

Visualizing causal hypotheses in environmental econometrics

Pierce Donovan[†]

March 2024

Abstract

Environmental economists have gravitated toward writing empirical papers with an emphasis on causal inference. Despite this development, there has not been much progress in adopting an explicit framework for communicating causal hypotheses based on prior beliefs about the structure of a data generating process. The shortfall reduces the transparency and accessibility of the assumptions underlying effect identification and limits the feasibility of causal hypotheses testing. This article explains why an explicit framework is worthwhile and demonstrates how Directed Acyclic Graphs can augment and standardize the communication of causal knowledge.

Keywords: causal models, identification, research design, specification testing, communication, directed acyclic graphs (JEL: A20, C12, C51, C52, Q50)

[†]Assistant Professor, Department of Economics, University of Nevada, Reno
email: pierce.donovan@unr.edu
web: piercedonovan.github.io

A framework for communicating causal knowledge

Most empirical papers in environmental economics today are structured around a single statistical inference of interest and written to emphasize how this measurement supports a causal interpretation (Segerson, 2019). Because causal inference is difficult, authors are rewarded for designing clever identification strategies. It is now the thoughtful communication of research design that defines “rigor” in empirical papers, rather than a barrage of regression specification tests (Angrist and Pischke, 2010).

This article considers the clarity, transparency, and testability of *causal hypotheses*—prior beliefs about the structure of a data generating process. This is thought to be addressed when testing a statistical inference’s sensitivity to varying empirical assumptions about the control variables or functional form used in an analysis, but the exercise in showing consistency across regression specifications is not undertaken with any explicit causal model in mind. While sensitivity analyses are likely to give researchers confidence in their results, they are unable to determine whether an approach estimates an effect of interest.

The spirit of specification testing is to determine the credibility of a statement such as “the effect of X on Y is likely positive.” A test of a result’s sensitivity to an assumption about a data generating process is suitable to bolster this claim because it tests whether a researcher is estimating the effect of interest. This change to a causal model will naturally motivate a specific change in an empirical strategy and highlight a critical prior belief. In contrast, a purely functional change to a regression specification fails to imply anything about the correctness of the underlying causal model. In essence, the typical “robustness test” assumes that the results of the preferred model are always causally interpretable. When causal models are not explicit, the testing of causal hypotheses becomes impossible.

Directed Acyclic Graphs (DAGs) (Pearl, 1995) systematize the creation and testing of causal hypotheses. DAGs visually display a researcher’s prior beliefs and assumptions, motivate identification strategies, and indicate the conditions under which a regression yields a causally interpretable measurement. They provide a way to choose conditioning variables, sources of data, and empirical methods in a manner that can be easily understood and validated by a wide audience of researchers, students, and stakeholders. This article introduces the visualization technique, shows how DAGs can facilitate causal hypothesis testing, assesses the costs and benefits of DAG adoption, and shares econometric insights unique to DAGs. This article seeks to lower the costs of adoption for an emerging method of broad applicability in environmental economics.

Several articles compare the DAG framework to competing foundations for presenting causal information. The potential outcomes (Neyman-Rubin) framework—which quantifies causal effects through comparisons of counterfactuals—is already widely used in applied economics to validate identification strategies and has the advantage over DAGs

in terms of the scope of empirical assumptions and conclusions that can be represented (e.g. monotonicity and the local average treatment effect) (Imbens, 2020). In contrast, the DAG framework is better at making the historically “ad hoc” facets of causal inference in economics more transparent and systematic (e.g. covariate selection and the identification of bias) (Schneider, 2020; Huntington-Klein, 2022). Heckman and Pinto (2022) claim that the scope of both frameworks is limited in non-empirical settings (e.g. general equilibrium) and promotes structural equation modeling to address causal questions.

The next two sections provide a primer on DAGs and several applications to illustrate their utility. The penultimate section demonstrates a review of a recent research article from the environmental economics literature with the aid of a DAG and describes the limitations of a purely graphical approach. The final section discusses the merits of integrating this innovation into future presentations, publications, peer reviews, and pedagogy.

A primer on Directed Acyclic Graphs

When constructing a DAG, the initial task is to center thinking around a particular relationship of interest, e.g. $T \rightarrow Y$ in Figure 1. Directed arrows like $T \rightarrow Y$ convey statements like “the outcome Y is in part determined by the status of some treatment T .” Additional causal relationships between T and Y may be mediated (facilitated) by other variables (e.g. M). A causal model should explicitly consider any variables which could distort the observed relationship between T and Y away from a causal interpretation. The complexity of a model is ultimately up to the researcher. Every inclusion or omission of a variable or arrow marks an explicit assumption about the underlying data generating process.




Figure 1: A representative DAG. Some variables may not be observable, and related links can be distinguished with dashed arrows. The use of color is a stylistic preference that aids identification.

Spurious relationships between T and Y are created by “confounding” variables—those which influence both T and Y —like X or Q . If the effects of these confounding variables are not mitigated, an estimate of the treatment effect will be biased. Conversely, variables like H will not influence the causal interpretation of a measured correlation between T and Y , since they do not contribute to a spurious relationship between T and Y .

Thus the DAG provides a way to differentiate between malignant and benign sources of variation in Y .

Ideally, there would exist some experimental source of variation for T , because this would sever any arrows leading towards T and yield a causal interpretation of the effect of T on Y . But this ideal isn't necessary. If data on X and Q are available, a simple matching strategy can control for these confounding variables by including them in a regression. Regression will automatically isolate the causal relationship between T and Y from the variation induced by changes in X or Q . In the language of DAGs, this "closes" the non-causal paths between T and Y . The DAG does not keep track of the paths that are closed by a researcher's analysis—or the paths that remain "open"—it only signals whether a path is a problem that an identification strategy needs to address.

Sometimes a confounding variable like Q isn't observable, but this may not be a problem if data on X are available. This is because of a restriction imposed on the model via the term "acyclic." By disallowing loops, DAGs remove the possibility to model bidirectional causal links, but enable statements about conditional independence, e.g. " Y is independent of Q conditional on X ." Since variation in Q can only impact Y through an intermediate effect on X , controlling for the variation in X halts the pass-through of information from Q to Y . Thus the requirement for closing a spurious path between T and Y can be relaxed to controlling for a confounding variable—or one of its "descendants"—along that path. Here, once X is controlled for, the status of Q becomes immaterial in the identification of the effect of T on Y . However, the decision to include Q (or H) as control variables may improve the *precision* of the $T \rightarrow Y$ estimate, since their inclusion reduces the unexplained variation in Y .

Because a DAG is a non-parametric representation of a causal process, it will not recommend a particular estimation procedure or functional form. For this step, economists should rely on traditional econometric knowledge. Regardless of specification, when reporting regression output, the marginal "effect" on Y attributable to T reflects all remaining open paths from T to Y . In the present example, the direct effect cannot be disentangled from the indirect effect, as the "mediator" M —which facilitates part of the effect of T on Y —is unobservable.

The researcher develops a model to determine whether the exploited variation generates a measurement of a causal effect, and the causal claim relies on an implicit assumption that the DAG accurately models the data generating process. Thus the addition of an explicit model does not remove the potential for misspecification. However, if someone disagrees with the model, that person may make changes to the DAG and review whether the new DAG admits the same identification strategy or a completely different one.

Applications of DAGs to environmental economics

This section provides evidence that discussions of research design can be made more clearly and succinctly with a DAG. Despite the simplicity of the following examples, they are difficult to explain without a DAG, which suggests that DAGs are valuable in augmenting understanding for both researchers and other stakeholders.

Bad controls are usually collider variables

Some research design strategies can inadvertently introduce bias to estimators. To illustrate, the following example concerns the long-run impact of cumulative wildfire smoke exposure on respiratory health, using hypothetical data from hospital admissions. Figure 2 provides a model of a data generating process that considers how the sample is being collected. Because smoke exposure also has negative long-run impacts on immune system health, it increases the likelihood of a hospital visit through a second causal channel that is unrelated to respiratory health. For clarity, it is assumed that no variables confound the relationship between smoke exposure and respiratory health.




Figure 2: Data collected using a selected sub-sample may fail to yield an internally-valid causal effect if the data generation itself creates spurious links between treatment and outcome.

Certain variables create spurious correlations between treatment and outcome only once they are included as controls—or in this case, held constant by the data collection approach. In the present example, weakened immune system health and weakened respiratory health are both sufficient conditions for being more likely to showing up in the hospital data. But the sufficiency of either condition implies that among hospital admits, those with poorer immune system health are more likely to have better respiratory health (and vice versa). This is not a causal claim, but a spurious correlation created by the data generating process. Since people impacted by wildfires are more likely to have worsened immune system health than the unaffected, the spurious link will contribute to an underestimation of the negative respiratory health effect.

DAGs alert researchers to potential identification pitfalls that would be difficult to explain without a graphical aid. Whenever a variable invites a collision of two arrows along a path from treatment to outcome—such as hospital admission here—the variable is called a “collider.” Controlling for a collider variable—through inclusion in the control set or the

data collection process—will open an otherwise closed path. The abstraction of the identification problem to a graphical one makes the detection and discussion of a collider bias simple. Without the concept of a collider, these biases are easily missed.

A common mistake is to reduce this example to an external validity concern—the question of whether the conclusion is valid beyond a particular experiment or sample. However, the measured relationship between smoke exposure and respiratory health won't even be *internally* valid here. The collider bias will reduce the measured smoke exposure effect for the hospitalized sub-sample. However, since the second mechanism through which smoke exposure increases the likelihood of hospital admission is known, an identification strategy that conditioned on immune system health would close the non-causal path from smoke exposure to respiratory health, even when the collider is in play. This allows the estimation of an unbiased respiratory health effect for the hospitalized population.

The role of a variable is important

Does adopting an electric vehicle decrease household emissions? The direct effect is likely positive due to increased mineral extraction and fossil fuel-derived electricity demand. However, the bulk of the adoption effect is likely indirect and negative, through the replacement of a gas-powered car. As a straw man, consider the economist who controls for the number of gas-powered cars in the household. This would clearly be a mistake, because, by controlling for the number of gas-powered cars, the researcher closed the replacement channel. That is, they have closed a causal path that should have been left open, as this strategy conditioned on a mediator, rather than a confounder.

But this economist wouldn't even measure the direct effect with their strategy, given the understanding of car buying preferences implied by Figure 3. Other relevant household characteristics will manifest a collider bias when controlling for the number of gas-powered cars. For example, households with car enthusiasts are drawn to gas-powered cars and driving more often, both for the sake of leisure. The control variable is a collider on the enthusiast path—conditioning on the number of gas-powered cars introduces a spurious correlation between electric vehicle adoption and enthusiasm.

To many environmental economists, the replacement story is probably obvious, although the collider story is likely not. But the recommended solution might be to control for the total number of cars instead. According to Figure 3, this is also a blunder. In that causal model, the total number of cars acts as another collider. By stratifying on the number of cars in the household, it is assumed that there is a one-to-one replacement of a gas-powered car for every electric vehicle adopted—which is unlikely. Some households may move from two to three total cars with their adoption, but they will be compared to

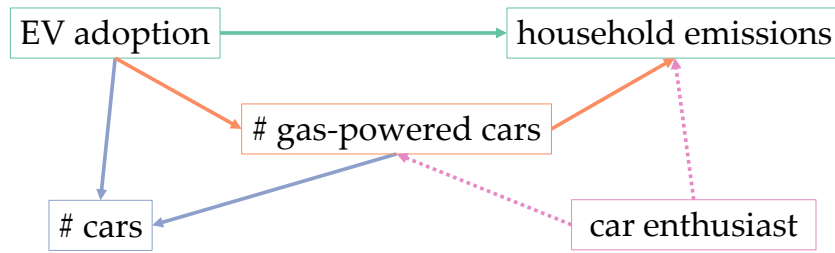


Figure 3: The true impact of electric vehicle adoption on household emissions is only revealed if controls for the number of cars in the household—gas-powered or otherwise—are omitted.

three [gas-powered] car households instead of two car households. The proposed solution moves the result from an underestimate to an overestimate of the decrease in emissions.

This model says that controls for the number of gas-powered cars and the number of total cars in the household should both be avoided. A different setting may imply a different causal structure than the one in Figure 3. In that case, the no-control strategy can be re-validated by checking the new DAG for any threats to identification.

Efficient communication of identification assumptions

Fishing quota systems have revolutionized the organization of fisheries, protecting the biological sustainability of targeted species and increasing the value of an authorized total allowable catch. In these systems, fishers explicitly own shares of this total allowable catch and are free to trade these fishing rights among themselves. The market for quota shares is intended to autonomously guide ownership into the hands of the most efficient fishers, therefore increasing the economic rents that accrue to each share.

How much of the variation in the price of quota is due to changes in fleet efficiency (i.e. the amount of fish caught per unit of effort made by the fisher), as opposed to changes in consumers’ demand for fish? A potential problem with separating these two impacts is the fact that any market development that increases the value of the quota may additionally incentivize investment that leads to the use of more efficient fishing methods. Figure 4 shows how uncontrolled demand-side factors create spurious relationships when interpreting the effect of changes in fleet efficiency on the price of quota.




Figure 4: Given the mental model above, the corresponding DAG reflects the assumptions required for the gear restriction variable to provide useful variation in fleet efficiency.

Consider a potential regulation enacted due to concerns about the bycatch (unintended catch) of an endangered species often found co-mingling with the targeted species, such as dolphins with tuna. This could manifest as a restriction on a certain type of gear or fishing location. The restriction reduces the efficiency of the fleet, impacts the value of fishing quota through this efficiency mechanism, and is independent of shifts in demand.

Figure 4 is consistent with this fishery story, and it suggests using the gear restriction variable as a solution to the identification problem. The regulation turns fleet efficiency into a collider variable along the spurious path from gear restriction to the quota price, and into a mediator on the causal path. Thus an unconditional regression of the price of quota on the gear restriction variable provides a causally interpretable estimate. A regression of fleet efficiency on the gear restriction is similarly unbiased. Dividing the two relevant regression coefficients yields the instrumental variables estimator for the effect of fleet efficiency on the price of quota—another causally interpretable estimate.

The DAG above reveals the familiar identification assumptions for the instrumental variables approach—relevance, independence, and validity—by checking for the presence or absence of three key arrows involving the gear restriction variable. The fishery story admits an arrow from gear restriction to fleet efficiency (relevance), but not to or from the price of quota or fish market development variables (validity and independence, respectively). Figure 4 transparently communicates what is believed to be true about the instrument, and readers can easily follow (or dispute) these assumptions.

Discovering new methods

The previous section used a DAG to identify multiple links between variables that could be estimated without bias. The resulting strategy emulated the familiar instrumental variables approach. But each new DAG structure provides an opportunity to discover a novel identification strategy that would have otherwise remained hidden. This last example illustrates the “random filter” identification strategy (Donovan, 2024) for deriving an unbiased estimate of a treatment effect amid selection into treatment.

How beneficial are climate resilience-motivated crop insurance programs in developing nations? The savvier farmers will likely have the most interest in an insurance program if made available, as well as the most sophisticated farming operations. Any naïve regression strategy would clearly pick up this selection bias. Figure 5 shows how to remedy this selection problem. Consider the mechanism through which insurance would lead to a benefit. If there is no adverse weather event (and thus no insurance claim), then crop insurance will have no positive impact on income. If the insurance claim variable isn’t correlated with business savvy, the two links in the chain from crop insurance to income can be estimated separately, without bias.




Figure 5: When an exogenous mediator exists between treatment and outcome, a researcher can separately identify the effects of treatment on the mediator and mediator on the outcome, then scale up the former effect by the latter.

The first link in the causal chain can be identified with a simple regression of insurance claim on crop insurance. The spurious path from crop insurance to insurance claim is not an issue, since income assumes the role of a collider and omitting income closes this path. The second link is identified by regressing income on insurance claim while controlling for crop insurance to close the other non-causal path. Multiplying these effects together recovers a causally interpretable estimate of the effect of crop insurance on income.

Figure 5 states three assumptions about this mediator. First, the mediator must intercept all causal paths from crop insurance to income. If another causal path existed outside of the mediator’s reach, the identification strategy will close that path. The other two assumptions require the insurance claim variable to have unconfounded relationships with crop insurance and income. This means the business savvy variable—or any other unobservable—may not create a spurious path between treatment and outcome that involves the mediator. If these assumptions hold, i.e. if one can genuinely draw the data generating process as in Figure 5, then the identification strategy above is valid.

The savviest of farmers could potentially mitigate their exposure to weather-related damages through siting or some other mechanism; this creates a confounding path involving the mediator. In this case, a modification to the DAG presents a more robust identification strategy. Conditional on the perceived risk of a disastrous weather event, the event itself (and thus the insurance claim) is now plausibly-exogenous with respect to the business savvy of individual farmers. Adding a measurable control variable like “risk of damage” in between the business savvy and insurance claim variables would signal that a control strategy is available to close the spurious path. A researcher should therefore run the two aforementioned regressions while controlling for exposure risk.

A DAG-enhanced review of a recent JAERE article

In early July, 2015, the water utility serving the city of Burbank, California sent warnings to households who would be violating upcoming summer irrigation restrictions based on their behavior in late June. The notices contained details about a new monitoring system that used real-time data to automatically inform the utility of irrigation outside of

permitted times. This treatment resulted in a tenfold increase in the number of households ever notified of an irrigation violation—and signaled a change in the likelihood of enforcement and threat of financial penalties should overuse continue.

West, Fairlie, Pratt, and Rose (2021) reported a substantial decrease in water consumption in response to receiving a notice—roughly 600 gallons per week, or 31% of mean household use—using a fuzzy regression discontinuity design. Treatment was predominantly determined by whether a household was found to be non-compliant with upcoming summer restrictions during a week in late June. The noncompliance algorithm counted the number of days with peak hourly water usage exceeding an arbitrary threshold (by which point irrigation was evident). Since the summer restrictions only allowed for irrigation two days per week, when the third-highest daily peak consumption hour exceeded 125 gallons (hereafter “peak water consumption”), a household would be deemed non-compliant.

The utility allowed for some leniency for historically efficient customers, so a household’s peak water consumption crossing the above threshold only implied a jump in the *probability* of receiving a warning. The lower-volume customers were instead evaluated based on their fifth-highest daily peak consumption hour, but this “consumption tier” information was not available to the research team. The imperfect compliance in treatment assignment using the stricter rule supported a fuzzy regression discontinuity approach, detailed below.

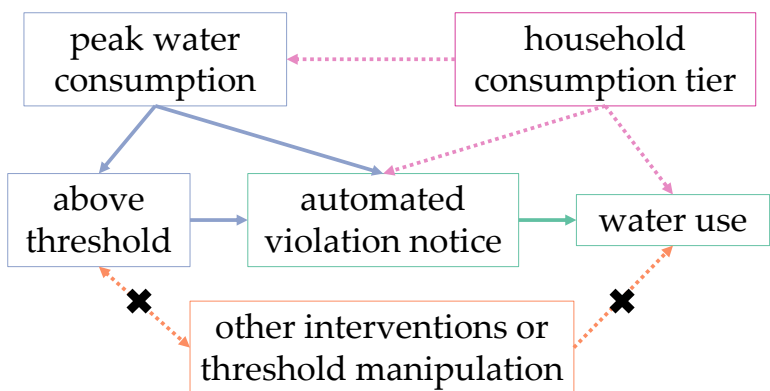


Figure 6: Sketching the DAG consistent with an article’s narrative highlights key assumptions, barriers to identification, sources of potential bias, and the validity of the identification method.

The nature of the omitted variable bias becomes clearer with the DAG in Figure 6—drawn to represent the authors’ understanding of the data generating process. Those who use less water historically would likely continue to do so and not receive a notice, while the higher baseline households would be treated. The unbalanced treatment and control groups biases the magnitude of the treatment effect downward.

Figure 6 aids a reviewer in running through the requirements for a satisfactory fuzzy

regression discontinuity design. The previously mentioned instrumental variables assumptions are conditionally satisfied. Clearly, the instrument—indicating if the stricter peak consumption threshold is exceeded—partly determines whether a household receives an irrigation restriction violation notice (it was a necessary precondition). And because the crossing of the peak consumption threshold is defined entirely by a household’s peak water consumption, the instrument is independent of the unobserved confounder—the household consumption tier—after conditioning on peak water consumption. Lastly, the instrument only impacts household water use through the delivery of a notice.

There are still the assumptions specific to the regression discontinuity approach to discuss. First, the goal of this regression discontinuity design is to separate the drop in water use due to receiving a notice from the [increasing] water use trend in the peak water consumption variable. Upon controlling for peak water consumption, any paths that utilize continuous variation in water use across the threshold are closed, leaving only the potential discontinuity to be measured—as desired.

Second, there shouldn’t be an additional cause for a discrete change in water use when comparing households just over the threshold to those just below it. For expositional clarity, the “local randomization” assumption is symbolized here by a non-path from the threshold to the outcome using the bold ‘X’s in Figure 6. Because water use will be regressed on the peak consumption threshold—and not the notice status itself—the authors hope to attribute the threshold-derived jump in outcomes to the notice, rather than some other factor. They make this assertion easy to believe, as no concurrent intervention utilizing the threshold existed (eliminating the arrow from the threshold to a potential intervention mediator), and no households were privy to the threshold or the automated detection algorithm (eliminating the arrow from a potential sorting confounder to the threshold).

From this discussion, the authors’ fuzzy regression discontinuity approach appears to generate an unbiased estimate of the effect of the irrigation restriction violation notices on summertime water use. But the above DAG cannot fully cover the scope of a peer review. This is discussed below.

DAGs are not the be-all and end-all

It should be clear that DAGs cannot provide a one-to-one mapping to everything commonly taught in econometrics classes, and for this reason their main role must be augmentative. For example, Figure 6 cannot suggest a particular implementation of fuzzy regression discontinuity (two stage least squares, local linear regression, bandwidth, etc.), and these decisions involve a separate set of assumptions from the causal ones shown by the DAG. The DAG does provide some guidance, however. As the introduction of this article establishes, the sensitivity of results to these additional assumptions should be an-

alyzed across specifications consistent with the same causal model.

Another shortcoming is that a DAG misses nuance relating to the estimate of interest. In the [common] case of heterogeneous treatment effects, an instrumental variables strategy typically makes use of a monotonicity assumption to show that the approach identifies a local average treatment effect. The additional reliance on the extrapolation of counterfactuals in regression discontinuity designs means that this treatment effect is also limited in scope to settings “close to” the threshold—for instance, in the case of West et al. (2021), where the peak consumption threshold for notice eligibility is around 125 gallons. The traditional potential outcomes logic embedded in these statements does not map to any feature in the graph.

An analysis aided by DAGs will still be subject to model misspecification. Where variations of functional form may instill a false sense of confidence in an inference’s robustness, a DAG may instill a similar confidence in a causal model. The interpretability of a result still depends on the accuracy of the causal model—regardless of whether it is explicitly represented by a DAG. While model correctness remains a strong assumption, the transparency provided by a DAG admits an opportunity to test it during peer review. A reviewer can start with a DAG that communicates the assumptions that must hold in order for the empirical strategy to be successful. If the setting in a paper admits an arrow inconsistent with the approach, the reviewer has a reason to be critical. If an author supplies a DAG that doesn’t seem to match their setting, the reviewer can point this out as well.

Even if environmental economists become more comfortable with the graphical approach, it is evident that additional discussion outside of the DAG framework will still be needed in empirical research. And as seen in the primer, certain questions involving non-recursive (bidirectional) model designs cannot be conveyed via DAGs. These two facts are occasionally taken as evidence that DAG adoption is impractical (e.g. Imbens (2020); Heckman and Pinto (2022)), but this extreme view ignores their potential for improving the communication of causal knowledge in a majority of settings. In reality, nothing is sacrificed by using a new tool in conjunction with the old ones.

Opportunities for implementation

Identification strategies are often nuanced and developed from a place of deep familiarity with specific data and settings. The inclusion of a DAG increases the comprehensibility of research because it makes the proposed identification strategy clear to others. This can lead to higher-quality feedback. It is easy to evaluate whether a researcher’s DAG represents their setting, which allows a reader to suggest whether a variation in the DAG still admits the use of the proposed identification strategy.

DAGs provide an opportunity to make the research process more transparent and responsive to criticism. When reviewers ask for some sort of sensitivity analysis, they can use the graphical language for raising concerns about causal hypotheses. A new causal hypothesis generates a modification to a DAG, which then organically motivates a new regression. This first-principles approach to testing model robustness is more academically rigorous than suggesting modifications to a regression with language concerning functional form alone. It matches the spirit of specification testing, which aims to build confidence that an effect of interest is being measured without bias.

The integration of DAGs in econometrics education creates significant value for pedagogy (e.g. Cunningham (2021); Huntington-Klein (2022)). Indeed, graduate studies in economics already highlight causal effect identification, yet several of the insights shared in this article are absent from applied econometrics instruction. At the undergraduate level, DAGs facilitate discussion of research design; for example, the introductory econometrics course at the University of Nevada, Reno attracts a diverse group of students to economics through its primary emphasis on discussing empirical research. The use of DAGs increases access to this class and builds on a desirable competency.

That accessibility can be extended to communication with policymakers and stakeholders. By simplifying explanations of causal effect identification while retaining rigor, DAGs can elevate a non-expert's understanding of an empirical approach. Perhaps the most promising consequence is that lawmakers would have an opportunity to engage with empirical work more critically, thus reducing reliance on trust as the key mediator for informing environmental policy.

References

- Angrist, J. D. and Pischke, J.-S. (2010). The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics. *Journal of Economic Perspectives*, 24(2):3–30.
- Cunningham, S. (2021). *Causal Inference: The Mixtape*. Yale University Press.
- Donovan, P. (2024). The random filter identification strategy. *Working Paper*. https://piercedonovan.github.io/files/research/donovan_random_filter.pdf.
- Heckman, J. J. and Pinto, R. (2022). The Econometric Model for Causal Policy Analysis. *Annual Review of Economics*, 14(1):893–923.
- Huntington-Klein, N. (2022). *The effect: an introduction to research design and causality*. CRC Press, Taylor & Francis Group.
- Imbens, G. W. (2020). Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics. *Journal of Economic Literature*, 58(4):1129–1179.
- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.

- Schneider, E. B. (2020). Collider bias in economic history research. *Explorations in Economic History*, 78:101356.
- Segerson, K. (2019). Reflections—On the Role of Theory in Contemporary Environmental and Natural Resource Economics. *Review of Environmental Economics and Policy*, 13(1):124–129.
- West, J., Fairlie, R. W., Pratt, B., and Rose, L. (2021). Automated Enforcement of Irrigation Regulations and Social Pressure for Water Conservation. *Journal of the Association of Environmental and Resource Economists*, 8(6):1179–1207.