trust-conditions framework shares Longino's insight that objectivity is a social achievement, not a metaphysical endowment. But it adds a dimension that Longino's account underspecifies: the *physical calibration* of the mediating layers. Critical interaction is more effective when it operates

on tokens produced by calibrated instruments—tokens that express functional relations—than when it operates on tokens produced by phenomenal impression. The physical dimension is what makes scientific critical interaction epistemically distinctive.

2.5 Scope of the claim

It is important to be explicit about what the trust-conditions framework is and is not claiming. The framework is *not* offering a replacement semantics for scientific language. Truth-conditional semantics—the Tarskian apparatus that specifies the conditions under which a sentence is true—remains indispensable for understanding predicative assertions, and predicative assertions remain a part of scientific discourse (taxonomy, classification, diagnosis). Nor is the framework a denial that truth matters to science. Scientists rightly care whether their claims are true, and the trust-conditions framework does not contest this.

What the framework *does* claim is that the *distinctive* form of modern scientific knowledge is functional, not predicative—it consists of claims about proportionalities, invariances, and lawful dependencies between variables—and that truth-conditional analysis, designed for predicative assertions, is the wrong level at which to understand the objectivity and warrant of functional knowledge. The trust-conditions framework provides the epistemology appropriate to functional relations: it specifies the conditions under which a claimed functional relation is *grounded* in the actual world through calibrated instruments.

The taxonomic/functional distinction is not a binary. Common-sense language, parts of biology, and diagnostic medicine still use predicative assertions, and truth conditions remain apt for evaluating them. The claim is about the direction of scientific development: as Cassirer documented, the trajectory of modern science has been consistently from substance-concepts toward function-concepts, from taxonomy toward invariance, from classification toward proportionality. It is this trajectory—and the epistemic demands it creates—that the trust-conditions framework addresses.

The scope of the framework extends beyond single instruments to the full spectrum of constructed micro-worlds through which trust is produced. A packaged instrument such as a thermometer is one kind of micro-world—a sealed, physically calibrated space in which a single functional relation is isolated and made readable. But an MRI scanner, with its multiple transducers and computational layers, is a more complex micro-world; a laboratory, with its instruments, people, and protocols, is a sociotechnical micro-world; and a social institution such as a court of law or a peer review system is a procedurally calibrated micro-world in which the behavior of reflexive agents is constrained to reduce interpretive variance. The trust-conditions framework provides a unified vocabulary for evaluating knowledge claims across this entire spectrum—from the geometry box to the AI-augmented data pipeline—by asking, at each point, how deep and how well-calibrated the mediating chain is, and how effectively reflexivity has been managed. The framework is thus not restricted to the natural sciences or to physical instruments; it applies wherever constructed micro-worlds produce knowledge claims that aspire to public stability and warrant.

The downstream claims about realism (Section 7) follow from this re-centering. Calibrational

realism is a position about what constitutes genuine epistemic progress in science, not a metaphysical thesis about the furniture of the world. It says that progress consists in trust-chain deepening—the progressive externalization of epistemic warrant from subjective experience to calibrated material process—and that this deepening is real, cumulative, and normatively evaluable.

3 Layers of Epistemic Mediation

Science is, at its core, *mediated cognition*. But a careful analysis reveals that *all* cognition is mediated—the question is by what, and how well calibrated. The agent's own sensory apparatus is the first mediating instrument, calibrated not by human design but through billions of years of organic evolution. The eye, the ear, the proprioceptive system—these are biologically calibrated instruments, shaped by natural selection to respond proportionally to features of the environment that were relevant to survival. They are not "unmediated" access to the world; they are the instrument we are born with.

What distinguishes scientific knowledge is that it interposes *additional* media between the agent and the phenomenon—language and inscriptions, physically fabricated instruments, procedures, theoretical models, mathematical formalisms—that overcome the limitations of the biological instrument. Science is constitutively Agent–Media–World: it extends the agent's mediating capacity beyond what biological calibration alone can achieve. The epistemic power of science lies precisely in the character of these additional media.

Gibson's concept of affordances—possibilities for action that the environment offers to an organism (Gibson, 1979)—applies naturally at the Agent–World level: surfaces afford walking, objects afford grasping, other organisms afford social interaction. But when media enter the picture, the affordance landscape is fundamentally transformed. As Rietveld and Kiverstein (2014) have argued, the landscape of affordances available to human agents is far richer than the motor possibilities emphasized in classical ecological psychology; it encompasses the entire range of sociocultural practices, material artefacts, and skilled engagements through which human forms of life are constituted. An instrument does not merely transmit information from world to agent; it *creates new affordances* that did not exist before. The thermometer affords temperature measurement; the spectrometer affords wavelength discrimination; the telescope affords observation of distant objects. These are not pre-existing possibilities waiting to be accessed; they are possibilities that come into existence with the instrument. Each layer of mediation expands the affordance landscape, and the history of trust-chain deepening is, from this perspective, a history of affordance expansion—the progressive construction of richer and more differentiated possibilities for epistemic action. I distinguish three broad strata.

3.1 Layer 0: Biological calibration and naming

At the most basic level, the agent maps a phenomenal encounter to a linguistic token using only the biologically calibrated instruments of perception and the socially negotiated conventions of language.

"That is a cup." "That patch is red." The semiotic triad is minimal: object, sign, interpreting agent. No fabricated instrument intervenes beyond the agent's own sensory apparatus—itself a product of evolutionary calibration.

The trust required at Layer 0 is trust in two calibrations: the biological calibration of the senses (shaped by natural selection, operating within species-typical parameters) and the social calibration of the naming convention (negotiated through shared use). Both are real forms of calibration, but both are limited. The biological instrument has a narrow range (the eye responds to a thin band of the electromagnetic spectrum), limited precision (colour categorization varies across individuals and languages), and no mechanism for self-correction beyond the slow process of evolutionary adaptation. Wilfrid Sellars' critique of the "Myth of the Given" applies directly: there is no Layer 0 encounter that is epistemically self-certifying (Sellars, 1956). Even the simplest naming act is shaped by the limitations of the biological instrument and the conventions of the linguistic community.

Note that Layer 0 knowledge is predominantly *predicative*: "That is a cup," "That patch is red," "Gold is a metal." These are class-membership assertions—the form of knowledge for which truth-conditional semantics was designed. Layer 0 is where truth conditions are at home.

Layer 0 is not epistemically worthless. It is where our negotiation begins—the baseline from which all deeper mediation departs. But its trust depth is constrained by the biological instrument's range and precision, and by the interpretive variance inherent in socially negotiated conventions.

3.2 Layer 1: Instrumented reading

The limitations of the biological instrument are overcome by interposing a physically fabricated device between phenomenon and token. A thermometer converts a thermal state—which the skin registers only as a vague continuum of "warm" to "hot"—into a precise numerical reading. A spectrometer converts reflected light—which the eye registers as a colour category—into a wavelength. The semiotic triad is extended into a semiotic *cascade*: the phenomenon produces a physical response in the instrument, and that response is then mapped to a linguistic or numerical token by the agent.

Crucially, the instrument is fabricated on the principle of *proportional material response*: devices are built whose physical states change proportionally with the quantity to be measured. A mercury column rises proportionally with temperature. A galvanometer deflects proportionally with current. A spring balance extends proportionally with applied force. This proportionality is the physical ground of trust in the reading (Hacking, 1983). The agent who reads "37.2 °C" trusts not only the convention (Celsius) but the physical calibration of the device.

The transition from Layer 0 to Layer 1 is the material enactment of the shift from predicative to functional knowledge described in Section 2. Where Layer 0 assigns an object to a class ("That is warm"), Layer 1 characterizes a functional relation between physical variables ($T = f(h)$). The instrument converts a predicative impression into a functional reading—and it is this conversion that makes scientific objectivity possible.

The additional mediating layer is a *physical regulation* that reduces the scope for interpretive variance. The spectrometer does not care about the agent's phenomenal experience of redness;

it reports a wavelength. The balance does not care about the agent's estimate of heaviness; it reports a mass. By interposing a physically calibrated mediator, the trust chain becomes partially independent of the interpreting agent's perceptual idiosyncrasies. This partial independence is what we ordinarily call "objectivity"—and the trust-conditions framework makes its mechanism explicit.

3.3 Layer N: Cascaded mediation

Multiple fabricated mediators are chained in series, each translating between physical states—often exploiting entirely different physical principles at each stage. The cascade is not merely "more of the same" but involves qualitatively different theoretical models and physical regularities, each adding calibration depth.

Consider temperature measurement. A clinical mercury thermometer measures body temperature through direct thermal contact and proportional expansion of mercury—a single Layer 1 instrument. But to measure the temperature of molten steel at 1500 °C, direct contact is impossible. Instead, the metallurgist uses a radiation pyrometer: an instrument that measures the intensity and spectral distribution of infrared radiation emitted by the molten surface, then computes temperature using Planck's radiation law. The pyrometer is calibrated against a black-body reference source, which is itself calibrated against the International Temperature Scale. Here the trust chain passes through multiple physical principles (thermal radiation, spectral analysis, black-body theory) and multiple calibration steps, each independently validated. The result is a measurement of far greater range and precision than any contact thermometer could achieve.

The measurement of time tells a similar story. A sundial tracks the shadow cast by the sun—a single proportional response to solar geometry (Layer 1). A pendulum clock exploits the proportionality between pendulum length and period. A quartz oscillator exploits the piezoelectric resonance of a crystal. An atomic clock exploits the frequency of hyperfine transitions in caesium-133 atoms. Each step in this cascade invokes a different physical regularity, a different theoretical model, and a more exacting calibration regime. The result is a progression from seconds-per-day accuracy to fractions of a nanosecond—not because later clocks are "more careful" versions of earlier ones, but because each cascaded layer adds a physically distinct and independently calibrated mediator.

The measurement of length provides a particularly instructive case. A wooden ruler is a Layer 1 instrument: a material artefact marked at known intervals, calibrated against a standard. A laser interferometer, by contrast, splits a laser beam, sends the two halves along different paths, recombines them, and counts the resulting interference fringes. Each fringe corresponds to half a wavelength of the laser light. The trust chain passes through the physics of laser coherence, the optics of beam-splitting, the wave theory of interference, and the metrological definition of the metre in terms of the speed of light. The interferometer achieves sub-nanometre precision not because it is a "better ruler" but because it cascades multiple physically distinct and independently calibrated mediating layers—each grounded in a different physical regularity—into a single measurement.

3.4 The gradient and its epistemological import

The three layers define a gradient from low-mediation knowledge (high phenomenological access, low portability and precision) to high-mediation knowledge (low phenomenological access, high portability and precision). As mediation depth increases, the agent's direct phenomenal contact with the measured phenomenon decreases, but the *precision, portability, and communicability* of the resulting token increases. The sundial is vivid and immediate; the atomic clock is phenomenologically opaque—no one "perceives" a hyperfine transition—but it is precise to fractions of a nanosecond and publicly checkable anywhere on Earth. The wooden ruler is tangible; the laser interferometer is abstract—but it resolves sub-nanometre displacements. In each case, cascaded mediation trades phenomenological immediacy for calibrated precision.

This gradient is not merely descriptive; it is normatively significant. It is also a gradient from *predicative* to *functional* knowledge: from Layer 0 assertions of class membership ("That is warm") through Layer 1 functional readings ($T = 37.2\,^{\circ}\mathrm{C}$) to Layer N cascaded functional chains that relate physical variables across multiple theoretical domains. It explains why STEM knowledge achieves a distinctive form of epistemic reliability—not because scientists are smarter or more rigorous as individuals, but because STEM practices systematically *deepen and calibrate the trust chain* at each mediating layer, interposing physically calibrated mediators between phenomenon and interpretation.

3.5 Procedural calibration: the general case

A crucial observation: physical calibration is not a separate category from procedural calibration. It is a *special case*. Every physically calibrated instrument is the product of a procedure—a sequence of operations (fabrication, testing, comparison against standards) that must be followed correctly for the calibration to hold. What makes a thermometer trustworthy is not the mercury in itself but the *procedure* by which the mercury column was enclosed, graduated, and checked against reference temperatures. Physical calibration is procedural calibration applied to material devices.

The asymmetry between these two species of calibration is the key to understanding why STEM knowledge has a distinctive epistemic profile. Physical calibration is especially powerful because the materially fabricated device does not change its behavior in response to being read: the mercury expands proportionally with temperature regardless of who holds the thermometer or what theory they endorse. This *non-reflexivity* is what makes physical calibration the most variance-reducing strategy available. Procedural calibration in domains where the measured system can react to being measured—social science, law, institutional design—must manage this reflexivity explicitly, through design strategies such as blinding, pre-registration, and institutional separation of researchers from subjects. Both are genuine forms of calibration; they differ not in kind but in the degree to which they must contend with reflexivity. The discussion below should be read with this asymmetry in mind: it is not that procedural calibration is weaker, but that it faces a distinctive challenge that physical calibration largely avoids.

If trust depth is generated by procedural calibration, and physical calibration is one species

of it, then the trust-conditions framework extends naturally to any domain in which disciplined procedures mediate between phenomena and knowledge claims—including the social sciences, law, institutional design, and the humanities.

Alva Noë's notion of "strange tools" (Noë, 2015) is illuminating: theories, conceptual frameworks, classificatory schemas, and institutional protocols are cognitive technologies that reorganize attention, constrain admissible inference, and make certain patterns visible. They function as *instruments of mediation*—not physically proportional, but procedurally calibrated to reduce interpretive variance.

Consider the range of procedurally calibrated mediators across disciplines:

- **Experimental design in social science.** Donald Campbell and Julian Stanley's framework of experimental and quasi-experimental designs (Campbell and Stanley, 1963) is itself a procedural instrument—a systematic methodology for controlling threats to validity (history, maturation, selection, regression to the mean) through design rather than through physical isolation. Randomized controlled trials, as developed in development economics by Duflo, Banerjee, and Kremer, achieve trust depth not through physical apparatus but through procedural innovations—randomization, blinding, pre-registration—that constrain interpretive variance within real-world institutional settings (Duflo et al., 2007).

- **Validated instruments in social science.** Psychometric instruments—standardized questionnaires, diagnostic scales (DSM, ICD), survey instruments—are explicitly calibrated through reliability testing (Cronbach's alpha, test-retest reliability) and validity testing (construct, criterion, content validity). These are instruments in every relevant sense: fabricated mediators, calibrated against standards, whose outputs are publicly checkable and independent of the individual researcher's phenomenal impressions.

- **Legal procedure.** Rules of evidence, cross-examination, appellate review, and the adversarial system itself are procedural mediators that subject claims to structured testing. A legal judgment produced through due process has greater trust depth than one produced by fiat—not because the courtroom contains physical instruments, but because the procedure constrains interpretive variance through institutionalized checks.

- **Ethical review as procedural calibration.** Institutional review boards (IRBs) and ethics committees, operating under principles articulated in the Belmont Report, function as trust-generating mechanisms: they ensure that research procedures respect informed consent, minimize harm, and distribute risks equitably. Ethical clearance is itself a mediating layer that adds procedural depth to the trust chain by subjecting the research design to independent scrutiny before data collection begins.

A deeper point emerges when we consider institutional design itself. Elinor Ostrom's analysis of institutions as structured "action situations" defined by rules—boundary rules, choice rules, information rules, pay-off rules—shows that institutions *are* procedural systems (Ostrom, 2005). When a government changes a policy, it is performing an intervention: altering the procedural

conditions under which social action occurs and observing the effects. When an educational system redesigns its curriculum, it is conducting an experiment in cognitive and social engineering. We tend to think that social sciences "cannot do experiments" because full isolation of variables is impossible in social settings. But this understates the situation. We are *already subjects* of institutional experiments: legal reforms, economic policies, public health interventions, educational programmes are all procedurally calibrated interventions whose effects are observed, measured, and evaluated. The difficulty is not that social experiments are impossible but that they are *ongoing and reflexive*—the experimenters are embedded in the systems they study.

This reflexivity—what Anthony Giddens called the "double hermeneutic," in which social-scientific knowledge feeds back into the social world it describes—is a genuine complication for trust chains in social science. But it is a complication, not a disqualification. The strategies mentioned above—blinding, pre-registration, independent replication, institutional separation—are precisely the procedural means by which reflexivity is managed.

To summarize the asymmetry stated at the outset of this subsection: the distinctive epistemic profile of the natural sciences is not that they alone possess trust chains—all disciplined inquiry does—but that their trust chains include physically calibrated instruments whose non-reflexive proportional responses make interpretive variance especially small. Social sciences, law, institutional design, and the humanities achieve trust depth through procedural calibration that must manage reflexivity explicitly. The difference is real and consequential, but it is a difference in the *character* of calibration, not in whether calibration is present. It does not license a hierarchy; it identifies a spectrum of procedural strategies for externalizing trust, with physical calibration as the most reflexivity-resistant and therefore the most variance-reducing.

Actor-network theory (Latour, 1987) provides a useful lens here: trust chains are sustained by networks of human and non-human actors—researchers, technicians, instruments, protocols, institutions, ethical review boards—all contributing to the stabilization of knowledge claims. The trust-conditions framework is compatible with this insight but adds a normative dimension that actor-network theory deliberately avoids: not all network-stabilizations are equally trustworthy, and the depth and character of procedural calibration provides criteria for evaluating them.

3.6 Negotiation, transparency, and the life of trust chains

Three further observations complete the picture.

First, trust chains are not discovered; they are *negotiated into existence*. Every calibration—biological, physical, or institutional—presupposes the capacity for adjustment: comparing the instrument's output against a standard, noticing discrepancies, and revising accordingly. Without this capacity for negotiation—the ability to say "this does not match; let us adjust"—there is no calibration, only fixed response. The cognitive basis of calibration is the agent's capacity for flexible, revisable coordination with the world and with other agents. A companion monograph develops the cognitive-architectural basis for this negotiability in detail (Nagarjuna and Karnam, 2026); here the point is philosophical: calibration is an inherently social and agentive achievement, not a mechanical

one.

Second, well-functioning instruments—both physical and institutional—become *transparent*. When the thermometer works, the scientist reads the temperature without thinking about the mercury. When the legal system functions, the citizen navigates contracts without thinking about the rules of evidence. When the conventions of a language are shared, speakers communicate without attending to the conventions themselves. Instruments, in this sense, are like well-fitting tools: they recede from attention precisely when they work (Noë, 2015). They become *visible* again when they fail—when the thermometer gives an unexpected reading, when the legal system produces an unjust outcome, when a word is misunderstood. Failure returns the instrument to the space of negotiation, where it can be inspected, recalibrated, or replaced. The life of a trust chain is this oscillation between transparent functioning and negotiated revision.

Third, STEM practices are not exempt from this social embeddedness. They are *sub-cultures*—specific institutional configurations in which the negotiation of trust chains has been stabilized through physically calibrated instruments, mathematically formalized models, and institutionally enforced protocols. The physicist's laboratory is a social institution as much as a courtroom is. What makes it epistemically distinctive is not that it transcends social negotiation but that its institutionalized practices *include* physically calibrated instruments whose non-reflexive proportional responses provide an anchor that is unavailable when all calibration is purely procedural. The natural sciences are socially embedded practices that have learned to externalize trust onto material devices—a specific and powerful strategy within the general space of procedural calibration, not a departure from it.

These observations reveal why the trust-conditions framework is superior to accounts that treat scientific knowledge as a finished product—whether that product is a true proposition (correspondence realism), an empirically adequate theory (constructive empiricism), or a socially stabilized consensus (social constructivism). Trust does not arrive automatically once a functional relation has been asserted or an instrument fabricated. Each layer of the trust chain is a *zone of negotiation*: a site where the claimed functional relation can be challenged, refined, or overthrown through further engagement with the world. The experimentalist who discovers that a thermometer gives inconsistent readings at extreme temperatures is *negotiating* with the functional relation $T = f(h)$—testing its limits, characterizing its deviations, and either correcting the instrument or revising the model. The metrologist who compares thermometers made of different materials and discovers discrepancies is negotiating at a deeper level: probing whether the assumed proportionality holds across physical substrates. Hasok Chang's reconstruction of the history of temperature measurement (Chang, 2004) is precisely a history of such negotiations—iterative cycles of instrument fabrication, comparison, and mutual calibration through which a coherent temperature scale was progressively constructed.

The consequence is that trust-chain deepening is not a one-time achievement but an *ongoing pursuit*. Instruments are revised; calibration protocols are updated; institutional standards are renegotiated in response to new evidence and new technologies. This makes the trust-conditions framework inherently *fallibilist*: no layer of the trust chain is permanently secure, and the very

mechanism that produces objectivity—calibrated negotiation with the world—is the same mechanism that can expose failure. Calibrational realism is therefore not a hard-headed claim that science has achieved objectivity; it is the claim that science is *pursuing* objectivity through the progressive deepening of calibrated trust chains, and that this pursuit is real, cumulative, and normatively evaluable even though it is never complete.

4 Instruments as Arbitrators of Objectivity

The layered framework places instruments at the heart of epistemology. This section develops that claim in full, arguing that instruments are not peripheral aids but the constitutive mechanism through which objectivity is achieved.

4.1 The standard view and its blindness

In standard philosophy of science, instruments appear under the heading of "observation" or "experiment" but are rarely given constitutive epistemic status. The theory-ladenness debate (Hanson, 1958; Kuhn, 1962) asks how theory shapes observation, but the question is framed in terms of the observer's conceptual framework, not in terms of the physical mediators through which observation occurs. Even logical positivism, which placed observation at the foundation of knowledge, treated instruments as transparent extensions of the senses rather than as epistemically constitutive devices (Carnap, 1928).

Ian Hacking comes closest to the position I am defending. His argument from intervention—that we believe in electrons because we can spray them, not because we have a good theory of them (Hacking, 1983)—relocates the ground of scientific realism from theory to practice, from representation to intervention. But even Hacking frames intervention as *evidence for the reality of entities*, not as the *constitutive mechanism of objectivity per se*. The instrument remains, in Hacking's account, a means by which we learn about a mind-independent world; it is not yet the place where objectivity is *produced*.

Davis Baird's concept of "thing knowledge" moves further: instruments carry knowledge in their material form, not merely as vehicles for propositional knowledge (Baird, 2004). A well-calibrated thermometer *knows* something about the relation between temperature and mercury expansion, in a sense that is irreducible to propositional expression. I build on Baird's insight but go further: instruments are not merely repositories of knowledge but the *arbitrators of objectivity*—the physical sites where trust is externalized from subjective experience to publicly checkable material process.

4.2 The instrument as a tiny isolated laboratory

Consider what a well-designed measuring instrument actually is. It is a *sufficiently isolated physical system* in which a controlled physical input produces a proportional physical output. The thermometer isolates a column of mercury from extraneous influences (vibration, air currents, ambient radiation) so that its expansion responds proportionally to one variable: temperature.

The galvanometer isolates a needle from everything except the magnetic field generated by the current to be measured. The spectroscope isolates the diffraction of light from everything except the wavelength distribution of the incoming signal.

Each instrument is, in this precise sense, a *tiny laboratory*: a physically bounded system in which external influences are controlled or excluded so that the phenomenon of interest can express itself through a proportional material response. This is exactly what a full-scale laboratory does—it controls the environment to isolate a phenomenon—but the instrument *miniaturizes and materializes* this control into a portable device.

The philosophical significance is this: the instrument does not passively receive information from the world and relay it to the observer. It *asks the physical world to respond to a physical change*—and the world's response is constrained by the instrument's design to be proportional, repeatable, and independent of who reads the result. The trust is externalized from the interpreting subject to the physical system. Two scientists reading the same thermometer get the same number not because they share a theory, a convention, or a conceptual framework (though they share all of these), but because the mercury's expansion is proportional to temperature and independent of who observes it. The *proportional invariant response of external material* replaces the subjective, publicly inaccessible phenomenal experience as the ground of the measurement.

This is the crucial step. Traditional epistemology locates the ground of knowledge in the subject's experience (empiricism) or in the subject's rational faculties (rationalism). Both grounds are subjective and publicly inaccessible. The instrument externalizes this ground: it converts a privately experienced phenomenon into a publicly readable physical state. In the language of Nagarjuna (1994), measurement *transforms arbitrariness into invariance*: the arbitrary, private, phenomenal impression is replaced by an invariant, public, proportional response. The externalization is not total—someone must still read the instrument, and reading involves convention and training—but it is decisive. The mercury column is not an inner experience; it is a physical object in the shared world, and its length is available for inspection by any competent observer.

4.3 From packaged instruments to extended laboratories

The "tiny isolated laboratory" model works well as a paradigm case—the thermometer, the galvanometer, the balance—but modern instrumentation has moved well beyond single-transducer devices. An MRI scanner, for instance, is not a tiny lab but an *integrated measurement system*: it employs multiple transducers (radio-frequency coils, gradient coils, receivers), each capturing a different aspect of the nuclear magnetic resonance signal. The outputs from these transducers do not pass directly to a human reader; they enter a *computational layer* where multiple rounds of Fourier transformation, spatial encoding, and image reconstruction occur before the final image appears on the screen. The result is not raw data from a single proportional response but a synthesized representation produced by cascading multiple functional relations through computational mediation. Modern mass spectrometers, gene sequencers, and satellite-based remote sensing systems exhibit the same architecture: multiple transducers, computational integration, and layered analysis-and-

synthesis before a terminal output is produced. Each such system is better understood not as a tiny lab but as an *engineered cascade of functional relations*—a trust chain materially packaged in a single device.

This progression from packaged instruments to complex measurement systems is part of a broader spectrum. A *laboratory* is not merely a room containing instruments; it is an extended system in which instruments, people, protocols, and media are organized into a coordinated trust-generating enterprise. As Latour and Pickering have documented in their ethnographic studies of laboratory practice (Latour, 1987), the working laboratory is a site where human agents actively modulate the inputs and outputs of instruments, interpret intermediate results, adjust protocols in response to unexpected findings, and negotiate with one another and with the material world. Reflexivity, which the packaged instrument largely eliminates through its sealed isolation, *re-enters* at the level of the extended laboratory, because intelligent agents are directly participating in the modulation of the measurement process.

This observation extends naturally beyond the science laboratory. Social institutions—peer review panels, regulatory agencies, legal courts, ethical review boards—are *extended laboratories* in this same sense: they organize multiple mediators (human and institutional) into trust-generating configurations. The difference between a physics laboratory and a legal court is not that one has trust chains and the other does not; it is a difference in the *character* of the mediators and the degree to which reflexivity is managed. The physics laboratory includes packaged instruments whose non-reflexive proportional responses provide powerful anchors; the legal court relies primarily on procedural calibration among reflexive agents. But both are sites of organized trust-chain production.

The spectrum thus runs from *packaged instrument* (sealed, non-reflexive, single functional relation) through *complex measurement system* (multiple transducers, computational layers) through *extended laboratory* (instruments plus people plus protocols) through *social institution* (primarily procedural calibration among reflexive agents). Reflexivity is never fully eliminated from science—it is managed, reduced, and constrained, but the human agents who design experiments, interpret data, and revise models are always part of the extended system. This is the point that the sociology of science has long insisted upon, and the trust-conditions framework concedes it fully. What the framework adds is a normative criterion: the degree to which reflexivity has been constrained by physically calibrated mediators determines the trust depth of the resulting knowledge claims. The natural sciences are not reflexivity-free; they are reflexivity-*managed* through an especially powerful combination of packaged instruments and procedural discipline.

A final extension concerns the introduction of *artificial agents* into the laboratory. As AI systems are increasingly embedded in measurement pipelines—automating data collection, performing pattern recognition, generating experimental designs—the trust chain deepens further, but so do the zones of possible failure. An AI system that processes telescope data to identify exoplanet candidates adds a computational layer to the trust chain, but it also introduces new questions of calibration: how was the algorithm trained? on what data? with what biases? The trust-conditions framework

treats AI-augmented measurement in the same terms as any other cascaded mediation: each layer must be evaluated for calibration, transparency, reproducibility, and institutional stabilization. The addition of artificial agents does not change the framework; it extends its application to a domain where the questions of trust are especially urgent and especially consequential.

4.4 Mathematics as the language of possible functional relations

The functional-relations framework developed in Section 2 yields a reinterpretation of the relationship between mathematics, instruments, and the physical world. When we write $T = f(h)$ for a mercury thermometer—relating temperature T to mercury column height h—what is the mathematical function *about*? The standard answer is that it describes a relation between temperature and mercury column height in the physical world. But this is ambiguous. It might mean: (a) a metaphysical relation between the property "temperature" and the property "mercury expansion" considered in themselves; or (b) a description of how a specific, physically isolated instrument converts one physical state into another—a functional relation physically instantiated in a calibrated device.

The trust-conditions framework favors interpretation (b), and the reason is now clear. Mathematics is the formal language of *possible* functional relations: it describes the space of all ways in which variables might be related—linear, quadratic, exponential, periodic, stochastic. This is what mathematicians have been building for centuries: an ever-expanding vocabulary for expressing possible functional worlds, with their internal coherence, symmetry, and beauty. The function f in $T = f(h)$ is one such possible relation. The instrument is where that possible relation is *physically instantiated*: the mercury thermometer converts f from a mathematical possibility into an actual, calibrated, physically operating device. Science is the enterprise of mapping possible functional relations (mathematics) to actual ones (through instruments), and the trust chain is what validates the mapping.

This reinterpretation deflates longstanding puzzles about the "unreasonable effectiveness of mathematics" in the natural sciences (Wigner, 1960). Mathematics is effective not because it mysteriously mirrors the deep structure of reality, but because it is the language of functional relations, and instruments are physically instantiated functional relations. The effectiveness is not unreasonable; it is the natural consequence of the fact that instruments are built to realize the very relations that mathematics describes. The mathematical model maps to the *instrument*—the tiny isolated laboratory—because the instrument was fabricated to embody a specific functional relation that the model expresses. Of course, instruments work because the physical regularities they exploit are real—mercury really does expand proportionally with temperature. But the mathematical model's domain of applicability is anchored to the instrument, and it extends to the world only insofar as the instrument's physical coupling to the phenomenon is well-understood and well-calibrated.

Whether this account extends to all of mathematical physics—to the differential equations of general relativity or the Hilbert-space formalism of quantum mechanics, which are not in any obvious sense models of specific instruments—remains an open question. But the direction is

suggestive. Even in theoretical physics, the equations express functional relations between variables (curvature and energy-momentum, state vectors and observables), and the empirical authority of these equations ultimately rests on trust chains that pass through calibrated instruments. The trust-conditions framework does not require that every mathematical model be a model of an instrument; it requires that the epistemic authority of mathematical models derives from the trust chains through which they are validated, not from a direct correspondence between the equation and a mind-independent process.

4.5 Calibration as trust externalization

How does an instrument become a node in a trust chain? Through *calibration*: the process by which an instrument's response is checked against a standard and its deviations are characterized, corrected, or bounded. Calibration is the institutional mechanism for ensuring that the physical proportionality holds within known tolerances.

The fabrication and calibration of a reliable measuring instrument involves three steps:

1. **Exploiting a physical regularity**: discovering that mercury expands proportionally with temperature, that electrical resistance changes proportionally with strain, that light bends proportionally with refractive index. This is an empirical engagement with the world—a discovery of physical proportionality.

2. **Packaging that regularity into a device**: enclosing the mercury in a calibrated tube, wiring the strain gauge into a circuit, grinding the lens to specification. The fabrication converts an empirically discovered physical sensitivity into a readable medium.

3. **Calibrating the device against a standard**: marking the tube at known temperatures, comparing the gauge against a reference load, checking the lens against a known spectrum. Calibration is the act that converts a physical artefact into a metrological instrument—an object whose readings are publicly meaningful within a community of practice.

The instrument, once calibrated, is *materialized trust*: its physical construction guarantees, within known tolerances, that the token it produces corresponds proportionally to the quantity it measures. The trust has been externalized from the interpreting subject to the physical device and the institutional standard against which it was calibrated. This is why Hacking was right to insist that *intervention*—not observation, not theory—is the wellspring of scientific realism. But the trust-conditions framework specifies *why* intervention is convincing: because instruments are physically calibrated mediators that externalize trust from subjective experience to proportional invariant material response.

Nancy Cartwright's concept of "nomological machines" is the closest existing notion to the "tiny isolated laboratory" proposed here (Cartwright, 1999). Cartwright argues that the laws of physics hold only within the shielded, controlled environments of laboratories and engineered systems—they are local, not universal. She uses this to cast doubt on the universality of physical laws and, by

extension, on scientific realism. I agree with Cartwright's description of instruments and laboratories as shielded, isolated systems. But I draw the opposite conclusion: precisely *because* instruments are isolated systems whose proportional responses are reliable within known tolerances, they provide the strongest available basis for epistemic trust. The locality of physical laws is not a defect but a feature: it means that the instrument's response is well-characterized, bounded, and trustworthy. Calibrational realism does not require universal laws; it requires locally reliable instruments.

Hasok Chang's meticulous reconstruction of how temperature scales were historically constructed—through iterative cycles of instrument fabrication, calibration, and mutual comparison—provides a concrete demonstration of how trust chains are built (Chang, 2004). Chang shows that there was no pre-existing "true temperature" against which thermometers could be checked. Instead, different thermometric substances (mercury, alcohol, air) were calibrated against each other and against fixed points (freezing and boiling of water) in a bootstrapping process that progressively stabilized a coherent temperature scale. This is trust-chain construction in action: not the discovery of a pre-existing fact, but the progressive externalization of trust through physically calibrated mediators.

5 Discovery and Invention Entangled

The trust-conditions framework dissolves one of the oldest dichotomies in the philosophy of science: the distinction between discovery and invention. If scientific knowledge is grounded through fabricated instruments that modulate reality under controlled conditions, then what we "discover" is inseparable from what we invent. This entanglement is not a defect of the framework but one of its most consequential implications.

5.1 Perception is action; observation is intervention

The enactive tradition in cognitive science holds that perception is not the passive reception of information but a form of skilled action: the organism explores its environment through movement, and what it perceives is constituted by the sensorimotor contingencies of that exploration (Noë, 2004). To see is to act; to perceive is to engage. If this is correct for biological perception, it applies *a fortiori* to instrumented observation. The scientist who reads a spectrometer is not passively receiving data from the world; she is actively intervening—preparing a sample, configuring the instrument, selecting operating parameters, calibrating against standards, interpreting the output. Instrumented observation is *engineered action*: a sophisticated modulation of the conditions under which phenomena manifest.

Calibration, in particular, is not a one-time setup but a *constant mediation process*. The instrument must be maintained, rechecked, adjusted; the protocol must be followed; the standards must be traced back through the metrological chain. The observation is not a snapshot of a pre-existing state of affairs but the product of an ongoing, actively maintained intervention. This is the enactive insight applied to scientific practice: just as the organism's perception is constituted by

its action, the scientist's observation is constituted by the instrumental intervention that makes it possible.

The consequence is that "discovery" through instruments is never purely passive. Every instrumented observation involves engineering the conditions under which a phenomenon can manifest. We must design and build the cloud chamber to "see" subatomic particles. We must fabricate the interferometer to "detect" gravitational waves. We must construct the gene sequencer to "read" a genome. In each case, the entities and processes are real—they are not conjured by the instrument—but they manifest only under the specific conditions that the instrument creates. Discovery and fabrication are *cognitively coupled*: we engineer our actions in increasingly sophisticated ways in order to perceive entities and processes that are otherwise inaccessible to the biological instrument.

5.2 Models as structures, instruments as tokens

The entanglement of discovery and invention becomes sharper when we consider the nature of scientific models. The non-statement or structuralist view of scientific theories—developed by Patrick Suppes, Frederick Suppe, and Wolfgang Stegmüller (Suppes, 1962; Suppe, 1977; Stegmüller, 1976)—holds that a scientific model is not a statement or proposition but a *structure*: a set-theoretic or mathematical entity that specifies relations among quantities. A model, on this view, is a complex predicate, not a sentence. It is neither true nor false in itself; it *applies* or *fails to apply* to a physical system.

This convergence with the functional-relations argument of Section 2 is not accidental: a model that specifies relations among quantities is precisely a specification of *functional relations*—and the structuralist view's insistence that models are not propositions but structures mirrors the distinction between predicative and functional knowledge. The instrument-as-minilab, mapped tightly to a model, functions as a *physical realization* of the model's structure—its functional relations physically instantiated. The thermometer realizes the model $T = f(h)$: the mathematical structure (proportional relation between temperature and mercury column height) is physically instantiated in the device. The instrument produces a *token*—a specific numerical reading—that describes the world, but only under the conditions that the instrument creates and the model specifies. The token is neither a pure discovery (it required fabrication) nor a pure invention (it is constrained by the world's response). It is a *procedurally constructed descriptor*: a product of the entanglement of fabrication and manifestation.

Consider the proportionality between temperature and the expansion of mercury. Is this proportionality discovered or invented? From one angle, it is discovered: mercury really does expand proportionally with temperature, and this is a fact about the physical world that the instrument-maker exploits. From another angle, it is invented: the instrument-maker *selects* this regularity, *isolates* it from confounding influences, *packages* it into a device, and *calibrates* it against a standard. The proportionality is discovered *within* a captured physical system—a system that was engineered precisely to make that proportionality manifest cleanly. The discovery required the

invention, and the invention exploited the discovery. From the lens of trust, they are inseparable: the trust chain runs through both.

5.3 Intervention as the condition of manifestation

The deeper point is that the entities and processes that science studies do not exist independently of the conditions under which they manifest. This is not an idealist claim that reality is mind-dependent; it is a claim about *access*. An electron does not appear in a cloud chamber because the observer wills it into existence; it appears because the cloud chamber creates the specific physical conditions (supersaturated vapour, ionization trail, photographic recording) under which the electron's passage becomes visible. The electron is real; but its manifestation as an observable track is an achievement of instrumental engineering.

To understand the conditions of manifestation is not different from understanding how to intervene. Hacking's insight that intervention grounds realism is, on this reading, an insight about *manifestation conditions*: we are realists about electrons because we can create the conditions under which electrons manifest in reproducible, calibrated ways. Intervention is modulating reality—setting up the conditions under which specific phenomena become accessible—and this modulation is itself an act of fabrication. The distinction between "discovering what exists" and "inventing the means of access" collapses: the means of access *are* the conditions of manifestation, and understanding those conditions *is* the scientific knowledge.

The entanglement of discovery and invention has a further consequence: the intellectual contribution of instrument-makers, metrologists, and experimental designers is not merely "technical"; it is epistemically constitutive. They do not merely facilitate discovery; they construct the conditions under which discovery becomes possible.

6 Simulations as Designing Instruments

The entanglement of discovery and invention finds its most explicit expression in computer simulations. A simulation instantiates a model's structure in a computational medium, runs it under specified conditions, and observes the outcome. It is, in a precise sense, a *virtual minilab*: a controlled environment in which the model's behavior can be tested, calibrated, and refined before the costly step of physical experimentation. The widespread use of simulations in contemporary science—from climate modelling to drug design to materials engineering—confirms and extends the trust-conditions framework.

6.1 From thought experiments to virtual laboratories

There is a continuum of mediated exploration that runs from the least to the most materially instantiated:

1. **Thought experiments.** The scientist manipulates a model entirely in imagination—Galileo's

falling bodies, Einstein's elevator, Maxwell's demon. The "instrument" is the thinker's capacity for structured counterfactual reasoning. Trust depth is minimal: it rests on the coherence of the reasoning and the plausibility of the idealizations.

2. **Paper-and-pencil calculation.** The model is instantiated in mathematical notation and solved analytically. The medium is inscription—marks on paper governed by formal rules. Trust depth increases: the calculation is inspectable, reproducible, and governed by mathematical conventions that are publicly checkable. But the model's behavior is accessible only where analytic solutions exist.

3. **Numerical simulation.** The model is instantiated in computer code and solved numerically. The medium is now a computational device—itself a physically fabricated instrument—executing algorithmic procedures. Trust depth increases further: the simulation can explore regions of parameter space inaccessible to analytic methods and can be run repeatedly under varied conditions.

4. **Virtual laboratory.** The simulation is embedded in an interactive environment that allows the scientist to manipulate variables, visualize outcomes, and iterate in real time. The virtual lab is a fully procedurally calibrated minilab: its behavior is governed by the model's structure, its interface is designed for systematic exploration, and its outputs are publicly inspectable.

5. **Physical experiment.** The model's predictions are tested against the behavior of a materially fabricated instrument—the physical minilab. This is where computational trust meets physical trust, and where the trust chain reaches its deepest calibration.

This continuum reveals that the transition from model to experiment is not a leap from theory to practice but a *gradual deepening of trust* through progressively more materially instantiated media. The simulation occupies a pivotal position: it is where the model's structure is tested under controlled conditions *before* the commitment to physical fabrication. In engineering practice, this is routine: aircraft are simulated before they are built; drug candidates are screened computationally before clinical trials; bridge designs are stress-tested in finite-element models before construction begins. The simulation is a *trust-depth test*—a way of assessing whether the model's structure is adequate before investing in the physical minilab.

6.2 Verification and validation as trust calibration

The epistemology of simulation is organized around two concepts that map directly onto the trust-conditions framework. *Verification* asks whether the computational implementation correctly instantiates the mathematical model: does the code do what the model says? *Validation* asks whether the model adequately represents the target phenomenon: does the model capture the relevant features of reality? (Winsberg, 2010).

In trust-conditions terms, verification is *procedural calibration of the computational instrument*: ensuring that the algorithm, the numerical methods, and the code correctly realize the model's

structure. Validation is *calibration of the model against the phenomenon*: ensuring that the model's structure maps adequately to the physical system it purports to represent. Together, verification and validation constitute the trust chain of the simulation, and their adequacy determines the trust depth of the simulation's outputs.

Wendy Parker has argued that simulations can be embedded in measurement practices—through data assimilation, for example, where observational data and model outputs are iteratively combined to produce estimates of quantities that neither data nor model could determine alone (Parker, 2020). In the trust-conditions framework, data assimilation is a cascaded mediation (Layer N) in which computational and physical trust chains are interwoven: the simulation calibrates the data, and the data calibrates the simulation, in a bootstrapping process analogous to Chang's account of how temperature scales were historically constructed (Chang, 2004).

6.3 Simulations confirm the minilab thesis

Paul Humphreys has argued that computer simulations constitute a genuinely new mode of scientific inquiry—neither pure theory nor pure experiment, but a "third way" that combines features of both (Humphreys, 2009). Margaret Morrison has shown that simulations and experiments often measure physical characteristics through the same model-based strategies (Morrison, 2015). These philosophical analyses converge on a conclusion that the trust-conditions framework makes explicit: the simulation is a *virtual minilab*.

Like the physical minilab, the simulation isolates a phenomenon (the model's structure) from extraneous influences (unmodelled variables, noise, physical limitations of the laboratory). Like the physical minilab, it produces tokens (numerical outputs) that describe the system's behavior under controlled conditions. Like the physical minilab, it requires calibration (verification and validation) to be epistemically trustworthy. The difference is that the simulation's medium is computational rather than material: it instantiates the model in code rather than in mercury, glass, or silicon.

This difference matters for trust depth. The simulation's trust chain passes through the model's mathematical structure, the algorithm's implementation, the hardware's physical computation, and the conventions of numerical analysis—each a calibrated mediating layer. But it does not pass through the physical regularity that the model purports to describe. The simulation tells us what the model *predicts*, not what the world *does*. The physical minilab is needed to close the trust chain: to check whether the model's structure, tested and refined in simulation, actually maps to the proportional material response of a physical system.

This is why simulations are best understood as *instruments for designing instruments*. They test the trust depth of a model—its coherence, its sensitivity to parameters, its behaviour under extreme conditions—before that model is instantiated in a physical device. The simulation is a trust-depth probe: it explores the model's structure in virtual space, identifying where the model is robust and where it is fragile, so that the physical minilab can be designed with informed confidence. The continuum from thought experiment to simulation to physical experiment is, in trust-conditions terms, a continuum of trust-chain deepening—from purely procedural (thought experiment) through

computationally calibrated (simulation) to physically calibrated (experiment). Each step adds a mediating layer; each layer adds trust depth; and the full chain, from imagination to instrumented measurement, is the epistemic infrastructure of modern science.

7 Calibrational Realism

The foregoing analysis yields a distinctive position in the realism debate. I call it *calibrational realism* and develop it here by stating the position, then defending it against the main alternatives.

7.1 The position stated

Calibrational realism holds that the epistemic success of science is explained by the *progressive externalization of trust conditions*: from subjective phenomenal experience (Layer 0) to simple instruments (Layer 1) to cascaded instrumented practice (Layer N). At each stage, trust is passed from the interpreting subject to a proportional invariant basis of external material. What we call "objective" knowledge is knowledge whose trust chain is deep enough and well-calibrated enough that individual interpretive variance becomes negligible.

This is a form of realism because it affirms that science makes *genuine epistemic progress*. Trust chains really do get deeper and more physically calibrated over the history of science. The progression from naked-eye observation to microscopy, from touch-based assessment to thermometry, from visual estimation to spectroscopy, from local measurement to satellite-mediated global monitoring—each step is a real deepening of the trust chain, not merely a change in fashion or convention. The claim is not that science converges on a "final theory" or mirrors a mind-independent reality; it is that science progressively externalizes the conditions of trust, making knowledge less dependent on individual phenomenal experience and more dependent on physically calibrated material processes.

This is not naive realism, which claims that scientific theories correspond to a mind-independent world *as it is in itself*. Calibrational realism makes no such claim. It claims that the *trust chains* through which knowledge is produced are progressively deepened and calibrated, and that this deepening constitutes genuine epistemic progress. The "reality" that calibrational realism is realist about is not a noumenal world behind appearances but the *reliability of physically calibrated trust chains*—a reliability that is itself empirically testable, improvable, and historically demonstrated.

7.2 Constructive empiricism and the observable/unobservable binary

Bas van Fraassen's constructive empiricism distinguishes the *observable* from the *unobservable* and recommends belief only in a theory's empirical adequacy—its agreement with observable phenomena—while remaining agnostic about unobservable entities and processes (van Fraassen, 1980). The observable/unobservable distinction is drawn by reference to the unaided human senses: what can be seen, heard, or felt without instruments is observable; everything else is unobservable.

The trust-conditions framework dissolves this binary. "Observable" is simply Layer 0: phenomenal encounter mediated only by convention and the agent's sensory apparatus. "Unobservable" is

everything accessed through Layer 1 or deeper: entities detected, measured, or inferred through instruments. Van Fraassen's recommendation to believe only in the observable is, in trust-conditions terms, a recommendation to trust only the shallowest layer of the mediation gradient. But why? The whole point of instrumented measurement is to *deepen* the trust chain beyond the unreliable vagaries of phenomenal impression. The spectrometer is more trustworthy than the naked eye precisely because it adds a physically calibrated mediating layer that reduces interpretive variance.

Van Fraassen might reply that instruments are theory-laden: we must already believe in the theory of optics to trust a spectrometer. But the trust-conditions framework turns this objection on its head. Theory-ladenness, far from being a source of epistemic vulnerability, is a source of trust depth. A radiation pyrometer that invokes Planck's law to measure the temperature of molten steel is *more* theory-laden than a mercury contact thermometer—and *more* trustworthy for that range. An atomic clock that exploits hyperfine transitions in caesium-133 is *more* theory-laden than a pendulum clock—and incomparably more precise. Each additional theoretical layer in the cascade is itself a calibrated mediator: a model that has been independently validated against other instruments and other physical regularities. Theory-ladenness adds mediation layers, and each such layer, when independently calibrated, adds trust depth.

Moreover, the objection applies equally to Layer 0. The eye is "theory-laden" by evolution: its response is shaped by millions of years of selective pressure, and its outputs are shaped by neural processing that the observer cannot inspect. If theory-ladenness is grounds for agnosticism, it applies to all layers equally; if it is not, it applies to none. The trust-conditions framework cuts through this impasse by replacing the binary observable/unobservable with a graduated assessment of trust depth. The question is not *whether* to trust instruments but *how deeply* each instrument has been calibrated and validated—a question that admits of empirical investigation rather than philosophical stipulation.

7.3 Entity realism and the scope of intervention

Hacking's entity realism vindicated: his argument that intervention grounds realism is explained by the trust-conditions framework as follows. When we "spray electrons" using an electron gun to investigate a target, we are using an instrument whose operation depends on the existence and behavior of electrons. The trust chain passes through physically calibrated devices whose design presupposes and exploits the very entities whose reality is in question. Intervention grounds realism because it *deepens the trust chain*: the instrument's operation is itself a physical test of the entity's existence.

But Hacking's entity realism is too narrow in one respect: it applies only to entities we can manipulate. Calibrational realism extends naturally to entities and processes we cannot directly manipulate but can access through cascaded instrumented observation—distant galaxies observed through telescopes, geological events inferred from isotope ratios in rock strata, gravitational waves detected through interferometry. In each case, the trust chain passes through physically calibrated instruments, and the reliability of the knowledge claim depends on the depth and calibration of

that chain. Manipulation is the strongest form of trust-chain deepening, but it is not the only form.

7.4 Structural realism and the grounding of structure

Structural realism holds that science is right about the *structure* of reality—the relational, mathematical structure expressed in successful theories—even if it is wrong about the nature of the entities that instantiate that structure (Worrall, 1989). In its ontic version, developed by Ladyman and French, structure is not merely epistemically privileged but *all there is*: the world is a structure of relations, and individual objects are derivative on those relations (Ladyman and Ross, 2007; French, 2014). The structural realist's core insight—that what survives theory change is relational structure, not entity-descriptions—aligns naturally with the trust-conditions framework's emphasis on functional relations over predicative assertions. Both views hold that the epistemically significant content of science is relational rather than classificatory.

But calibrational realism departs from structural realism in a crucial respect: it asks how relational structure is *grounded*. For the structural realist, the structure that matters is the mathematical structure expressed in theoretical equations—the group-theoretic symmetries, the state-space topologies, the Hamiltonian formalism. The question of how we come to *know* this structure is typically left to a general theory of scientific inference. Calibrational realism makes this question central. The relational structures that science is realist about are not, in the first instance, the structures of theoretical equations considered abstractly; they are the *physically calibrated functional relations* that instruments exploit and that calibration validates. The proportional invariants—mercury expansion with temperature, galvanometer deflection with current, spectral dispersion with wavelength—are relational structures, but they are structures that are *materially instantiated* in calibrated devices, not merely represented in mathematical formalism.

This grounding claim addresses a persistent challenge for structural realism: the problem of identifying which mathematical structures correspond to physical structure rather than to artefacts of representation (French, 2014). The trust-conditions framework offers a criterion: the structures that deserve realist commitment are those that are physically instantiated in calibrated instruments and validated through reproducible trust chains. A mathematical symmetry that appears in a theoretical formalism but cannot be linked to any calibrated instrumental response has a weaker claim to realist significance than one that is embodied in the proportional response of a physical device. The instrument provides an *empirical anchor* for structural realism—a way of distinguishing physically significant structure from merely formal structure.

A further point of contact concerns what structural realists call the "Newman problem": that purely structural descriptions of the world are trivially satisfied by any domain of the right cardinality, so structural realism threatens to become vacuous unless the structures are constrained by more than their mathematical form (Worrall, 1989). Calibrational realism provides exactly such a constraint: the structures are constrained by the *physical calibration* of the instruments that instantiate them. The proportional invariants that instruments exploit are not arbitrary structures; they are selected, isolated, and calibrated against standards. The Newman problem dissolves when structure is

grounded in physical calibration rather than abstracted from theoretical representation.

The dialogue between structural realism and calibrational realism is thus potentially fruitful in both directions. Structural realism gives calibrational realism a richer vocabulary for discussing the relational content of science. Calibrational realism gives structural realism the grounding it needs—a material, instrumental, calibrated basis for the structures it identifies as epistemically significant.

7.5 Instrumentalism and operationalism revisited

The instrumentalist says theories are useful tools for organizing experience, not descriptions of reality (Duhem, 1954). The operationalist says the meaning of a concept is exhausted by the operations that measure it (Bridgman, 1927). Calibrational realism shares with both the emphasis on practice, intervention, and measurement. But it denies the instrumentalist's conclusion that theories are "merely" useful: the progressive deepening of trust chains constitutes genuine epistemic progress, not merely increasing pragmatic convenience. And it goes beyond operationalism by explaining *why* operational definitions work: because the operations involve physically calibrated instruments whose proportional material responses externalize trust from subjective experience. Operationalism correctly identifies the importance of measurement operations but does not explain what makes some operations epistemically better than others. The trust-conditions framework provides the answer: operations that involve deeper, more physically calibrated trust chains are epistemically superior.

Percy Bridgman was right that the meaning of "length" is tied to the operations by which length is measured. But Bridgman did not explain why a laser interferometer is a better length-measuring operation than a wooden ruler, or why a wooden ruler is better than a paced-off estimate. The trust-conditions framework does: each instrument adds a physically calibrated layer that reduces interpretive variance, and the laser interferometer represents a deeper trust chain than the ruler, which represents a deeper chain than the estimate. The graduated character of trust depth is what operationalism misses and calibrational realism provides.

8 Objections and Replies

Three objections merit explicit consideration.

Objection 1: Is the predicative/functional distinction really as sharp as claimed? A critic may object that the distinction between predicative and functional knowledge is overdrawn. Many scientific claims *are* predicative: "water is H_2O," "the electron has spin $\frac{1}{2}$," "the patient has influenza." If science still uses predicative assertions, why claim that functional relations are its distinctive form?

The reply is that the distinction is one of direction and emphasis, not a binary. Predicative assertions remain a part of scientific discourse—taxonomy, classification, and diagnosis are real

forms of scientific knowledge. But consider what underlies even the predicative-sounding claims just cited. "Water is H_2O" is grounded in functional relations: the mass ratios of hydrogen and oxygen in electrolysis, the spectral signatures of O-H bonds, the stoichiometric proportionalities of chemical reactions. "The electron has spin $\frac{1}{2}$" is grounded in functional relations: the proportionality between angular momentum and magnetic moment measured in Stern-Gerlach experiments. The predicative form ("X is Y") is the surface; the functional relations that ground it are the epistemic substance. As Cassirer documented, the trajectory of modern science has been consistently from substance-concepts toward function-concepts (Cassirer, 1923). The trust-conditions framework does not deny that predicative assertions occur in science; it argues that the *grounding* of those assertions—the process by which they become warranted—passes through functional relations validated by calibrated instruments. It is this grounding process, not the surface grammar of scientific statements, that the framework addresses.

Objection 2: Are the four dimensions jointly sufficient? The paper specifies four dimensions of trust conditions: physical calibration, transparency, reproducibility, and institutional stabilization. A critic may ask whether these are jointly sufficient for epistemic adequacy, or merely a useful but incomplete heuristic. What about cases of reproducible but systematically biased measurement? What about opaque but highly reliable instruments?

The reply is that the four dimensions are *evaluative criteria*, not a formal sufficiency condition. They provide the vocabulary for assessing trust chains, not an algorithm that outputs a binary verdict. This is by design: the trust-conditions framework replaces binary distinctions (objective/subjective, true/false, scientific/unscientific) with graduated assessments along multiple dimensions. These dimensions can conflict—an instrument may be highly reproducible but poorly calibrated, or well-calibrated but institutionally unsupported—and resolving such conflicts requires judgment informed by context. The framework provides the dimensions along which that judgment operates; it does not eliminate the need for judgment. The analogy is with Kuhn's criteria for theory choice (accuracy, consistency, scope, simplicity, fruitfulness): individually insufficient, collectively illuminating, and requiring interpretive application (Kuhn, 1977).

Objection 3: Is calibrational realism genuinely new? Calibrational realism draws on entity realism (Hacking), social objectivity (Longino), measurement theory (Chang, Tal), nomological machines (Cartwright), and thing knowledge (Baird). A critic might see the position as a novel recombination of existing insights rather than a genuinely distinct contribution.

The reply is that the recombination is itself the contribution, and it yields claims that none of the component positions make individually. Specifically, calibrational realism is distinctive in four respects that form an integrated package: (i) it takes trust-chain depth as the *unit* of epistemic evaluation, not merely as evidence for or against realism; (ii) it identifies physical calibration as a special case of procedural calibration, unified by the criterion of reflexivity-resistance; (iii) it defines objectivity as negligible interpretive variance produced by calibrated mediation, rather than as correspondence, consensus, or empirical adequacy; and (iv) it grounds scientific progress

in trust-chain deepening rather than in theory change, empirical accumulation, or puzzle-solving capacity. No existing position in the realism debate combines all four. The value of a synthesis lies not in whether each component is novel but in whether the combination illuminates something that the components individually do not.

9 Implications

9.1 Objectivity reframed

The traditional framing places objectivity on the side of agent-independent truth—propositions that hold regardless of who entertains them—and subjectivity on the side of agent-dependent experience. On this view, science is objective because it discovers truths about a mind-independent world; perception is subjective because it depends on the perceiver's constitution.

The trust-conditions framework shows why this framing is misleading. *All* meaning passes through interpreting agents: no knowledge claim is agent-free. Even the most rigorous scientific measurement terminates in an agent's interpretive act—a human reads the instrument, records the number, enters it into a model. The question is not whether an agent is present (one always is) but how deep and how physically calibrated the trust chain is between the phenomenon and the agent's interpretive act.

"Objectivity," on this account, is not the absence of a subject but the *depth and physical calibration of the trust chain* that makes individual interpretive variance negligible. When two scientists independently measure the melting point of aspirin and agree on 135 °C, their agreement is not evidence that they have accessed a subject-independent truth. It is evidence that the trust chain—fabricated instrument, calibration protocol, shared convention, physically proportional response—is deep enough and well-calibrated enough that the scientists' phenomenological differences do not matter. The "objectivity" of the measurement is a property of the trust chain, not a property of the proposition considered in isolation.

Lorraine Daston and Peter Galison's historical study of objectivity (Daston and Galison, 2007) shows that the concept has meant different things in different periods: truth-to-nature, mechanical objectivity, trained judgment. The trust-conditions framework provides a normative account that can evaluate these historical ideals. Mechanical objectivity—the nineteenth-century insistence on letting instruments record phenomena without human intervention—is, in trust-conditions terms, an attempt to deepen the trust chain by removing a Layer 0 interpretive step. Trained judgment—the twentieth-century acknowledgment that human expertise is needed to interpret instrument outputs—is, in trust-conditions terms, a recognition that trust chains always terminate in an agent's interpretive act. The framework does not privilege one historical ideal over another but provides the criterion by which all can be assessed: how deep and how well-calibrated is the trust chain?

9.2 The history of science as trust-chain deepening

The trust-conditions framework suggests a reinterpretation of the history of STEM as the progressive deepening of trust chains through instrument fabrication. The trajectory is clear: from naked-eye observation to optical microscopy, from touch-based temperature assessment to mercury thermometry to platinum resistance thermometry to infrared spectroscopy, from visual estimation of distance to triangulation to radar to laser ranging to satellite-based geodesy. Each step is a material innovation—an act of engineering—that converts embodied sensitivities into publicly checkable, physically calibrated measurements.

This is not a replacement for existing historiography. Kuhn's paradigm shifts (Kuhn, 1962), Lakatos's research programmes (Lakatos, 1978), and Laudan's research traditions (Laudan, 1977) all capture important dimensions of scientific change. But the trust-conditions framework offers a complementary lens that foregrounds the material and instrumental dimension—a dimension that is often acknowledged but rarely given systematic philosophical treatment. The growth of STEM is, in large part, a history of trust-chain deepening: the progressive externalization of trust from subjective phenomenal experience to proportional invariant bases of external material.

This reinterpretation also offers a new lens on the role of mathematics in science. Mathematics, on the trust-conditions account, is the formal language of *possible* functional relations—differential equations for continuous dependencies, probability for stochastic regularities, geometry for spatial structure, algebra and group theory for relational invariants. What mathematicians have built over centuries is a vast and beautiful space of possible functional worlds; what scientists do is map these possible worlds to the actual one, through calibrated instruments that physically instantiate functional relations. Logic complements mathematics as the organizer of consistency within the system of trust chains, detecting contradictions and preserving inferential coherence. Together, mathematics and logic are the formal infrastructure of the trust-chain system, not themselves trust chains but the languages in which trust chains are described, evaluated, and extended.

9.3 Engineers and technologists as epistemic agents

The trust-conditions framework has a striking consequence for the sociology of knowledge production: it brings laboratory staff, instrumentation engineers, and technologists to the centre of the epistemic enterprise. In the conventional picture, the scientist—the theorist, the principal investigator—is the epistemic agent; the technician who builds and maintains the instruments is a support worker, epistemically peripheral. The trust-conditions framework inverts this hierarchy, or at least flattens it. If trust is externalized through physically calibrated instruments, then the people who design, build, calibrate, and maintain those instruments—the engineers, the metrologists, the laboratory technicians—are doing the epistemic heavy lifting. They are the ones who construct the "tiny isolated laboratories" through which trust is materialized. Every precisely ground lens, every carefully calibrated sensor, every meticulously maintained reference standard is a contribution to the epistemic strength of the science that relies on it.

This revaluation is not merely symbolic. It has practical implications for how we fund, credit,

and organize scientific work. If the trust chain is where the epistemic action is, then the material infrastructure of science—the instruments, the calibration facilities, the technical workforce—deserves the same intellectual attention and institutional support as the theoretical superstructure. The persistent tendency to valorize theoretical over experimental work, and to treat engineering and material fabrication as epistemically subordinate to "pure" science, is, on the trust-conditions account, a systematic misidentification of where scientific objectivity actually comes from.

9.4 Markets and consumers as epistemic participants

The trust-conditions framework reveals an often-overlooked dimension of epistemic validation: the role of markets, trade, and consumers in testing and deepening trust chains. STEM practices do not produce a single instrument; they produce *many copies*. Companies fabricate, package, price, and distribute instruments to a world of consumers who put them to use in contexts that no laboratory can fully anticipate. A clinical thermometer is tested not only in the metrological laboratory but in thousands of clinics and households; a GPS receiver is validated not only by the satellite engineers but by millions of navigators; a pharmaceutical compound is assessed not only in controlled trials but in the post-market experience of patients and physicians.

Consumers participate in this epistemic process, often without explicit awareness. By purchasing and using instruments, they subject those instruments to a distributed, large-scale, long-duration reliability test. The feedback loop—through product returns, warranty claims, adverse-event reports, customer reviews, and the simple fact that reliable instruments generate repeat purchases while unreliable ones do not—constitutes a form of epistemic validation that complements laboratory calibration. Markets are, in this sense, extended testing grounds: they provide wider and more varied conditions of use than any single laboratory can replicate, and the commercial success or failure of an instrument over time reflects, among other things, the robustness of the functional relations it embodies.

This observation reframes the relationship between industry and scientific research. The university-industry connection is not merely an institutional policy or a funding mechanism; it reflects a deeper structural fact about how trust chains are deepened and validated. Industry manufactures the micro-worlds that science invents; it distributes them to populations of users whose collective experience tests those micro-worlds under conditions of real-world diversity; and the resulting feedback—financial, reputational, and technical—funds further cycles of instrument revision and trust-chain deepening. Trade practices participate indirectly but substantively in scientific validation. The full epistemic circuit runs not only from laboratory to publication but from laboratory to factory to market to user and back—a cycle in which the world at large participates in the ongoing negotiation of trust.

9.5 Demarcation revisited

The trust-conditions framework offers a fresh approach to the demarcation problem—the question of what distinguishes science from non-science (Popper, 1959; Laudan, 1983). On the trust-conditions

account, what distinguishes science from pseudoscience is not a single criterion (falsifiability, progressiveness, puzzle-solving capacity) but the *depth and calibration of trust chains*. Pseudoscience characteristically involves shallow trust chains—anecdotal evidence (Layer 0), unvalidated instruments, unreproducible results, minimal institutional oversight—presented as if they were deep. Homeopathy relies on phenomenal reports of symptom change (Layer 0) without physically calibrated mediators; astrology relies on conventional associations (Layer 0) without any instrumentally mediated link between celestial positions and human affairs. The trust-conditions framework does not dismiss these practices as "irrational"; it diagnoses them as operating with shallow trust chains that are inadequate for the claims they make.

A caveat is in order. Shallow trust chains are a necessary but not sufficient marker of pseudoscience: some pseudosciences develop elaborate institutional and procedural infrastructures (credentialing bodies, journals, conferences) that mimic the surface features of deep trust chains without the physical calibration that grounds them. The trust-conditions framework's diagnostic power lies not in a single criterion but in the capacity to examine the *character* of calibration at each layer—whether the mediating chain includes physically calibrated instruments whose proportional responses are independently validated, or whether the chain, however institutionally elaborate, rests ultimately on Layer 0 phenomenal reports.

9.6 Contemporary implications

Three contemporary issues illustrate the framework's diagnostic power.

First, the *replication crisis* in psychology, medicine, and other sciences is, in trust-conditions terms, a crisis of trust-chain reproducibility. When a study fails to replicate, it means that the trust chain—from experimental setup through measurement to data analysis—is not robust enough to produce the same terminal token across different instantiations. The framework suggests that the depth of calibration—both physical and procedural—is a key variable. Disciplines that combine physically calibrated instruments with procedurally calibrated protocols (physics, chemistry) tend to have fewer replication problems. Disciplines that rely primarily on procedural calibration (psychology, social science) are more vulnerable, not because they lack trust chains, but because procedural calibration alone is more susceptible to uncontrolled variation than the combination of physical and procedural calibration.

Second, *AI-generated knowledge claims* present a distinctive trust-chain challenge. When a large language model asserts "the melting point of aspirin is 135 °C," it produces a terminal token that is indistinguishable from the terminal token of a deep trust chain (calibrated instrument, validated protocol, peer-reviewed measurement). But the AI's token was produced through distributional pattern-matching over text, not through a physically calibrated trust chain. It is a token from the shallowest possible source—statistical co-occurrence—dressed in the clothing of deep instrumented measurement. The trust-conditions framework provides the diagnostic: ask not whether the token is "true" but how deep the trust chain that produced it is.

Third, *algorithmic decision-making* in governance, medicine, and finance raises questions about

trust in computational processes. The trust-conditions framework suggests evaluating these systems by the same criteria applied to instruments: how physically calibrated are the inputs? how transparent is the processing? how reproducible are the outputs? how institutionally stabilized are the validation procedures? An algorithm trained on calibrated medical data and validated through clinical trials has a deeper trust chain than one trained on unvalidated internet text, even if both produce confident-sounding outputs.

10 Conclusion

This paper has argued that truth-conditional analysis, while adequate for predicative assertions, does not reach the functional relations that constitute the distinctive form of modern scientific knowledge. The epistemology appropriate to functional relations is supplied by *trust conditions*—the depth and calibration of the mediating chains through which claimed functional relations are grounded in the actual world. An instrument is possible when a functional relation between variables exists; it physically instantiates that relation; calibration checks that the relation holds within tolerances. Mathematics is the formal language of possible functional relations; science maps these to the actual world through calibrated instruments. The resulting position—calibrational realism—holds that scientific objectivity consists not in the absence of interpreting subjects but in the depth and physical calibration of the trust chains through which functional relations are validated. The progressive deepening of these trust chains—from predicative naming through instrumented measurement to cascaded experimental practice—is the mechanism of scientific progress.

Calibrational realism resolves persistent tensions in the philosophy of science. It explains why intervention grounds realism (Hacking) without requiring a metaphysics of mind-independent correspondence. It dissolves the observable/unobservable binary (van Fraassen) by replacing it with a graduated trust-depth metric. It explains why operational definitions work (Bridgman) without reducing meaning to operations. It agrees with social accounts of objectivity (Longino) that objectivity is an achievement, while adding the physical dimension that makes scientific achievement epistemically distinctive. And it provides a normative account of the historical ideals of objectivity (Daston and Galison) by specifying the criterion—trust-chain depth and calibration—against which those ideals can be evaluated.

Crucially, calibrational realism is not a hard-headed claim that science has achieved objectivity once and for all. Every layer of the trust chain is a zone of negotiation—a site where claimed functional relations can be challenged, refined, or overthrown through further engagement with the world. The spectrum from packaged instruments through complex measurement systems to extended laboratories and social institutions shows that reflexivity is never fully eliminated but only progressively managed. Objectivity, on this account, is not a state but a *pursuit*: the ongoing, cumulative, fallible deepening of calibrated trust chains. That the pursuit is never complete does not make it any less real or any less worth undertaking.

Instruments deserve a central place in epistemology. They are not the servants of observation

but the arbitrators of objectivity—the physical sites where functional relations are instantiated, calibrated, and made publicly available. A philosophy of science that treats them as peripheral misses where the epistemic action actually is.

References

Baird, Davis (2004). *Thing Knowledge: A Philosophy of Scientific Instruments*. Berkeley: University of California Press.

Bridgman, Percy W. (1927). *The Logic of Modern Physics*. New York: Macmillan.

Campbell, Donald T. and Julian C. Stanley (1963). *Experimental and Quasi-Experimental Designs for Research*. Boston: Houghton Mifflin.

Carnap, Rudolf (1928). *Der logische Aufbau der Welt*. English translation: *The Logical Structure of the World*, University of California Press, 1967. Berlin: Weltkreis-Verlag.

Cartwright, Nancy (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Cassirer, Ernst (1923). *Substance and Function and Einstein's Theory of Relativity*. Originally published as *Substanzbegriff und Funktionsbegriff*, Berlin, 1910. Chicago: Open Court.

Chang, Hasok (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford: Oxford University Press.

Daston, Lorraine and Peter Galison (2007). *Objectivity*. New York: Zone Books.

Duflo, Esther, Rachel Glennerster, and Michael Kremer (2007). "Using Randomization in Development Economics Research: A Toolkit". In: *Handbook of Development Economics* 4, pp. 3895–3962.

Duhem, Pierre (1954). *The Aim and Structure of Physical Theory*. Originally published in French, 1906. Princeton, NJ: Princeton University Press.

French, Steven (2014). *The Structure of the World: Metaphysics and Representation*. Oxford: Oxford University Press.

Gibson, James J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Goldman, Alvin I. (1979). "What Is Justified Belief?" In: *Justification and Knowledge*. Ed. by George S. Pappas, pp. 1–23.

Hacking, Ian (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.

Hanson, Norwood Russell (1958). *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*. Cambridge: Cambridge University Press.

Humphreys, Paul (2009). "The Philosophical Novelty of Computer Simulation Methods". In: *Synthese* 169.3, pp. 615–626.

Hutchins, Edwin (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Knorr-Cetina, Karin (1999). *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.

Kuhn, Thomas S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

— (1977). "Objectivity, Value Judgment, and Theory Choice". In: *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Chicago: University of Chicago Press, pp. 320–339.

Ladyman, James and Don Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Lakatos, Imre (1978). *The Methodology of Scientific Research Programmes*. Vol. 1. Philosophical Papers. Cambridge: Cambridge University Press.

Latour, Bruno (1987). *Science in Action: How to Follow Scientists and Engineers Through Society*. Cambridge, MA: Harvard University Press.

Laudan, Larry (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.

— (1983). "The Demise of the Demarcation Problem". In: *Physics, Philosophy and Psychoanalysis*. Ed. by R. S. Cohen and L. Laudan, pp. 111–127.

Longino, Helen E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, NJ: Princeton University Press.

Morrison, Margaret (2015). *Reconstructing Reality: Models, Mathematics, and Simulations*. Oxford: Oxford University Press.

Nagarjuna, G. (1994). "The Role of Inversion in the Genesis, Development and the Structure of Scientific Knowledge". Available at https://web-archive.southampton.ac.uk/cogprints.org/4340/1/th.pdf. PhD thesis. Indian Institute of Technology, Kanpur.

— (2026). "Snapshots and Streams: A Lens on Natural and Artificial Meaning-Making". Preprint.

Nagarjuna, G. and Durga Prasad Karnam (2026). "Sensation Modulating Network: An Architecture of a Cognitive Agent Transforming a Habitat into Memetat". Monograph in preparation.

Noë, Alva (2004). *Action in Perception*. Cambridge, MA: MIT Press.

— (2015). *Strange Tools: Art and Human Nature*. New York: Hill and Wang.

Odling-Smee, F. John, Kevin N. Laland, and Marcus W. Feldman (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton, NJ: Princeton University Press.

Ostrom, Elinor (2005). *Understanding Institutional Diversity*. Princeton, NJ: Princeton University Press.

Parker, Wendy S. (2020). "Evidence and Knowledge from Computer Simulation". In: *Erkenntnis* 87.4, pp. 1781–1799.

Popper, Karl R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

Psillos, Stathis (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.

Rietveld, Erik and Julian Kiverstein (2014). "A Rich Landscape of Affordances". In: *Ecological Psychology* 26.4, pp. 325–352.

Sellars, Wilfrid (1956). "Empiricism and the Philosophy of Mind". In: *Minnesota Studies in the Philosophy of Science*. Ed. by Herbert Feigl and Michael Scriven. Vol. 1. Minneapolis: University of Minnesota Press, pp. 253–329.

Shapin, Steven and Simon Schaffer (1985). *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton, NJ: Princeton University Press.

Stegmüller, Wolfgang (1976). *The Structure and Dynamics of Theories*. Originally published in German, 1973. New York: Springer.

Suppe, Frederick (1977). *The Structure of Scientific Theories*. 2nd. Urbana: University of Illinois Press.

Suppes, Patrick (1962). "Models of Data". In: *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*. Ed. by Ernest Nagel, Patrick Suppes, and Alfred Tarski, pp. 252–261.

Tal, Eran (2017). "Measurement in Science". In: *Stanford Encyclopedia of Philosophy*. Fall 2017 Edition, Edward N. Zalta (ed.) URL: https://plato.stanford.edu/archives/fall2017/entries/measurement-science/.

Tarski, Alfred (1944). "The Semantic Conception of Truth and the Foundations of Semantics". In: *Philosophy and Phenomenological Research* 4.3, pp. 341–376.

van Fraassen, Bas C. (1980). *The Scientific Image*. Oxford: Clarendon Press.

Wigner, Eugene P. (1960). "The Unreasonable Effectiveness of Mathematics in the Natural Sciences". In: *Communications on Pure and Applied Mathematics* 13.1, pp. 1–14.

Winsberg, Eric (2010). *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press.

Worrall, John (1989). "Structural Realism: The Best of Both Worlds?" In: *Dialectica* 43.1–2, pp. 99–124.