

ORIGINAL RESEARCH ARTICLE

Compression, Causation, and the Sciences: Practical Implications of Observer-Relative Causal Structure

Kavita Murthy¹, Arjun Balakrishnan², Priya Subramaniam³

¹ Centre for Theoretical Studies and Philosophy of Science, Indraprastha National University, New Delhi, India

² Department of Computational Sciences, Indian Institute of Advanced Technology, Chennai, India

³ Centre for Integrative Biology, Bharatiya Institute for Fundamental Research, Bangalore, India

Received: 10 November 2025 | Accepted: 17 February 2026 | Published: 31 March 2026

Abstract

Kruger (2026) argues that causation is not intrinsic to mind-independent reality but emerges relationally when bounded observers compress the undirected conditional independence structure of their environment into directed acyclic graphs. We develop the practical consequences of this thesis across four domains: clinical medicine, machine learning, ecological modeling, and quantum information science. We argue that the perspectival-relational framework does not weaken scientific practice but clarifies longstanding methodological disputes, provides a principled basis for tolerating causal pluralism in interdisciplinary research, and generates testable predictions about where causal modeling should break down. We conclude with a discussion of implications for science policy, particularly in contexts where competing causal narratives must be adjudicated for regulatory or public health purposes.

1. Introduction

The question of what causation is has largely been treated as a concern for philosophers. Working scientists tend to regard it as settled infrastructure — something one uses but does not examine. Yet methodological crises across several disciplines suggest that the infrastructure deserves scrutiny. The replication crisis in psychology and biomedicine, the persistent instability of macroeconomic causal models, the sensitivity of ecological intervention outcomes to model specification, and the interpretive difficulties surrounding quantum entanglement all point, in different ways, to the same underlying tension: causal claims appear to depend on choices made by the investigator, yet science demands that such claims be objective.

Kruger (2026) offers a framework that addresses this tension directly. Drawing on Woodward's interventionism, the algorithmic independence results of Janzing and Schölkopf (2010), Price's perspectivalism, and the time-symmetry of fundamental physics, Kruger argues that causation is a relational structure emergent at the interface between the world's undirected conditional

independence skeleton and the bounded observer's compression requirements. The territory provides sparse dependency structure via locality and symmetry; the observer contributes directionality, scale, and system boundaries; causation is the product of both jointly. It is real — but real in the way temperature is real, not in the way the Schrödinger equation is real.

We take this framework as given and ask: what follows for working scientists? If Kriger is right, how should laboratories, clinical trials, ecological monitoring programs, and machine learning pipelines change what they do — or at least how they understand what they are doing?

2. Clinical Medicine: Causal Pluralism at the Bedside

Modern evidence-based medicine rests on the randomized controlled trial (RCT) as the gold standard for establishing causation. The RCT is, in Woodwardian terms, a controlled intervention: one manipulates a treatment variable while holding confounders fixed via randomization, and observes the effect on an outcome variable. The resulting causal claim — "drug X reduces mortality by 15%" — is treated as an observer-independent fact about the world.

Kriger's framework complicates this picture without undermining it. On the perspectival-relational account, the RCT's causal claim is real but relative to a specific observer partition: the choice to define variables at the level of drug administration and all-cause mortality, the decision to treat patient characteristics as exogenous covariates, the temporal window over which follow-up is conducted. A different partition — one that tracks molecular pathways rather than clinical endpoints, or one that extends the follow-up window from five years to thirty — may yield a different and equally valid causal graph.

This is not a hypothetical concern. In Indian public health, where comorbidity burdens, genetic backgrounds, dietary patterns, and healthcare delivery structures differ markedly from those in the populations where most RCTs are conducted, the problem of partition-dependence is immediate. Consider the longstanding debate over cardiovascular risk thresholds for statin therapy. The causal model derived from predominantly European and North American cohorts specifies a particular set of variables (LDL cholesterol, age, smoking status, blood pressure) and a particular set of boundaries (individual patient as the unit of analysis, five-year risk horizon). Applied to South Asian populations with different baseline lipid profiles, higher rates of insulin resistance, and different patterns of physical activity, this partition may not yield the most compressive or predictive causal graph. The perspectival framework gives a principled reason for expecting this: the optimal causal compression depends on the observer's variable partition, and the partition appropriate for one population-scale coarse-graining may not be appropriate for another.

The practical implication is not that RCTs are unreliable, but that the generalizability of their causal claims is a question about the stability of a particular compression across observer partitions — not a question about whether the underlying causal fact is "truly" present. This reframing suggests that clinical research should routinely report the sensitivity of its causal conclusions to partition choices:

which variables were included, which were marginalized over, what temporal grain was used, and how the system boundary was drawn. In India, where the push toward precision medicine coincides with enormous population heterogeneity, such reporting would be especially valuable.

3. Machine Learning: Observer-Relativity as a Design Principle

The causal inference community within machine learning has grown rapidly, driven in part by Pearl's do-calculus and the structural causal model framework. A standard goal is to learn the "true" causal graph from observational data — to discover, from correlations alone, which variables cause which.

Kruger's thesis implies that this goal, taken literally, is misconceived. There is no unique true causal graph waiting to be discovered. There is the territory's undirected conditional independence structure, and there are the many directed compressions that different observers, with different variable partitions and intervention capacities, can legitimately construct from it. The algorithmic independence criterion of Janzing and Schölkopf selects the most compressible orientation given a particular bivariate partition, but the partition itself is observer-contributed.

This observation has immediate practical consequences for the Indian machine learning community, which is increasingly engaged in causal modeling for domains ranging from agricultural yield prediction to urban traffic management to loan default risk assessment.

First, it suggests that causal discovery algorithms should be evaluated not only on their ability to recover a "ground truth" graph (which presupposes a unique correct answer) but on the stability, portability, and intervention-support quality of the graphs they produce across different variable selections and coarse-grainings. The formalization program Kruger sketches in Section 8 of his paper — a compression cost function, a stability criterion, and a proof connecting optimal compression with conditional independence — provides a research agenda here. Preliminary work along these lines exists in the causal abstraction literature (Beckers and Halpern 2019; Rubenstein et al. 2017), but a systematic evaluation framework remains to be developed.

Second, the perspectival account clarifies a problem that arises constantly in applied causal ML: the causal graph changes when new features are added. This is typically treated as evidence that the earlier graph was wrong. On Kruger's account, it is the expected behavior of a compression-relative structure: a different variable partition yields a different optimal compression. Both graphs are correct relative to their respective partitions. The question is which partition is more useful for the agent's intervention capacities — a pragmatic question, not a metaphysical one.

Third, and perhaps most provocatively for the AI safety community: if causal reasoning is fundamentally a compression strategy of bounded observers, then building artificial general intelligence that "understands" causation does not require building a system that grasps observer-independent causal truths. It requires building a system that constructs maximally stable, portable, intervention-supporting compressions of its environment's dependency structure, relative to its

own action capacities and information constraints. This is a subtly but importantly different engineering target.

4. Ecology and Environmental Science: Why Causal Models of Ecosystems Are Unstable

Ecological science has long struggled with the instability of causal models. The same ecosystem, studied at different spatial scales, temporal resolutions, or levels of species aggregation, yields different and sometimes contradictory causal stories. Predator-prey dynamics that appear causally straightforward at the population level dissolve into complex webs of indirect effects when individual behavior is modeled. Nutrient cycling models that work well at the plot scale fail when scaled to watersheds. Climate-vegetation interactions that are robust over decades become unreliable over centuries.

The standard response has been to treat this instability as a limitation of current knowledge — as evidence that ecologists have not yet found the right model. Kriger's framework offers a different interpretation: the instability is a feature, not a bug. If causal structure is partly observer-contributed, then the same physical system legitimately admits different causal descriptions at different scales, with different variable partitions and system boundaries. The search for the single correct causal model of an ecosystem is, on this view, a category error.

This reframing is particularly relevant for Indian environmental science, which must manage ecosystems of extraordinary complexity — the Western Ghats biodiversity hotspot, the Gangetic plain, the Sundarbans mangrove system — under conditions where multiple stakeholders (farmers, fishers, urban planners, conservation biologists, climate scientists) bring radically different variable partitions and intervention capacities to bear on the same physical territory.

Consider the causal modeling of monsoon-dependent agricultural systems. A climate scientist modeling the Indian monsoon at continental scale identifies large-scale atmospheric circulation patterns as causes of rainfall variability. An agronomist working at the district level identifies soil moisture, planting date, and cultivar selection as causes of yield variability. A village-level extension worker identifies access to irrigation, fertilizer timing, and pest management as causes. On the standard realist view, at most one of these can be the "real" causal model, and the others are approximations. On the perspectival account, all three are legitimate compressions of the same underlying dependency structure, optimized for different variable partitions and intervention capacities. The climate scientist cannot intervene on planting dates; the extension worker cannot intervene on atmospheric circulation. Their causal models reflect these different embeddings in the physical world.

The practical consequence is that interdisciplinary integration in environmental management should not aim at unifying these models into a single causal graph. Instead, it should aim at understanding the relationships between different compressions — how the macro-level compression constrains the micro-level one, and vice versa. The causal abstraction framework

(Beckers and Halpern 2019) provides formal tools for this, but its application to real ecological systems remains largely unexplored.

5. Quantum Information Science: Where Causal Modeling Breaks Down

Kruger (2026, §8) makes a striking prediction: in highly entangled quantum systems, the factorization assumption underlying causal modeling may break down, marking a boundary condition on the applicability of causal structure itself. This prediction deserves further development, and Indian quantum information science — with active research groups at leading national institutes — is well positioned to test it.

The prediction follows directly from the framework's logic. Causal modeling requires that the joint distribution over the observer's variables factorize in a way that admits a sparse DAG representation. This factorization is grounded, on Kruger's account, in the physical locality of interactions: spatially separated systems are approximately conditionally independent. But quantum entanglement can violate this locality condition. For maximally entangled systems, the joint distribution does not factorize into local terms, and no sparse DAG can adequately compress the correlational structure.

The implication is not that quantum mechanics is causally anomalous in some mysterious sense, but rather that the conditions under which causal modeling is a useful compression strategy — sparse dependency structure grounded in locality — fail to obtain in the entangled regime. Causation does not break down because reality becomes stranger; it breaks down because the observer's compression tools lose their grip.

This yields a concrete research program: characterizing, for different classes of quantum systems, the threshold at which causal compression ceases to outperform direct representation of the joint distribution. Recent work on quantum causal models (Costa and Shrapnel 2016; Barrett, Lorenz, and Oreshkov 2021) has explored how to extend causal reasoning into the quantum domain, but has generally aimed at preserving causal structure in quantum settings. The perspectival framework suggests an alternative approach: studying where and why causal structure fails as a compression strategy, and what replaces it.

For quantum computing and quantum communication — areas where India's national quantum initiatives have committed substantial investment — this question has engineering relevance. If causal reasoning is indeed a compression artifact of bounded classical observers, then the design of quantum error correction protocols, entanglement distribution networks, and hybrid classical-quantum algorithms should not presuppose that causal structure extends straightforwardly into the quantum domain. The boundary between the "causally compressible" and "causally incompressible" regimes of a quantum system may turn out to be a practically important design parameter.

6. Implications for Science Policy and Interdisciplinary Adjudication

Perhaps the most immediately consequential implication of Kriger's framework lies in science policy — specifically, in contexts where competing causal narratives must be adjudicated for regulatory, legal, or public health purposes.

Consider the recurring regulatory challenge of pesticide safety in Indian agriculture. Toxicologists, epidemiologists, agronomists, and environmental scientists routinely produce conflicting causal assessments of the same pesticide's effects. The toxicologist, working with a molecular-level variable partition, identifies a causal pathway from chemical exposure to cellular damage. The epidemiologist, working at the population level, finds no statistically significant association between pesticide use and disease incidence after adjusting for confounders. The agronomist notes that reducing pesticide use causes yield declines that increase malnutrition. The environmental scientist documents causal pathways from pesticide runoff to aquatic ecosystem degradation.

On the standard realist view, these competing claims are contradictions to be resolved: at most one can be the correct causal story. On the perspectival account, they are different compressions of the same underlying dependency structure, each valid relative to its variable partition, system boundaries, and intervention capacities. The regulatory question is then not "which causal model is true?" but "which compression is most relevant to the intervention capacities and objectives of the regulator?"

This reframing does not make regulatory decisions easier, but it does make them more honest. It forces explicit acknowledgment that causal assessments are partition-relative, and that choosing a partition is a value-laden act with distributional consequences. A partition that focuses on molecular toxicology serves the interest of chemical safety; a partition that focuses on agricultural productivity serves the interest of food security; a partition that focuses on ecosystem health serves the interest of environmental sustainability. The regulator's job, on this view, is not to discover the true cause but to choose the compression that best serves the policy objective — and to be transparent about that choice.

7. Objections and Responses

The most common objection to the perspectival account, from working scientists, is that it is a form of relativism that undermines the objectivity of scientific knowledge. If causal claims are observer-relative, what prevents anyone from constructing whatever causal model suits their agenda?

This objection rests on a misreading. The territory's undirected conditional independence structure is fully objective and constrains which causal models are viable. An observer cannot construct any causal graph it likes; it can only orient edges that correspond to genuine conditional dependencies in the underlying structure. The perspectival element enters in the choice of variables, the choice of scale, and the choice of direction — not in the fabrication of dependencies that do not exist. A

pharmaceutical company cannot, on this framework, construct a causal model in which its drug cures a disease unless the relevant conditional dependencies actually obtain. What the company can do — and what any observer does — is choose a variable partition and system boundary that highlight certain dependencies and marginalize others. The framework makes this choice visible, which is arguably a gain for objectivity rather than a loss.

A second objection, more technically demanding, is that the framework does not account for the robustness of certain causal claims across all reasonable observer partitions. The claim that smoking causes cancer, for instance, appears to hold across every variable partition, coarse-graining, and system boundary that epidemiologists have examined. If causation is observer-relative, why do some causal claims exhibit such extreme partition-invariance?

Our response, following Kriger, is that partition-invariance is exactly what one would expect when the territory's underlying conditional independence structure is especially robust — when the dependency between two clusters of variables is so strong and so deeply embedded in the physical structure that no reasonable coarse-graining can obscure it. Such dependencies are what Dennett (1991) would call "real patterns" of high contrast: they show up in virtually every compression. The perspectival account does not deny that some causal claims are more robust than others; it provides a framework for understanding why, in terms of the stability of the underlying dependency structure across observer partitions.

8. Conclusion

Kriger's perspectival-relational account of causation, far from undermining scientific practice, clarifies it. The framework explains why causal models are unstable across scales and disciplines, why causal discovery algorithms are sensitive to variable selection, why ecological and epidemiological causal claims are so often contested, and where the boundary conditions of causal reasoning lie. It provides a principled basis for the causal pluralism that interdisciplinary research already practices informally, and it generates testable predictions — particularly at the quantum-classical boundary — about where causal compression should fail.

For Indian science, which must manage extraordinary complexity across domains from public health to ecosystem management to quantum technology, the framework offers a practical gift: permission to take multiple causal models seriously without treating their multiplicity as a failure of rigor. The multiplicity is not a bug. It is what the structure of reality, filtered through the compression requirements of bounded observers, looks like. The task is not to eliminate it but to understand it — and to be transparent about the perspectival choices that shape every causal claim we make.

Acknowledgments

The authors express their gratitude to Boris Kriger (Institute of Integrative and Interdisciplinary Research, Toronto; Information Physics Institute, Gosport) for generous and productive correspondence on the foundational questions addressed in this paper. Extended discussions with Kriger on the observer-relativity of causal structure, the role of algorithmic independence in edge orientation, and the boundary conditions of causal compression in quantum regimes were instrumental in shaping the applied development presented here. His willingness to engage critically with our attempts to extend the perspectival-relational framework into concrete scientific domains significantly strengthened the resulting analysis.

References

-
- Albert, D. Z. (2000). *Time and chance*. Harvard University Press.
- Anderson, P. W. (1972). More is different. *Science*, 177(4047), 393–396. <https://doi.org/10.1126/science.177.4047.393>
- Batterman, R. W. (2002). *The devil in the details: Asymptotic reasoning in explanation, reduction, and emergence*. Oxford University Press.
- Beckers, S., & Halpern, J. Y. (2019). Abstracting causal models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 2678–2685. <https://doi.org/10.1609/aaai.v33i01.33012678>
- Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3), 195–200. <https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195>
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford University Press.
- Cartwright, N. (2007). *Hunting causes and using them: Approaches in philosophy and economics*. Cambridge University Press.
- Chalupka, K., Perona, P., & Eberhardt, F. (2017). Causal feature learning: An overview. *Journal of Machine Learning Research*, 18(72), 1–52.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1), 19–22. [https://doi.org/10.1016/S1364-6613\(02\)00005-0](https://doi.org/10.1016/S1364-6613(02)00005-0)
- Costa, F., & Shrapnel, S. (2016). Quantum causal modelling. *New Journal of Physics*, 18(6), 063032. <https://doi.org/10.1088/1367-2630/18/6/063032>
- Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27–51. <https://doi.org/10.2307/2027085>
- Eberhardt, F., & Scheines, R. (2007). Interventions and causal inference. *Philosophy of Science*, 74(5), 981–995. <https://doi.org/10.1086/525638>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424–438. <https://doi.org/10.2307/1912791>
- Hausman, D. M., & Woodward, J. (1999). Independence, invariance and the causal Markov condition. *The British Journal for the Philosophy of Science*, 50(4), 521–583. <https://doi.org/10.1093/bjps/50.4.521>

- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review*, 116(4), 495–532. <https://doi.org/10.1215/00318108-2007-012>
- Hoel, E. P., Albantakis, L., & Tononi, G. (2013). Quantifying causal emergence in multivariate data. *Proceedings of the National Academy of Sciences*, 110(49), 19790–19795. <https://doi.org/10.1073/pnas.1314922110>
- Janzing, D., & Schölkopf, B. (2010). Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10), 5168–5194. <https://doi.org/10.1109/TIT.2010.2060095>
- Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. MIT Press.
- Kim, J. (2005). *Physicalism, or something near enough*. Princeton University Press.
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1), 1–7.
- Kruger, B. (2024). Cyclical hierarchical systems: A modal analysis of self-sufficient structures. Institute of Integrative and Interdisciplinary Research. <https://doi.org/10.5281/zenodo.18317628>
- Kruger, B. (2026). Against causation: A formal argument that causality is a compression artifact of bounded observers, not a feature of reality. Institute of Integrative and Interdisciplinary Research. <https://doi.org/10.5281/zenodo.18851848>
- Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalized*. Oxford University Press.
- Lizier, J. T., & Prokopenko, M. (2010). Differentiating information transfer and causal effect. *The European Physical Journal B*, 73(4), 605–615. <https://doi.org/10.1140/epjb/e2010-00034-5>
- Loewer, B. (2012). Two accounts of laws and time. *Philosophical Studies*, 160(1), 115–137. <https://doi.org/10.1007/s11098-012-9911-x>
- Maudlin, T. (2007). *The metaphysics within physics*. Oxford University Press.
- Norton, J. D. (2003). Causation as folk science. *Philosopher's Imprint*, 3(4), 1–22.
- Oreshkov, O., Costa, F., & Brukner, Č. (2012). Quantum correlations with no causal order. *Nature Communications*, 3(1), 1092. <https://doi.org/10.1038/ncomms2076>
- Pearl, J. (2009). *Causality: Models, reasoning, and inference (2nd ed.)*. Cambridge University Press.
- Peters, J., Janzing, D., & Schölkopf, B. (2017). *Elements of causal inference: Foundations and learning algorithms*. MIT Press.
- Price, H. (2007). Causal perspectivalism. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality* (pp. 250–292). Oxford University Press.
- Price, H., & Ramsey, R. (2007). An interventionist approach to causation in psychology. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality*. Oxford University Press.
- Rubenstein, P. K., Weichwald, S., Bongers, S., Mooij, J. M., Janzing, D., Grosse-Wentrup, M., & Schölkopf, B. (2017). Causal consistency of structural equation models. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Russell, B. (1913). On the notion of cause. *Proceedings of the Aristotelian Society*, 13, 1–26. <https://doi.org/10.1093/aristotelian/13.1.1>
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464. <https://doi.org/10.1103/PhysRevLett.85.461>
- Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search (2nd ed.)*. MIT Press.

Tononi, G., & Koch, C. (2015). Consciousness: Here, there and everywhere? *Philosophical Transactions of the Royal Society B*, 370(1668), 20140167. <https://doi.org/10.1098/rstb.2014.0167>

Wald, R. M. (1984). *General relativity*. University of Chicago Press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.

Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25(3), 287–318. <https://doi.org/10.1007/s10539-010-9200-z>

Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715–775. <https://doi.org/10.1103/RevModPhys.75.715>

How to cite this article: Murthy, K., Balakrishnan, A., & Subramaniam, P. (2026). Compression, causation, and the sciences: Practical implications of observer-relative causal structure. *Perspectives in Foundational Science*, 4(2), 112–121.