

KNO-014 · Causality and Correlation

KNO · Ways of Knowing

Juan Segurola

v0.7 · EN

All Rights Reserved

Ways of Knowing

Epistemic profile: Current scientific interpretation **Licence at publication:** All Rights Reserved

Article Abstract

When two variables move together, what additional evidence is needed before the pattern can support a claim that changing one would change the other? This article examines that question through temporal order, confounding, reverse causation, randomisation, causal diagrams, adjustment, propensity scores, instrumental variables, negative controls and target-trial emulation. It asks how each method changes the inferential problem, which assumptions remain hidden when an association is treated as self-explanatory, and why different designs identify different causal contrasts. It also asks where observational evidence can be genuinely causal without pretending to reproduce a randomised experiment, and where formal methods merely make assumptions more explicit. The goal is to separate the observable pattern from the intervention question and to show what a defensible causal claim must make inspectable.

In 60 Seconds

Association alone does not tell us which causal story generated a pattern. A cause may produce an outcome, a common cause may produce both variables, the outcome may influence the supposed exposure, or selection and measurement may create the dependence. Causal inference therefore begins by defining the intervention or contrast of interest and asking what design makes that contrast identifiable. Randomisation can protect against baseline confounding when properly implemented. Observational studies can also support causal claims, but only through explicit assumptions about time, causal structure, adjustment, measurement and selection. Directed acyclic graphs help state those assumptions; propensity scores address balance in measured covariates; instrumental variables substitute different assumptions; negative controls test some bias pathways; target-trial emulation aligns timing and design. None is a causality machine. The practical discipline is to expose the assumptions that connect an observed association to a counterfactual or interventional claim, then stress-test those assumptions rather than treating statistical precision as causal proof.

Why This Question Matters

Medicine, engineering, economics, climate science and public policy often need to know what would happen if an exposure, treatment, process variable or policy were changed. An association can motivate that question but cannot answer it by itself. The difference matters because prediction and intervention are not the same task. Understanding the inferential bridge from correlation to causation helps readers see why design, timing, causal assumptions and bias diagnostics are part of the evidence rather than technical details added after the result.

1. Association Is a Pattern; Causation Is a Counterfactual Question

A statistical association describes how variables co-vary in observed data. A causal effect compares outcomes under different possible exposure or treatment conditions. Hernán's formal epidemiologic presentation makes the distinction explicit: causal effects are defined through contrasts between potential outcomes, while observed associations can differ from those causal contrasts because the groups receiving different exposures may already differ in relevant ways [2]. Holland's account of the Rubin causal model makes the same separation central to statistical reasoning about causation [3].

The counterfactual formulation also exposes a fundamental obstacle. For the same unit at the same time, we do not simultaneously observe the outcome both under treatment and under no treatment. The missing counterfactual is not a technical nuisance that a larger dataset automatically fills. Causal inference therefore depends on comparisons across units, times or regimes together with assumptions that make those comparisons informative.

Claim KNO014-C01. An observed association and a causal effect are different quantities; causal effects require a contrast between potential outcomes or intervention regimes, not merely a covariance or regression coefficient. [2,3]

Claim KNO014-C02. The fundamental causal problem is that, for one unit, only one of the relevant potential outcomes is observed. [3]

Where That Leaves Us

A statistical association describes how variables co-vary in observed data. Hernán's formal epidemiologic presentation makes the distinction explicit: causal effects are defined through contrasts between potential outcomes, while observed associations can differ from those causal contrasts because the groups receiving different exposures may already differ in relevant ways [2]. The next section asks why temporal order helps but does not settle the question.

2. Why Temporal Order Helps but Does Not Settle the Question

A cause must precede its effect in the causal relation under study. This makes temporal ordering an essential constraint. If an alleged cause occurs only after the outcome, that claim fails for that temporal relation.

But precedence is not enough. An exposure may precede an outcome while both are driven by an earlier common cause. A biomarker may change before diagnosis while already reflecting an occult disease process. A policy may be adopted before an economic change because governments anticipate that change. Temporal order removes some causal stories, not all competing stories.

Bradford Hill treated temporality as one consideration in evaluating whether an observed association might be causal, while also warning against treating his considerations as rigid criteria that mechanically prove causation [1]. His framework remains historically important precisely because it addressed inference from imperfect observational evidence, not because it supplied a universal checklist.

Claim KNO014-C03. Temporal precedence is necessary for the cause to precede the effect in the causal ordering, but temporal order alone is not sufficient to identify a causal effect. [2]

Where That Leaves Us

A cause must precede its effect in the causal relation under study. If an alleged cause occurs only after the outcome, that claim fails for that temporal relation. But precedence is not enough. An exposure may precede an outcome while both are driven by an earlier common cause. The next section examines confounding.

3. Confounding: When a Common Cause Produces the Association

Suppose smoking is associated with an outcome, but age influences both smoking history and the outcome. If age is not appropriately handled, some of the observed association between smoking and outcome may reflect age rather than the causal effect of smoking.

This is confounding in its simplest form: a common cause of exposure and outcome opens a non-causal path between them. Modern causal diagrams make this structure explicit. Greenland, Pearl and Robins showed how directed acyclic graphs can encode assumed causal relations and help identify variables that must be controlled to estimate an unconfounded effect [4].

The important word is assumed. A graph is not discovered merely by fitting the data. Its arrows represent causal commitments supplied by scientific knowledge, design information and temporal constraints. Data can contradict some consequences of a graph, but a compatible graph is not thereby proven true.

Claim KNO014-C04. Confounding can produce an exposure-outcome association even when part or all of that association is not the causal effect of the exposure. [4]

Claim KNO014-C05. Causal diagrams can make adjustment assumptions explicit, but the causal structure encoded by their arrows requires substantive assumptions beyond the observed joint distribution. [4,5]

Where That Leaves Us

Suppose smoking is associated with an outcome, but age influences both smoking history and the outcome. If age is not appropriately handled, some of the observed association between smoking and outcome may reflect age rather than the causal effect of smoking. Data can contradict some consequences of a graph, but a compatible graph is not thereby proven true. The next section examines reverse causation.

4. Reverse Causation: When the Arrow Points the Other Way

An association between X and Y does not tell us whether X causes Y, Y causes X, both occur, or a third structure explains the dependence. Reverse causation is especially important when disease processes begin before diagnosis. Weight loss may predict later illness because preclinical illness caused the weight loss, not because weight loss caused the illness.

Longitudinal data can help because they establish sequence, but they do not automatically eliminate reverse causation. The relevant causal process may begin before the first recorded measurement. The correct response is to define time zero carefully, use knowledge of latency and disease history, and test whether plausible lag structures alter the inference.

Claim KNO014-C06. Longitudinal ordering can constrain reverse-causation explanations but cannot by itself exclude causal processes that began before measurement. [2,7]

Where That Leaves Us

An association between X and Y does not tell us whether X causes Y, Y causes X, both occur, or a third structure explains the dependence. Weight loss may predict later illness because preclinical illness caused the weight loss, not because weight loss caused the illness. The next section examines randomisation.

5. Randomisation: Why an Experiment Can Change the Evidential Problem

In an ideal randomised experiment, treatment assignment is generated independently of baseline causes of the outcome. With adequate sample size, concealment, adherence to the protocol and valid outcome measurement, randomisation can make treatment groups comparable in expectation and support estimation of causal effects without the same confounding assumptions required in ordinary observational comparisons [2].

This does not mean every randomised trial is unbiased. Loss to follow-up, non-adherence, outcome misclassification, post-randomisation selection and analysis choices can still create problems. Nor does randomisation answer every causal question: some interventions are unethical, impossible, too rare or too long-term to randomise.

The key epistemic point is narrower. Randomisation changes the assignment mechanism. It does not make causality visible; it creates a design in which a specific causal contrast can be identified under a smaller and more transparent set of assumptions.

Claim KNO014-C07. Proper randomisation can identify causal effects by breaking systematic baseline causes of treatment assignment, but trial validity still depends on implementation, follow-up, measurement and analysis. [2]

Where That Leaves Us

In an ideal randomised experiment, treatment assignment is generated independently of baseline causes of the outcome. This does not mean every randomised trial is unbiased. Loss to follow-up, non-adherence, outcome misclassification, post-randomisation selection and analysis choices can still create problems. The next section examines the proposition that adjustment is not “control for everything”.

6. Adjustment Is Not “Control for Everything”

A common intuition says that observational studies become more causal if they adjust for more variables. Causal diagrams show why this is false.

If a variable is a confounder, adjustment may block a non-causal backdoor path. If a variable lies on the causal pathway, adjusting for it may remove part of the total effect one intended to estimate. If a variable is a collider, meaning it is caused by two variables on a path, conditioning on it can open an association that was absent before conditioning. In such cases, “controlling for” a variable manufactures bias.

Greenland, Pearl and Robins formalised how graph structure guides adjustment [4]. Later methodological work has emphasised that collider control, mediator control and indiscriminate pre-treatment adjustment can each distort causal effect estimates in the wrong structure [6].

Claim KNO014-C08. More statistical adjustment is not monotonically better: adjusting for a confounder may reduce bias, while adjusting for a mediator or collider can alter or introduce bias depending on the estimand and causal structure. [4,6]

Where That Leaves Us

A common intuition says that observational studies become more causal if they adjust for more variables. If a variable is a confounder, adjustment may block a non-causal backdoor path. If a variable lies on the causal pathway, adjusting for it may remove part of the total effect one intended to estimate. The next section examines the proposition that potential outcomes and causal diagrams are complementary languages.

7. Potential Outcomes and Causal Diagrams Are Complementary Languages

Two major modern traditions describe causal inference using different but overlapping languages. Potential-outcome frameworks define effects by comparing outcomes under alternative treatment states. Structural causal models and DAGs represent causal relations and intervention structure graphically and algebraically.

These are not simply rival vocabularies where one must be chosen and the other rejected. In many practical analyses they illuminate different parts of the same problem. Potential outcomes force precision about the causal contrast and target population. DAGs force precision about assumed pathways, confounding, mediation and selection. Both make assumptions more visible than a regression coefficient presented without a causal model.

Pearl emphasised that causal claims require assumptions not contained in ordinary probability distributions and that causal languages are needed to express those assumptions [5]. Hernán and Robins similarly organise causal inference around explicitly specified treatment strategies and counterfactual outcomes [8].

Claim KNO014-C09. Potential-outcome and graphical approaches provide complementary ways to state causal estimands and assumptions; neither reduces causal inference to ordinary association alone. [5,8]

Where That Leaves Us

Two major modern traditions describe causal inference using different but overlapping languages. These are not simply rival vocabularies where one must be chosen and the other rejected. Structural causal models and DAGs represent causal relations and intervention structure graphically and algebraically. Hernán and Robins similarly organise causal inference around explicitly specified treatment strategies and counterfactual outcomes [8]. The next section examines propensity scores.

8. Propensity Scores: Balance Measured Covariates, Not Unknown Reality

Rosenbaum and Rubin introduced the propensity score as the probability of treatment assignment conditional on observed covariates [9]. Under their assumptions, balancing on the propensity score can balance those observed covariates between treatment groups and support causal effect estimation.

But a propensity score does not automatically solve unmeasured confounding. If an important common cause is absent from the measured covariates, balancing the measured variables does not balance the missing cause. The method is therefore powerful within its identification conditions and powerless against violations it cannot see.

Claim KNO014-C10. Propensity-score methods can balance observed covariates under specified assumptions but do not by themselves eliminate bias from unmeasured confounding. [9]

Where That Leaves Us

Rosenbaum and Rubin introduced the propensity score as the probability of treatment assignment conditional on observed covariates [9]. But a propensity score does not automatically solve unmeasured confounding. If an important common cause is absent from the measured covariates, balancing the measured variables does not balance the missing cause. The next section examines instrumental variables.

9. Instrumental Variables: A Different Route with Different Assumptions

Sometimes researchers use an instrumental variable: a variable that shifts treatment or exposure but, under the identification model, affects the outcome only through that treatment and satisfies additional independence and relevance conditions.

Angrist, Imbens and Rubin showed how instrumental-variable reasoning can identify a local average treatment effect for compliers under explicit assumptions, including monotonicity in their framework [10]. This result illustrates a recurring lesson: methods do not produce causal effects in the abstract. They identify particular estimands for particular populations under particular assumptions.

Claim KNO014-C11. Instrumental-variable methods can identify causal effects without ordinary no-unmeasured-confounding assumptions, but only by substituting other strong identification assumptions and often a more specific estimand. [10]

Where That Leaves Us

Sometimes researchers use an instrumental variable: a variable that shifts treatment or exposure but, under the identification model, affects the outcome only through that treatment and satisfies additional independence and relevance conditions. This result illustrates a recurring lesson: methods do not produce causal effects in the abstract. The next section examines negative controls.

10. Negative Controls: Look for Bias Where No Effect Should Exist

A negative-control exposure or outcome is chosen so that, under the causal hypothesis being tested, the target causal effect should be absent. If an association nevertheless appears, it can reveal residual confounding, selection bias or other systematic error.

Lipsitch, Tchetgen Tchetgen and Cohen proposed negative controls as a routine diagnostic for observational causal studies [11]. A negative control is not a universal falsifier. Its interpretation depends on why the control is expected to have no causal effect and on whether it shares the bias mechanism one is trying to detect. Used carefully, however, it turns background assumptions into an empirical stress test.

Claim KNO014-C12. Negative controls can detect some forms of residual confounding or bias when the control is causally null for the target relation but shares relevant bias pathways. [11]

Where That Leaves Us

A negative-control exposure or outcome is chosen so that, under the causal hypothesis being tested, the target causal effect should be absent. A negative control is not a universal falsifier. Lipsitch, Tchetgen Tchetgen and Cohen proposed negative controls as a routine diagnostic for observational causal studies [11]. The next section examines target trials.

11. Target Trials: Design the Observational Question Before Analysing the Data

Many observational studies are implicitly trying to answer the question a randomised trial would have answered if that trial were feasible. Hernán and Robins proposed making that hypothetical trial explicit: eligibility criteria, treatment strategies, assignment procedure, time zero, follow-up, outcome, causal contrast and analysis plan are specified first, and the observational analysis is then designed to emulate those elements as closely as possible [7].

This approach can reveal avoidable design defects. If eligibility is assessed at one time, treatment later, and follow-up begins at a third time, the analysis can introduce immortal-time or selection biases. Target-trial thinking therefore shifts causal inference upstream: design precedes modelling.

Claim KNO014-C13. Explicit target-trial emulation can expose timing and design mismatches in observational analyses before statistical modelling begins. [7,12]

Where That Leaves Us

Many observational studies are implicitly trying to answer the question a randomised trial would have answered if that trial were feasible. Target-trial thinking therefore shifts causal inference upstream: design precedes modelling. If eligibility is assessed at one time, treatment later, and follow-up begins at a third time, the analysis can introduce immortal-time or selection biases. The next section revisits Bradford Hill's evidence-integration framework.

12. Bradford Hill Revisited: Evidence Integration, Not a Causal Scorecard

Hill's 1965 discussion remains influential because it addressed how to reason when controlled experiments are unavailable. Strength, consistency, specificity, temporality, biological gradient, plausibility, coherence, experiment and analogy were offered as viewpoints to consider, not as a validated additive score and not as necessary-and-sufficient conditions [1].

Modern causal inference adds formal tools that Hill did not have: counterfactual estimands, DAGs, g-methods, propensity scores, instrumental-variable identification, negative controls and target-trial emulation. Yet Hill's broader lesson remains compatible with them: causal judgment integrates multiple kinds of evidence while resisting automatic inference from association alone.

The important update is methodological. A strong association can still be confounded. A weak association can still represent an important causal effect. Biological plausibility evolves with background knowledge. Consistency across studies can reflect shared bias. The evidential value of each feature depends on the causal question and the design that generated the evidence.

Claim KNO014-C14. Bradford Hill's considerations are aids to causal reasoning, not a mechanical checklist that proves or disproves causation. [1]

Where That Leaves Us

Hill's 1965 discussion remains influential because it addressed how to reason when controlled experiments are unavailable. Strength, consistency, specificity, temporality, biological gradient, plausibility, coherence, experiment and analogy were offered as viewpoints to consider, not as a validated additive score and not as necessary-and-sufficient conditions [1]. The next section asks what data alone cannot tell you.

13. What Data Alone Cannot Tell You

A sufficiently rich dataset can estimate many conditional associations with great precision. Precision does not identify the causal structure that generated them. Distinct causal models may imply the same observed distribution. The difference between those models appears only when we ask intervention or counterfactual questions.

This is why causal inference cannot be reduced to “use a more powerful model”. Machine learning may estimate prediction functions exceptionally well while leaving the causal estimand unidentified. Conversely, a simple analysis can support a strong causal claim when design and assumptions make the contrast identifiable.

Claim KNO014-C15. Predictive accuracy and causal identification are distinct goals; a model can predict well without establishing the effect of an intervention. [5]

Where That Leaves Us

A sufficiently rich dataset can estimate many conditional associations with great precision. Precision does not identify the causal structure that generated them. Distinct causal models may imply the same observed distribution. The difference between those models appears only when we ask intervention or counterfactual questions. The next section turns these points into a practical sequence for analysis.

14. A Practical Causal-Inference Sequence

A disciplined analysis can be organised as a sequence rather than a single statistical test:

1. State the causal question and target population.
2. Define the exposure or intervention and the comparison regime.
3. Define the outcome and time horizon.
4. Specify time zero and the ordering of relevant events.
5. Draw or otherwise state the assumed causal structure.
6. Decide which variables are confounders, mediators, colliders or consequences of selection for the estimand at hand.
7. Choose a design and estimator whose identification assumptions match that structure.
8. Check balance, overlap, measurement quality and missingness.
9. Use negative controls, alternative specifications or sensitivity analyses where appropriate.
10. Interpret the estimate as causal only to the extent that the identification and measurement assumptions remain defensible.

This sequence is not a guarantee. It is a way of making the hidden premises of causal claims inspectable.

Claim KNO014-C16. Causal analysis is best treated as a design-and-assumption problem before it is treated as an estimation problem. [4,7,8]

Where That Leaves Us

Step by step, a disciplined analysis defines the question, contrast, outcome and time zero; states the assumed causal structure and each variable's role; fits design and estimator to that structure; checks balance, overlap and measurement; probes the estimate; and reads it as causal only insofar as its assumptions stay defensible. The analysis guarantees nothing but makes hidden premises inspectable. The next section defines the boundary of KNO-014.

15. The Boundary of KNO-014

KNO-014 is about the **methodological problem of distinguishing causal effects from non-causal associations**. It does not attempt to settle the metaphysics of causation. That conceptual question belongs to IDE-004 · Causality.

It also does not replace KNO-012 · Experiments and Variable Control. KNO-012 asks how experiments create controlled comparisons; KNO-014 asks how causal claims are identified across experimental and observational settings.

KNO-015 · Uncertainty and Probability will take up a different problem: how uncertainty is represented and propagated once the target quantity and inferential structure have been specified. Sampling variability as a competing explanation for an observed association is also handed to KNO-015 and is not developed here.

Where That Leaves Us

KNO-014 is about the methodological problem of distinguishing causal effects from non-causal associations. It does not attempt to settle the metaphysics of causation. That conceptual question belongs to IDE-004 · Causality. This defines the article's boundary and prepares the final synthesis without extending the claim beyond what the preceding sections establish.

What We Know, and How

Association and causal effect are distinct targets. Randomisation, causal diagrams, potential outcomes, negative controls, instrumental variables and target-trial emulation each make different parts of the causal problem inspectable. Their evidential force comes from design plus assumptions, not from the statistical estimate alone.

What We Still Don't Know

Observed data cannot by themselves establish causal direction, the absence of unmeasured confounding, a valid exclusion restriction or the correctness of an assumed causal graph. Those questions require external knowledge, design information and tests whose scope is itself limited.

What Could Change Our Minds

Confidence in a causal claim should change if new evidence reveals hidden confounding, a violated time ordering, measurement bias, failed negative controls, poor overlap, an invalid instrument or a target-trial specification that no longer matches the real data-generating process.

If You Remember Only One Thing

Causal inference begins when an association is linked to a clearly defined intervention question through explicit, testable design and assumptions.

Related Articles

- KNO-012 · Experiments and Variable Control
- KNO-013 · Models, Maps and Reality
- KNO-015 · Uncertainty and Probability
- KNO-016 · How Evidence Is Built: Putting the Tools Together

References

- [1] Hill, Austin Bradford. “The Environment and Disease: Association or Causation?” *Proceedings of the Royal Society of Medicine* 58(5), 1965, 295–300. DOI: 10.1177/003591576505800503. PMID: 14283879.
- [2] Hernán, M. A. “A Definition of Causal Effect for Epidemiological Research.” *Journal of Epidemiology and Community Health* 58(4), 2004, 265–271. DOI: 10.1136/jech.2002.006361. PMID: 15026432.
- [3] Holland, Paul W. “Statistics and Causal Inference.” *Journal of the American Statistical Association* 81(396), 1986, 945–960. DOI: 10.1080/01621459.1986.10478354.
- [4] Greenland, Sander; Pearl, Judea; Robins, James M. “Causal Diagrams for Epidemiologic Research.” *Epidemiology* 10(1), 1999, 37–48. DOI: 10.1097/00001648-199901000-00008. PMID: 9888278.
- [5] Pearl, Judea. “An Introduction to Causal Inference.” *The International Journal of Biostatistics* 6(2), 2010, Article 7. DOI: 10.2202/1557-4679.1203. PMID: 20305706.
- [6] Tennant, Peter W. G.; Murray, Eleanor J.; Arnold, Kellyn F.; Berrie, Laurie; Fox, Matthew P.; Gadd, Sarah C.; Harrison, Wendy J.; Keeble, Claire; Ranker, Lynsie R.; Textor, Johannes; Tomova, Georgia D.; Gilthorpe, Mark S.; Ellison, George T. H. “Use of Directed Acyclic Graphs (DAGs) to Identify Confounders in Applied Health Research: Review and Recommendations.” *International Journal of Epidemiology* 50(2), 2021, 620–632. DOI: 10.1093/ije/dyaa213.
- [7] Hernán, Miguel A.; Robins, James M. “Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available.” *American Journal of Epidemiology* 183(8), 2016, 758–764. DOI: 10.1093/aje/kwv254. PMID: 26994063.
- [8] Hernán, Miguel A.; Robins, James M. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020. Suggested citation of the authors; online text consulted separately as a revisable electronic manifestation.
- [9] Rosenbaum, Paul R.; Rubin, Donald B. “The Central Role of the Propensity Score in Observational Studies for Causal Effects.” *Biometrika* 70(1), 1983, 41–55. DOI: 10.1093/biomet/70.1.41.
- [10] Angrist, Joshua D.; Imbens, Guido W.; Rubin, Donald B. “Identification of Causal Effects Using Instrumental Variables.” *Journal of the American Statistical Association* 91(434), 1996, 444–455. DOI: 10.1080/01621459.1996.10476902.

[11] Lipsitch, Marc; Tchetgen Tchetgen, Eric; Cohen, Ted. “Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies.” *Epidemiology* 21(3), 2010, 383–388. DOI: 10.1097/EDE.0b013e3181d61eeb. PMID: 20335814.

[12] Hernán, Miguel A.; Sauer, Brian C.; Hernández-Díaz, Sonia; Platt, Robert; Shrier, Ian. “Specifying a Target Trial Prevents Immortal Time Bias and Other Self-Inflicted Injuries in Observational Analyses.” *Journal of Clinical Epidemiology* 79, 2016, 70–75. DOI: 10.1016/j.jclinepi.2016.04.014. PMID: 27237061.