Deterministic by design: causal inference challenges for ‘biodiversity change’ syntheses

Rebecca Spake[1], Richard McElreath[2], Luke Christopher Evans[3,4]

1 Environmental Change and Sustainability, School of Geography and Environmental Science, University of Southampton, Southampton, UK, SO17 1BJ, UK. r.spake@soton.ac.uk

2 Max Planck Institute for Evolutionary Anthropology, Deutscher Platz 6, 04103 Leipzig, Germany. richard_mcelreath@eva.mpg.de

3 UK Centre for Ecology & Hydrology, Wallingford, OX10 8BB, UK.

4 Biodiversity, Land Use and Investment directorate, Defra, London, United Kingdom of Great Britain and Northern Ireland. lukechristopherevans@gmail.com

Abstract

Understanding biodiversity change over space and time is a central goal of ecology. A common approach to asking why biodiversity is changing over time, is to collate biodiversity time series from multiple sites and undertake a two-step analysis: First, site-level estimates of ‘biodiversity change’ are calculated. Second, those change estimates are regressed on putative environmental drivers, such as climate, biome, and land-use change. Such change-driver relationships are typically interpreted without reference to an explicit estimand and due consideration of the data-generating process, risking causal misinterpretation. We demonstrate this first logically, using simple examples and directed acyclic graphs (DAGs). We show that ‘biodiversity change’ calculations are deterministic variables, and change-driver associations thus represent an aggregate effect of a driver across historical and contemporary causal paths. Using simulations of species richness across space and time, we then demonstrate that multiple distinct data-generating processes, including confounding, species-pool constraints that impose ceiling/floor effects, and detection bias, can produce similar change-driver associations, highlighting the importance of understanding the data-generating process. We caution against using biodiversity-change calculations without consideration of the underlying data-generating process, and present DAG-centric guidance for researchers asking questions about biodiversity change across time and space, thereby making assumptions and estimands explicit.

Keywords: Biodiversity trend; Causal diagram; Change score; Effect modification; Meta-analysis; Meta-regression; Regression to the Mean; Synthesis

1. Introduction

Understanding how and why biodiversity is changing is a fundamental goal in ecology (1,2). Going beyond change detection and attributing changes to causal drivers is essential for informing targeted conservation interventions. For example, identifying drivers that accelerate or buffer biodiversity loss can guide where restoration or protection efforts might be most effective. Biodiversity change studies can be highly influential yet can generate considerable scientific debate (3). Robust inference is therefore essential, yet numerous analytical challenges exist, many of which are well recognised, including sampling biases (4) and non-independence arising from phylogenetic, spatial, and temporal structure (5). The challenges concerning causal inference have been less appreciated until more recently (e.g., (2)), and so ecologists do not usually embed causal reasoning into their research workflows to explicitly justify analytical choices. Instead, choices are based on highly variable heuristics (6), so that research teams apparently asking the same question of the same data reach different conclusions (7). Biodiversity-change research is particularly vulnerable to this problem, which we aim to highlight here.

A common approach to asking how and why biodiversity is changing uses biodiversity time series collated from multiple sites distributed across large, heterogeneous spatial extents (8,9). It typically proceeds in two steps (Figure 1). First, to ask *how* biodiversity (Y) is changing, researchers estimate biodiversity change (or 'trends') (ΔY) for each site, for example, calculating biodiversity differences between two time points, or for a longer time series, by regressing biodiversity on time and estimating the slope. The frequency distribution of change estimates within the sample is interrogated to ask, for example, how variable biodiversity change is across sites, and how it is changing on average (10). Second, to ask *why* biodiversity is changing, these change estimates are typically regressed on putative 'drivers' (Z) such as climate, biome, conservation status, or taxon (or regressions specify interactions between time and Z). Drivers may change estimates themselves, for example, temperature change over time. Figures of change-driver (ΔY ~ Z) relationships are common in the literature. They are a seemingly intuitive means of conveying how the magnitude and direction of change vary with such drivers, e.g., yielding interpretations such as 'warmer sites are declining faster' or 'species gains are greater at lower latitudes'.

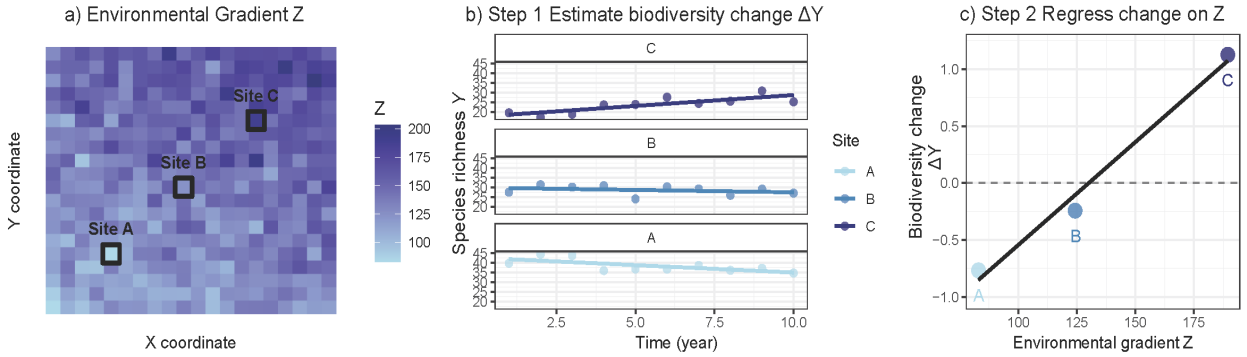

Figure 1. Demonstration of the common two-step approach for asking how biodiversity change varies across environmental gradients. a) Three study sites (A-C) distributed across a spatial environmental gradient Z. b) Biodiversity time series for the three sites showing species richness Y over ten years. In Step 1, site-level biodiversity change (ΔY) is estimated for each site as the fitted slope. c) In Step 2, site-specific slopes from Step 1 are regressed against the environmental driver Z. The dashed horizontal line indicates zero change.

However, analyses of change variables (such as biodiversity change) are not straightforward with observational data, because change constitutes a 'deterministic variable' (11–13). Consider biodiversity Y (e.g., species richness) measured at two time points, Y_t and Y_{t+1}. Any measure of change over time (e.g., $\Delta Y = Y_{t+1} / Y_t$) is completely determined by Y_t and Y_{t+1}. Consequently, estimates of change~driver relationships (i.e., $\Delta Y \sim Z$) reflect some composite of multiple, possible causal paths, including $Y_t \rightarrow Y_{t+1}$, $Z \rightarrow Y_t$ and $Z \rightarrow Y_{t+1}$ relationships (12). The same issue applies to regression slopes from longer time series: whether estimated separately for each site (Y ~ time) or estimated from interaction models (Y ~ time × Z); slopes remain derived variables that are mathematical functions of all observed biodiversity values across time. It follows that causal inference is challenging when the outcome variable of interest is deterministic, and risks leading to misinterpretations of the causes of biodiversity change over time.

Ecology is undergoing a causal revolution (14,15). Increasing calls to adopt causal frameworks in analyses of biodiversity change (2,16–18) promote the use of directed acyclic graphs (DAGs): tools that graphically describe the causal relationships within a system of interest. DAGs help identify appropriate analytical strategies for given questions and transparently document causal assumptions (14,19). DAGs are particularly useful for illuminating the challenges specific to analyses involving change variables, which are especially prone to misinterpretation (11,12,20). In this paper, we apply causal frameworks and DAGs to biodiversity-change questions to demonstrate how change-driver relationships can be misinterpreted. Using simulations of biodiversity time-series, we demonstrate that multiple, distinct causal processes can yield similar

change-driver relationships, highlighting the importance of causal thinking about the data-generating processes at play. We conclude by proposing that researchers clearly state their target estimand of interest and identify an appropriate analytical strategy for its estimation using DAGs. While our focus is on biodiversity change, the workflow we propose is relevant wherever deterministic variables are used as outcomes in observational research, including meta-analyses.

1.2 Depicting biodiversity change with Directed Acyclic Graphs (DAGs)

Most scientific data analyses have the logical status of conjectures: the assumptions that make them valid for causal inference are not stated and are usually unknown. Instead of using logic to decide how to process data, researchers rely upon heuristics and conformity. For centuries, that was the best we could do. But in the 1970s, statisticians began to make causal inference axiomatic (21,22). It became possible to prove which minimal assumptions are required for an analytical procedure to be valid for a formally defined goal. In the 1990s, this framework was extended to structural causal models represented by graphs of causal relationships (23), including DAGs. Many fields, including ecology, have begun to incorporate this “causal revolution” into their practice and training. We provide a brief introduction to this approach and then use it to explain why analysis of change variables usually requires unrealistic assumptions.

A causal DAG is a graphical representation of the hypothesised causal relationships between a set of variables. Any two variables in the DAG may be connected by a unidirectional arrow, which signifies that the first ‘parent’ variable exerts a causal effect on the second ‘child’ node. By depicting the data-generating processes in a system of interest, DAGs illustrate potential sources of confounding and selection bias and enable researchers to algorithmically identify strategies to address them.

Although DAGs can clarify the assumptions required for valid causal inference, many ecological analyses continue to rely on heuristics. Ecologists asking questions about biodiversity change commonly assert that some reference to a baseline is required to detect and attribute biodiversity change to drivers. This intuition perhaps explains the widespread use of change calculations, ΔY, in biodiversity-change analyses, as a means of somehow incorporating or controlling for the baseline. Yet recent studies have used DAGs to clarify challenges inherent to interpreting analyses that use ΔY (11,12,20). The challenges that concern us here are first that change variables are not measurements, but rather calculations, and second that these calculations do not automatically statistically control for baseline outcome values.

1.2.1 ΔY is entirely determined by its parents

Change variables are not measurements, but rather calculations. Any measure of change, whether calculated as an arithmetic difference $Y_{t+1}-Y_t$, a ratio Y_{t+1}/Y_t, a per cent change or a regression slope (Y~time), is a deterministic variable that is derived, or functionally created from, parent variables (DAGs in Figure 2). Determinism means that Y_t, Y_{t+1}, and ΔY are tautologically associated purely by an arithmetic procedure, a data transformation (24). And this transformation is destructive: it is not possible to reverse the transformation to recover the original values Y_{t+1} and Y_t, and the same ΔY can represent ecologically distinct trajectories. A change of ΔY=1 species may be ecologically quite different if the change is from $Y_{t=1}$ to Y_{t+1}=2 or instead from $Y_{t=10}$ to Y_{t+1}=11. For simplicity, we represent richness change in our DAGs as the difference between two time points (Y_t and Y_{t+1}), although the same logic applies to regression slopes from longer time series.

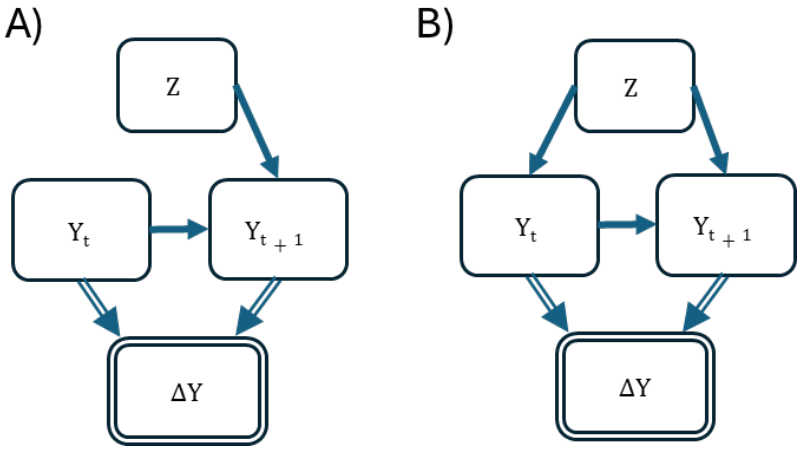

Figure 2. Directed acyclic graphs (DAGs) depicting the deterministic nature of change variable, showing the relationships between a biodiversity variable at baseline (Y_t) and follow-up (Y_{t+1}), and the change variable ΔY; and a potential driver of interest, *Z*. We present change variables by a double-outlined node and double lines for all arrows entering a change variable to denote that they are part of a functional, not probabilistic, relationship (12). A) Driver Z influences biodiversity at time t+1 only. B) Driver Z affects both Y_t and Y_{t+1} (baseline biodiversity Y_t is a mediator of Z's direct effect on recent biodiversity Y_{t+1}).

1.2.2 Calculating ΔY does not necessarily 'control' for baseline values

By showing how ΔY is constructed, we can use DAGs to see that the imposed arithmetic coupling does not necessarily 'control for' or 'remove' the influence of

baseline Y_t on ΔY. Instead, ΔY will inherit causal paths from its parents Y_t and Y_{t+1}. To demonstrate this point, we next consider how ΔY behaves in two simple causal systems involving a single driver variable Z (Figure 2), and consider whether effects of interest are identifiable or not.

Before considering statistical identification, we must carefully consider precisely which causal effect we want to estimate. In causal inference, this quantity is called the estimand (25), and corresponds to a specific question. In the context of biodiversity-change studies, two estimands are particularly relevant. The first is the total causal effect of driver Z on ΔY, which captures how Z influences the difference between Y_t and Y_{t+1}, through all of the causal pathways involved. The second is the effect of Z on Y_{t+1}, conditional on Y_t (which, for a DAG with mediation by Y_t, corresponds to the direct effect of Z on contemporary biodiversity). In ecological terms, this effect captures the effect of *Z* on contemporary biodiversity, independent of its influence on past richness. These estimands coincide only under the rare case that Z does not influence Y_t (Case A below).

Case A: Z affects Y_{t+1} only (Z does not influence baseline)

In Case A (Figure 2A), recent biodiversity Y_{t+1} is determined by both baseline biodiversity Y_t and a driver Z that influences Y_{t+1} only. Here, Z might represent a recent environmental change that occurs after baseline, such as conservation interventions or disturbances that directly affect colonisation-extinction dynamics. Under the strong assumption that Y_t is not affected directly or indirectly by Z (i.e., $Y_t \perp Z$), an observed $\Delta Y \sim Z$ relationship will represent the causal effect of Z on biodiversity change, unconfounded by past effects on richness.

In this special and unrealistic case, the total causal effect of Z on ΔY is equivalent to the total effect of Z on Y_{t+1}, because Z has no causal path to Y_t. This is the only case where ΔY can be used to isolate the effect of Z on biodiversity change. However the use of ΔY is unnecessary, because the baseline Y_t is not a confound. It is not necessary to control for it. Additionally, this scenario is rare in observational biodiversity settings, because most environmental gradients that directly or indirectly influence current species persistence (such as climatic variables) will have also shaped past richness, either directly or indirectly through shared causes. This scenario would therefore occur under controlled conditions, such as experimental manipulations and novel environmental pressures absent during baseline community assembly. Yet, this scenario most closely aligns with what researchers are implicitly asking when they enquire what 'drives' biodiversity change; identifying variables that influence ongoing ecological dynamics, rather than historical ones.

Case B: Z affects both Y_t and Y_{t+1} (baseline biodiversity is a mediator)

Here (Figure 2B), an environmental driver Z causes biodiversity at both points in time, and later biodiversity Y_{t+1} is also influenced by its historical values Y_t. In Figure 2B, we see that three causal paths exist between Z and ΔY: i) $Z \rightarrow Y_t \rightarrow \Delta Y$; ii) $Z \rightarrow Y_{t+1} \rightarrow \Delta Y$; and iii) $Z \rightarrow Y_t \rightarrow Y_{t+1} \rightarrow \Delta Y$. In this example, any association between ΔY and Z (i.e., a ΔY ~ driver relationship) consists of the direct effects of Z on Y at two time points (i and ii), as well as the indirect effect of Z on Y_{t+1}, mediated by Y_t (iii).

One might think that regressing the change ΔY on Z would suffice to identify the causal effect of Z on change from baseline. But this is incorrect. The reason is that the baseline Y_t contributes twice to the association between Z and the change ΔY: through the influence on Yt+1 and through the direct influence on ΔY. If we don't adjust for Y_t, then the estimate of the effect of Z on Yt+1 can be biased even with infinite data. If all effects are additive, then by the rules of path analysis, the covariance between ΔY and Z is:[1]

$$\mathrm{Cov}(\Delta Y,Z) = \mathrm{Cov}(Y_{t+1},Z) - \mathrm{Cov}(Y_t,Z) = P_{Z,Yt+1} + P_{Z,Yt}\,(\,P_{Yt,Yt+1} - 1\,)$$

Where the $P_{A,B}$ indicates the path coefficient from A to B. Only the term $P_{Z,Yt+1}$ is relevant for estimating the causal effect of Z on Yt+1. The rest is bias through Z's influence on Y_t. For these confounding terms to equal zero, we require that the effect of Y_t on Y_{t+1} is exactly 1, i.e., exact proportionality. But if there is any other trend in the outcome or a ceiling or floor effect, where Y is constrained to minimum or maximum values, $P_{Yt,Yt+1} \neq 1$. Thus, in Case B, ΔY does not isolate the effect of Z on Yt+1, unless we make unrealistic assumptions.

But all is not lost. To determine a valid statistical strategy for estimating the causal effect of Z, without assuming that $P_{Yt,Yt+1} = 1$, we can use an algorithm called the back-door criterion (26). Given a DAG and a cause of interest Z, the back-door criterion determines which sets of covariates are logically sufficient to identify an unbiased estimate of the causal effect of Z. Applied to this example, it is possible to block all "back-door" non-causal paths between Z and Yt+1 by conditioning on Y_t. So a regression of ΔY on both Z and Y_t can identify the causal effect of Z on Y_{t+1}. But so can a regression of Y_{t+1} on Z and Y_t. An important insight is that the use of the change variable ΔY does nothing to control for influence through the baseline.

---

[1] Translating the DAG in Figure 2B to structural equations, we have $Y_t = aZ$, $Y_{t+1} = bZ + cY_t$, $\Delta Y = Y_{t+1} - Y_t$. Set all variances to 1 for simplicity. Then the covariance between two variables. the total covariance between Z and ΔY is $\mathrm{cov}(Z,\Delta Y) = \mathrm{cov}(Z,Y_{t+1}) - \mathrm{cov}(Z,Y_t) = [b+ac] - [a] = b + a(c-1)$. Only when c=1 is the true direct effect b recovered without bias. If c<1, then the analysis will underestimate the influence of Z on Y_{t+1}.

1.2.4 What causal estimands should biodiversity-change studies target?

We have shown that while the total causal effect of Z on ΔY can be technically identifiable, it is not necessarily what researchers intend to know. This brings us to the question: What causal estimands should studies target? The goals of biodiversity-change studies are broad, ranging from informing macro-ecological theories to guiding where restoration or protection efforts might be most needed. Which estimand to target depends on the exact question. We have introduced two relevant causal estimands: 1) the total causal effect of Z on ΔY, and 2) the direct effect of Z on contemporary biodiversity Y_{t+1}, conditional on Y_t. In most cases, we argue that the former is most likely to mislead, and the latter estimand is likely to be more informative.

For researchers interested in how communities are tracking climate change, such as leading-edge theories that species are shifting their ranges poleward and producing apparent richness gains at high latitudes and losses at low latitudes, the total causal effect of *Z* on ΔY is not useful. The total effect does not isolate the biodiversity change dynamics occurring between time t and t+1. The more appropriate estimand here is the direct effect of Z on Y_{t+1}, conditional on Y_t, which captures Z's effect on contemporary diversity independent of its past effect. For understanding drivers of contemporary biodiversity, i.e, where biodiversity is being lost now, the direct effect of Z on Y_{t+1} conditional on Y_t is the most informative estimand. Identifying contemporary drivers of loss requires us to separate the influence of current pressures from past conditions and isolate the contribution of present-day drivers to present-day biodiversity.

The key insight here is that the estimand that researchers target will depend on the question. Change variables are appropriate in some contexts (12,20), but they are not universally appropriate and are never necessary. Some methodologists flatly recommend against analyses of change variables, arguing they are 'not of causal interest' (11,13). What is needed, in every case, is an explicit estimand and a logical derivation of an analysis strategy that does not simply assume that using a change calculation for an outcome correctly controls for effects of the baseline.

2 Methods

The examples above demonstrate the basic problems with change variables, but they are too simple to represent the diversity of inferential settings ecologists face. So here we further demonstrate challenges to causal inference about biodiversity change by simulating four scenarios and applying the common two-stage analytical approaches to each. Despite different underlying mechanisms, ranging from statistical artefacts to 'real' ecological processes, we show that all scenarios can produce similar change-driver (ΔY

~ Z) associations, and are therefore observationally indistinguishable using conventional two-step approaches, highlighting the need for explicit causal thinking in biodiversity change analyses.

We present each scenario using DAGs that show the relationships between an environmental gradient Z and species richness Y at different time points t. For each scenario, we simulate biodiversity at $N = 200$ sites distributed along a spatial environmental gradient Z. For each site, we generate a time series of richness over 20 years with evenly spaced, annual observations. We presume perfect detection of species richness in all scenarios, except for *Detection bias* (described below), which includes varying detection probabilities over time.

We refer to our scenarios as: 1) *Process effects*, wherein *Z* influences recent biodiversity dynamics only; 2) *Confounding*, where *Z* shapes both historical and recent colonisation-extinction dynamics via unobserved variables; 3) *Compression*, where species pools impose ceiling and floor effects on richness values; and 4) *Detection bias*, where observed ΔY~Z patterns arise due to detection biases that depend on both *Z* and time. Each scenario is summarised in Table 1 and represented by DAGs. While our simulations generate twenty-year time series for each site, for simplicity/conceptual clarity, our DAGs represent biodiversity at two time points, Y_t and Y_{t+1}. As we use a twenty-year time series, our change metric used in the two-step approach is the rate of log-biodiversity change per unit time (a proportional change metric). In all four scenarios, the total causal effect of Z on biodiversity change is identifiable, but identifiable does not imply interpretable: the total effect has a different causal interpretation across scenarios (summarised in Table 1). All analyses were conducted using R 4.5.1 (27). The R script to reproduce our simulations is available [**https://github.com/LukeChrisEvans/deterministic_by_design_code**].

Table 1. Four different causal scenarios for biodiversity change, shown with species richness Y observed at two time points (Y_t and Y_{t+1}). Perfect species detection is assumed for all scenarios except Scenario 4 *Detection bias*, in which the observed, imperfectly measured richness is denoted with asterisks. Unobserved variables are shown in light gray.

Scenario	**Interpretation of the total causal effect of Z on ΔY, and alternative estimands**
Process effects 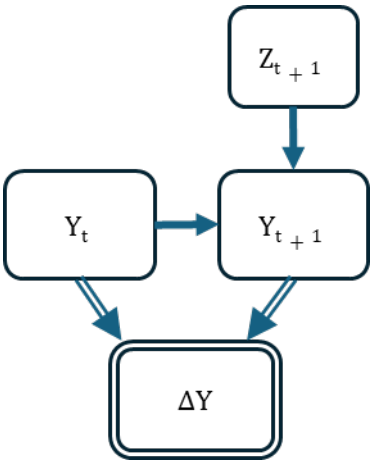 Perfect species detection assumed	The total causal effect of Z on ΔY is identifiable. To estimate the direct effect of Z on Y_{t+1} or ΔY, no adjustment is necessary. In this case, ΔY ~ Z and Y_{t+1} ~ Z estimate the same causal effect if, and only if, Y_t is independent of Z ($Y_t \perp Z$), i.e., Y_t and Y_{t+1} have no shared causes. The assumption that $Y_t \perp Z$ is unrealistically strong. In real settings, Y_t and Y_{t+1} will share common causes, directly or indirectly. When Y_t and Z are associated, ΔY ~ Z and Y_{t+1} ~ Z will estimate different causal quantities.
Confounding 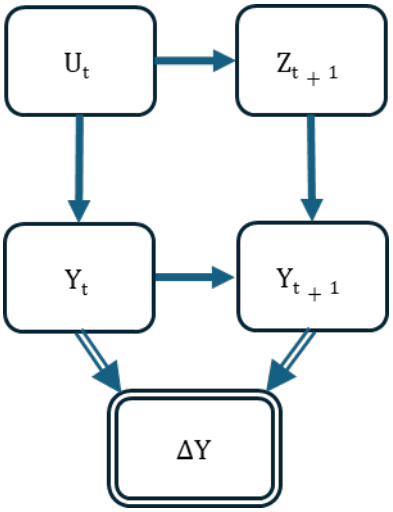 	The total causal effect of Z on ΔY is identifiable. This effect (ΔY~Z association) aggregates multiple causal pathways and, because these pathways cannot be separated using ΔY alone, the regression conflates historical and contemporary effects. The direct effect of Z on Y_{t+1} is not identifiable without adjustment, because a back-door path is open: $Z_{t+1} \leftarrow U_t \rightarrow Y_t \rightarrow Y_{t+1}$, unless we adjust for U_t. However, U_t is unmeasured or unconsidered. By adjusting for Y_t however, the effect of Z on Y_{t+1} can be identified.

U_t is an unobserved or unconsidered site-level confounder Perfect species detection assumed U_t causes sites to have both higher Z and higher Y	
Compression 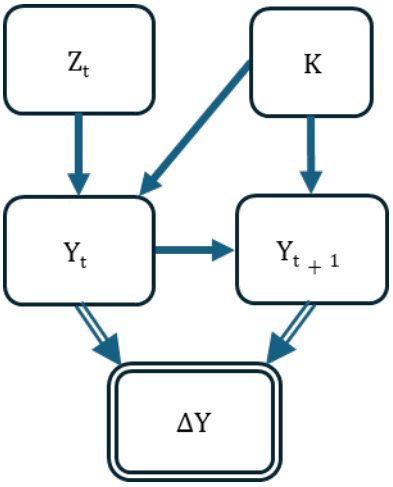 K denotes a maximum, ceiling species richness set by the regional species pool Perfect species detection assumed	The total causal effect of Z on ΔY is identifiable. The ΔY~Z association is driven only by Z setting the baseline, not recent biodiversity. Z has no causal effect on the process generating recent biodiversity, beyond its historical influence (via Y_t). There is no direct effect of Z on Y_{t+1} to estimate for this DAG. The association ΔY ~ Z is therefore dominated by ceiling and/or floor effects, not by any causal effect of Z on biodiversity change in an ecological sense.
Detection bias Z environmental gradient D detection bias (unobserved) Y true species richness Y* observed species richness t time ΔY* observed richness change	There is no effect of Z on Y_t, Y_{t+1} or ΔY to estimate for this DAG. Any ΔY~Z association is driven by measurement error. Variation in Z causes variation in the discrepancy between ΔY* and true ΔY, not ecological change itself. The issue here is that observed Y* is not a good

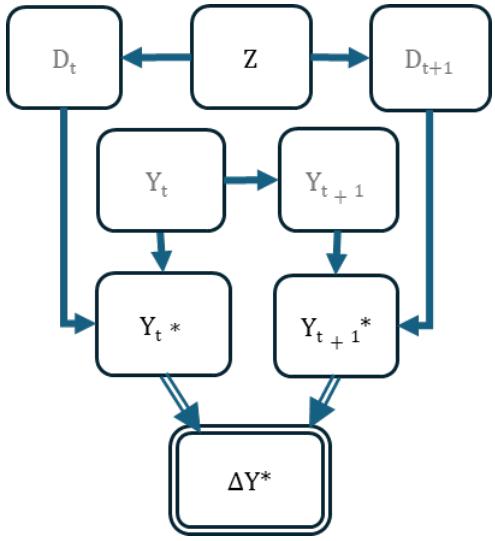	estimate of true Y.

2.1 Process effect

We first consider a simple scenario ('Process effect'; Table 1), where richness Y_{t+1} is determined by Y_t and also a driver Z, which has occurred after baseline richness Y_t (analogous to Scenario A). Baseline richness Y_t is independent of Z ($Yt \perp Z$), meaning that the gradient has not shaped historical community richness.

We simulated this as follows:

$Y_{i,1} \sim Poisson(\lambda)$

Provides a random latent abundance at each site that is independent of Z. Z_2 is a post-baseline gradient (e.g., disturbance intensity, conservation intervention) that drives exogenous change C_i, which represents the site-level annual rate of change; the balance of colonisations and extinctions occurring throughout the observation period is modelled on the log scale with process error.

$C_i = \gamma_0 + \gamma_Z Z_{2,i} + \varepsilon_{C,i}, \qquad \varepsilon_{C,i} \sim N(0, \sigma_C^2)$

Richness at each subsequent time point is determined by the richness of the previous time point, plus the site-specific rate of change and demographic stochasticity, with the observed count impacted by detection error:

$\log(\lambda_{i,t+1}) = log(Y_{i,t}) + C_i,$ t=1,…,T

$Y_{i,t+1} \sim Poisson(\lambda_{i,t+1})$

$$Y*_{i,t+1} \sim Binomial(Y_i, 1)$$

Over *T* years, richness follows a linear trend with slope C_i plus noise. In this scenario, the observed richness (Y*) is equal to the true latent abundance (i.e. no detection error), but we introduce this common structure as it is applied to all our scenarios, with the detection error varying in the final scenario.

2.2 Confounding

In the next scenario ('Confounding'; Table 1), an unmeasured factor U_t influences both baseline richness Y_t and Z_{t+1}. This generates two paths for the effect of U_t on Y_{t+1}, one path mediated by Z_{t+1} and one mediated through $Y_{t..}$ This scenario might occur when factors such as historic land use change impact the baseline biodiversity level but also impact processes that lead to varying rates of contemporary biodiversity change. For example, if land use history determines contemporary disturbance intensity or the likelihood of extreme events. In such scenarios, the meaning of ΔY~Z is conflating U_t's historical effects on richness through Y_t with the contemporary effects through Z that influence biodiversity dynamics. In our simulation for illustrative purposes, U was parameterised to positively influence Z but negatively influence initial richness, inducing a positive association between Z and temporal change that partially offset a negative direct effect of Z. U was generated but withheld from the analyst or simply not considered, replicating the observational setting in which only Z, Y_t,* Y_{t+1}*, and ΔY are considered in the analysis.

We simulate this as follows. Baseline richness is determined by U, an unobserved (or not considered) factor.

$\log(\lambda_{i,1}) = \mu_0 + \beta_Z U_i + \varepsilon_{i,1}, \qquad \varepsilon_{i,1} \sim N(0, \sigma_0^2)$
$Y_{i,1} \sim Poisson(\lambda i, 1)$

The same gradient *U* generates an unobserved post-baseline variable *Z2*, representing an unmeasured contemporary environmental process:

$z_{2,i} = \alpha 0 + \alpha_U U_i + \varepsilon_{U,i}, \qquad \varepsilon_{Z,i} \sim N(0, \sigma_z^2)$

Z2 drives change *C*, which represents the balance of colonisations and extinctions occurring after baseline:

$Ci = \gamma 0 + \gamma_z z_{2,i} + \varepsilon_{C,i}, \qquad \varepsilon_{C,i} \sim N(0, \sigma_C^2)$

Latent richness changes over time as a linear trend on the log scale with slope $C_i$plus noise

$log(\lambda_{i,t+1}) = log(Y_{i,t}) + C_i + \eta_i, t,$ t=1,…,T

$$\eta_i, t \sim N(0, \sigma\eta 2)$$

The latent richness is then translated into observed richness, as in scenario 1.

2.3 Compression

In a third scenario ('Compression'; Figure 2B), species richness, like many ecological variables, is subject to ceiling and floor effects. Species richness cannot fall below zero (the floor) and is constrained by regional species pools K (the maximum, ceiling amount). Sites with initially low richness can only increase or remain constant, while sites with initially high richness can only decrease or remain constant. If baseline richness Y_t varies systematically along an environmental gradient Z_t, compression will generate apparent ΔY~Z relationships.

Baseline richness is set by Z, and rather than representing community dynamics more mechanistically for simplicity richness is just distributed between a fixed floor (0) and ceiling *K* shared for all sites:

$\lambda_{i,1} = \mu_0 + \beta_Z Z_i + \varepsilon_{i,1}$ $\varepsilon_{i,1} \sim N(0, \sigma_0^2)$

$Yi, 1 = clamp(Poisson(\lambda i, 1), 0, K),$

where λ is constrained to be positive and $clamp(x, 0, K) = max(0, min(K, x))$.

Richness changes were modelled using a global growth term (G) on the additive scale, with all sites increasing in richness at the same rate but with random-walk noise. There is no effect of *Z* on the dynamics, and the global increments are identical across sites:

$$\lambda_{i,t+1} = Yi, t + G + \eta i,$$

$Yi, t + 1 = clamp(Poisson(\lambda_{i,t+1}), 0, K),$ t=1,…,T−1

$$\eta_{i,t} \sim N(0, \sigma^2_{proc})$$

Where λ is again constrained to be positive. The latent richness $Yi, t + 1$ is then translated into observed richness $Y^*_{i,t+1}$ as in scenario 1.

2.4 Detection bias

In our final scenario ('Detection bias'; Figure 3A), true species richness Y remains constant across both time and the environmental gradient Z, but the gradient affects detection probability D_t, with the detection probability increasing over time (e.g., survey coverage/effort improving across time). True richness and detection probability jointly determine observed species richness, denoted by an asterisk (Y_t* and Y_{t+1}*). For example, low D_t in low-Z sites precludes species detection during early surveys, while greater D_t means observers detect a larger proportion of the true richness at high-Z sites. So while true richness does not depend on Z, there are systematic differences in observed baseline richness Yt*. Improved technology or participation during more recent surveys could lead to high detectability at all sites.

Apparent ΔY~Z relationships are observed because sites that appeared species-poor at low Z increase toward the true richness over time. This pattern is a form of regression to the mean (29): extreme baseline values induced by Z-dependent detection bias tend to move back toward the common underlying mean when re-measured under less biased conditions (30).

We simulated this causal structure as follows:

True (latent) richness *Y* at site *i* is constant across time *t* informed by a fixed density λ and is independent of Z, with variation between years deriving from only demographic stochasticity:

$Y_{i,t} \approx Y_i$ for all $t = 1, \ldots, T$

$$Y_{it} \sim Poisson(\lambda)$$

Detection at baseline (t=1) strongly depends on Z:

$W_i = 1 - (t_i / t_{20})$

$$logit(p_{i,1}) = \alpha_0 + ag\ ti + \alpha_Z Z_i\, Wi$$

Detection at all follow-up surveys increases globally with time, and the effect of Z on detection is reduced with Wi such that it becomes independent of *Z* in the final survey year:

So observed species richness Y* at each time point:

$$Y^*_{i,t} \sim Binomial(Y_i, p_{i,t})$$

2.5 Application of two-step approaches

Using our simulated time series, we apply three common two-step approaches to ask the question ‘how does spatial gradient Z drive biodiversity change?’. While an explicit causal estimand is not usually stated in papers, for our scenarios, applying common analytic strategies represents estimating the total causal effect of Z on the rate of log-biodiversity change over time (or more accurately, the effect of Z on log *observed* biodiversity change over time). All approaches first estimate biodiversity change at each site and then relate these estimates to drivers. In our first analysis, we fit a single ordinary least squares (OLS) model for each site using OLS regression (function lm() in R), with log species richness (+1) as a response variable, and time as a predictor. Coefficients were subsequently regressed on Z using a second OLS regression. In the second approach, we fitted a generalised linear multilevel model to the entire dataset using the R package lme4 (29(30) with a Poisson error and log link function. We specified site-level random intercepts and random slopes for each time series [Y ~ Time + (1+Time/Site)]. Site-level random slope estimates were then regressed on Z using OLS regression. In our third approach, we fit a generalised linear multilevel Poisson model that additionally specified an interaction between Time and Z [Y ~ Time * Z (1+Time/Site)]. We subsequently estimated and visualised the coefficients for Time conditional on values across the range of Z using the package interplot(32). We visualised all other relationships using the ggplot2 package(33).

3 Results

All of the two-step approaches that we applied to ask ‘how does Z drive biodiversity change?’ yielded qualitatively similar results across all four data-generating scenarios. Here we present the results of the two-step approach that fitted separate OLS regressions to log(species richness) values and regressed the coefficients against Z. We provide the results of the other two approaches using generalised linear multilevel models in Appendix S, which are qualitatively similar. For all scenarios, the estimated coefficient for Z represents the total causal effect of Z on the estimated rate of change of observed log-biodiversity over time.

The four distinct data-generating scenarios yielded qualitatively different intercepts and biodiversity trajectories over time (Figure 3, top row). Yet, across all scenarios, a negative relationship between biodiversity change and Z was observed (Figure 3, bottom row), leading to the conclusion that biodiversity change is most strongly positive at low Z, the effect declines with increasing Z and becomes negative at high Z (Figure 3). However, the direct effect of Z on contemporary biodiversity change and the mechanism underlying Z's effect varied significantly between scenarios.

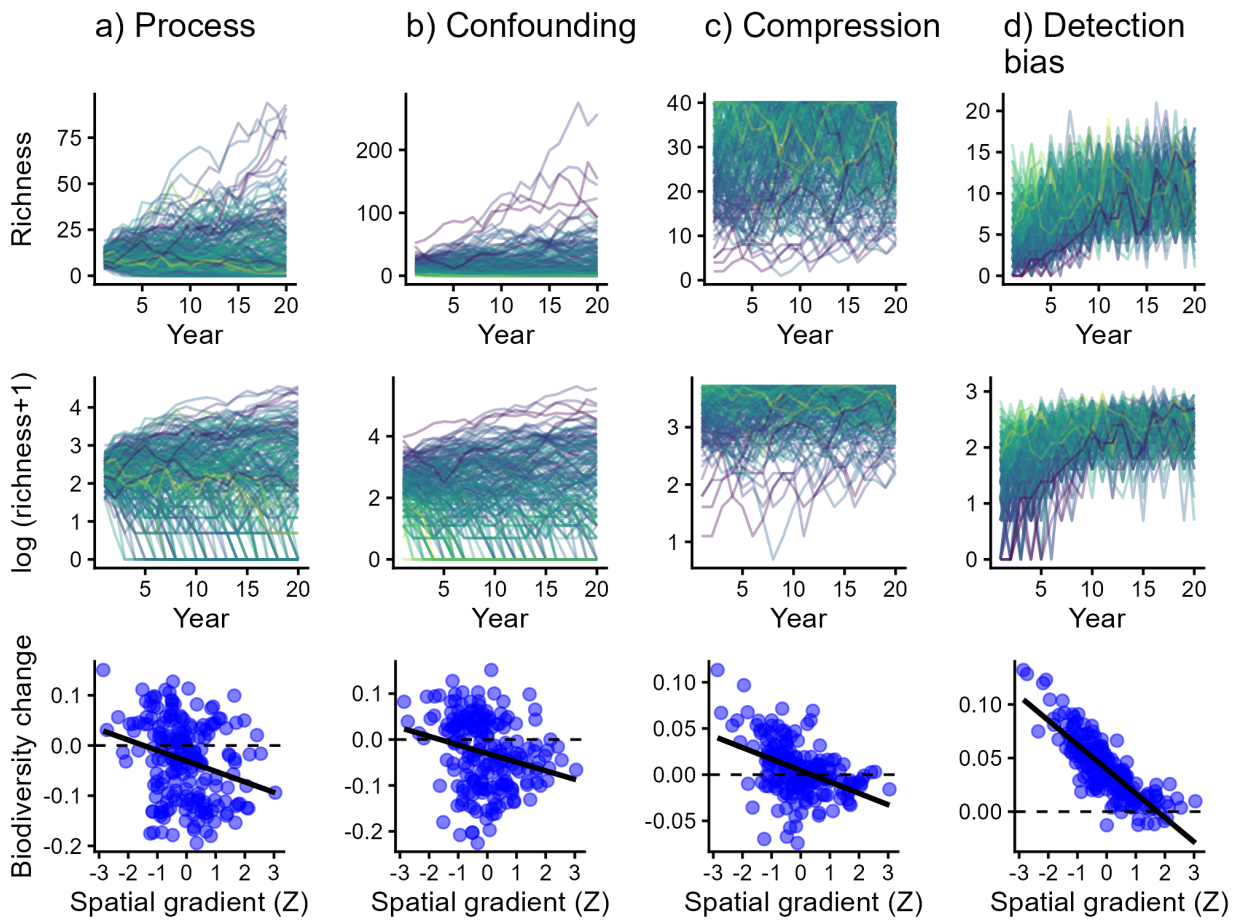

Figure 3. Four distinct data-generating scenarios yield distinct, observed biodiversity time series (first row) that lead to similar patterns in change~Z patterns (third row). Horizontal dashed lines at zero indicate no change over time. Measures of biodiversity change (bottom row) were estimated from separate, site-level OLS regression models fitted to log(species richness) for each site. Fitted black lines represent biodiversity change~Z trends estimated from OLS regression. Scenarios include a) Process effect, b) Confounding, c) Compression and d) Detection bias (corresponding causal structures in Table 1).

Table 2: Two-step approaches yield similar regression coefficients from models regressing biodiversity change (ΔY) on the environmental gradient (Z), representing the total causal effect of *Z* on ΔY. Coefficients (and standard errors) from the same statistical model (ΔY~Z) fitted to simulated datasets generated from four different data-generating processes. Coefficient magnitudes and directions are comparable. We additionally present the true direct effect of Z on contemporary (yearly) change specified in the simulations for comparison.

term	Process effect (log)	Confounding (log)	Compression (log)	Detection bias (log)
(Intercept)	-0.031 (0.005)	-0.030 (0.005)	0.005 (0.002)	0.040 (0.001)
Z	-0.021 (0.005)	-0.019 (0.005)	-0.012 (0.002)	-0.023 (0.001)
True direct effect of Z	-0.025	-0.05	0	0

4 Discussion

We have used DAGs and simulations to show that questions about 'drivers of biodiversity change' are causally ambiguous without an explicitly defined estimand and knowledge about the data-generating process. When analysts adopt the common two-step approach of first calculating biodiversity change and then regressing it on driver Z, the causal interpretation of the estimated slope is obscure because a ΔY~Z coefficient represents a total causal effect that aggregates across multiple causal pathways. Moreover, multiple plausible mechanisms that can cause historic and current levels of biodiversity can generate similar estimates of total effects and patterns of change across gradients. Here, we provide guidance for researchers seeking answers to questions about biodiversity change in future studies and most importantly, call for wider consideration of the data-generating processes when evaluating patterns of biodiversity change across gradients.

4.1 Identifiability does not ensure interpretability: we need knowledge of the data-generating process

Our study adds to the growing calls for researchers in ecological research to use DAGs and state their target estimand (14,15). Importantly, our Detection Bias scenario

underscores the need to include processes that influence species' detection (33,34). Additionally, including a selection node, S, is necessary to convey assumptions about the sampling process when the available data are not a random sample of the target population of inference (33). Automated software such as DAGitty, available as a browser-based interface (dagitty.net), makes it relatively simple to automate identification queries and identify minimally sufficient adjustment sets for closing confounding paths. However, it is important to note that identifiability is not sufficient for an estimand to be interpretable.

In our four scenarios, the $\Delta Y \sim Z$ association represents the total causal effect of driver Z on the proportional rate of change of biodiversity. Although this total effect is technically identifiable (it is not confounded), its interpretability is questionable. Its meaning differs across our scenarios and requires knowledge of the data-generating process. In each scenario, the total effect is made up of varying contributions from multiple causal pathways, including Z's influence on historical and contemporary biodiversity. Differences in both the trajectories (slopes) and baselines among our scenarios are apparent from looking at Figure 3A, which can help with interpretation to an extent, but this information gets lost in biodiversity-change studies that adopt the two-step approach, because the calculation of change in step 1 is destructive; it is not possible to recover information about actual biodiversity values at a given point in time. That is, summarising biodiversity trajectories using the slope of log-richness over time produces a proportional rate of change, but it discards information about the intercept, which encodes baseline richness and important information about causal processes.

Of our four scenarios, only our Process Effect scenario yielded a total effect that is causally interpretable, under strong assumptions of no compression and that baseline richness is independent of the driver Z ($Y_t \perp Z$), conditions rarely met in observational settings. In our Compression scenario, the total effect reflects Z's effect on the biodiversity trend through its influence on baseline richness alone, which in turn determines how strongly ceiling/floor effects will constrain the trajectory. In other words, the total causal effect reflects the effect of Z on historical richness only, and not contemporary biodiversity dynamics. The key point here is that identifiability is not enough to ensure interpretability. In addition to an identifiable estimand, we require knowledge of the data-generating process, which can be achieved by producing generative models that simulate data (34).

We note that some studies investigating drivers of biodiversity change have regressed biodiversity change on measures of climate change and baseline climate (i.e., ΔZ), with or without baseline values of biodiversity (e.g., $\Delta Y \sim \Delta Z + Z_t$, or $\Delta Y \sim Y_t + \Delta Z + Z_t$). In either case, the core issue remains: the estimand is typically obscure without an explicit DAG.

Moreover, a model of $\Delta Y \sim \Delta Z + Z_t$ is structurally equivalent to $Y_{t+1} \sim Y_t + Z_t + Z_{t+1}$. The change variables here do not add information, and obscure what is being estimated. Whether such a model estimates a causal effect of contemporary climate on contemporary biodiversity, a direct effect of climate, or something else, depends entirely on the DAG, and the corresponding adjustment set required for a particular estimand.

4.2 Implications for basic and applied ecology

Researchers often aim to identify drivers that accelerate or buffer biodiversity loss, to inform macro-ecological theories or guide where restoration or protection efforts might be most needed. We have argued that the total causal effect of a driver on biodiversity change is unlikely to be informative for these goals, and likely provide 'insufficiently deep explanations' for biodiversity change, as they aggregate across multiple causal pathways (38). We therefore emphasise caution; their use can be justified given an explicit estimand, assumptions of the data-generation process and calculation of an appropriate adjustment set.

We note that these issues are relevant to meta-analyses that synthesise studies of spatial comparisons of biodiversity (and other ecological variables) between alternative conditions or 'treatments', for example, comparisons of richness under different land uses or conservation designations. Meta-analysis is widely used in ecology to evaluate whether such differences systematically vary across factors that serve as 'effect modifiers' of interest (39). In such studies, an 'effect size' is estimated for each study, commonly a log response ratio ($\ln RR = \log(Y_{t+1}/Y_t)$, to represent study-level proportional differences in richness between treatment and control groups. As a change variable, effect-sizes require a destructive calculation that loses information about actual values (40). Effect sizes are subsequently meta-regressed on study-level variables such as biome, taxon, and latitude, termed 'effect modifiers', to explain variation in observed effects. Yet effect sizes are fully determined by their parents, and therefore, total causal effects of effect modifiers on effect sizes will aggregate multiple causal pathways between modifiers and units in control and treatment groups, leaving the association vulnerable to misinterpretation.

Acknowledgements

RS and LE are grateful to Common Ground Oxford. We are grateful to informative discussions about our results with D Bowler, C Callaghan, R Boyd and R Cooke.

Data availability statement

All code to re-run the simulations is available on the github repository**:**
https://github.com/LukeChrisEvans/deterministic_by_design_code

References

1. Estimates of local biodiversity change over time stand up to scrutiny - Vellend - 2017 - Ecology - Wiley Online Library [Internet]. [cited 2026 Apr 21]. Available from: https://esajournals.onlinelibrary.wiley.com/doi/full/10.1002/ecy.1660

2. Gonzalez A, Chase JM, O'Connor MI. A framework for the detection and attribution of biodiversity change. Philos Trans R Soc Lond B Biol Sci. 378(1881):20220182. doi:10.1098/rstb.2022.0182 PubMed PMID: 37246383; PubMed Central PMCID: PMC10225858.

3. Dornelas M, Chase JM, Gotelli NJ, Magurran AE, McGill BJ, Antão LH, et al. Looking back on biodiversity change: lessons for the road ahead. Philos Trans R Soc Lond B Biol Sci. 2023 May 29;378(1881):20220199. doi:10.1098/rstb.2022.0199

4. Boyd RJ, Powney GD, Burns F, Danet A, Duchenne F, Grainger MJ, et al. ROBITT: A tool for assessing the risk-of-bias in studies of temporal trends in ecology. Methods in Ecology and Evolution. 2022;13(7):1497–507. doi:10.1111/2041-210X.13857

5. Johnson TF, Beckerman AP, Childs DZ, Webb TJ, Evans KL, Griffiths CA, et al. Revealing uncertainty in the status of biodiversity change. Nature. 2024 Apr 1;628(8009):788–94. doi:10.1038/s41586-024-07236-z

6. McElreath R. Chapter 1: The Golem of Prague. In: Statistical Rethinking: A Bayesian Course with Examples in R and Stan (2nd ed.). Chapman and Hall/CRC.; 2020. p. 17. doi:https://doi.org/10.1201/9780429029608

7. Gould E, Fraser HS, Parker TH, Nakagawa S, Griffith SC, Vesk PA, et al. Same data, different analysts: variation in effect sizes due to analytical decisions in ecology and evolutionary biology. BMC Biol. 2025 Feb 6;23(1):35. doi:10.1186/s12915-024-02101-x

8. Dornelas M, Chase JM, Gotelli NJ, Magurran AE, McGill BJ, Antão LH, et al. Looking back on biodiversity change: lessons for the road ahead. Philos Trans R Soc Lond B Biol Sci. 2023 May 29;378(1881):20220199. doi:10.1098/rstb.2022.0199

9. Dornelas M, Antão LH, Bates AE, Brambilla V, Chase JM, Chow CFY, et al. BioTIME 2.0: Expanding and Improving a Database of Biodiversity Time Series. Global Ecology and Biogeography. 2025;34(5):e70003. doi:10.1111/geb.70003

10.McGill BJ, Dornelas M, Gotelli NJ, Magurran AE. Fifteen forms of biodiversity trend in the Anthropocene. Trends in Ecology & Evolution. 2015 Feb 1;30(2):104–13. doi:10.1016/j.tree.2014.11.006

11.Tennant PWG, Arnold KF, Ellison GTH, Gilthorpe MS. Analyses of 'change scores' do not estimate causal effects in observational data. Int J Epidemiol. 2022 Oct 1;51(5):1604–15. doi:10.1093/ije/dyab050

12.Berrie L, Arnold KF, Tomova GD, Gilthorpe MS, Tennant PWG. Depicting deterministic variables within directed acyclic graphs: an aid for identifying and interpreting causal effects involving derived variables and compositional data. Am J Epidemiol. 2024 Jun 24;194(2):469–79. doi:10.1093/aje/kwae153 PubMed PMID: 38918044; PubMed Central PMCID: PMC11815499.

13.Shahar E, Shahar DJ. Causal diagrams and change variables. Journal of Evaluation in Clinical Practice. 2012;18(1):143–8. doi:10.1111/j.1365-2753.2010.01540.x

14.Byrnes JEK, Dee LE. Causal Inference With Observational Data and Unobserved Confounding Variables. Ecology Letters. 2025;28(1):e70023. doi:10.1111/ele.70023

15.Franks DW, Ruxton GD, Sherratt T. Ecology needs a causal overhaul. Biological Reviews. 2025;100(5):1950–69. doi:10.1111/brv.70029

16.Tekwa E, Gonzalez A, Zurell D, O'Connor M. Detecting and attributing the causes of biodiversity change: needs, gaps and solutions. Philos Trans R Soc Lond B Biol Sci. 378(1881):20220181. doi:10.1098/rstb.2022.0181 PubMed PMID: 37246389; PubMed Central PMCID: PMC10225850.

17.Schrodt F, Beck M, Estopinan J, Bowler DE, Fontaine C, Gaüzère P, et al. Advancing causal inference in ecology: Pathways for biodiversity change detection and attribution. Methods in Ecology and Evolution. 2025;16(10):2276–304. doi:10.1111/2041-210X.70131

18.Bonebrake TC, Yau EYH. Quantifying drivers of biodiversity change through increasing data availability and improved analytical frameworks. Journal of Animal Ecology. 2025;94(10):1900–3. doi:10.1111/1365-2656.70138

19.McElreath R. Statistical Rethinking: A Bayesian Course with Examples in R and Stan (2nd ed.). Chapman and Hall/CRC.; 2020. doi:https://doi.org/10.1201/9780429029608

20.Glymour MM. Commentary: Modelling change in a causal framework. Int J Epidemiol. 2022 Jul 28;51(5):1615–21. doi:10.1093/ije/dyac151 PubMed PMID: 35900266; PubMed Central PMCID: PMC9558068.

21.Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology. 1974;66(5):688–701. doi:10.1037/h0037350

22.Holland PW. Statistics and Causal Inference. Journal of the American Statistical Association. 1986;81(396):945–60. doi:10.2307/2289064

23.Pearl J. Causal Diagrams for Empirical Research. Biometrika. 1995;82(4):669–88. doi:10.2307/2337329

24.Oldham PD. A note on the analysis of repeated measurements of the same subjects. Journal of Chronic Diseases. 1962 Oct 1;15(10):969–77. doi:10.1016/0021-9681(62)90116-9

25.Lundberg I, Johnson R, Stewart BM. What Is Your Estimand? Defining the Target Quantity Connects Statistical Evidence to Theory. Am Sociol Rev. 2021 Jun;86(3):532–65. doi:10.1177/00031224211004187

26.Pearl J. Comment: Graphical Models, Causality and Intervention. Statistical Science. 1993;8(3):266–9.

27.R Core Team. R: A Language and Environment for Statistical Computing [Internet]. Vienna, Austria.: R Foundation for Statistical Computing; 2026. Available from: https://www.R-project.org/

28.McClure CJW, Rolek BW. Pitfalls arising from site selection bias in population monitoring defy simple heuristics. Methods in Ecology and Evolution. 2023;14(6):1489–99. doi:10.1111/2041-210X.14120

29.Regression to the mean: what it is and how to deal with it | International Journal of Epidemiology | Oxford Academic [Internet]. [cited 2026 Apr 21]. Available from: https://academic.oup.com/ije/article/34/1/215/638499?guestAccessKey=

30.Bates D, Mächler M, Bolker B, Walker S. Fitting Linear Mixed-Effects Models Using lme4. J Stat Soft. 2015 Oct 7;67(1):1–48. doi:10.18637/jss.v067.i01

31.Solt F, Hu Y, Kenkel B. interplot: Plot the Effects of Variables in Interaction Terms [Internet]. 2021 [cited 2026 Apr 21]. Available from: https://cran.r-project.org/web/packages/interplot/index.html

32.Wickham H. ggplot2: Elegant Graphics for Data Analysis. New York, NY, USA: Springer-Verlag; 2016.

33.Boyd RJ, Cooke R, Powney GD, Pescott OL. Why your causal diagram should probably include sampling and measurement processes [Internet]. 2026 Mar 11 [cited 2026 Apr 21]. Available from: https://ecoevorxiv.org/repository/view/12104/

34.Priyadarshana TS, Slade EM. Predictive entomology: A causal framework for detecting and attributing insect population change. Current Opinion in Insect Science. 2026 May 19;101538. doi:10.1016/j.cois.2026.101538

35.Deffner D, Fedorova N, Andrews J, McElreath R. Bridging theory and data: A computational workflow for cultural evolution. Proc Natl Acad Sci U S A. 2024 Nov 26;121(48):e2322887121. doi:10.1073/pnas.2322887121 PubMed PMID: 39556723; PubMed Central PMCID: PMC11621747.

36.Lord FM. A paradox in the interpretation of group comparisons. Psychological Bulletin. 1967;68(5):304–5. doi:10.1037/h0025105

37.Pearl J. Lord's Paradox Revisited – (Oh Lord! Kumbaya!). Journal of Causal Inference, TECHNICAL REPORT. 2016;4(2):R436.

38.Winship C, Morgan SL, editors. Mechanisms and Causal Explanation. In: Counterfactuals and Causal Inference: Methods and Principles for Social Research [Internet]. 2nd edn. Cambridge: Cambridge University Press; 2014 [cited 2026 May 11]. p. 325–53. (Analytical Methods for Social Research). Available from: https://www.cambridge.org/core/books/counterfactuals-and-causal-inference/mechanisms-and-causal-explanation/36ECACBB8D707C58A5A5FCEDA532161B doi:10.1017/CBO9781107587991.011

39.Meta-analysis and the science of research synthesis - PubMed [Internet]. [cited 2026 Apr 21]. Available from: https://pubmed.ncbi.nlm.nih.gov/29517004/

40.Spake R, Bowler DE, Callaghan CT, Blowes SA, Doncaster CP, Antão LH, et al. Understanding 'it depends' in ecology: a guide to hypothesising, visualising and interpreting statistical interactions. Biological Reviews. 2023;98(4):983–1002. doi:10.1111/brv.12939

Appendix S1 - Results of multilevel models fitted to the simulated data

Here, we present analyses of two additional common approaches to biodiversity change analyses that ask how biodiversity change depends on driver *Z*, using multilevel models with the R package lme4 (Bates et al. 2015)(1), which are common in biodiversity change analyses, to demonstrate equivalent challenges for causal inference.

We fitted a generalised linear multilevel model to the entire dataset using the R package lme4 (Bates et al. 2015)(1) with a log link function. We specified site-level random intercepts and random slopes for each time series [Y ~ Time + (1+Time/Site)]. Site-level random slope estimates were then regressed on Z using OLS regression.

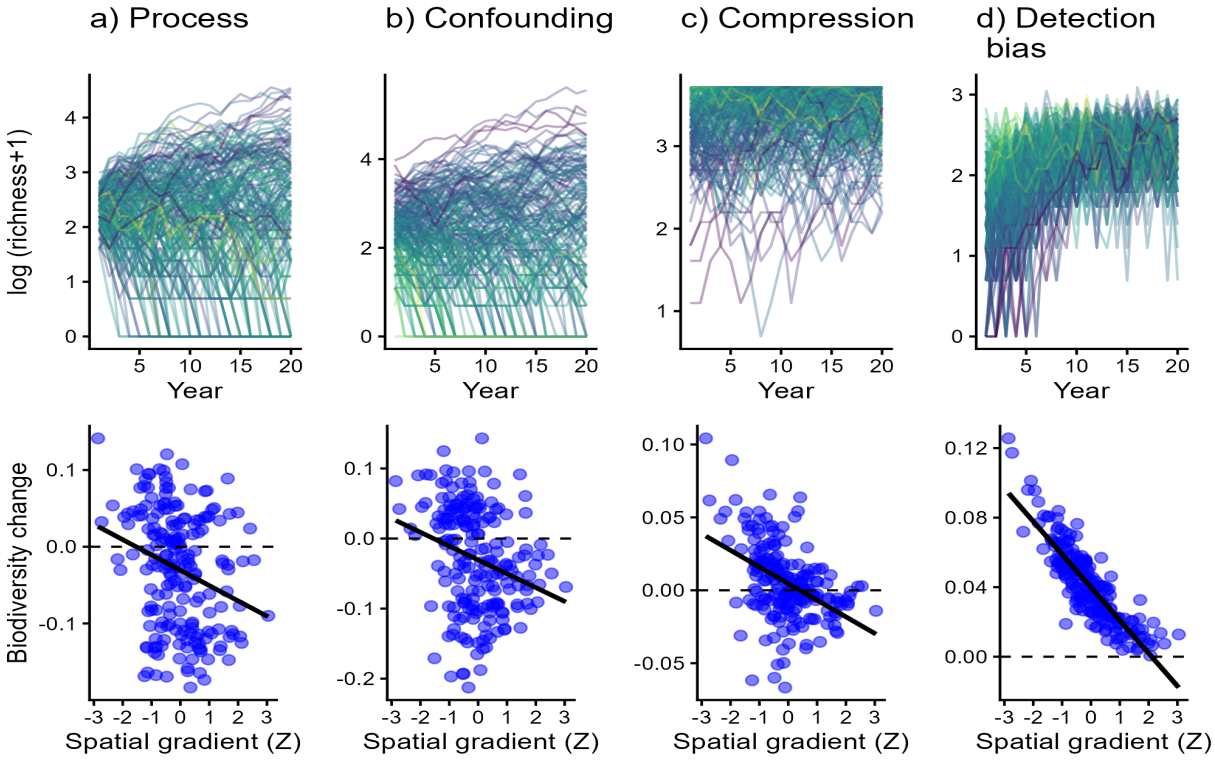

Figure S1.1 Four distinct data-generating scenarios (a-d) generate similar change~Z patterns. Horizontal dashed lines at zero indicate no change over time. Estimates of biodiversity change are site-level random slopes for the effect of time (step 1). Fitted black lines represent biodiversity change~*Z* trends estimated from OLS regression of these slopes on *Z* (step 2). Scenarios include a) Process effect, b) Confounding, c) Compression, d) Detection bias (corresponding causal structures in Table 1).

Table S1.1 Two-step approaches yield similar regression coefficients from multi-level models regressing biodiversity change (ΔY) on the environmental gradient (Z), representing the total causal effect of *Z* on ΔY. Coefficients (and standard errors) from the same statistical model (ΔY~Z) fitted to simulated datasets generated from four different data-generating processes. Coefficient magnitudes and directions are comparable. We additionally present the true direct effect of Z on contemporary (yearly) change specified in the simulations for comparison.

term	Process effect (log)	Confounding (log)	Compression (log)	Detection bias (log)
(Intercept)	-0.031 (0.005)	-0.031 (0.005)	0.005 (0.002)	0.040 (0.001)
Z	-0.020 (0.005)	-0.020 (0.005)	-0.011 (0.002)	-0.019 (0.001)
True direct effect of Z	-0.025	-0.05	0	0

We also fit a generalised linear multilevel Poisson model that additionally specified an interaction between Time and *Z* [Y ~ Time * Z (1+Time/Site)]. We subsequently estimated the coefficients for Time conditional on values across the range of Z using the package interplot (Solt and Hu 2018)(2). We then estimated the site-level random slopes and regressed them on Z.

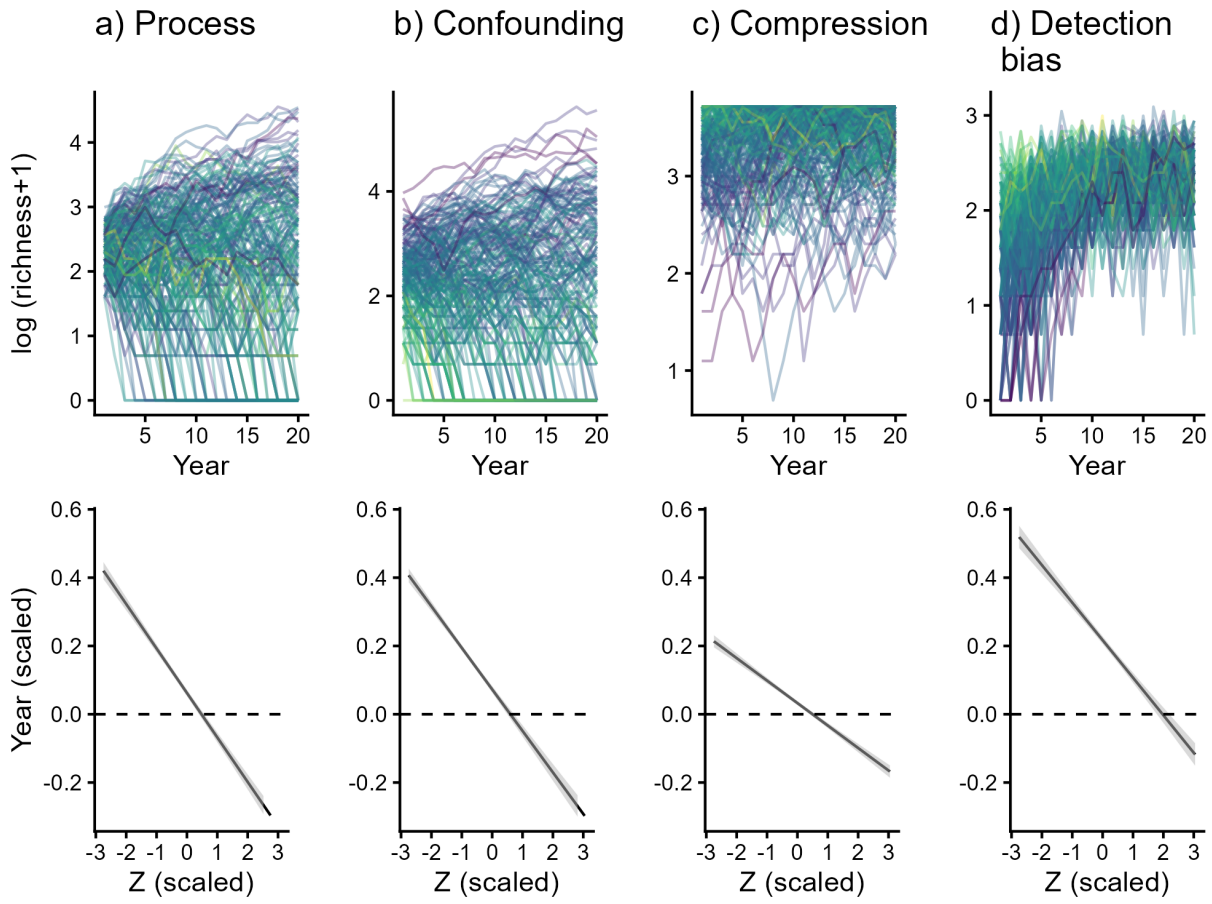

Figure S1.2 Four distinct data-generating scenarios (a-d) generate similar change~Z patterns. Horizontal dashed lines at zero indicate no change over time. Estimates of biodiversity change represent the estimated coefficient for time, conditional on *Z,* using R package interplot. Scenarios include a) Process effect, b) Confounding, c) Compression, d) Detection bias (corresponding causal structures in Table 1).

Table S1.2 One-step approaches fitting Year-by-Z interactions yield similar relationships across all causal structures. Values show coefficients (and standard errors) from the same statistical model fitted to simulated datasets generated from four different data-generating processes. Coefficient magnitudes and directions are comparable.

term	Process effect	Confounding	Compression	Detection bias
Intercept	2.063 (0.069)	1.947 (0.082)	3.348 (0.017)	2.078 (0.006)
Year (scaled)	0.062 (0.005)	0.073 (0.005)	0.033 (0.003)	0.217 (0.006)
Z (scaled)	-0.149 (0.069)	-0.942 (0.082)	0.134 (0.017)	0.110 (0.006)
Year × Z (scaled)	-0.130 (0.005)	-0.122 (0.004)	-0.066 (0.003)	-0.110 (0.005)

1. Bates D, Mächler M, Bolker B, Walker S. Fitting Linear Mixed-Effects Models Using lme4. J Stat Soft. 2015 Oct 7;67(1):1–48. doi:10.18637/jss.v067.i01

2. Solt F, Hu Y, Kenkel B. interplot: Plot the Effects of Variables in Interaction Terms [Internet]. 2018 [cited 2026 Apr 21]. Available from: https://cran.r-project.org/web/packages/interplot/index.html