Implementing Causal Inference in Ecology Through the Structural Causal Model (SCM) Framework

by

Suchinta Arif

Submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
December 2022

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Ecologists are often interested in understanding causal relationships from ecological data. However, developed methods for causal inference, particularly for observational-based studies are often not taught or applied in ecology. This thesis overviews how the structural causal model (SCM) framework can be employed to increase both observational and experimental causal inference in ecology, increasing the validity of causal conclusions drawn from ecological data. First, this thesis presents a comprehensive review of the SCM framework geared towards a general ecology audience interested in observational causal inference. Next, the SCM framework is applied to a local observational coral reef dataset, to determine the causal drivers of coral-reef regime shifts in Seychelles. This framework is also applied to a global observational coral reef dataset, to understand the global drivers of reef fish biomass. The cumulative results from the above two studies provide practical guidelines on how to apply the SCM to both local and global ecological data, each highlighting that novel, reliable causal conclusions can be drawn from this approach. Using theory and simulated data, this thesis further explains how the SCM framework can be used to ensure proper study design and analysis across quasi-experimental (e.g., matching methods, before control impact, regression discontinuity design, instrumental variables) approaches. Last, using key ecological examples, this thesis explores how the SCM framework can also be employed to help ensure valid causal conclusions are drawn from experimental data. Ultimately, the increased uptake of the SCM framework across ecology can increase the depth and pace at which we understand causal relationships in nature.

LIST OF ABBREVIATIONS AND SYMBOLS

BACI	Before after control impact
BEF	Biodiversity Ecosystem Function
DAG	Directed Acyclic Graph
DHW	Degree Heating Week
DiD	Difference-in-Difference
GLM	Generalized Linear Model
HDI	Human Development Index
IV	Instrumental Variable
MPA	Marine Protected Area
NPP	Net Primary Productivity
PS	Propensity Score
RCT	Randomized Control Trail
RDD	Regression Discontinuity Design
SCM	Structural Causal Model
SST	Sea Surface Temperature

ACKNOWLEDGEMENTS

I am so grateful to have had such a positive and fulfilling PhD experience, which would not be possible without the guidance and help of several faculty members, colleagues, and friends at Dalhousie University.

First, a huge thank you to my advisor and mentor, Dr. Aaron MacNeil. Thank you for always remaining positive, allowing me an abundance of autonomy, guiding me to become a better person and researcher, and always choosing to view all your students in the best possible light. Thank you for encouraging me to pursue initiatives outside the scope of my thesis, which has made my PhD journey more holistic and rewarding. Simply put, I could not have asked for a better student-supervisor fit for this phase of my career, and for that I will be forever grateful.

My biggest thank you to my work wife, Melanie Duc Bo Massey, whose research I still do not quite understand. Nonetheless, we have worked on numerous non-thesis related projects together, which has enriched my PhD experience 10-folds. I simply would not have the confidence to tackle the following very rewarding initiatives without you: co-creating a scientific outreach organization from scratch, co-teaching a graduate module, co-publishing pedagogical papers, and more.

A huge thank you to Dr. Matthew McLean for being an outstanding human and scientist. You have made me a kinder person just by association. A special thank you to Aaron Judah for being the most ambitious mentee I will likely ever have. Thank you to everyone in the MacNeil Lab for an abundance of positive interactions.

Last, but not least, I would like to thank my committee members Dr. Nancy Shackell, Dr. Nick Graham and Dr. Derek Tittensor for your ongoing support and guidance. A special thank you to Dr. Nick Graham for hosting me for a Mitacs placement at Lancaster University, which was one of the absolute highlights of my PhD journey.

CHAPTER 1 INTRODUCTION

1.1 WHAT IS CAUSAL INFERENCE?

Causal inference refers to the transdisciplinary practice of integrating assumptions, study designs, and estimation strategies to allow researchers to draw causal conclusions from data (Imbens and Rubin 2015; Morgan and Winship 2015). It often requires the leveraging of theory and deep knowledge to accurately estimate the effect of events, choices, or other factors on a given outcome of interest (Cunningham 2021). As such, causal inference aims to provide evidence of causality guided by causal reasoning. The study and application of causal inference has roots in philosophy, statistics, social science, economics, and computer science (Imbens and Rubin 2015; Morgan and Winship 2015). Causal inference methods for both observational and experimental data are broadly applicable across disciplines, with many methods that were designed in one discipline finding broadscale use across others (e.g., Pearl 2009). Researchers are often interested in the causal mechanisms that govern natural and human processes, to use this gained knowledge to make decisions that benefit us and our environment.

1.2 A BRIEF HISTORY OF CAUSAL INFERENCE

Although ideas of causality have existed throughout history (e.g., see Aristotles' four causes in *Physics* and *Metaphysics*), David Hume is often credited with providing the first modern definition of causality. Hume attempted to define what a cause is as well as how humans come to possess causal knowledge, and what is required to infer it from observations. Hume posited three essential requirements to determine cause: (1) *contiguity*, meaning that cause and effect must be nearby in time and space, (2) temporal priority, meaning that a cause must precede its effect, and (3) *necessary connection*, meaning that causes must always produce the effect, and that the effect is not produced without the cause (Hume 1739). His definition of a cause "We may define a cause to be an object, followed by another, […], where, if the first object had not been, the second had never existed" is still widely applied today (Hume 1748).

In the 1920s, two major advances defined alternative approaches to causal inference, the first being the popularization of randomized control experiments (RCTs). In 1923, Jerzy Neyman introduced randomized experiments in agriculture by discussing the potential yield to be gained from agricultural plots under different experimental exposures (Neyman 1923). This basis of his work was later formulized statistically and is now widely known as the potential outcome framework (Rubin 1974; Holland 1986). The potential outcome framework defines the causal effect for an individual unit as the difference between the potential outcomes that would be observed for that individual unit with and without a given treatment. Imagine we are interested in quantifying the causal effect of a treatment X_i on an outcome Y_i, for an individual *i*. For each individual, there are two potential outcomes:

$Y_i(1)$ is the outcome of individual *i* if they take treatment X ($X_i = 1$)
$Y_i(0)$ is the outcome of individual *i* if they do not take treatment X ($X_i = 0$)

The causal effect of the treatment on the outcome for individual *i* is defined as:

$$\text{Causal effect}_i = Y_i(1) - Y_i(0) \quad [1.1]$$

In words, the causal effect for individual *i* is defined as the difference between their potential outcomes. However, in reality, only one potential outcome, either $Y_i(1)$ or $Y_i(0)$ can be realized. An individual *i* can either take the treatment ($X_i = 1$) resulting in the observed $Yi(1)$ or not take the treatment ($X_i = 0$) resulting in the observed $Y_i(0)$, but not both. This is referred to as the fundamental problem of causal inference (Rubin 1974; Holland 1986). To work around this, the potential outcome framework formally articulates how to quantify an average treatment effect (ATE) for RCTs (Rubin 1974; Holland 1986) as well as non-experimental (Rosenbaum and Rubin 1983) data. Under the potential outcome framework ATE can be defined as:

$$\text{ATE} = E[Y_i(1)-Y_i(0)] \quad [1.2]$$

In words, ATE is defined as the difference between the average potential outcomes $Y_i(1)$ and $Y_i(0)$ for a population of *n* units where $i = 1 \ldots n$. Again, this cannot be directly quantified as we do not

observe $Y_i(1)$ and $Y_i(0)$ simultaneously. However, under RCTs, if the treatment X_i is randomly assigned, then treatment X_i is independent of the potential outcomes that would be realized under each treatment. This is referred to as ignorability (Rosenbaum 2002), and written as:

$$\{Y_i(1), Y_i(0)\} \perp X_i \quad [1.3]$$

Given this, we can estimate the ATE of a population by taking the difference in mean outcomes of the individuals in the treated group and control group, written as:

$$ATE = E[Y_i \mid X_i = 1] - E[Y_i \mid X_i = 0] \quad [1.4]$$

In words, under a randomized experiment where ignorability is met, ATE = (average treatment outcome) – (average control outcome). The average difference in outcome is attributed to the treatment effect, and not to confounding that can arise from non-random treatment assignment. The potential outcome framework is arguably the most widely used framework for causal inference and provides the underlying theory behind RCTs.

Around the same time as Neyman, Ronald Fisher also published on the importance of randomization in experimental designs (Fisher 1925). Throughout his life, Fisher remained a strong proponent of randomized experiments. He noted that the randomization process eliminated bias, for example, by removing all uncontrolled differences between treatment groups. According to Fisher, "randomization…relieves the experimenter from the anxiety of considering and estimating the magnitude of the innumerable causes by which data may be distributed." (Fisher 1935). Although randomized experiments first appeared in the 1800s (Peirce and Jastrow 1885; Stolberg 2006), and Neyman re-introduced randomized experiments a few years prior, Fisher is often credited with the "invention" or "discovery" of RCTs due to his advocacy and popularization of the concept (Rubin 2005; Hall 2007). RCTs are now prevalent across disciplines and is often regarded as "the gold standard" for causal inference (Hariton 2018). However, it is important to note that RCTs, much like any causal inference tool, rests on a set of causal assumptions, that when not met, can lead to erroneous causal conclusions (e.g., see Kimmel et al. 2021).

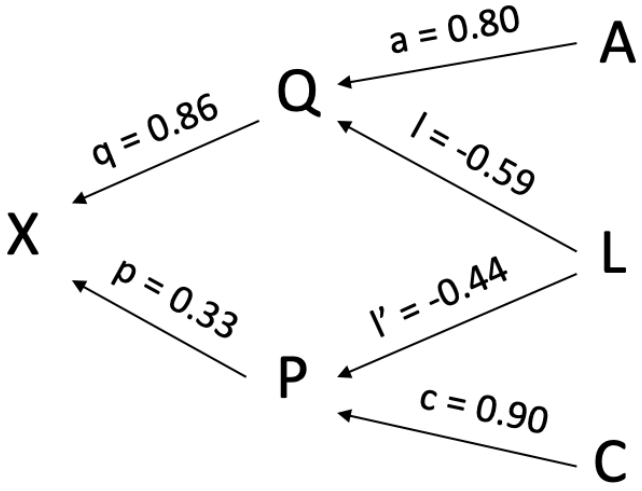

Figure 1.1 A path diagram and associated path coefficients representing the causal relationships between birth weight (X), growth rate (Q), gestation period (P), size of litter (L), and other factors (A, C) for guinea pigs. Adapted from Wright (1921).

A second advance in causal inference that occurred in the 1920s came with Sewall Wright's path analysis, which uses path diagrams to encode a researcher's causal assumptions about a system or process under study (Wright 1920; Wright 1921). Figure 1.1 shows a path diagram adapted from Wright (1921), representing how different variables are hypothesized to effect birth weight in guinea pigs. Directed arrows point from cause to effect; for example, size of litter (L) directly effects gestation period (P), which in turn effects birth weight (X). Based on the path diagram and the observed correlations between paired variables, a system of equations can be constructed to solve for "path coefficients", which represent the direct effects of variables on each other (Wright 1921). First, equations can be constructed based on the principle that "correlation between two variables is equal to the sum of the products of the chains of path coefficients along all the paths by which they are connected (Wright 1921)" For example, in Figure 1.1, X and P are connected by the path P → X (associated with path coefficient p) as well as the path P ← L → Q → X (associated with path coefficients qll'). The correlation between X and P can thus be represented by the equation:

$$r_{XP} = p + qll' \qquad [1.5]$$

A second set of equations can also be constructed based on the principle that the complete determination of an effect in a path diagram must equal 1 (Wright 1921). For example, in Fig 1, the complete determination of P by factors L and C can be expressed by the equation:

$$l'^2 + c^2 = 1 \quad [1.6]$$

Given a path diagram and observed correlations between paired variables, these two types of equations can be used to solve for path coefficients. For example, if we know the correlation between P and L is 0.44444, then we can solve for the path coefficient l' since $r_{PL} = l' = 0.4444$. Given that $l'^2 + c^2 = 1$, we can solve for l'^2 ($0.4444^2 = 0.1975$), c^2 (1-0.1975 = 0.8025) and the path coefficient c ($\sqrt{0.8025} = 0.8958$). Figure 1.1 presents the calculated path coefficients for our path diagram (see Wright 1921; p 568-570 for full calculations using this example). Path coefficients show the strength and direction of direct effects; for example, assuming Figure 1.1 accurately represents the causal relationships among variables, rate of growth (Q) has a stronger positive effect on birth weight (X) relative to gestation period (P). A generalization of path analysis known as structural equation modeling (SEM) was subsequently developed by Otis Dunley Duncan in 1975 (Duncan 1975) and continues to be widely employed across disciplines, including ecology and the social sciences (Hershberget 2003; Grace et al. 2010). SEM is meant to be used as a mathematical tool for drawing causal conclusions from observational data, given a researchers' causal assumptions, which are encoded in SEM diagrams (Pearl 2012). These causal assumptions must ultimately be justified based on domain knowledge, such as scientific consensus, prior studies and experiments, scientific judgement, and other informed sources (Pearl 2012). Conditional on these causal assumptions, the SEM approach allows researchers to draw causal conclusions from observational data.

SEMs offers great potential for quantifying causal relationships in nature, particularly when working with data across a wide range of questions and study systems that may otherwise be impossible or impractical to explore under experimental design. However, the application of SEM through the years has led to many misunderstandings about its original causal intent. Bollen and Pearl (2013) highlight eight core myths about causality and SEMs, including critiques by opponents that claim SEMs aim to establish causal relationships from correlation alone. Examples

are perpetuated both in historical (e.g., Guttman 1977, De Leeuw 1985; Baumrind 1983) and recent (Goldthorpe 2001; Freedmand 2004; Sobel 2008) literature. This gives the general impression that SEMs are used to derive causal conclusions from correlations that arise from observational data and complicated models of partial association alone, and thus such results should not be trusted as valid causal conclusions (Pearl 2012). However, SEM diagrams represent the causal assumptions of the researcher; the credibility of conclusions therefore depends on the credibility of the causal assumptions in each application. Unfortunately, the misconceptions around SEM have led to many researchers explicitly avoiding causal language and interpretation (e.g., Muthen 1987) and relying extensively on the data (instead of a researcher's knowledge) to guide SEM diagrams and subsequent interpretations (Pearl 2012).

It is worth noting that RCTs quickly became a widely applied and trusted tool for causal inference, whereas observational causal inference approaches such as SEM remain relatively underutilized. The opposition against observational approaches to causal inference can largely be attributed to Karl Pearson, whom many consider to be the founder of modern statistics. Pearson, inventor of the correlation coefficient, believed that ideas around causality were outdated and unscientific compared to the mathematically clear and precise concept of his correlation coefficient (Pearson 1911). In his words, causality was just a "fetish amidst the inscrutable arena of modern science (Pearson 1911)." Pearson also pointed out non-causal spurious correlations that can often arise in observational data, leading to the often-repeated phrase "correlation does not imply causation" (F.A.D. 1900). This led to a general misconception that correlation *can't* equal causation, even when valid observational causal inference is employed. For his efforts, Pearson has been regarded as "causality's worst adversary" (Pearl 2009). The notion that causal conclusions can only be drawn from experiments has also been perpetuated across disciplines by many researchers through time. For example, in his influential article, Paul Holland (1986), noted "no causation without manipulation" as one of his mottos, explaining that variables that cannot be manipulated, such as race and sex, cannot act as causal factors. He then criticized SEMs for representing these factors (race and sex) as causal variables in SEM diagrams (Holland 1986).

In recent years, a "causal revolution" has largely been attributed to computer scientist Judea Pearl, who has integrated the potential outcome framework with nonparametric SEM and other theories

of causation to formulate a comprehensive causal inference methodology known as the structural causal model (SCM; Pearl 2009). Similar to the original intent of SEMs, this framework encodes a researcher's causal assumptions in graphical diagrams in the form of directed acyclic graphs (DAGs). DAGs are causal graphs comprised of a set of nodes (*i.e.,* variables) and *directed* arrows pointing from cause to effect; they are also *acyclic* meaning they cannot contain bi-directional relationships or a feedback loop where a variable causes itself (Glymour and Greenland 2008, Elwert 2014). For example, the path diagram presented in Figure 1.1 is also a DAG. Given a DAG, a set of graphical rules can be applied to determine which variables must be controlled for (e.g., through statistical adjustment, stratification, experimental control) to eliminate bias and quantify causal relationships. Importantly, this framework highlights that causal assumptions encoded in DAGs can be used to guide valid causal inference across both observational and experimental data, helping remove the long-standing bias against observational causal inference. The SCM framework forms the basis of this thesis and is expanded upon in subsequent sections and chapters; however, it is worth noting that is one of many causal inference frameworks that exist. Other landmark developments in observational causal inference include the development of quasi-experimental approaches including instrumental variables (Wright 1928), regression discontinuity design (Thistlethwaite et al. 1960), and propensity score matching (Rosenbaum and Rubin 1983), as well as early developments in time-series causal analysis such as Granger causality (Granger 1969). The history and ongoing development of causal inference frameworks is complex and varied, and researchers interested in drawing causal conclusions from data should remain open to varied approaches available for doing so.

1.3 Causal Inference in Ecology

In ecology, RCTs remain the primary tool used for causal inference. RCTs rely on the randomization and control of a treatment variable to understand its effect of an outcome variable. Indeed, RCTs have led to significant advances in our understanding of ecological phenomena, such as Gauses' laboratory experiments on yeast confirming the theory of competitive exclusion (Gause 1932), Paines' starfish removal experiment creating the concept of keystone species (Paine 1966), and Simberloff's field experiment on mangroves, testing the theory of island biogeography (Simberloff 1976). However, although RCTs are seen as the gold standard in ecology and the sciences, with conclusions drawn from RCTs often not questioned, they are nonetheless

susceptible to potential biases that can lead to erroneous causal conclusions. For example, Kimmel et al. (2021) discuss four core causal assumptions required for valid causal inference in experimental biology. These include excludability, which is the assumption that the process by which treatments are assigned has no effect on the outcome. For example, an experimental set-up investigating the effect of drought on plant growth can lead to the excludability criterion being violated if the drought treatment also alters temperature, humidity, and light, which in turn also influence the outcome plant growth. In this situation, the violation of the excludability criterion will lead to confounding bias, a commonly acknowledged bias in observational studies (e.g., Gray et al. 2016). Other causal assumptions for RCTs include no interference between units, no multiple versions of treatment, and no compliance, meaning that all units receive the treatment they were assigned. Kimmel et al. (2021) highlight how each of these biases can arise in experimental designs, explaining that valid causal inference always rests on a set of core causal assumptions that must be carefully communicated and justified by the researcher, even under experimental set-ups. The general assumption that RCTs can lead to valid causal conclusions without the careful consideration of how the study design and/or statistical analysis of an experiment may bias causal estimates represents a current limitation of many ecological RCTs (Kimmel et al. 2021). Further, many pertinent causal queries in ecology cannot be answered through perfectly designed RCTs, and instead rely on observational data collected across broad scales of time and space, and in situations that are not easily manipulated. For example, ecologists continue to rely on observational data to understand how anthropogenic disturbances such as climate-induced bleaching events, pollution, or overfishing effect ecosystem health and services. The increased availability of observational data through citizen science and technological advances (Sagarin and Pauchard 2010) further highlights the importance of observational ecological studies.

Across observational settings, some ecologists have advocated for the use of quasi-experimental approaches to estimate causal effects (Butsic et al. 2017; Larsen et al. 2019; Wauchope et al. 2021). These include matching methods, which aim to balance the distribution of covariates between treatment and control groups to remove confounding bias within observational studies (e.g., see Adnam et al. 2008; Herrera et al. 2019 for ecological examples); before-after control impact (BACI) studies, which measure a response both before and after an intervention for both treatment and control sites, with differences in the rate of change between treatment and control

being attributed to the effect of the intervention (e.g., see Kadye and Booth 2012 and Bousquin and Colee 2014 for ecological examples); regression discontinuity design (RDD), which uses observational data near a discontinuity in either space, time, or policy to separate observations into treatment and control groups, with the assumption that at or near this discontinuity, confounding variables are equal between treated and control groups (e.g., see Perez et al. 2017 for ecological example); and instrumental variables (IV), which uses an instrument (i.e. a third variable) that is correlated with the predictor but not the outcome to overcome confounding and measurement errors (e.g., see Bush and Cullen 2009 and Amin et al. 2014 as ecological examples). Although these causal tools are valuable, like RCTs, they too require researchers to critically examine the causal assumptions required for each method. For example, although past ecological studies employing matching methods state the presumed confounding variables used in their matching procedure, they do not state how these variables interact with one another within the broader causal structure of a study system. Without this knowledge, it is unclear whether there are other confounding variables that need to be included in the matching analysis or whether the inclusion of all presumed confounding variables may lead to other forms of bias (Mansournia et al. 2013). As another example, past reviews of BACIs have noted the prevalence of improper study design, whereby key variables required to ensure proper BACI design across ecological studies are often neglected (Ferraro et al. 2019; Adams et al. 2019). Collectively, ecological causal inferences drawn from quasi-experimental approaches can be strengthened by ensuring that causal assumptions required for each approach are carefully examined and met.

Other notable observational causal inference techniques in ecology include SEM, which has been employed to disentangle causal relationships across varied ecological systems (Grace et al. 2010), as well as convergent cross mapping (CCM; Sugihara et al. 2012), a time-series analysis specially created for understanding complex ecological relationships in nature. Although observational causal inference methodologies exist and have been employed across several ecological studies, it is important to note that most observational studies aiming to answer causal questions in ecology do not employ *any available causal inference methodology*. Instead, a commonly applied statistical approach is to generate a 'causal salad' model (Bhalla 2018 McElreath 2020), whereby predictor variables of interest are placed under one statistical model and subsequently interpreted for causal effects. Such an approach will generate a correlation between a predictor and response

variable, controlling for all other covariates entered in the model. The 'causal salad' model is only partially guided by causal reasoning, and often reflects only which variables a researcher may be interested in and/or have available. For example, both Cinner et al. (2016) and Cinner et al. (2020) investigated the effect of socio-environmental drivers on global reef fish biomass using the Reef Life Survey data (Reef Life Survey Foundation 2019). The former included 18 predictor variables, and the latter included 17 predictor variables, with 11 overlapping variables. Although the two studies used the same dataset, the effect sizes of the 11 overlapping predictor variables varied across studies due to differences between the two 'causal salad' models. This illustrates that correlative conclusions drawn from such approaches should not be trusted as valid causal estimates (although the authors did not explicitly set out to conduct a causal analysis). Ultimately, the 'causal salad' approach often leads to biased estimates for all predictor variables of interest (McElreath 2020).

Another commonly applied approach across observational ecological studies is to apply predictive model selection techniques such as Akaike's Information Criterion (AIC; e.g., see Millard et al. 2021 and Lu et al. 2021 for ecological examples). These approaches select the 'best' model among a candidate set and subsequently make inferences from parameters that are of ecological interest within the top-ranked model. However, such approaches are meant to measure out of sample predictive accuracy, and not for generating causal effects (Arif and MacNeil 2022). Recently, the advance of machine learning (ML) techniques combined with the availability of big data have also led to their misuse for drawing causal conclusions (e.g., see Guy et al. 2020 and Alkhamis et al. 2021 for ecological examples). The substantive use of predictive approaches for drawing causal conclusions from ecological data signal that most ecologists are not equipped with the theory and tools required for observational causal inference.

In addition to applying non-causal approaches, ecologists often avoid explicitly acknowledging the causal nature of their research, and instead use coded language that implies causality without explicitly saying so (Hernan 2018). Collectively, this has led to a culture whereby ecologists continue to depend on observational data to answer fundamental causal queries in ecology but do so under the prescription that "correlation does not imply causation" (attributed to Karl Pearson; F.A.D. 1900). However, correlation can equal causation, given the application of valid causal

inference approaches. What is needed then, is a fundamental shift in culture, whereby ecologists are equipped with the theory and tools required for widespread observational causal inference. This is turn can strengthen the pace and depth at which we understand our natural world.

1.4 THE STRUCTURAL CAUSAL MODEL (SCM) FRAMEWORK

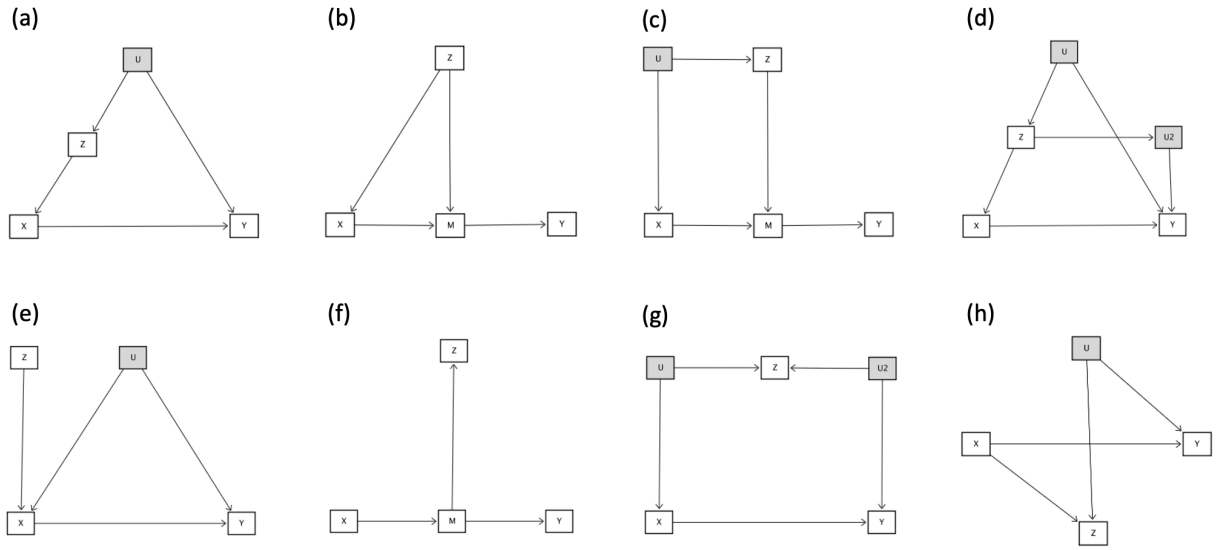

Figure 1.2 Simplified directed acyclic graphs (DAGs). Variables are represented by nodes, with directed arrows pointing from cause to effect. Unobserved variables are depicted in grey. In scenarios a-d, the effect of X on Y can only be determined if Z is controlled for. In scenarios e-h, the effect of X on Y will be biased if Z is controlled for.

Judea Pearl's SCM framework, draws on directed acyclic graphs (DAGs) to visualize the hypothesized causal structure of a system or process under study, based on a researchers' domain knowledge. DAGs comprise of a set of nodes (*i.e.,* variables) and directed arrows pointing from cause to effect, with causes preceding their effects. For example, in Figure 1.2a, X effects Y, U effects Y and U effects X indirectly through Z. DAGs are also non-parametric, making no assumptions about the functional form (e.g., linear, nonlinear, stepwise) or effect size of direct effects (Glymour and Greenland 2008). This non-parametric nature of DAGs makes them compatible with a wide range of ecological systems. Given a DAG, graphical rules can be applied to determine variables that need to be controlled for (*e.g.,* through statistical adjustment, stratification, experimental control) to answer causal queries from both observational and

experimental data. These graphical rules are detailed in later chapters, but the main takeaway is that given a DAG, they guide researchers on how to create causal models that can be used to answer a specific causal query at hand.

To demonstrate the importance of the SCM framework, Figure 1.2 shows a series of DAGs representing the causal structure of simplified systems. In each instance, imagine we are interested in determining the effect of X on Y. Figure 1.2 a-d represent scenarios where controlling for Z is required to determine the effect of X on Y, whereas Figure 1.2 e-h represent scenarios where controlling for Z will lead to non-causal correlative associations between X and Y. In each case, the application of graphical rules within the SCM framework can confirm when Z (as well as any other variable) does and does not need to be controlled for. Having a formulized causal inference framework for determining covariate selection stands in contrast with commonly applied statistical techniques used in ecology. For example, if we are interested in the effects of both X and Z on Y, then a 'causal salad' approach would include both X and Z as predictor variables, leading to biased estimates under the scenarios presented in Figure 1.2 e-h. Particularly when dealing with complex real-world systems, knowing which variables need to be controlled for (or not controlled for) for a given causal query can quickly become unimaginable without the formulized guidelines provided by the SCM framework.

The SCM framework has been implemented across a wide array of disciplines, leading to key insights in epidemiology (e.g., Griffith et al. 2020), paediatrics (e.g., Williams et al. 2018), psychology (e.g., Jiang et al. 2021), and more. Within ecology, this framework has already been utilized to understand the drivers of species-level trait variation (Cronin and Schoolmaster 2018), re-define biodiversity-ecosystem functioning relationships (Schoolmaster et al. 2020), as well as conceptualize wildlife recovery dynamics (Wilson et al. 2021). As a widely applicable causal inference method, the SCM framework holds tremendous potential for shaping the way ecologists understand and apply causal inference. However, several limitations remain. First the SCM framework has only been applied to a few ecological studies (Cronin and Schoolmaster 2018; Schoolmaster et al. 2020; Wilson et al. 2021), and within these, only one (Schoolmaster et al. 2020) included real-world ecological data. Second, these studies focus on niche ecological topics and do not facilitate a comprehensive understanding of how to apply the SCM framework geared

towards an ecological audience. Further, only one study (Schoolmaster et al. 2020) has illuded to the use of the SCM framework for experimental set-ups. However, it is important to note that this framework can be employed across observational, as well as quasi-experimental and experimental settings to guide effective study design and statistical analysis required for drawing causal conclusions from ecological data.

1.5 THESIS OBJECTIVES

A major objective of this thesis is to effectively communicate the SCM framework to a general ecology audience, with the aim of increasing its utility and uptake across future ecological studies. To this end, **Chapter 2** provides a user-friendly review of the SCM framework, with an emphasis on its application for observational causal inference in ecology. This chapter uses simulated ecological examples to detail the key steps required for observational causal inference, overviews common statistical biases that can occur in non-causal correlative studies, as well as highlights the recent application of the SCM framework within the ecological literature. **Chapter 3** applies the SCM framework to determine the drivers of climate-induced coral reef regime shifts in Seychelles, providing the first application of the SCM framework in the coral reef literature. **Chapter 4** applies the SCM framework to understand the socio-environmental drivers of reef fish biomass across temperate and tropical reefs, providing the first application of this framework to a largescale ecological dataset. Whereas Chapter 2 can be referenced to provide a detailed overview of the SCM framework, Chapters 3 and 4 can be used as a practical guide for ecologists aiming to answer causal questions from ecological data.

A second objective of this thesis is to show that the SCM framework can be employed to strengthen causal inference across both quasi-experimental and experimental settings. **Chapter 5** applies the key principles of SCM framework to strengthen causal conclusions drawn from quasi-experimental approaches including matching methods, difference-in-difference, instrumental variables, and regression discontinuity design. **Chapter 6** further highlights the utility of the SCM framework for strengthening causal inference in experimental ecological studies. Both chapters show that DAGs can be used to visualize the overall causal structure of a system, and guide ecologists on how to remove any biases (e.g., confounding) that may still be present under quasi-experimental or experimental settings.

CHAPTER 2 APPLYING THE STRUCTURAL CAUSAL MODEL (SCM) FRAMEWORK FOR OBSERVATIONAL CAUSAL INFERENCE IN ECOLOGY

A version of this work is currently accepted as Arif S, M MacNeil A. *Ecological Monographs.*

2.1 ABSTRACT

Ecologists are often interested in answering causal questions from observational data but generally lack the training to appropriately infer causation. When applying statistical analysis (e.g., generalized linear model) on observational data, common statistical adjustments can often lead to biased estimates between variables of interest due to processes such as confounding, overcontrol, and collider bias. To overcome these limitations, we overview the structural causal model (SCM), an emerging causal inference framework that can be used to determine cause and effect relationships from observational data. The SCM framework uses directed acyclic graphs (DAGs) to visualize a researchers' assumptions about the causal structure of a system or process under study. Following this, a DAG-based graphical rule known as the backdoor criterion can be applied to determine statistical adjustments (or lack thereof) required to determine causal relationships from observational data. In the presence of unobserved confounding variables, an additional rule called the frontdoor criterion can be employed to determine causal effects. Here, we use simulated ecological examples to review how the backdoor and frontdoor criteria can return accurate causal estimates between variables of interest, as well as how biases can arise when they are not employed. We further provide an overview of studies that have applied the SCM framework in ecology. SCM and its application of DAGs have been broadly employed in other disciplines to make valid causal inference from observational data. Their use in ecology holds tremendous potential for quantifying causal relationships and investigating a range of ecological questions without randomized experiments.

2.2 INTRODUCTION

Observational studies in ecology rely on data that have not been experimentally manipulated and are commonly used to understand ecological patterns and processes seen in nature (Carmel et al. 2013). Observational approaches are increasing in relevance due to the emergence of large-scale ecological questions that are not easily manipulated or controlled, such as invasive species and the consequences of climate change. New advances in technology, such as remote-sensing, environmental genetics, and animal-borne sensors, as well as increased availability of data online and from citizen science, have enhanced opportunities to answer previously intractable ecological questions using observational data (Sagarin and Pauchard 2010).

Many observational studies in ecology are aimed at answering causal questions, such as the impact of marine protected areas on fishing communities (Mascia et al. 2010) or the effect of forest fragmentation on species richness (Sam et al. 2014). However, causal inference – the leveraging of theory and deep knowledge to estimate the impact of events, choices or other factors on a given outcome of interest (Cunningham 2021) – is rare. Yet without the consideration of causal relationships, statistical analysis can frequently lead to biased estimates (i.e., estimates that differ from the true parameter being estimated) that undermine ecological inferences by providing non-causal correlations among variables of interest. This is the basis of the often-repeated phrase "correlation does not imply causation" (F.A.D. 1900). We believe that increasing the use of causal inference methods in observational ecology will reduce bias throughout the discipline and lead to more accurate assessments across a range of ecological questions, especially when experimental approaches are unfeasible.

Structural causal modelling (SCM, Pearl 2009) is an emerging causal inference framework, which unifies the strong features of structural equation modeling (SEM; Wright 1921, Shipley 2016) and Rubin's potential outcome framework (Rubin 2005) among others, to create a powerful theory of causation and framework for causal inference. Importantly, this framework can be used to determine cause and effect relationships from observational data, without needing to set up randomized control experiments (Pearl 2009). SCM has been widely employed across other disciplines, including econometrics (Imbens 2020), epidemiology (Pearce and Lawlor, 2016),

paediatrics (Williams et al. 2018) and psychology (Rohrer, 2018), as well as a few ecological studies (Cronin and Schoolmaster 2018; Schoolmaster et al. 2020; Schoolmaster et al. 2022; Arif et al. 2022; Arif and MacNeil 2022). It holds tremendous potential for increasing the use of causal inference across observational ecological studies.

Under the SCM framework, the derivation of causal effects rests on a set of causal assumptions about the data generating process (e.g., X effects Y and not the other way around). These causal assumptions are visualized using directed acyclic graphs (DAGs), which represent a researchers' assumptions about the causal structure of a system or process under study (Pearl 2009; Morgan and Winship 2014). Given a DAG, a graphical rule known as the backdoor criterion determines the sufficient sets of variables for adjustment required to determine causal effects from observational data. When the backdoor criterion cannot be employed – due to the presence of an unobserved confounding variable – a second graphical rule called the frontdoor criterion can be employed. Using simulated ecological examples with specified (i.e., known) causal effects, we define these criteria and review how they can be employed to determine causal effects between variables of interest.

To date, the few ecological studies that have employed the SCM framework have identified key causal relationships across study systems (Cronin and Schoolmaster 2018; Schoolmaster et al. 2020), outlined general steps required for observational causal inference (Cronin and Schoolmaster 2018), as well as clarified SCM theory (Schoolmaster et al. 2022). However, these studies can be niche topics and theoretically complex. Here, we provide an easily accessible overview of the SCM framework, highlighting two key tools – the backdoor and frontdoor criteria – that can be used for causal inference across observational ecological studies.

2.3 Directed Acyclic Graphs (DAGs)

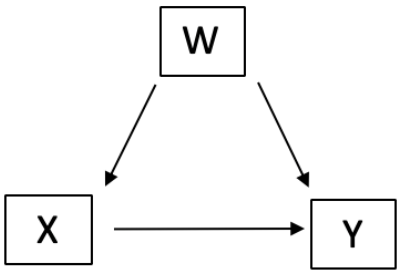

Figure 2.1 A directed acyclic graph (DAG) representing the causal structure between three variables, X, Y and W.

DAGs are used to represent causal relationships within a given system. A DAG consists of a set of nodes (variables) that are connected to each other by edges (arrows). These arrows represent causal relationships between variables, pointing from cause to effect, with causes preceding their effects. For example, the DAG in Figure 2.1 shows that X directly effects Y (X → Y), W directly affects both X (W → X) and Y (W → Y), and W indirectly effects Y through X (W → X → Y). It is important to note that the arrows between nodes (variables) represent hypothesized causal relationships (i.e., a lack of causal relationship can be found following a SCM analysis). On the other hand, a lack of arrow between two nodes assumes no causal relationship between variables, representing strong *a priori* causal assumptions. Therefore, missing arrows encode strong causal assumptions, whereas arrows between nodes represent the possibility of an effect (Elwert, 2013).

A key characteristic of DAGs is that they must be *acyclic*, meaning that they cannot contain bi-directional relationships (i.e., arrows need to be unidirectional) or a feedback loop where a variable causes itself (Glymour and Greenland 2008, Elwert 2014). This limits the application of DAGs to ecological systems that do not contain bi-directionality and or feed-back loops. However, one way to resolve this issue is to articulate the temporal sequence of events more finely (Greenland et al. 1999). For example, if temperature at time one ($Temp_{t1}$) effects ice cover, which then influences temperature at time two ($Temp_{t2}$), $Temp_{t1}$ and $Temp_{t2}$ can be represented as separate nodes within a DAG, without violating acyclic requirements. For interested readers, Schoolmaster et al. (2020) provide a published ecological DAG that incorporates the temporal sequence of events (see their Appendix S2).

DAGs are also non-parametric, meaning that they do not make any assumptions about the stochastic nature of variables or their observation, or the functional form of direct effects (e.g., linear, nonlinear, stepwise) and their effect size (Glymour and Greenland 2008). In this sense, a DAG is qualitative: X → Y only communicates that X causally affects Y in some way, without specifying any other restrictions. This non-parametric nature of DAGs makes them compatible with a wide range of ecological systems.

2.4 THE SCM FRAMEWORK

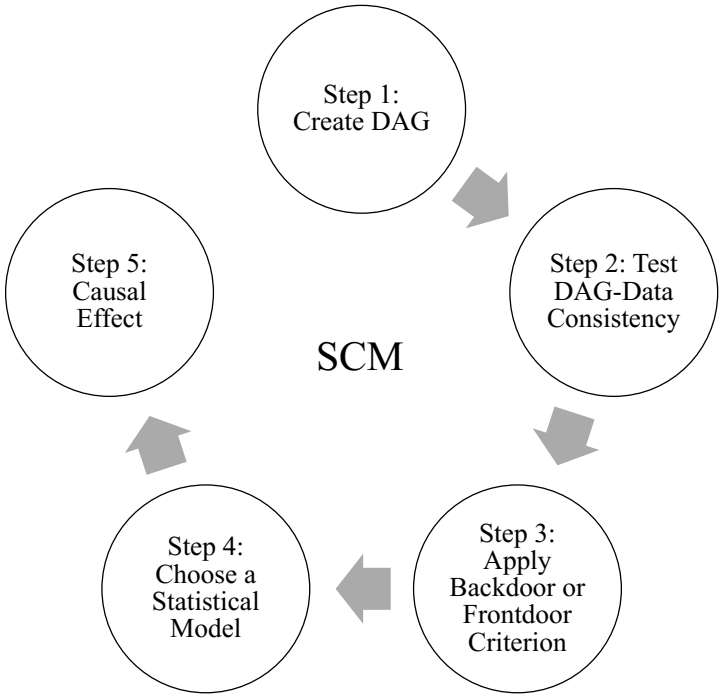

Figure 2.2 A workflow for going from DAGs to causal inference under the SCM framework.

DAGs are central to the SCM framework as they are used to visualize and quantify causal relationships from observational data (Pearl 2009). Figure 2.2 summarizes the SCM framework which includes creating a DAG (step 1), testing a DAG for DAG-data consistency (step 2), applying either the backdoor or frontdoor criterion (step 3), choosing an appropriate statistical model (step 4), and making inference by quantifying a causal effect (step 5). As we walk through our review, we will follow the workflow in Figure 2.2 using simulated ecological examples interspersed with relevant theory and background information.

2.4.1 Step 1: Creating a DAG

DAGs represent a researcher's causal assumptions about the data generating process of a system of process under study (Pearl 2009, Morgan and Winship 2014). As such, researchers should ensure that their DAG represents the complete causal structure of the system or process, including all relevant measured and unmeasured variables, as well as all common causes of any pair of variables included in the DAG (Sprites et al. 2001, Glymour and Greenland 2008). DAGs should also be rigorously justified based on domain knowledge, theory, and research. A combination of background information including experimental data, past literature, and expert knowledge can be used to create DAGs of ecological systems. For example, Ethier and Nudds (2017) gathered information from published literature and local stakeholder knowledge to create DAGs depicting factors affecting population dynamics of bobolink (*Dolichonyx oryzivorus*). In another study, Cronin and Schoolmaster (2018) synthesized past literature to create a DAG representing the causes of trait covariation. Expert opinion can also be elicited to generate DAGs. To ensure credibility and transparency, researchers should apply formal methods for surveying experts, which has been developed within the ecological literature (e.g., Choy et al. 2009, Kuhnert et al. 2010, Martin et al. 2012, Drescher et al. 2013), including for the development of causal diagrams (e.g., Marcot et al. 2006, McNay et al. 2006). For example, Marcot et al. (2006) show how to use expert review to create their DAG on the probability of capture of northern flying squirrels.

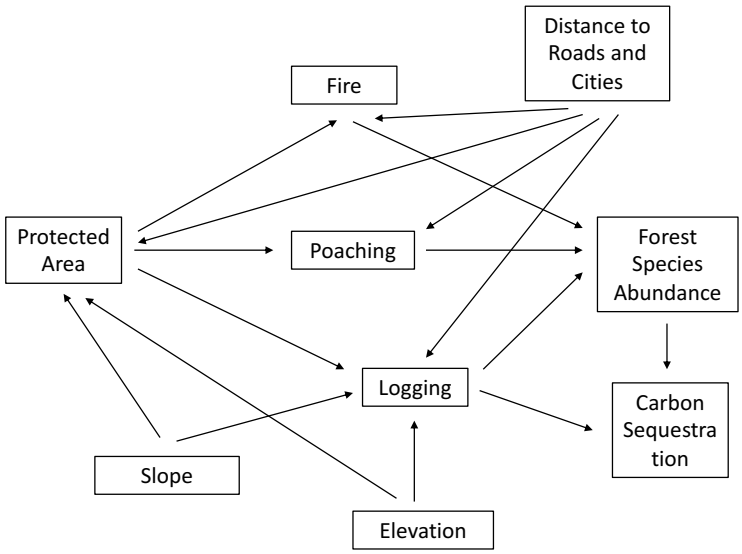

Figure 2.3 A DAG representing how different factors may influence forest species abundance.

As a general ecological example, Figure 2.3 presents a DAG adapted from Adams et al. (2015), showing how different factors are expected to influence forest species abundance across a hypothetical region (Step 1, Figure 2.2). Here, protected areas are shown to effect forest species abundance through three intermediate processes: fire, poaching and logging (Adams et al. 2015). Other variables including distance to roads and cities, slope, and elevation effect both protected areas placement (protected areas are often placed in high and far places; Joppa et al. 2009) and forest species abundance through their effects on fire, poaching and/or logging (Adams et al. 2015). We have created a simulated dataset, matching the causal structure of this DAG (see https://doi.org/10.6084/m9.figshare.19541059 for R code). We will use this DAG and simulated dataset to work through the rest of the SCM workflow (Steps 2-5, Figure 2.2). Specifically, we will aim to answer how protected areas, fire, logging and poaching each effect forest species abundance. Because our simulated data was created with specified (i.e., known) causal effects, we can use it to show how the SCM framework can return accurate causal estimates.

2.4.2 Step 2: Test DAG-Data Consistency

Once a DAG has been created, it can be tested against observational data to check for DAG-data consistency. Simply put, a DAG often asserts multiple independencies that should hold in the observational data, given that both the DAG and observational data are representative of the data generating process. Given a DAG, a pair of variables can be independent of each other (e.g., X is independent of Y) if there are no paths (i.e., a sequence of nodes and arrows) connecting them. As well, a pair of variables can be conditionally independent. Conditional independencies emerge from **d-separation** (dependency separation; Pearl 1988), a graphical rule for deciding whether a variable X is independent of another variable Y, given a set of variable(s), Z in a path.

> **d-separation** (Pearl, 1988): A set of variables, Z, is said to block (or d-separate) a path from one variable to another if either
>
> (i) the path contains at least one arrow-emitting variable that is in Z, or
> (ii) the path contains at least one collider variable (variable with two incoming arrows) that is outside Z and does not cause any variables in Z
>
> If all paths between X and Y are blocked (or d-separated) by Z, then X and Y are independent given Z, written $X \perp Y|Z$. For a more detailed discussion of d-separation, readers can reference Shipley (2000) and Shipley (2016) which discuss d-separation within an ecological context.

DAG-data consistency requires that all implied independencies for a given DAG (including conditional independencies based on d-separation rules) are consistent with the observational dataset. For example, in a simplified DAG, X → Z → Y, X is independent of Y, given Z (an arrow-emitting variable that d-separates the path from X to Y). Therefore, the associated observational data should show that X is independent of Y when Z is adjusted for. Often a DAG will hold many independencies and these independencies can be tested against a dataset to ensure DAG-data consistency. If all implied independencies within a DAG coincide with the dataset, then this supports DAG accuracy. However, if at least one implied independency is refuted (i.e., does not match the data), then the DAG is not consistent with the data and would need to be altered.

For our DAG (Figure 2.3), there are 28 independencies that can be tested against our simulated data to ensure DAG-data consistency (Supplemental Material 2.8.1). In an observational study, we would test these independencies against observational data. Here, we proceed by testing DAG-data consistency using our simulated dataset, to walk readers through the process. Specifically, we use the R package 'dagitty', which provides a user-friendly way to evaluate whether a DAG is consistent with a dataset, even when DAGs become increasingly complex and include many variables (Textor et al., 2016). Dagitty uses a formal test of zero correlation to test whether each identified independency of a specified DAG is consistent with a given dataset (see Textor et al. 2016 for details). Using dagitty, we tested DAG-data consistency and found that all 28 independencies were consistent with our simulated dataset (Step 2, Figure 2.2; see https://doi.org/10.6084/m9.figshare.19541059 for R code). This is expected as our simulated data was created to match the causal structure of our DAG.

In real world applications, a DAG may require a series of adjustments until DAG-data consistency is reached. As an ecological example, Schoolmaster et al. (2020) provide a real-world example of a DAG used to understand the relationship between tree species composition and canopy cover. Their initial DAG failed DAG-data consistency and was subsequently updated using a combination of domain knowledge and results from failed independence tests (Schoolmaster et al. 2020). Anken et al. (2021) further provide general examples and guidelines on updating DAGs based on DAG-data consistency, using the R package 'dagitty'. Importantly, they note that this process should be handled with care and always supported by domain knowledge. Failed independence tests are not necessarily proof that a DAG is incorrect; they can also indicate problems with the data (e.g., if the collected data does not represent the data-generating process). Ultimately, there should be a firm theoretical basis for creating and revising DAGs.

Once a DAG has been sufficiently justified and tested and updated based on DAG-data consistency, the backdoor (or frontdoor) criterion can be employed (Step 3, Figure 2.2). Before moving on to application of backdoor and frontdoor criteria, we briefly review why they can be applied to DAGs to determine causal effects from observational data.

DAGs for causal effects

Causal effects describe to what extent a predictor variable X (i.e., the cause) influences a response variable Y (i.e., an effect). The SCM framework uses counterfactual reasoning to determine the causal effect of X on Y (Pearl 2009). A counterfactual represents the potential outcome that would be realized if a predictor variable X was set to a different value, i.e., X=x. Specifically, a counterfactual for response variable Y is noted as Yx(u), which represents the value of (outcome) Y, had (predictor) X been x in unit (or situation) U = u (Rubin 2005, Morgan and Winship 2014). This counterfactual Yx(u) is represented by the equation:

$$Y_x(\mathrm{u}) \triangleq Y_{Mx}(\mathrm{u}) \qquad [2.1]$$

Under the SCM framework, a DAG represents a structural model, M. In equation 1, M_x stands for a modified version of a model M, where X is intervened upon (*i.e.,* "if X had been x", X=x). Graphically M_x is represented by a modified DAG, where the arrows pointing into X are eliminated. Equation 1 states that the counterfactual Yx(u) is the solution for Y in the modified model M_x (see Galles and Pearl 1998 for axiom of Eq 1).

This definition of counterfactuals can be used to predict the effect of interventions from observational data alone. Under the SCM framework, interventions are denoted by what's known as the do-operator, written do() (Pearl 1995, 2009). For example, the query Q = P(y|do(x)) asks what the distribution of Y would be, if X is set to a particular value of x (*i.e.,* the causal effect of X on Y). Related to Eq 1, this can be defined as

$$P(y|do(x)) \triangleq P_{Mx}(y) \qquad [2.2]$$

showing that the distribution of outcome Y (if X is set to a particular value of x) is equal to the distribution of Y in the modified model M_x (Pearl 1995, 2009).

Given that we do not have post-interventional data (following the distribution of M_x), the question becomes whether the query Q = P(y|do(x)) can be estimable from observational data (following

the distribution of M) and the set of causal assumptions represented by its associated DAG. When a query includes a do-expression, an algebraic procedure known as do-calculus (Pearl 1995) can be used to *equate* post-interventional distributions (those represented in M_x) to pre-interventional (or observational) distributions (those represented in M). To identify an interventional query, *e.g.,* Q = P(y|do(x), the inference rules of do-calculus (outlined in Pearl 1995) need to be repeatedly applied until an expression is obtained that no longer contains a do-operator. If this can be done, then the post interventional query is estimable from observational data. While the application of do-calculus makes for challenging reading, based on its derived inference rules, Pearl created the *backdoor criterion* and the *frontdoor criterion*, which are two DAG-based graphical rules that can be applied to estimate interventional queries from observational data (i.e., the causal effect of X on Y), without the need for do-calculus operations.

2.4.3 Step 3: (Option 1): Apply Backdoor Criterion

The backdoor criterion (Pearl 1993, Pearl 2009) is used to identify a set of variables, Z, that when controlled for, allows the post-interventional query Q = P(y|do(x)) to be accurately estimated from observational data. The backdoor criterion states that a set of variables, Z, is sufficient for estimating the causal effect of X on Y under two conditions:

1. The variables in Z block all *backdoor paths* from X to Y. A *path* within a DAG is any sequence of arrows and nodes connecting two variables of interest, X and Y, regardless of direction. A *backdoor path* is a path between X and Y with an arrow pointing from Y and an arrow pointing into X. Backdoor paths create bias by providing one or more indirect, non-causal pathways through which information can leak from one variable through another, leading to spurious correlation. To block a backdoor path from X to Y, the backdoor path from X to Y must be d-separated. Again, the rules for d-separation are:

   d-separation (Pearl 1988): A set of variables, Z, is said to block (or d-separate) a path from one variable to another if either

(i) the path contains at least one arrow-emitting variable that is in Z, or
(ii) the path contains at least one collider variable (variable with two incoming arrows) that is outside Z and does not cause any variables in Z

2. No element of Z is a descendant of (*i.e.,* caused by) X.

When applied, the backdoor criterion blocks all non-causal pathways between a predictor and response variable of interest, while leaving all causal paths open. As such, the application of backdoor criterion eliminates common statistical biases that can otherwise plague observational studies, including confounding, overcontrol, and collider bias. Supplementary Material 2.8.2 defines each of these biases and shows how the backdoor criterion removes each of them. The main takeaway is that given a DAG, the application of the backdoor criterion will avoid all three biases, allowing for causal estimates to be made.

Given our DAG (Figure 2.3), we can use the backdoor criterion to determine the sufficient set for adjustment required to answer our causal questions (Step 3, Figure 2.2). For example, if we want to quantify the causal effect of protected area on forest species abundance, there are nine backdoor paths that need to be blocked (i.e., d-separated):

1. Forest Species Abundance → Carbon Sequestration ← Logging ← Elevation → Protected Area

2. Forest Species Abundance → Carbon Sequestration ← Logging ← Slope → Protected Area

3. Forest Species Abundance ← Fire ← Distance to Roads and Cities → Logging ← Elevation → Protected Area

4. Forest Species Abundance ← Fire ← Distance to Roads and Cities → Logging ← Slope → Protected Area

5. Forest Species Abundance ← Logging ← Elevation → Protected Area

6. Forest Species Abundance ← Logging ← Slope → Protected Area

7. Forest Species Abundance ← Poaching ← Distance to Roads and Cities → Protected Area

8. Forest Species Abundance ← Logging ← Distance to Roads and Cities → Protected Area

9. Forest Species Abundance ← Fire ← Distance to Roads and Cities → Protected Area

The first four backdoor paths are already blocked because we have not adjusted for a collider variable (i.e., a variable with two incoming arrows: → X ←) in each of these four paths. Specifically, carbon sequestration acts as a collider variable in backdoor paths 1 and 2, and logging acts as a collider in backdoor paths 3 and 4. The remaining backdoor paths do not contain collider variables and must be blocked by adjusting for an arrow-emitting variable that isn't a descendent of (i.e., caused by) protected area, our predictor variable. As such, path 5 can be blocked by adjusting for elevation, path 6 can be blocked by adjusting for slope, and paths 7-9 can all be blocked by adjusting for distance to roads and cities. Collectively, the causal effect of protected area on forest species richness, given this DAG can be quantified by adjusting for slope, elevation and distance to roads and cities.

Given that application of the backdoor criterion can rapidly become difficult to keep track of for increasingly complex DAGs, researchers are encouraged to draw out their DAG on www.daggity.net (instructions within site), which will apply the backdoor criterion and generate the minimal sufficient adjustment set(s) required to determine causal effects, given a specified DAG and causal question. As an example, readers can visit dagitty.net/m18S_bV to work with our protected area DAG. Using this website (see Supplementary Material 2.8.3 for quick steps), to determine the causal effect of fire on forest species abundance, we can adjust for either (distance to roads and cities and protected area) or (logging and poaching). To determine the causal effect of poaching on forest species abundance we can adjust for either (distance to roads and cities and protected area) or (fire and logging). Last, to determine the causal effect of logging on forest

species richness we can adjust for either (distance to roads and cities and protected area) or (fire and poaching). When there are multiple options for a sufficient adjustment set based on the backdoor criterion, researchers can choose a set based on data availability and measurement error. If known, it is best to select the set where variables are measured most accurately.

We note that given our DAG and linear simulated data, causal effects between variables of interest could also be determined using alternative methods such as SEM. However, a strength of the backdoor criterion is that it can allow causal estimation without requiring the availability of all variables in a DAG (Pearl 2009). For example, the effect of protected area on forest species abundance requires observational data on only variables for protected area, forest species abundance, slope, elevation, and distance to roads and cities. By only including variables necessary for answering specific causal queries, this can further enhance estimation accuracy by reducing researchers' reliance on noisy and irrelevant data (MacDonald 2004). In addition, the application of the backdoor criterion does not require lengthy algebraic manipulations, isn't computationally taxing and is compatible across linear and non-parametric statistical approaches (Pearl 2009). Ultimately, it provides ecologists with a widely applicable method for covariate selection across observational studies.

2.4.4 Step 4: Choose a Statistical Model

Once the backdoor criterion is used to determine the sufficient set(s) for adjustment, researchers must decide on an appropriate statistical model to carry out their causal analysis. Since our simulated data was created with a linear causal structure, we have chosen linear regression models for analysis (Step 4, Figure 2.2). However, it is up to each researcher to decide what form of analysis will best suit their data. As DAGs are non-parametric, they make no assumptions about the distribution of variables (e.g., normal) or the functional form of effects (e.g., linear, nonlinear, stepwise), making them compatible with a wide range of statistical methods. DAGs are also compatible with both frequentist and Bayesian statistical approaches since they are used to determine the sufficient set(s) for adjustment, and not the analysis itself. Statistical models developed under the SCM framework are still beholden to the same issues of sample size and

measurement error in terms of the precision of resulting estimates; however, they are based on causal reasoning.

2.4.5 Step 5: Causal Effect

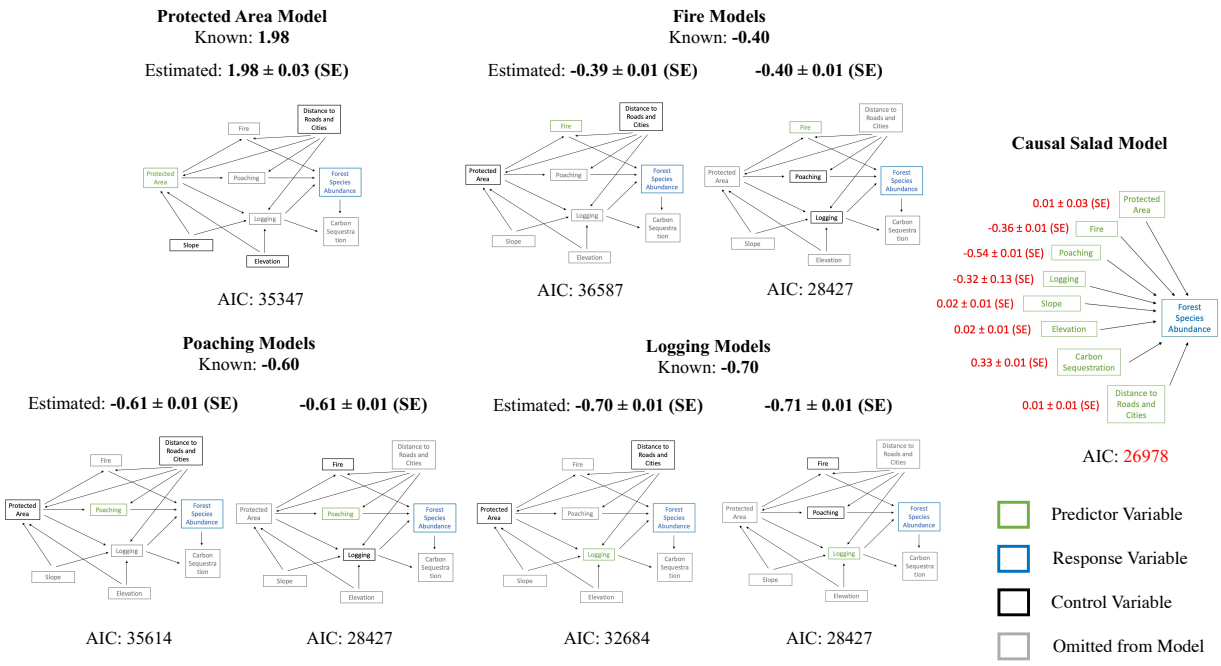

Figure 2.4 Results from linear regression models that employed the backdoor criterion to determine the causal effect of different predictor variables on forest species abundance, using our simulated dataset with specified (i.e., known) causal effects. Predictor, response, and control variables are highlighted in green, blue, and black, respectively; omitted variables are shaded in grey. We chose generalized linear regression as our statistical models; for example, the protected area model is represented by the linear regression equation: $Forest\ Species\ Abundance_i = \alpha + \beta 1 Protected\ Area_i + \beta 2 Slope_i + \beta 3 Elevation_i + \beta 4 Distance\ to\ Roads\ and\ Cities_i + \varepsilon_i$. The known and estimated causal effects, along with AIC values are noted for each model. Lastly, the results from a causal salad model (where all variables are placed under one model) are shown as a contrast, with estimated effects for each included variable noted in red.

Figure 2.4 shows that when the backdoor criterion was used to determine the sufficient set for adjustment, our linear regression models were able to correctly estimate the causal effect between selected predictor variables and forest species abundance, our response of interest (Step 5, Figure 2.2; see https://doi.org/10.6084/m9.figshare.19541059 for R code). This is achieved because the

backdoor criterion blocks all non-causal pathways (i.e., backdoor paths) between our predictor and response variable of interest, while leaving all causal paths open. By adjusting for specific variables (if necessary) to answer specific causal questions, the backdoor criterion can guide causal inference in observational settings.

Importantly, in performing a causal analysis we are not trying to find a 'best model' of the data according to criteria of model fit such as AIC, which seek to find the model with the greatest predictive support, regardless of potential biases present in estimated effect sizes (Arif and MacNeil 2022). For example, in Figure 5 we include a 'causal salad' model (Bhalla 2018; McElreath 2020) typical of ecological observational studies (including our own past work), whereby all available variables thought to affect a response are thrown into one statistical model and subsequently interpreted, without directly addressing the causal structure of the system. In our simulated example, the 'causal salad' model (Figure 2.4) is strongly favored over all other models by AIC, yet it provides an entirely inaccurate picture of the causal structure in the system. Under this approach, we obtain inaccurate estimates of our predictor variables of interest (Figure 2.4). For example, the estimated effect of protected area on forest species abundance is negligible due to overcontrol bias (see Supplementary Material 2.8.2) occurring from the inclusion of fire, poaching and logging, which are intermediate variables between the predictor and response variable of interest. Effect sizes for fire, poaching and logging are also biased due to the inappropriate inclusion of carbon sequestration, which is not a predictor variable but is instead influenced by our response variable of interest. Collectively, these results demonstrate the general principle that the models used for causal inference must be carefully built to consider relevant causal relationships within a system prior to analysis. It also directly undermines 'variance explained' as a modelling objective or arbiter of truth – without causal thinking to support modelling decisions, it is easy to add variables that seem to represent a better model according to a range of widely-used statistical criteria. In this, the backdoor criterion can play a critical role in model development that stands apart from typical model-selection methods, by determining the sufficient set(s) for adjustment required for causal inference.

2.4.6 The Frontdoor Criterion

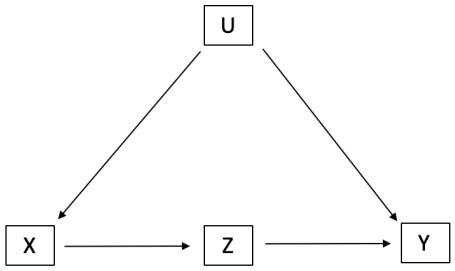

Figure 2.5 A DAG where the effect of X on Y cannot be estimated (due to an unobserved confounding variable U) without the use of the front door criterion.

The DAG-based approach to causal models up to this point has assumed we have observational data on all variables needed to satisfy the backdoor criterion. However, in some circumstances, there may be a known but unobserved variable that confounds our results, preventing application of the backdoor criterion for determining causal effects. For example, if we want to determine the causal effect of X on Y for the DAG in Figure 2.5, the backdoor criterion instructs us to adjust for U. However, U is unobserved, and therefore cannot be used as a covariate in our final model. In such cases, an approach called the frontdoor criterion can be employed for causal inference (Pearl 1995, Pearl 2009). To quantify the effect of X on Y in the presence of unobserved confounders, a variable Z satisfies the frontdoor criterion if:

1. Z blocks all directed paths from X to Y
2. There are no unblocked paths from X to Z
3. X blocks all backdoor paths from Z to Y

Once a Z variable is identified, the causal effect of X on Y can be determined by first employing the backdoor criterion to separately determine the effect of X on Z and Z on Y (Figure 2.5). The product of these two causal effects (i.e., point estimates) then becomes the effect of X on Y (Pearl 1995; 2009). Below we show how to apply the front door criterion to determine the effect of sharks on rays based on a hypothetical ecological example.

Step 1: Create a DAG

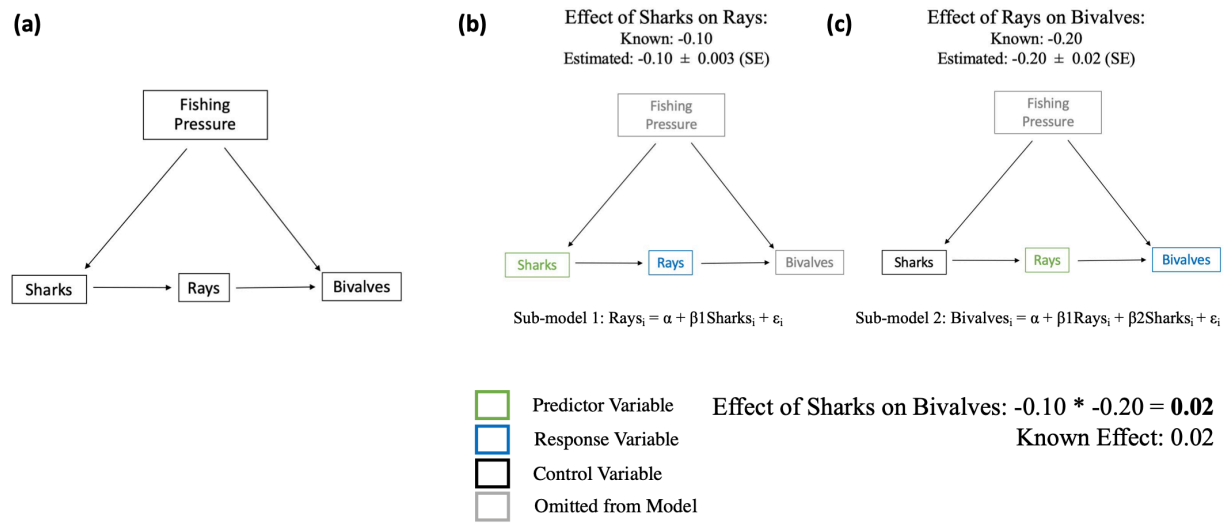

Figure 2.6 Employing the frontdoor criterion. (a) A DAG representing the causal structure between sharks and bivalves. Here, fishing pressure is an unobserved variable, and the frontdoor criterion needs to be employed to determine the effect of sharks on bivalves. (b, c) Employing the frontdoor criterion to determine the effect of sharks on bivalves from our simulated shark-bivalve dataset. Linear regression models were used to first determine the effect of sharks on rays and the effect of rays on bivalves, using the backdoor criterion to determine the sufficient set for adjustment. The product of these two effects gives us the effect of sharks on bivalves. Known causal effects (from our simulated data) are noted for comparison.

The DAG in Figure 2.6a asserts that sharks effect rays, which in turn effect bivalves, through a top-down trophic cascade which has previously been supported (Myers et al. 2007, Buam and Worm 2009) and refuted (Grubbs et al. 2016) in the literature. In our hypothetical scenario, we also assert that fishing pressure effects both sharks and bivalves, but not rays. Here, observational data on fishing pressure isn't available, making it an unobserved variable. Like our prior example, we created a simulated dataset (with known causal effects) matching our DAG (see https://doi.org/10.6084/m9.figshare.19541059 for R code) to demonstrate the use of the frontdoor criterion. Specifically, we will show how to employ the frontdoor criterion to return the causal effect of sharks on bivalves, which we have set to 0.02.

Step 2: Test DAG-data Consistency

Given the DAG in Figure 2.6a, there are two independencies that can be tested based on d-separation rules: 1) fishing pressure is independent of rays, given sharks and 2) sharks are independent of bivalves, given fishing pressure and rays. However, testing either independency requires observational data on fishing pressure (our unobserved variable). Therefore, due to our unobserved confounding variable, DAG-data consistency cannot be tested based on d-separation rules in this case. However, we can still apply the frontdoor criterion for causal estimates with our asserted DAG (unchecked for DAG-data consistency).

Step 3 (Option 2): Apply Frontdoor Criterion

The frontdoor criterion can be employed to find the effect of sharks on bivalves. Rays satisfy the frontdoor criterion since (1) they block all directed paths from sharks to bivalves, (2) there are no unblocked backdoor paths from sharks to rays, and (3) all backdoor paths from rays to bivalves are blocked by sharks (see rules for frontdoor criterion above). To determine the effect of sharks on bivalves, we first need to apply the backdoor criterion to determine the effect of sharks on rays (which can be estimated without any adjustments), and the effect of rays on bivalves (which can be estimated by adjusting for sharks). Both sub-models can employ the backdoor criterion without needing to adjust for fishing pressure (our unobserved variable). The causal effect of sharks on bivalves can then be estimated by multiplying the effect of sharks on rays by the effect of rays on bivalves.

Step 4: Choose a Statistical Model

We use linear regression models because our simulated data was created using linear relationships.

Step 5: Causal Effect

Figure 2.6 shows that when the frontdoor criterion is employed, we were able to accurately determine the causal effect of sharks on bivalves (see https://doi.org/10.6084/m9.figshare.19541059 for R code). Specifically, the product of the effect of sharks on rays (Figure 2.6b) and the effect of rays on bivalves (Figure 2.6c) gave us an accurate causal estimate of sharks on bivalves (0.02), without having to adjust for fishing pressure, our unobserved confounding variable. In contrast, a model with just rays regressed on sharks gives a misleading estimate of 0.99. Here, the correlation between sharks and rays is spurious due to the confounding effect of our unobserved fishing pressure variable.

The front door criterion is not as widely applicable to ecological data as the backdoor criterion, given that it requires a specific causal structure, specified by its three rules (see above). However, in cases where these rules are met, the frontdoor criterion can provide causal estimates, regardless of the strength of unobserved confounding. As well, it can be employed in the presence of multiple unobserved confounding variables.

2.5 Examples of SCM in Ecology

Although currently underutilized, the SCM framework and its application of DAGs has been used to understand the causal structure of ecological systems. Here, we provide an overview of two recent applications of Pearl's SCM framework in ecology.

2.5.1 What Causes Species-level Trait Covariation?

Ecological theory suggests that there may be several causes of species-level trait covariation including size, pace of life, evolutionary history, and ecological condition (Cronin and Schoolmaster 2018). Although numerous studies have attempted to quantify the causal effect of these factors on trait covariation, these studies do not explicitly consider the causal structure driving trait variation, which in turn can lead to inappropriate statistical adjustments and biased estimates. To resolve this, Cronin and Schoolmaster 2018 synthesized relevant literature and

domain knowledge to create a DAG representing the causes of species-level trait covariation that can be applied to across multiple kingdoms.

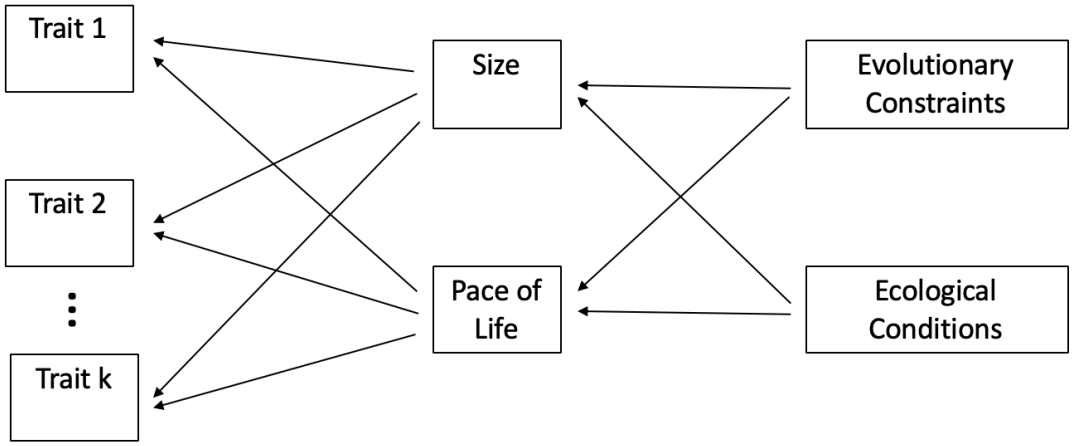

Figure 2.7 A DAG representing how different factors influence species-level trait covariation, from Cronin and Schoolmaster 2018.

As their Figure 2.7 DAG suggests, size and pace-of-life may be two direct causes of trait covariation, and their influence on traits are confounded by evolutionary history and ecological conditions. To determine how size and pace of life effect trait covariation, they first had to accurately quantify their causal effect on each trait, as this information was subsequently used to determine their influence on trait covariation. One way to do this is to employ the backdoor criterion. For example, to determine the effect of size on a trait (e.g., Trait 1 in Figure 2.7), the backdoor criterion instructs us to adjust for either pace of life or evolutionary constraints and ecological condition to remove the confound of evolutionary history and ecological condition. In contrast, previous studies have estimated the effect of either size or pace of life on traits without first controlling for these confounding variables (e.g., Brown et al. 2004; Johnson et al. 2012). Another widely accepted approach has been to first account for evolutionary constraints and then analyze the residuals (e.g., Bielby et al. 2007). However, Cronin and Schoolmaster 2018 show that these approaches lead to erroneous estimates about the causes of trait covariation. They also showed that methods including principle component analysis (PCA) and exploratory factor

analysis (EFA) are not able to partition trait covariance when the direct causes (size and pace of life) are correlated due to shared drivers (evolutionary history and ecological conditions). This is concerning as several high-profile studies have used these techniques to reach their conclusions (e.g., Wright et al. 2004 concluded from a PCA that size is the only causes of lead trait covariance). Taken together, a well-considered DAG guides ecologists on the sufficient set(s) for adjustment required to quantify the causes of trait-covariation and further highlights the utility of Pearl's SCM framework for observational causal inference.

2.5.2 Is Biodiversity a Cause of Ecosystem Functioning?

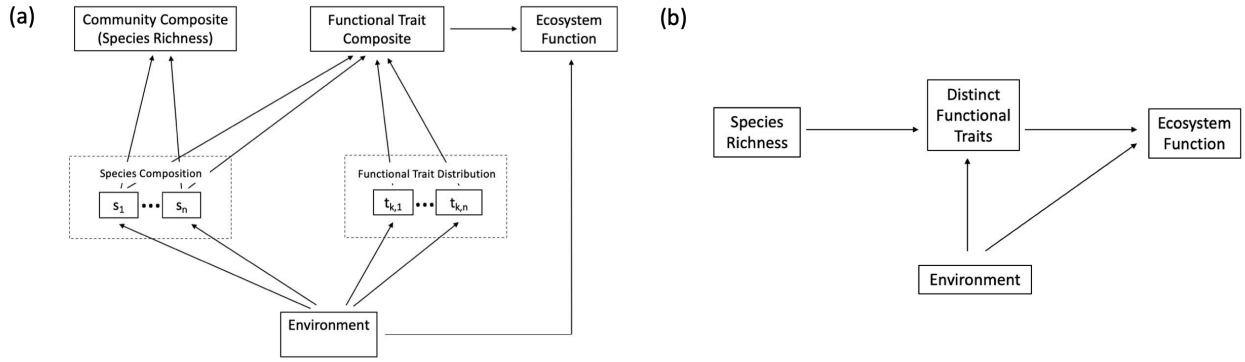

Figure 2.8 Two DAGs representing the causal relationship between biodiversity and ecosystem function. The DAG in (a) is from Schoolmaster et al. 2020 and the DAG in (b) is from Grace et al. 2021.

A central goal of ecology is to understand the causes of ecosystem functioning (Mittelbach 2012); however, correctly identifying these causes has been difficult because there are numerous hypothesized drivers that are often interrelated. A widespread belief among ecologists is that biodiversity is a prominent cause of ecosystem functioning (Tilman et al. 2014). Hundreds of papers have published Biodiversity-Ecosystem Function (BEF) correlations across various ecological systems, with conflicting theories and conclusions (Schoolmaster et al. 2020). To better understand whether biodiversity causally effects ecosystem functioning, Schoolmaster et al. 2020

created a DAG by synthesizing BEF literature and logic (Figure 2.8a). Their DAG deviates from the standard model whereby species richness is assumed to effect ecosystem functioning through functional trait diversity (Loreau 2001), and instead posits that species composition effect both species richness and functional trait diversity, with functional trait diversity driving ecosystem functioning (Figure 2.8a).

Given their DAG (Figure 2.8a), the backdoor criterion states that functional trait distribution and the environment needs to be adjusted for to determine the causal effect (or lack thereof) of biodiversity on ecosystem function. Using simulated and empirical data, Schoolmaster et al. 2020 show that when this is done, there is no causal relationship found between biodiversity and ecosystem functioning. Instead, they argue that previous observational studies that have found an association between biodiversity and ecosystem function arise from model misspecification (i.e., having an incomplete or incorrect set of predictors). For example, given their DAG, confounding bias from failing to condition on environmental factors can lead to spurious (i.e., non-causal) associations between biodiversity and ecosystem functioning. Given their DAG, Schoolmaster et al. 2020 conclude that BEF correlations are non-causal associations. Instead, their model suggests that it is species composition and not biodiversity that drives ecosystem functioning.

Recently, a comment on Schoolmaster et al. (2020) was published by Grace et al. (2021), criticizing their DAG and conclusions, asserting that biodiversity causally effects ecosystem functioning. They provide an alternative DAG, which maintains that biodiversity can causally effect ecosystem functioning indirectly through its effect on trait diversity (i.e., 'distinct functional trait'; Figure 2.8b). This aligns with the standard model (Loreau, 2001) on BEF correlations being causal. Schoolmaster et al. 2022 responded with a comprehensive reply, addressing critiques of their DAG, clarifying the SCM framework, and showing that the standard model and past interpretations of BEF experiments are not supported by causal analyses. Interestingly, Schoolmaster et al. 2022 note that the simulations provided by Grace et al. (2021) do not represent the standard model DAG they defend, but instead map onto the DAG presented by Schoolmaster et al. (2020).

Although the issue of BEF correlation versus causation has yet to be resolved, there now exist two contradictory DAGs that can be used to focus critical debate and deepen our understanding of this potential process. As noted by Grace et al. (2021), DAGs allow researchers to state their causal assumptions explicitly and transparently. Ultimately, this allows other researchers to examine those causal assumptions and subsequent interpretations critically, as was done by Grace et al. (2021) and Schoolmaster et al. 2022. Ultimately, communicating and critiquing researchers' causal assumptions through DAGs may lead to a deeper understanding of BEF correlations, as well as for other ecological phenomena.

2.6 Additional Considerations

2.6.1 Inaccurate or Unknown Causal Structure

One of the potential limitations of DAGs is that they may not accurately represent the true causal nature of an ecological system. Simply put, inaccurate DAGs will lead to inaccurate causal inference. This can arise when using incorrect theory and background information, or by creating DAGs based on available data, rather than incorporating all relevant variables (such as omitted or unobserved variables). However, as a researcher's causal assumptions are explicitly stated through graphical representation, DAGs allow reviewers to explicitly critique and correct potential problems with far more transparency than is typical (Pearl 2009). Further, the ability to test DAG-data consistency via d-separation rules facilitates more reliable conclusions (Textor et al. 2016).

We believe that SCM should be used whenever researchers have causal objectives and sufficient background knowledge to create and justify the assertions made in their DAG. If, however, the causal structure between the predictor and response variables of interest are not fully known, but there exists enough background knowledge and support to create several plausible DAGs (each of which support DAG-data consistency), it may be advantageous to present all DAGs as plausible alternatives, reflecting this epistemic uncertainty. This should provide more accurate estimates, especially when predictor variables have the same covariate adjustments across a range of

plausible DAGs. We emphasize that since several DAGs can pass DAG-data consistency, it is always imperative to first justify a DAG (or set of DAGs) based on theory, instead of relying solely on DAG-data consistency.

2.6.2 Application within Quasi-Experimental and Experimental Approaches

In recent years, ecologists have promoted the use of quasi-experimental methods for causal inference, including propensity score matching, before-after-control-impact (BACI) studies, regression discontinuity design (RDD), and instrumental variables (IV; Butsic et al. 2017, Larsen et al. 2019). Here, DAGs and the principles of the SCM framework (e.g., the backdoor criterion) can be used to create more robust study designs as well as explicitly communicate assumptions required for quasi-experimental approaches (see Chapter 5 for details). For example, propensity score matching is employed to remove confounding bias associated with ecological observational studies (e.g., Ramsey et al. 2019). However, although past ecological studies assume that all confounding variables enter a propensity score analysis, it is unclear how these variables relate to one another and within the broader causal structure of a study system. Without this knowledge, it is unclear whether there are unmeasured confounding variables that need to be included in the propensity score (leading to confounding bias) or whether the accidental inclusion of non-confounding variables will lead to other forms of bias (e.g., overcontrol and collider bias; Shrier 2009, Sjolander 2009, Mansournia et al. 2013). As noted by Pearl, for a propensity score analysis to be valid, the selected variables that enter a propensity score must satisfy the backdoor criterion to remove bias (Pearl 2009). In other words, the variables that enter a propensity score should be the sufficient set for adjustment based on the backdoor criterion. For an overview of how the SCM framework can guide quasi-experimental study designs, we refer readers to Chapter 5. By utilizing DAGs and the principles of the SCM framework, ecologists can design more robust quasi-experimental approaches, while explicitly communicating their causal assumptions to their audience.

DAGs and the SCM framework can also guide causal inference in experimental studies (see Chapter 6 for details). Like observational studies, experimental studies rely on causal assumptions

that must be ensured by the researcher (Kimmel et al. 2021). Here, DAGs can be used to understand if data collected from an experimental set up (e.g., RTCs) can be used for causal inference or if there are sources of bias that need to be accounted for (e.g., Williams et al. 2018; Schoolmaster et al. 2020; Schoolmaster et al. 2022). For example, Williams et al. (2018) overview a RCT investigating the effect of an intervention promoting breastfeeding on cognitive development during childhood. A DAG of this study clarifies that only using data from individuals who attend a follow-up session can lead to collider bias because both the intervention and outcome can affect the likelihood of individuals following up; therefore, follow-up data should not be distinctly analyzed (Williams et al. 2018). As an ecological example, Schoolmaster et al. (2020) use their biodiversity-ecosystem function (BEF) DAG to argue that BEF experiments do not directly manipulate biodiversity, but rather manipulate community structure, failing to isolate for the biodiversity effect.

2.7 Conclusion

Ecology has relied on observational data from its inception (Elton 1927), yet use of causal logic has typically been limited to RCTs. Our ongoing reliance on observational data to understand fundamental questions in ecology requires the increased use of valid causal inference methodologies. Here we have introduced Pearl's SCM framework, which allows causal inference to be made in a wide range of observational contexts. The SCM framework uses DAGs to visualize the hypothesized causal structure of a system or process under study, allowing researchers to explicitly communicate their causal assumptions. Once a DAG has been built that is sufficient to characterize a system or process under study, the backdoor or frontdoor criterion can be employed to guide appropriate statistical adjustments required for causal inference. Doing so can improve the validity of causal conclusions drawn from observation-based research.

2.8 Supplemental Material

2.8.1 Testing DAG-data Consistency

Once a DAG has been created, it can be directly tested against observational data to ensure DAG-data consistency, which is when data are consistent with all independencies implied by a given DAG. For example, there are 28 independencies that are implied by our forest species abundance DAG (Figure 3):

1. Carbon Sequestration ⊥ Distance to Roads and Cities | Fire, Logging, Poaching
2. Carbon Sequestration ⊥ Distance to Roads and Cities | Forest Species Abundance, Logging
3. Carbon Sequestration ⊥ Protected Area | Fire, Logging, Poaching
4. Carbon Sequestration ⊥ Protected Area | Forest Species Abundance, Logging
5. Carbon Sequestration ⊥ Elevation | Distance to Roads and Cities, Logging, Protected Area
6. Carbon Sequestration ⊥ Elevation | Fire, Logging, Poaching
7. Carbon Sequestration ⊥ Elevation | Forest Species Abundance, Logging
8. Carbon Sequestration ⊥ Fire | Forest Species Abundance, Logging
9. Carbon Sequestration ⊥ Poaching | Forest Species Abundance, Logging
10. Carbon Sequestration ⊥ Slope | Distance to Roads and Cities, Logging, Protected Area
11. Carbon Sequestration ⊥ Slope | Fire, Logging, Poaching
12. Carbon Sequestration ⊥ Slope | Forest Species Abundance, Logging
13. Distance to Roads and Cities ⊥ Forest Species Abundance | Fire, Logging, Poaching
14. Distance to Roads and Cities ⊥ Elevation
15. Distance to Roads and Cities ⊥ Slope
16. Forest Species Abundance ⊥ Protected Area | Fire, Logging, Poaching
17. Forest Species Abundance ⊥ Elevation | Distance to Roads and Cities, Logging, Protected Area
18. Forest Species Abundance ⊥ Elevation | Fire, Logging, Poaching
19. Forest Species Abundance ⊥ Slope | Distance to Roads and Cities, Logging, Protected Area
20. Forest Species Abundance ⊥ Slope | Fire, Logging, Poaching
21. Elevation ⊥ Fire | Distance to Roads and Cities, Protected Area
22. Elevation ⊥ Poaching | Distance to Roads and Cities, Protected Area
23. Elevation ⊥ Slope
24. Fire ⊥ Logging | Distance to Roads and Cities, Protected Area
25. Fire ⊥ Poaching | Distance to Roads and Cities, Protected Area

26. Fire ⊥ Slope | Distance to Roads and Cities, Protected Area
27. Logging ⊥ Poaching | Distance to Roads and Cities, Protected Area
28. Poaching ⊥ Slope | Distance to Roads and Cities, Protected Area

We use the R package 'dagitty' to test if these 28 independencies coincide with our simulated data (see https://doi.org/10.6084/m9.figshare.19541059 for R code). All 28 independencies are consistent with our simulated data, ensuring DAG-data consistency.

2.8.2 Preventing Overcontrol, Confounding and Collider Bias:

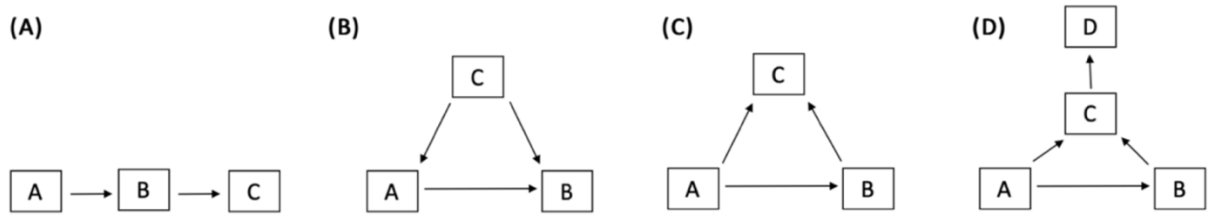

Supplementary Figure 2.9 DAGs depicting a (a) chain, (b) fork, (c) collider, and (d) descendant of a collider. To estimate the effect of A on B, biases associated with these structures, which include confounding, overcontrol, and collider bias, must be avoided.

The components of a DAG can be broken down into three types of causal structures: chains, A → C → B (Supplementary Figure 2.9a); forks, A ← C → B (Supplemtary Figure 2.9b); and colliders A → C ← B (Supplementary 2.9c). When appropriate statistical adjustments are not made, these three causal structures can each lead to a specific type of bias: overcontrol, confounding, and collider bias (e.g., see also Elwert, 2014). Here, we review how the backdoor criterion directs us to avoid these biases, which can otherwise plague observational correlative studies.

Overcontrol bias: In a chain, A → C → B, two variables may be associated because one variable, A, indirectly causes the other, B (Supplementary Figure 2.9a). If we want to find the effect of A on B, conditioning on C would block the association flowing from A to B. This is known as 'overcontrol bias' and can be resolved by not conditioning on an intermediate variable between predictor and response variable. Here, the application of the backdoor criterion instructs us to not condition on variable C in order to determine the effect of A on B.

Confounding bias: In a fork, A ← C → B, two variables, A and B, may be associated because of a common cause, C (Supplementary Figure 2.9b). If we want to find the effect of A on B, then not conditioning on C would cause a spurious, or biased association between A and B. This is known as 'confounding bias' and can be resolved by conditioning on the common cause. Here, the application of the backdoor criterion instructs us to condition on variable C in order to determine the effect of A on B.

Collider bias: In a collider, A → C ← B, two variables, A and B, may be associated because they have a common outcome, C (i.e., a collider, Figure 2.9c). Conditioning on a collider variable, C, creates a spurious, or biased association between A and B. This is known as 'collider bias' and can be resolved by not conditioning on a collider (e.g., C in Figure 2.9c and Figure 2.9d) or any descendant of a collider (e.g. D in Figure 2.9d). Here, the application of the backdoor criterion instructs us not to condition on variable C (Fig S1c and S1d) or D (Fig S1d) in order to determine the effect of A on B.

While examples of confounding variables can be found throughout the ecology literature (*e.g.,* land use change acts as a confound for determining the effect of climate on elevational species redistribution, Guo et al., 2018), the same cannot be said for overcontrol and collider bias. A literature search using Web of Science with the search terms "ecology" and "overcontrol" or "collider" resulted in no papers that mentioned these biases. This does not mean that these biases do not occur in observational ecological studies, but rather that they are not knowingly adjusted for. It is also noted that while ecologists are aware of confounding, it does not mean they are accurately being adjusted for across observational studies. For example, covariate adjustments resulting in the inclusion of multiple potential confounders can lead to bias, instead of reducing it (Shrier and Platt, 2008).

Given a DAG, the application of the backdoor criterion will eliminate overcontrol, confounding, and collider bias, allowing for more reliable causal estimates from observational data.

2.8.3 Applying the Backdoor Criterion Using www.dagitty.net

The backdoor criterion can be employed to determine the sufficient set(s) for adjustment required to quantify causal estimates between variables of interest. Since the application of the backdoor criterion can become complex and time consuming, users can use www.dagitty.net as a resource to draw their DAG (instructions within site). Given a DAG, and specified predictor and response variables, this website will automatically generate the backdoor adjustment set without needing to do it by hand. A saved version of our forest species abundance DAG (Figure 2.3) can be found here: dagitty.net/m18S_bV. Given this DAG, let's use this website to determine the effect of protected area on forest species abundance:

1. To select the predictor variable, protected areas, click on this variable, then select "exposure" under the Variable tab on the top left of the screen.
2. To select the response variable, forest species abundance, click on this variable, then select "outcome" under the Variable tab on the top left of the screen.
3. Under the 'Causal effect identification' tab on the top right of the screen, select 'Adjustment (total effect)'. The backdoor adjustment set will be noted here. Given our DAG, to determine the effect of protected are of forest species abundance, distance to roads and cities, elevation, and slope must be adjusted for.

2.8.4 Literature Cited

Elwert, F. 2014. Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology,* 40, 31-53.

Guo, F., Lenoir, J., Bonebrake, T. 2018. Land-use change interacts with climate to determine elevation species redistribution. *Nature Communications,* 9, 1315.

Shrier, I., Platt, R. 2008. Reducing bias through directed acyclic graphs. *BMC Medical Research Methodology,* 8, 70.

CHAPTER 3 CAUSAL DRIVERS OF CLIMATE-MEDIATED CORAL REEF REGIME SHIFTS

A version of this work has been published: as Arif S, Graham NAJ, Wilson S, M MacNeil A. 2022. *Ecosphere* 13(3): e3956.

3.1 ABSTRACT

Climate-induced coral bleaching events are a leading threat to coral reef ecosystems and can result in coral-macroalgal regime shifts that are difficult to reverse. It is unclear how different factors causally influence regime shift or recovery trajectories after a bleaching event. Here, we use structural causal modeling (SCM) and its application of directed acyclic graphs (DAGs) to determine how key factors affect regime shift vs recovery potential across coral reefs in Seychelles, which were severely impacted by bleaching events in 1998 and 2016. Our causal models reveal additional causal drivers of regime shifts, including initial macroalgae cover, wave exposure, and branching coral cover. We also find that reduced depth and structural complexity and increased nutrients increase the likelihood of regime shifting. Further, we use a DAG-informed predictive model to show how recovering reefs are expected to change after a recent 2016 bleaching event, suggesting that three out of twelve recovering reefs are expected to regime shift given their pre-disturbance conditions. Collectively, our results provide the first causally-grounded analysis of how different factors influence post-bleaching regime shift vs recovery potential on coral reefs. More broadly, SCM stands apart from previous observational analysis and provides a strong framework for causal inference across other observational ecological studies.

3.2 INTRODUCTION

Climate-induced coral bleaching is currently one of the leading threats to coral reef ecosystems and is expected to be an increasingly frequent stressor for coral reefs in the future (Hughes et al. 2018). A potential long-term consequence of climate-induced bleaching events is that they can lead to a coral-macroalgal regime shift, whereby the benthic composition abruptly transitions from a coral dominated reef to one dominated by macroalgae (Graham et al. 2015). Regime shifts have become a key concern for coral reef conservation as they represent substantial change and degradation of coral reefs worldwide, which are often difficult to reverse (Bellwood et al. 2004). For example, climate-driven regime shifts have led to altered trophic structure, diversity, and species composition of reef fish communities (Hempson et al. 2017; Robinson et al. 2019a), as well as increased catch instability and fisheries dependence on herbivorous fish (Robinson et al. 2019b). It is important to note that not all coral reefs shift towards algal domination after a bleaching event (e.g., Gilmour et al. 2013), and past research has found correlations between key predictor variables and regime shift vs recovery trajectory (Graham et al. 2015). However, these findings were not grounded in causal inference, the deliberate use of specific methods to infer causation (Pearl 2009).

A literature review of causal inference in coral reef ecology (Supplementary Material 3.6.7) shows that no observational studies to date have employed causal inference methods to determine relationships for reef regime shifts; however, most studies used causal language to communicate their results (e.g., “the effect of X on Y”). With the development of structural causal modelling (SCM; Pearl 2009), there is an opportunity to revisit these analyses to understand causal effects of factors influencing regime vs recovery trajectories on coral reef ecosystems. SCM is a causal inference method that can be used to determine causal relationships from observational data. It uses directed acyclic graphs (DAGs) to visualize the causal structure of a system under study, which is then used to guide covariate selection required for observational causal inference (see Methods for details). Already, DAGs have been applied across several ecological

studies, leading to more informed insights across study systems (Cronin and Schoolmaster 2018; Schoolmaster et al. 2020).

Here, we employ SCM to determine how key factors have influenced recovery vs regime shift trajectories after a widespread bleaching event in Seychelles. The mass coral bleaching event of 1998 reduced coral cover by over 90% across 21 coral reef sites in the inner Seychelles (Graham et al. 2015). Post-disturbance trajectories in cover of coral and macroalgae resulted in approximately half of these sites recovering live coral, while the other half shifted towards macroalgal domination. In addition, the 2016 bleaching event further impacted reefs across Seychelles, reducing coral cover by 70% on those reefs that had recovered from the 1998 disturbance (Wilson et al. 2019). It is currently unclear whether reefs that recovered from the 1998 bleaching event will recover a second time, or undergo regime shifts to macroalgal dominance. Applying SCM to this unique and well-studied system, our study addresses two research questions: (1) does SCM lead to additional insights on the causal drivers of regime shifts following the 1998 bleaching event in Seychelles and (2) which of the reef sites that recovered following 1998 are expected to regime shift as a result of the 2016 bleaching event? By employing SCM, our study aims to better understand the causal factors influencing coral-algal regime shift dynamics.

3.3 METHODS

3.3.1 Ecological Surveys

Seychelles were impacted by a widespread climate-induced coral bleaching event in 1998 (Goreau et al. 2000). Pre- (1994) and post-bleaching (2005, 2008, 2011, 2014) surveys of 21 coral reefs throughout the inner Seychelles Islands were conducted using identical methods (see Graham et al. 2015 for details). Coral reefs were categorized as either regime shifting or recovering based on data collected before and 16 years after the 1998 bleaching event. Regime shifting reefs had post-disturbance macroalgae cover greater than coral cover and patterns through time showing high and/or increasing cover of

macroalgae over time. In contrast, recovering reefs had post-disturbance coral cover greater than macroalgal cover and patterns through time showing high and/or increasing levels of coral cover over time. In total, 12 reefs were classified as recovering and 9 were classified as regime shifting (see Graham et al. 2015 for details).

3.3.2 Causal Framework

Pearl's structural causal model (SCM; Pearl 2009) framework uses directed acyclic graphs (DAGs) to visually represent the causal structure of a system under study. Specifically, nodes within a DAG represent variables, with *directed* arrows between nodes representing possible causal effects (e.g., X → Y shows that X affects Y). A lack of arrow assumes no causal relationship between variables, and these represent our priori assumptions about where causality cannot occur (Elwert 2014). DAGs must be acyclic, meaning that they cannot contain bi-directional relationships or a feedback loop where a variable either directly or indirectly causes itself (Elwert 2014). However, DAGs may still represent ecological systems with bi-directional relationships by more finely articulating the temporal sequence of events (Greenland et al. 1999). DAGs are also non-parametric, making them compatible with a wide range of statistical analyses (Glymour and Greenland, 2008).

Figure 3.1 A workflow for going from DAGs to causal inference under the SCM framework.

The first step of a SCM is to create a DAG (Step 1, Figure 3.1). DAGs should be created and justified based on the accumulation of domain knowledge, which can include expert opinion and past and ongoing research. DAGs should include all measured and unmeasured variables required to depict the system or process under study, as well as all

common causes of any pair of variables included in the DAG (Sprites et al. 2001; Glymour and Greenland, 2008). Here, we have created a DAG representing how variables may influence regime shift vs recovery trajectories in Seychelles (Figure 3.2) based on ecological knowledge, and our own past and ongoing research.

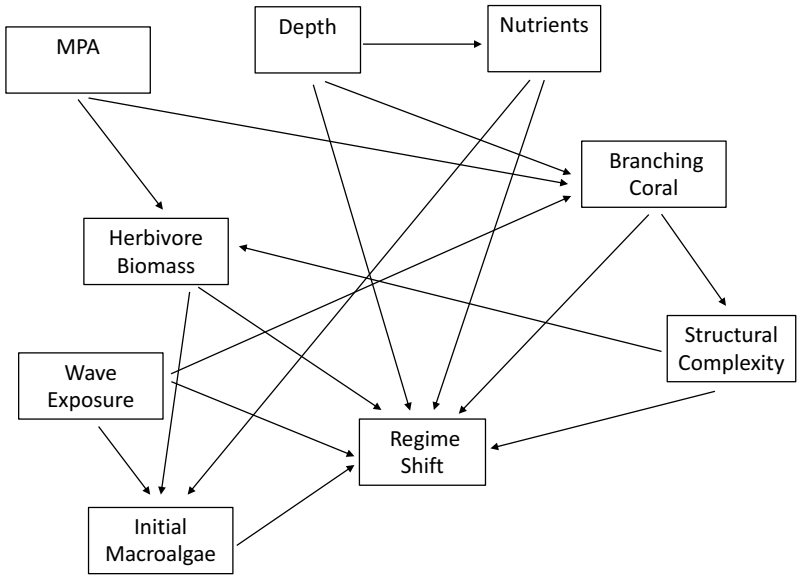

Figure 3.2 DAG representing the causal structure of factors influencing regime shift vs recovery trajectories in Seychelles coral reefs. We note that herbivorous fish biomass, branching coral, macroalgae cover, and structural complexity represent pre-disturbance observational data (1994), whereas depth, MPA status, and wave exposure data represent values that are assumed to stay stable across years. Although nutrient data was collected in 2004, they are expected to capture pre-disturbance nutrient levels across reef sites (see Graham et al. 2015).

Several key factors are assumed to influence regime shift trajectory after a climate-induced bleaching event: MPA, depth, nutrients, branching coral, structural complexity, herbivorous fish biomass, wave exposure and pre-disturbance macroalgal cover (Figure 3.2). We note that we used 1994 (pre-disturbance) data for branching coral, macroalgae cover (not included in Graham et al. 2015 analysis), herbivorous fish biomass, and structural complexity because we wanted to know how their condition prior to the 1998 bleaching event would influence regime shift vs recovery trajectory and to resolve any bi-

directional relationships that may exist between our predictor variables and response (Glymour and Greendale 2008). A detailed rational of each directed arrow in our DAG is presented in our Supplementary Table 3.1.

Once a DAG is created, it can be checked for DAG-data consistency (Step 2, Figure 3.1). Simply put, a specified DAG will have (often many) independencies between variables (e.g., A is independent – or *d-separated* – from B, if C is adjusted for) that should be compatible with the observational dataset, given that both the DAG and data are representative of the data-generating process. If *all* implied independencies are compatible with the data, it provides overall support for a DAG. We tested our DAG for DAG-data consistency using the R package 'dagitty' (Textor et al. 2016) which confirmed that all 32 independencies implied by our DAG were consistent with our observational data (Supplementary Material 3.6.3).

A finalized DAG is then used to determine which variables need to be controlled (e.g., through covariate adjustment) for to determine a causal effect for a specific causal query. This is the critical step that separates SCM from correlative observational studies. Specifically, a graphical procedure known as the *backdoor criterion* guides covariate selection required to determine the causal effect of X on Y (Pearl 2009; Supplementary Material 3.6.4). In short, the backdoor criterion instructs us to block all non-causal pathways (i.e., backdoor paths) between our predictor and response variable of interest, while leaving all causal paths open. As such, the application of backdoor criterion eliminates common statistical biases that can otherwise plague observational studies, including confounding, overcontrol, and collider bias (see Chapter 2.8.2 for definitions) . Here, we employ the backdoor criterion to guide covariate selection for *each predictor variable* expected to influence recovery vs. regime shift trajectory (Step 3, Figure 3.1; Supplementary Material 3.6.4). We note that only the effect size of the predictor variable of interest is interpreted for its associated model, with additional covariates acting as required controls. This approach differs from the 'causal salad' model (Bhalla 2018; McElreath 2020) commonly used throughout ecology – including in our own work –

where all assumed predictor variables are placed into one model and subsequently interpreted.

Once the backdoor criterion is applied for covariate selection, researchers must choose an appropriate statistical model (Step 4, Figure 3.1). As DAGs are non-parametric, they make no assumptions about the distribution of variables (e.g., normal) or the functional form of effects (e.g., linear, nonlinear, stepwise), making them compatible with a wide range of statistical methods. Here, we applied a Bayesian logistic regression analysis to each of our causal models, where the response variable was 0 for recovering sites, and 1 for regime shifting sites. We standardized our data by subtracting the mean of each variable and dividing by 2 standard deviations in order to assess relative effect sizes of our predicter variables (Gelman and Hill, 2007). We ran our models using the 'rethinking' package on R, using weakly informative priors. Our final Bayesian logistic regression models (one for each predictor variable) were:

$$Y_i \sim \text{Bernoulli}(p_i)$$

$$MPA\ model: logit(p_i) = \beta_0 + \beta_1 MPA$$

$$Nutrient\ model: logit(p_i) = \beta_0 + \beta_1 Nutrient\ + \beta_2 Depth$$

$$Herbivore\ biomass\ model: logit(p_i) = \beta_0 + \beta_1 Herbivore\ biomass\ + \beta_2 MPA\ + \beta_3 Structural\ complexity$$

$$Branching\ coral\ model: logit(p_i) = \beta_0 + \beta_1 Branching\ coral\ + \beta_2 MPA\ + \beta_3 Depth\ + \beta_3 Wave$$

$$Depth\ model: logit(p_i) = \beta_0 + \beta_1 Depth$$

$$Structural\ complexity\ model: logit(p_i) = \beta_0 + \beta_1 Structural\ complexity\ + \beta_2 Branching\ coral$$

$$Wave\ model: logit(p_i) = \beta_0 + \beta_1 Wave$$

$$Initial\ macroalgae\ model: logit(p_i) = \beta_0 + \beta_1 Macroalgae\ + \beta_2 Herbivore\ biomass\ + \beta_3 Depth\ + \beta_3 Nutrients + \beta_4 Wave$$

Priors (for standardized data):

$$\beta_0 \sim Cauchy(0,10)$$

$$\beta_{1,\ldots,N} \sim Student\_t(4,0,2.5)$$

3.3.3 DAG-Informed Predictive Model

We further created a predictive model to determine how recovering reef sites would be expected to respond to subsequent bleaching events in 2016 (n=12). Predictor variables included all factors that were found to directly influence this response based on our previous DAG-based analysis, which were: depth, nutrient, branching coral cover, structural complexity, wave exposure and initial macroalgal cover (see results). This approach captures all relevant variables assumed to influence a response variable of interest and is expected to lead to both high in-sample and out-of-sample predictive accuracy. We excluded herbivorous fish biomass because our results suggested that our coarse biomass metric may not be representative of herbivore grazing effects (see discussion). We employed a Bayesian logistic regression model, with 0 for recovering sites, and 1 for regime shifting sites. Our final predictive Bayesian logistic model is specified as:

$$\mathrm{Y}_i \sim \mathrm{Bernoulli}(p_i)$$

$$Predictive\ model\text{: } logit(p_i) = \beta_0 + \beta_1 Depth + \beta_2 Nutrient + \beta_3 Branching\ coral + \beta_4 Structural\ complexity + \beta_5 Macroalgae + \beta_6 Wave$$

Priors (for unstandardized data):

$$\beta_{0,\ldots,N} \sim N(0,10)$$

We used data from the 1998 bleaching event to train our model. A posterior predictive check (McElreath, 2020) showed that our predictive model was able to correctly identify the trajectory of 90% (19/21) of sites after the 1998 bleaching event. Our trained model was then used to predict recovery vs. regime shift trajectory following the 2016 bleaching event using 2014 data.

3.4 RESULTS

3.4.1 Causes of Regime Shifts

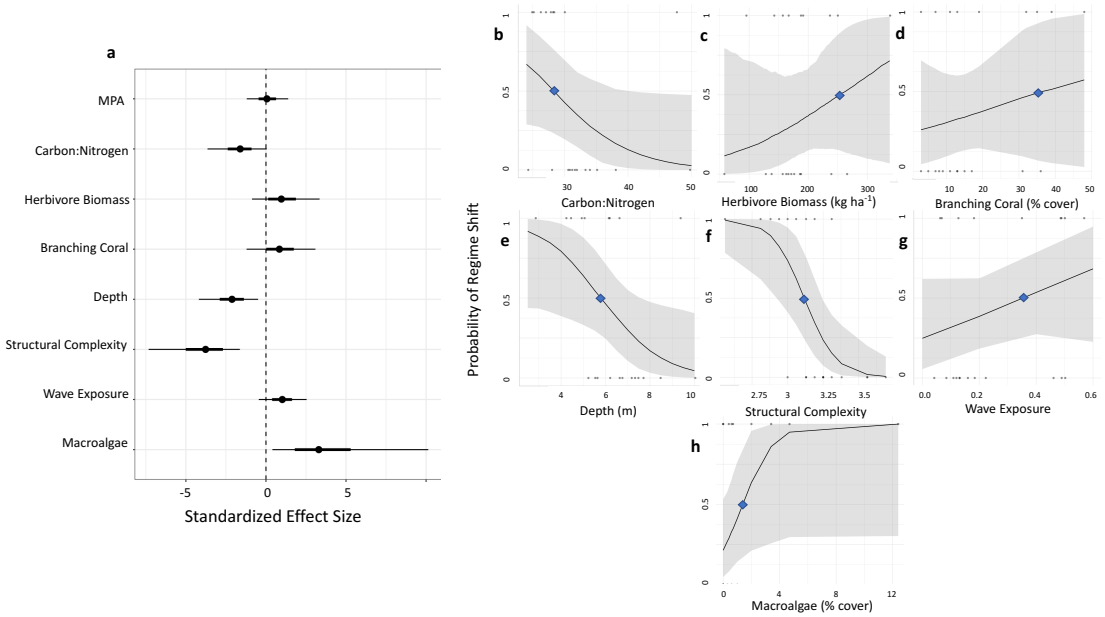

Figure 3.3 Causal effect of factors on regime shift trajectory. (a) Standardized effect size of factors influencing regime shift trajectory. Parameter estimates are posterior median values (dot) 50% percentile interval (PI; thick lines), and 89% highest percentile interval (PI; thin lines). (b-h) Marginal plots of predictor variables (for continuous variables) affecting regime shift trajectory: solid black line represents the predicted median value of all drawn posterior predictive samples; grey shading represents the 95% Bayesian predictive intervals; blue dot represents the point at which regime shifts and recovery are equally likely.

Our causal models show that depth and structural complexity decreased the likelihood of a climate-induced regime shift following the 1998 bleaching event in Seychelles (Figure 3.3). Similar to Graham et al. 2015, we find that deeper and structurally complex reefs are more resilient against climate-induced bleaching events. Our nutrient causal model also shows that high nutrients (low carbon:nitrogen ratios) increase the likelihood of regime shifts (Figure 3.3). In addition to these insights, our causal models revealed several factors that influenced regime shift trajectory, which were not evident in our past correlative study (Graham et al. 2015). Importantly, higher initial macroalgal cover

increased the likelihood of regime shifting and had the strongest effect size (and largest variation) of all predictor variables (Figure 3.3). High wave exposure was also shown to increase the likelihood of regime shifting (Figure 3.3). To a lesser extent, both higher herbivorous fish biomass and branching coral cover increased the likelihood of regime shifting (Figure 3.3).

3.4.2 Predictions for Future Regime Shifts

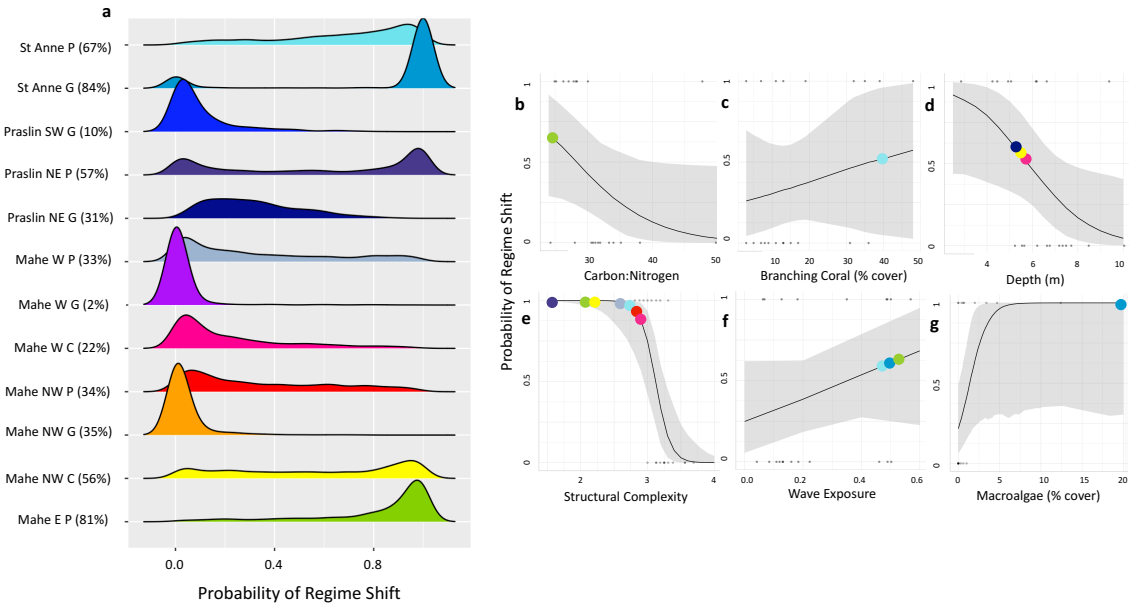

Figure 3.4 Predictions of future regime shifts in Seychelles. (a) Probability of regime shift for recovering sites across Seychelles after a subsequent bleaching event in 2016, based on our predictive model. Estimates are based on 1000 samples drawn from the posterior predictive distribution, with mean probability (%) noted in brackets. (b) Marginal plots of predictor variables affecting regime shift trajectory, highlighting 2014 values from recovering reefs that that fall above the 0.5 probability for regime shifting.

Our predictive model suggests that out of the 12 sites that recovered from the 1998 bleaching event in Seychelles, five reefs have a greater than 50% probability of regime shifting following the second bleaching event in 2016 (Figure 3.4a). These reefs include St Anne P, St. Anne G, Praslin NE P, Mahe NW C, and Mahe E P (Figure 3.4a). Of

these, three reefs are particularly vulnerable, with over 60% probability of regime shifting: St Anne P (67%), St. Anne G (84%), Mahe E P (81%, Figure 3.4a).

3.5 DISCUSSION

3.5.1 Causal vs. Correlative Analysis

Our causal models indicate multiple additional factors that influenced regime shift vs recovery trajectory, including initial macroalgae cover, wave exposure, and branching coral cover, which were not evident from our past correlative analysis on the same study system (Graham et al. 2015), with most of the same variables. The main difference between a causal vs correlative analysis lies in covariate selection. Our previous analyses placed all assumed predictor variables into one model (i.e., a 'causal salad' model), and subsequently removed all variables which did not show an effect to arrive at a final model that included: herbivore biomass, depth, structural complexity, and nutrients, as well as juvenile coral density, which was not considered in the current analysis (see next paragraph; Graham et al. 2015). In comparison, here we used our ecological knowledge to create our DAG and subsequently applied the backdoor criterion to build a causal model for each of our predictor variables. The backdoor criterion guides covariate selection required to determine causal relationships from observational data, theoretically equivalent to what would be expected under a perfectly executed randomized controlled experiment (Pearl 2009). Given a DAG, the application of the backdoor criterion removes spurious correlations that may otherwise plague observational studies.

DAGs also allow us to carefully think about which observational data may best suite our causal questions, reflecting the causal structure of our DAG (Figure 3.2). In our prior study, we used post-disturbance 2005 data for herbivorous fish biomass, to allow the short-term disturbance effects on herbivore biomass to be accounted for (Graham et al. 2015). However, this could be bi-directional, since reef sites that are in the process of regime shifting can lead to increased post-disturbance herbivorous fish biomass. Such bi-directionality is prohibited in a DAG, and the presence of bi-directional relationships in a

regression analysis can lead to erroneous results due to the presence of simultaneity bias (Merton, 1968). To remedy this, we used pre-disturbance 1994 data for herbivore biomass, allowing for a directed arrow pointing from herbivorous fish biomass to regime shift trajectory. We previously used post-disturbance juvenile coral density (2011 data) as a predictor variable (Graham et al. 2015). While post-disturbance juvenile coral density data may influence regime shift trajectory, it may also be influenced by the regime shift vs recovery process itself. Indeed, higher post-disturbance juvenile coral density may be more a process of recovery (Hughes et al. 2010; Gilmour et al. 2013) than a factor that influences regime shift vs recovery trajectories. Given this, we excluded post-juvenile coral density as a predictor variable in our DAG and analysis. In addition, we included pre-disturbance macroalgal cover as a predictor variable in our DAG, which was not previously considered, but ended up having the strongest effect size (Figure 3.3).

In our review of all coral reef regime shifts, we found that observational studies did not apply causal inference techniques (Supplementary Material 3.6.7). Whereas some studies did not include any covariate adjustments, and others employed a 'causal salad' approach, and across all studies, causal analysis and the consideration of the overall causal structure between variables of interest was missing (Supplementary Material 3.6.7). Here, the application of DAGs and the backdoor criterion provides of formal causal framework that can guide future covariate selection on coral reef regime shift studies. To further demonstrate how the backdoor criterion can lead to improved causal estimates, we compare our results guided by the backdoor criterion to other statistical models, including those that don't include any covariate adjustments as well as a 'causal salad' model (Supplementary Table 3.2). Our estimates vary significantly across different models, highlighting the potential for more theory-driven and reliable causal models across observational studies (Supplementary Table 3.2).

3.5.2 Additional Insights

Our causal models reveal several key additional factors that affected regime shift vs recovery trajectory that were not evident in our previous study (Figure 3.3). In particular,

we find that higher initial (pre-disturbance) macroalgal cover, a variable not included in Graham et al. (2015) increases the likelihood of regime shifting (Figure 3.3). Following a disturbance event, additional space made available through coral mortality can lead to macroalgae expansion and subsequent inhibition of coral recruitment (McCook et al. 2001). In Seychelles, the rate of coral recovery has been strongly negatively associated with the rate of macroalgae cover increase (Wilson et al. 2012). Initial macroalgal cover may lead to post-disturbance coral-algal shifts in two ways. First, reef macroalgae may have higher thermal tolerance than coral species, with some macroalgae species experiencing no mortality under elevated temperatures (Anderson 2006). Therefore, established macroalgae may remain intact following a bleaching event, creating a competitive starting point for macroalgal expansion. Second, various macroalgae species exhibit limited dispersal ability, with propagule settlement and recruitment remaining close to the source population (Capdevila et al. 2018). As such, having higher levels of macroalgae already established within a site may create a strong basis for macroalgae recruitment following a disturbance event. To our knowledge, the impact of initial macroalgal cover prior to disturbance on reef recovery dynamic has yet to be investigated elsewhere. Future research should examine whether this pattern is generalizable across other reef ecosystems, as well as the mechanisms that underlie this process.

Our results also indicate that wave exposure increased the likelihood of regime shifting (Figure 3.3). Wave-exposed reefs with higher water flow can favor macroalgae growth through increased exposure to and uptake of inorganic carbon and nutrients, which can lead to higher photosynthesis and growth rates of macroalgae (Hurd 2000). Wave action can also limit coral growth and larval settlement (Gove et al. 2015), as well as remove coral through colony dislodgement and abrasive damage (Madin and Connolly 2006). On the other hand, higher flow rates associated with greater wave exposure have also been attributed to reduced bleaching susceptibility and faster recovery of corals through the passive diffusion of harmful oxygen radicals that can accumulate in corals under high SST and irradiance (Nakamura and Woesik 2001; Nakamura et al. 2003). However, McClanahan et al (2005) show that increased flow rate is correlated with increased bleaching intensity in Mauritius, reasoning that higher flow rates occurred in regions with

lower variation in water temperature, which in turn can increase bleaching susceptibility (Safaie 2018). Collectively, these factors may promote post-disturbance coral-algal regime shifts at wave-exposed reef sites. For example, wave exposure was found to be the main determinant of a coral-algal regime shift following a catastrophic typhoon disturbance in Micronesia (Ross et al. 2015). In Seychelles, recovery rate for recovering reef sites has also been negatively associated with increased wave exposure (Robinson et al. 2019c).

We found a positive association between herbivorous reef fish biomass and regime shift trajectory (Figure 3.3). Following a climatic disturbance, herbivorous reef fish is expected to limit coral-algal shifts through grazing pressure, which limits the growth of macroalgae and enhances coral recruitment through creating space for larval settlement (McCook et al. 2001; Hughes 2007; Mumby and Harborne 2010). Yet in Seychelles, the pre-disturbance biomass of herbivorous fish seem to be positively correlated with regime shift occurrence. Previous studies in Seychelles have shown that higher herbivore biomass post-disturbance was correlated with a reduced likelihood of regime shifting (Graham et al. 2015), but a slower recovery rate on recovering reefs (Robinson et al. 2019c). Coarse biomass metrics of herbivorous fish biomass, which combine distinct functional groups, may not be representative of true grazing effects on coral reefs, which are tightly linked to size structure and functional composition of herbivore assemblage (Nash et al. 2015; Steneck et al. 2018; Robinson et al. 2019c). Future studies in Seychelles can address this gap by looking at more accurate grazing metrics, which have been shown to clarify the effect of herbivorous reef fish across other coral reef systems (Steneck et al. 2018). In general, and particularly along reefs where fishing has drastically reduced herbivorous fish biomass, it is expected that lower levels of herbivory will limit recovery and lead to more coral-algal shifts following climatic disturbance (Hughes 2007; Mumby and Harborne 2010).

Higher initial branching coral cover was also shown to increase the likelihood of regime shifting (Figure 3.3). Branching corals are often vulnerable to bleaching given their lower

heat tolerance (Loya et al. 2001). Following a climate-induced bleaching event, this can lead to higher loss of branching coral; for example, surveys following the 2016 bleaching event in Seychelles found that 95% of *Acropora* and *Pocillopora* colonies were either bleached or recently dead (Wilson et al. 2019). Ultimately, reefs with high pre-disturbance branching coral cover may result in a large coral mortality and the subsequent availability of open space, creating favourable conditions for macroalgae growth and dominance, which may overwhelm the effect of grazing pressure from herbivorous reef fish (Williams et al. 2001). Moreover, once dead, the structures provided by branching corals erode rapidly, impacting reef fish and other organisms (Sheppard et al. 2002). On the other hand, branching corals are often fast growing (Darling et al. 2012) and can play a critical role in coral recovery following disturbance, as demonstrated by the recovering reefs in Seychelles (Robinson et al. 2019c; Wilson et al. 2019) and the Great Barrier Reef (Linares et al. 2011). How this trade-off between heat tolerance and growth rates relates to longer term patterns and predictions of coral recovery may depend on the frequency and intensity of disturbances. Given that severe bleaching events are now expected every six years (Hughes et al. 2018), branching coral cover may ultimately be at a disadvantage over heat-tolerant corals (Kubicek et al. 2012; Kubicek et al. 2019), with reefs with higher branching coral cover potentially being more vulnerable to coral-algal regime shifts.

3.5.3 Predicting Recovery vs Regime Shift Trajectory Post 2016 Bleaching Event

Our predictive model suggests that out of the 12 sites that recovered from the 1998 bleaching event in Seychelles, five reefs show a greater than 50% probability of regime shifting following the second bleaching event in 2016, with three reefs having a probability above 60% (Figure 3.4a). These reef sites are vulnerable to coral-algal regime shifts due to a combination of factors. For example, St Anne P, which shows a 67% chance of regime shifting had (according to 2014 data) very high branching coral, low structural complexity, and high wave exposure (Figure 3.4c,e,f). In comparison, St Anne G, showing a 84% change of regime shifting had high wave exposure and very high macroalgae cover (Figure 3.4f,g). Last, Mahe E P, which had a 81% chance of regime

shifting had high nutrients (low carbon:nitrogen), low structural complexity and high wave exposure (Figure 3.4 b,e,f). We note that our predictions may underestimate reef vulnerability to the 2016 bleaching event as our predictive model was 'trained' using results from the first bleaching event in 1998. Indeed, Seychelles reefs may now be under unstable equilibria (May 1997; Scheffer and Carpenter 2003), essentially requiring less cumulative stress to drive coral-algal shifts (Mumby and Hastings 2009).

Our predictions suggest that several factors can come together to influence regime shift vs recovery trajectories on coral reefs impacted by subsequent bleaching events. Given that severe bleaching events are now expected every 6 years (Hughes et al. 2018), conservation and management efforts may benefit from prioritizing locations where the rate of warming and threat of frequent bleaching is lowest (Van Hooidonk et al. 2016) and where recovery from climatic disturbances is most likely (Cote et al. 2010; Graham et al. 2020). Specifically, reefs with increased depth, resilient coral species and structural complexity, and low macroalgae, nutrients, and wave exposure may be more resilient against future climatic disturbances in Seychelles. Deeper, structurally complex granitic reefs with higher cover of heat-tolerant massive corals and low macroalgae may be important areas for future conservation efforts in this region (Graham et al. 2006; Dajka et al. 2019). Collectively, incorporating these findings into management efforts may aid in prioritizing potentially resilient coral reefs amidst our current environmental and climate change crisis.

3.6 Supplementary Material

3.6.1 Application of SCM

Applying SCM to determine causal relationships from observational data require 4 key steps outlined in Figure 3.1. Here we detail each step that was taken to determine the causal drivers of coral-algal regime shifts following the 1998 bleaching event in Seychelles.

3.6.2 Step 1: Create DAG

We created our DAG by consolidating our domain knowledge with past and ongoing research, as well as the expert opinion of NAJG, SKW and MAM, who have extensive knowledge of Seychelles coral reefs in particular. We included relevant variables required to show how factors may influence climate-induced regime shifts across Seychelles coral reefs and included common causes of any pair of variables included in the DAG (Sprites et al. 2001, Glymour and Greenland, 2008). The justification of the causal links between each variable in our DAG is detailed in Supplementary Table 3.1.

Supplementary Table 3.1**.** Justification for the causal links in our coral reef DAG (Figure 3.2).

Causal Assumptions	Rational
MPA → Herbivorous Fish Biomass	It is widely accepted among coral reef scientists that MPAs generally lead to higher herbivorous fish biomass through the reduction of fishing pressure (e.g. Edwards et al. 2014; Soler et al. 2015). In Seychelles, MPAs have led to higher herbivorous biomass under stable conditions (Jennings et al. 1996), as well as after bleaching and the subsequent loss of corals (Graham et al. 2020).
Structural Complexity → Herbivorous Fish Biomass	Structural complexity provides a wider range of habitat variability for reef fish species, which can increase herbivorous fish biomass (Graham and Nash, 2013; Rogers et al. 2014; Verges et al. 2011).

Structural Complexity → Regime Shift	Structural complexity can play a critical role in reef recovery by providing niche space for coral settlement and survival (Connell et al. 1997; Victor 2008). It can also increase the abundance and diversity of reef fish through habitat availability and refuge (Verge´s and Vanderklift, 2011; Graham and Nash 2013), which in turn can limit algal growth and regime shifts (Bellwood et al. 2004; Mumby and Harborne 2010). Structural complexity has played an important role in coral recovery in Guam (Colgan 1987) and the eastern Pacific (Guzman and Cortes 2007) and has previously been correlated with post-bleaching recovery in Seychelles (Graham et al. 2015).
MPA → Branching Coral	MPAs had higher branching coral cover relative to fished sites across Seychelles coral reefs prior to the 1998 coral bleaching event (Graham et al. 2020). There is some uncertainty regarding why higher branching coral cover was found in MPAs pre-bleaching, and MPA placement in areas with higher branching coral may be a presumed factor. However, one of the marine reserves representing three of our MPA sites was chosen based on bird conservation and not coral cover (Graham et al. 2020). We suspect that MPAs may

	have more pre-disturbance branching coral relative to fished sites as branching corals are easily damaged by fishing gear and other recreational activities (McManus et al. 1997; Cros and McClanahan 2003; Strain et al. 2019), particularly trap fishery (Stevens 2020).
Branching Coral → Regime Shift	Branching coral are generally more susceptible to bleaching and mortality from heat stress (Loya et al. 2001; Sheppard et al. 2002). Moreover, once dead, the structures provided by branching corals erode rapidly, impacting reef fish and other organisms (Sheppard et al. 2002; Graham and Nash 2013; Wilson et al. 2019). This in turn can leave reefs with higher pre-disturbance branching coral cover more vulnerable to regime shifts following a climate-induced bleaching event.
Wave Exposure → Regime Shift	Wave exposure can favor macroalgae growth through increased exposure and uptake of inorganic carbon and nutrients (Larned and Atkinson, 1997; Hurd 2000). Wave action can also limit coral growth and larval settlement (Gove et al. 2015), as well as remove coral through colony dislodgement and abrasive damage (Madin and Connolly 2006). These factors are expected to increase the chances of

	regime shifting at wave-exposed sites. However, we note that higher flow rates have also been attributed to reduced bleaching susceptibility and faster recovery of corals (Nakamura and Woesik 2001; Nakamura et al. 2003), which in turn may limit coral-algal shifts following a climate-induced bleaching event.
Depth → Regime Shift	Many threats, such as sediment input and fishing pressure, are less pronounced in deeper waters, making them less vulnerable to regime shifts (Bridge et al. 2013). Reduced light penetration and temperature may also decrease the likelihood of heat stress and subsequent coral bleaching (Slattery et al. 2011; Bridge et al. 2013). However, we note that light is also required for recovery of corals and on turbid reefs, coral may recovery more quickly in shallow water (Evans et al. 2020). Overall, we expect that depth will likely lower the changes of a regime shifting, which has previously been correlated with decreased likelihood of regime shifts in Seychelles (Graham et al. 2015).
Nutrient → Regime Shift	Increased nutrient loads from human sources can increase macroalgae growth and abundance (Schaffelke and Klumpp 1997; Burkepile and Hay, 2006), limit

	coral growth rates (Koop et al. 2001), reproductive success and settlement (Harrison and Wallace 1990; Ward and Harrison 1997), as well as increase bleaching susceptibility of corals (Wooldrige 2009). Increased nutrient loading in combination with reduced herbivory has also been shown to make coral reefs less resilient against disturbances (Hughes et al. 2003; Burkepile and Hay 2006). Collectively, these factors are expected to increase the likelihood of regime shifting for reefs with higher nutrients, a correlation that has previously been shown in Seychelles (Graham et al. 2015).
Depth → Nutrient	Nutrients tend to decrease with depth because shallow reefs are more exposed to nutrient run-off (Bridge et al. 2013).
Depth → Branching Coral	Branching coral species inhabit shallow regions and are therefore expected to decrease with depth (e.g. Tamir et al. 2019).
Herbivore Fish Biomass → Regime Shift	Herbivorous reef fish can limit coral-algal shifts through increased grazing pressure, which limits the growth of macroalgae and enhances coral recruitment through creating space for larval settlement (McCook et al. 2001; Bellwood et al.

	2004; Mumby et al. 2006; Hughes 2007; Mumby et al. 2013).
Wave Exposure → Branching Coral	Wave exposure influences branching coral distribution patterns and colony size, with higher wave exposure reducing branching coral cover (Done, 1983; Madin and Connolly, 2006).
Branching coral → Structural Complexity	Branching coral is positively correlated with structural complexity (Graham and Nash, 2013). Branching coral may be particularly likely to contribute to fine-scale structural complexity on reefs (Chabanet et al. 1997).
Initial Macroalgae → Regime Shift	Higher cover of pre-disturbance macroalgae may create a stronger basis for macroalgae expansion post-disturbance given that macroalgae species tend to have limited dispersal ability, with propagule settlement and recruitment remaining close to the source population (e.g., Kendrick and Walker, 1991; Kendrick and Walker 1995; Capdevila et al. 2018).
Herbivore Biomass → Initial Macroalgae	Herbivore fish can decrease macroalgal cover through increased grazing pressure (McCook et al. 2001; Bellwood et al. 2004; Mumby et al. 2006; Hughes 2007; Mumby et al. 2013).

Wave Exposure → Initial Macroalgae	Wave exposure can increase macroalgae cover through increased exposure and uptake of inorganic carbon and nutrients (Larned and Atkinson, 1997; Hurd 2000).
Nutrients → Initial Macroalgae	Nutrients can increase macroalgae cover by simulating macroalgal growth (Schaffelke and Klumpp 1997; Burkepile and Hay, 2006).

3.6.3 Step 2: Test DAG-Data Consistency

Once a DAG has been created, it can be directly tested against observational data, to ensure DAG-data consistency. If an initial DAG does not pass DAG-data consistency it can be altered until DAG-data consistency is reached. Testing DAG-data consistency emerges from **d-separation** rules, which is a DAG-based rule for deciding whether a variable X is independent of another variable Y given a set of variable(s), Z in a *path*. A path is a sequence of arrows and nodes connecting two variables of interest within a DAG.

> **d-separation** (Pearl, 1988): A set of variables, Z, is said to block (or d-separate) a path from one variable to another if either
>
> (i) the path contains at least one arrow-emitting variable that is in Z, or
> (ii) the path contains at least one collider variable (variable with two incoming arrows) that is outside Z and does not cause any variables in Z. A collider variable is a variable that is influenced by two or more variables in a given path. Collider variables block the association between the variables that influence it.

There are 32 independencies (based on d-separation rules) that are implied by our DAG (see below). These independencies can be tested against our observational data to ensure DAG-data consistency. Our observational data should be consistent with these

conditional independencies, given that both our DAG and data are representative of the data generating process. For example, given our DAG, MPA should be independent of wave exposure (see 1 below), and MPA should also be independent of macroalgae when herbivore biomass, nutrients, and wave exposure are adjusted for (see 2 below). Collectively, all 32 conditional independencies implied by our DAG should be consistent with our observational data to provide support for our overall DAG structure.

1. MPA ⊥ Wave Exposure *(meaning MPA is independent of wave exposure)*
2. MPA ⊥ Macroalgae | Herbivore Biomass, Nutrients, Wave Exposure *(meaning MPA is independent of macroalgae given that herbivore biomass, nutrients, and wave exposure are adjusted for)*
3. MPA ⊥ Macroalgae | Depth, Herbivore Biomass, Wave Exposure
4. MPA ⊥ Depth
5. MPA ⊥ Nutrients
6. MPA ⊥ Structural Complexity | Branching Coral
7. MPA ⊥ Regime Shift | Branching Coral, Depth, Herbivore Biomass, Structural Complexity, Wave Exposure
8. Herbivore Biomass ⊥ Wave Exposure | Branching Coral, MPA
9. Herbivore Biomass ⊥ Wave Exposure | MPA, Structural Complexity
10. Herbivore Biomass ⊥ Depth | Branching Coral, MPA
11. Herbivore Biomass ⊥ Depth | MPA, Structural Complexity
12. Herbivore Biomass ⊥ Nutrients | Depth
13. Herbivore Biomass ⊥ Nutrients | Branching Coral, MPA
14. Herbivore Biomass ⊥ Nutrients | MPA, Structural Complexity
15. Herbivore Biomass ⊥ Branching Coral | MPA, Structural Complexity
16. Wave Exposure ⊥ Depth

17. Wave Exposure ⊥ Nutrients
18. Wave Exposure ⊥ Structural Complexity | Branching Coral
19. Macroalgae ⊥ Depth | Branching Coral, MPA, Nutrients, Wave Exposure
20. Macroalgae ⊥ Depth | MPA, Nutrients, Structural Complexity, Wave Exposure
21. Macroalgae ⊥ Depth | Herbivore Biomass, Nutrients, Wave Exposure
22. Macroalgae ⊥ Branching Coral | Depth, MPA, Structural Complexity, Wave Exposure
23. Macroalgae ⊥ Branching Coral | MPA, Nutrients, Structural Complexity, Wave Exposure
24. Macroalgae ⊥ Branching Coral | Depth, Herbivore Biomass, Wave Exposure
25. Macroalgae ⊥ Branching Coral | Herbivore Biomass, Nutrients, Wave Exposure
26. Macroalgae ⊥ Structural Complexity | Branching Coral, Herbivore Biomass, MPA
27. Macroalgae ⊥ Structural Complexity | Depth, Herbivore Biomass, Wave Exposure
28. Macroalgae ⊥ Structural Complexity | Herbivore Biomass, Nutrients, Wave Exposure
29. Depth ⊥ Structural Complexity | Branching Coral
30. Nutrients ⊥ Branching Coral | Depth
31. Nutrients ⊥ Structural Complexity | Branching Coral
32. Nutrients ⊥ Structural Complexity | Depth

We used the R package 'dagitty' (Textor et al. 2016) to test DAG-data consistency between our observational data and specified DAG. Dagitty uses a formal test of zero (partial) correlation for each identified independency based on d-separation rules. To mitigate problems around multiple testing (e.g., for complex DAGs with many testable

implications), the p-values obtained are then corrected using the Holm-Bonferroni method (see Textor et al. 2016 for details). These tests assume linearity and multivariate normality, though other options are available (see Ankur et al. 2021). Our observational data was found to be consistent with all 32 conditional independencies implied by our DAG providing support for our overall DAG. Once a DAG is finalized (i.e., ensures DAG-data consistency), we can employ the backdoor criterion to guide model selection required to determine the effect of a predictor variable on a response variable of interest. We note that since several DAGs can pass DAG-data consistency, it is critical to ensure that a finalized DAG can be justified (or set of DAGs) based on theory, instead of relying solely on DAG-data consistency.

3.6.4 Step 3: Apply of the Backdoor Criterion for Model Selection

The backdoor criterion is used to identify a set of variables, Z, that when adjusted for, allows the causal effect of X n Y to be accurately estimated from observational data, given a DAG. The backdoor criterion (Pearl, 1993, 2009) states that a set of variables, Z, is sufficient for estimating the causal effect of X on Y under two conditions:

1. The variables in Z block all backdoor paths from X to Y. A *path* within a DAG is any sequence of arrows and nodes connecting two variables of interest, X and Y. A *backdoor path* is a path between X and Y with an arrow pointing into X. The backdoor paths are the paths that create bias by providing an indirect, non-causal path along which information can flow. To block a backdoor path from X to Y, X and Y must be d-separated by adjusting for Z (see requirements for d-separation above).
2. No element of Z is a descendant of (i.e., caused by) X.

The backdoor criterion blocks all non-causal pathways (i.e., backdoor paths) between our predictor and response variable of interest, while leaving all causal paths open. Note that

while causality flows in the direction of DAG arrows, information can flow in both directions, inducing various sources of potential bias. As such, the application of backdoor criterion eliminates common statistical biases that can otherwise plague observational studies, including confounding, overcontrol, and collider bias. Collectively, given a DAG, the backdoor criterion guides model selection required to determine causal relationships without biased spurious correlations (Pearl 2009). Below we apply the backdoor criterion to guide covariate selection for determining the causal effect of each of our predictor variables. We note that only the effect size of the predictor variable should be interpreted for its associated model, as additional covariates are there as required controls.

MPA, depth and wave exposure models:

To determine the effect of MPA on regime shift trajectory, there are no backdoor paths that need to be blocked. Similarly, the effect of depth and wave exposure on regime shift also do not have any backdoor paths that need to be blocked. Therefore, no covariate adjustments are required for all three models.

Nutrient model:

To determine the effect of nutrient on regime shift, there are two backdoor paths that need to be blocked or d-separated:

1. Regime Shift ← Depth → Nutrients
2. Regime Shift ← Branching Coral ← Depth → Nutrients

Adjusting for depth will block both backdoor paths. This satisfies condition 1 of the backdoor criterion; condition 2 is also satisfied because depth is not a descendant of

nutrients (i.e., it is not affected by nutrients). Therefore, the effect of nutrients on regime shift can be quantified by adjusting for depth (i.e., adding depth as a covariate) in our statistical model.

Herbivore biomass model:

To determine the effect of herbivores fish biomass on regime shift, there are two backdoor paths that need to be blocked or d-separated:

1. Regime Shift ← Structural Complexity → Herbivore Biomass
2. Regime Shift ← Structural Complexity ← Branching Coral ← MPA → Herbivore Biomass
3. Regime Shift ← Branching Coral ← MPA → Herbivorous Fish Biomass

Adjusting for structural complexity (blocking path 1 and path 2) and branching coral (blocking path 3) will allow us to estimate the causal effect of herbivorous fish biomass on regime shift.

To determine the effect branching coral on regime shift, there are five backdoor paths that must be blocked or d-separated:

1. Regime Shift ← Wave Exposure → Branching Coral
2. Regime Shift ← Depth → Branching Coral
3. Regime Shift ← Nutrients ← Depth → Branching Coral
4. Regime Shift ← Herbivorous Fish Biomass ← MPA → Branching Coral

5. Regime Shift ← Structural Complexity → Herbivorous Fish Biomass ← MPA → Branching Coral

Adjusting for wave exposure (blocking backdoor path 1), depth (blocking backdoor paths 2 and 3), and MPA (blocking backdoor path 4) will allow us to estimate the causal effect of branching coral on regime shift. We note that backdoor path 5 is already blocked because herbivorous fish biomass – a collider variable – has *not* been adjusted for.

Structural complexity model:

To determine the effect of structural complexity on regime shift, there are six backdoor paths that must be blocked or d-separated, including:

1. Regime Shift ← Branching Coral → Structural Complexity
2. Regime Shift ← Wave → Branching Coral → Structural Complexity
3. Regime Shift ← Herbivorous Fish Biomass ← MPA → Branching Coral → Structural Complexity
4. Regime Shift ← Depth → Branching Coral → Structural Complexity
5. Regime Shift ← Nutrients → Initial Macroalgae ← Depth → Branching Coral → Structural Complexity
6. Regime Shift ← Initial Macroalgae ← Wave → Branching Coral → Structural Complexity

All backdoor paths between regime shift and structural complexity can be blocked or d-separated by adjusting for branching coral cover.

Initial macroalgae model:

Lastly, to determine the effect of initial macroalgae cover on regime shift, there are ten backdoor paths that must be blocked of d-separated:

1. Regime Shift ← Herbivorous Fish Biomass → Initial Macroalgae
2. Regime Shift ← Structural Complexity → Herbivorous Fish Biomass → Initial Macroalgae
3. Regime Shift ← Structural Complexity ← Branching Coral Cover ← MPA → Herbivorous Fish Biomass → Initial Macroalgae
4. Regime Shift ← Herbivorous Fish Biomass ← MPA → Branching Coral
5. Regime Shift ← Depth → Initial Macroalgae
6. Regime Shift ← Nutrients ← Depth → Initial Macroalgae
7. Regime Shift ← Branching Coral ← Depth → Initial Macroalgae
8. Regime Shift ← Structural Complexity ← Branching Coral ← Depth → Initial Macroalgae
9. Regime Shift ← Wave → Branching Coral ← Depth → Initial Macroalgae
10. Regime Shift ← Nutrients → Initial Macroalgae
11. Regime Shift ← Wave → Initial Macroalgae
12. Regime Shift ← Branching Coral ← Wave → Initial Macroalgae

Adjusting for herbivorous fish biomass will block or d-separate backdoor paths 1-4, adjusting for depth will block or d-separate backdoor paths 5-9, adjusting for nutrients will block backdoor path 10 and adjusting for wave will block backdoor paths 11 and 12.

3.6.5 Step 4: Statistical Model

The final step in SCM is to choose our statistical model, to estimate effect sizes. As DAGs are non-parametric, they make no assumptions about the distribution of variables (e.g., normal) or the functional form of effects (e.g., linear, nonlinear, stepwise), making them compatible with a wide range of statistical methods. Any non-linear relationships between variables must come from domain knowledge or have a sound theoretical justification to avoid chasing noise in the observed data. DAGs are also compatible with both frequentist and Bayesian statistical approaches since they are used to guide model selection, and not the analysis itself. We note that statistical models developed under the SCM framework are still beholden to the same issues of sample size and measurement error in terms of the precision of resulting estimates; however, they will be more causally accurate. Here, we have employed the same point estimation method used by our past correlative analysis (Graham et al. 2015), a Bayesian logistic regression.

Specifically, we applied a Bayesian logistic regression analysis to each of our causal models, where the response variable was 0 for recovering sites, and 1 for regime shifting sites. We standardized our data by subtracting the mean of each variable and dividing by 2 standard deviations in order to assess relative effect sizes of our predicter variables (Gelman and Hill, 2007). We ran our models using the 'rethinking' package on R, using weakly informative prior. For logistic regression using standardized data, Gelmen et al. (2008) recommend a Cauchy(0,10) prior for B_0 and a student_t(1,0,2.5) priors for $B_{1...N}$ as weakly informed priors. Recently, others have recommended a student_t(nu,0,2.5) for $B_{1...N}$ where nu is 3<nu<7 (Glosh et al. 2015). We have chosen nu of 4 (student_t(4,0,2.5)) but have compared our results using a nu of 1-7, which all gave similar results.

3.6.6 Causal vs. Correlative Models

Supplementary Table 3.2 Results from causal (backdoor criterion) vs correlative models (causal salad; no covariates).

Predictor Variable	**Backdoor Criterion**	**Causal Salad**	**No Covariates**
MPA	0.06 [-1.20, 1.39]	-1.57 [-5.02, 1.01]	NA
Nutrient	-1.61 [-3.65, 0.05]	-0.33 [-2.62, 1.81]	-1.30 [-3.44, 0.23]
Herbivore Biomass	1.76 [-0.44, 4.04]	0.08 [-1.75, 3.68]	0.90 [-0.65, 2.62]
Branching Coral	0.84 [-1.21, 3.09]	-0.04 [-2.68, 2.71]	1.41 [-0.04, 3.12]
Depth	-2.12 [-4.19, -0.49]	-4.76 [-9.67, -1.46]	NA
Structural Complexity	-3.77 [-7.33, -1.63]	-4.66 [-9.33, -1.77]	-3.81 [-7.06, 1.62]
Wave Exposure	1.02 [-0.45, 2.54]	1.83 [-0.89, 5.16]	NA
Initial Macroalgae	3.3 [0.40, 10.14]	1.44 [-1.67, 6.92]	4.44 [1.11, 10.24]

Supplementary Table 3.2 contrasts causal estimates of our predictor variables from our causal models (where we used our DAG and the application of the backdoor criterion to guide model selection) with a 'causal salad' model (where all variables were placed into one model and subsequently interpreted) and no covariate models (where additional variables were not adjusted for). "NA" is noted in instances where the backdoor models coincided with our no covariate model (i.e., no covariate adjustments were required by the backdoor criterion). The results show that causal estimates can vary greatly depending on model selection, highlighting the importance of applying DAGs to visualize causal structure and guide appropriate causal analysis.

3.6.7 Lack of Causal Inference Among Other Regime Shift Studies

We conducted a literature review on observational studies investigating factors that cause coral reef regime shifts. A systematic literature search was conducted using the Web of Science database, with the search terms “coral reef + (phase OR regime) shift”, which identified 1,262 results. We read through each abstract and selected studies that (1) used observational data and (2) focused on a form of coral reef phase shift (e.g., coral reef to macroalgae phase shift; recovery back to coral reef state). We excluded studies that looked at coral cover but did not consider this specifically under a phase shift context. The summary of our 14 selected studies is highlighted in Table S3.

The questions motivating each study are causal (e.g., factors influencing coral-macroalgal phase shifts, Graham et al, 2015), and causal language such as “effect”, “influence”, or “cause” is used in all but three studies (Supplementary Table 3.3). However, these studies have not employed any causal inference methods (Supplementary Table 3.3). We note that while we have emphasized the use of SCM, there are other approaches to causal inference that are available to researchers. These include quasi-experimental methods employed under the potential outcomes (PO) framework (Butsic et al. 2017; Larsen et al. 2019), as well as convergent cross mapping (CCM; Sugihara 2012) and granger causality (Detto et al. 2011), which can be used on time series data.

Our study stands apart from past observational work on coral reef regime shifts because we have employed the backdoor criterion to guide covariate selection. In doing so, we have explicitly stated the causal structure of our system via our DAG and have adjusted for specific variables to answer specific causal questions. In contrast, while the studies listed under Supplementary Table 3.3 may explain why a predictor variable is expected to influence the response, none (including our own) state their causal assumptions about how all relevant variables are causally connected. For example, Jouffray et al. 2019 used 20 anthropogenic and biophysical predictors thought to influence regime shifts without stating how variables may be causally connected. Without such consideration, issues of over-adjustment can lead to biased estimates between variables of interest due to processes such as overcontrol bias and collider bias (Pearl 2009; Elwert 2014). On the

other hand, some studies (Mumby and Harborne, 2010; Arias-Gonzalez et al., 2017) used only one predictor variable, without considering the need (or lack thereof) for covariate adjustment, which in turn can lead to confounding bias (Pearl 2009; Elwert 2014). Collectively, past observational studies on coral reef regime shifts have not employed causal inference methods and may therefore be prone to spurious and potentially misleading conclusions. In contrast, SCM offers a strong causal inference framework for coral reef ecologists interested in cause-and-effect relationships. It allows researchers to utilize their domain ecological knowledge to create DAGs and guide appropriate model selection for causal analysis.

Supplementary Table 3.3 Summary of observational studies that look at factors associated with coral reef regime shifts. * denotes when causal language was used in the title.

Paper	**Predictor Variables**	**Response Variable**	**Conclusion**	**Causal Language**	**Causal Inference**
Graham et al. 2015	Herbivorous biomass, juvenile coral density, depth, initial structural complexity, carbon nitrogen ratio of sampled algae	Recovery of coral reefs	Post-bleaching coral reef recovery favored when reefs were structurally complex, in deep water, with high density of juvenile corals and herbivorous fish and low nutrient loads	Affect; influence	No

Jouffray et al. 2015	Various; E.g. grazer biomass, caper biomass, brewer biomass, large predator biomass, latitude, depth, reef zone, etc.	Reef regimes (calcifying regime; turf regime; Macroalgal/sand regime)	Various; E.g. macroalgal regime decreased with increase in grazer, scrapper, and browser biomass	Drivers*; influence; effect, impact, explain	No
Cheal et al. 2010	herbivore diversity, scrapper/excavator abundance, grazer/dentritivore abundance, algal browser abundance	Phase 1 (low macroalgae), Phase 2 (transitional), Phase 3 (established macroalgae)	Low fish herbivore diversity and low abundance of algal browsers and grazers/detritivores correlated with coral reef recovery	None	No
Aronson and Precht 2000	Year	Urchin density, fish count, percent cover macroalgae, percent cover CTB, percent cover hard corals, percent cover Halimeda spp (separate ANOVA for each variable)	Macroalgae cover correlated with reduced urchin herbivory	None	No

Kumagai et al. 2018	Sea surface temperature, surface current, estimated herbivorous fishes occurrence	Macroalgae-coral shift	Ocean current and herbivory drive macroalgae-to-coral shift under climate warming	Drive*, effect, influence, promote	No
Jouffray et al. 2019	20 anthropogenic and biophysical variables (e.g. effluent, habitat modification, non-commercial boat fishing)	Occurrence of 4 distinct regime shifts	Various: E.g. Regime 2 was best predicted by a strong positive relationship with maximum monthly climatological mean of wave power	Drivers*, effect, influence, impact	No
Johns et al. 2018	Habitat	Coral recovery (coral recruitment, recruit survival, juvenile coral persistence)	Negative impact of macroalgae on recruitment and recruit survival of corals maintain macroalgae state	Maintain*, effect, increase, influence, impact, create	No
Mumby and Harborne 2010	Marine reserve status	Rate of change of coral cover	Marine reserves enhance the recovery of corals	Enhance, effect, affect, higher, lower, increase	No

Cruz et al. 2018	Anthropogenic and environmental factors	Zoanthid and macroalgal phase shift indices (PSI)	Marginal reefs: correlation between coral-zoanthid shifts and local human impacts; coral-algal shifts and ports, urbanized surfaces, higher latitudes and shore proximity.	Effect; increase, decrease,	No
Bennett et al. 2015	Year	Fish trophic biomass, cover of benthic functional groups (separate ANOVA for each response variable)	Tropical herbivores maintain coral reef state	Maintain, increase, reduce, cause, impact, drive, lower	No
Ledlie et al. 2007	Year and monitoring site	Percent coral cover, percent macroalgae, changes in abundance and biomass of herbivorous fish (separate ANOVA for	Absence of macroalgal consumers and current dominance of macroalgae reduces a shift back to coral-dominated state	None	No

		each response variable)			
Arias-Gonzalez et al. 2017	Year (associated with coral-algal phase shift)	Length, abundance, biomass for herbivorous fish	Coral-algal phase shift not driven by herbivorous fish	Driven*, effect, affect, precipitate	No
Huges, 1994	NA: interpretation based on graphs showing changes in coral cover, human population, number of Diadema, etc. with time		Overfishing, hurricane damage and disease have combined to cause coral-macroalgal shift	Cause; effect;	No
Lapointe 1997	NA: nutrient-related variables were recorded for sites experiencing phase shifts, which exceeded nutrient thresholds noted to sustain macroalgal blooms		Nutrient enrichment is a causal factor for coral-macroalgae phase shift in Jamaica and southeast Florida	Causal factor	No

3.6.8 Literature Cited

Ankan, A., Wortel, I., Textor, J. 2021. Testing graphical causal models using the R package “dagitty”. *Current Protocols*, 1, e45.

Arias-Gonzalez, J., Fung, T., Seymour, R., Garza-Perez, J., Acosta-Gonzalez, G., Bozec Y. et al. 2017. A coral-algal phase shift in Mesoamerica not driven by changes in herbivorous fish abundance. *PLoS ONE,* 12(4), e0174855.

Aronson, R., Precht, W. 2000. Herbivory and algal dynamics on the coral reef at Discovery Bay, Jamaica. *Limnology and Oceanography,* 45(1), 251-255.

Bellwood, D., Hughes, T., Folke, C., Nystrom, M. 2004. Confronting the coral reef crisis. *Nature,* 429, 827-833.

Bennett, S., Wernberg, T., Harvey, E., Santana-Garcon, J., Saunders, B. 2015. Tropical herbivores provide resilience to a climate-mediated phase shift on temperate reefs. *Ecology Letters,* 18, 714-723.

Bridge, T., Hughes, T., Guinotte, J., Bongaerts, P. 2013. Call to protect all coral reefs. *Nature Climate Change,* 3, 528-530.

Burkepile, S., Hay, M. 2006. Herbivore vs. nutrient control of marine primary producers: context-dependent effects. *Ecology,* 87, 3128-3139.

Butsic, C., Lewis, D., Radeloff, V., Baumann, M., Kuemmerle, T. 2017. Quasi-experimental methods enable stronger inferences from observational data in ecology. *Basic and Applied Ecology*, 19, 1-10.

Capdevila, P., Linares, C., Aspillaga, E., Riera, J., Hereu, B. 2018. Effective dispersal and density-dependence in mesophotic macroalgal forests: Insights from the Mediterranean species *Cystoseira zosteroides*. *PLoS ONE*, 13(1), e0191346.

Chabanet, P., Ralambondrainy, H., Amanieu, M., Faure, G., Galzin, R. 1997. Relationships between coral reef substrata and fish. *Coral Reefs*, 16, 93-102.

Cheal, A., Aaron, M., Cripps, E., Emslie, M., Jonker, M., Schaffelke, B., Sweatman, H. 2010. Coral-macroalgal phase shifts or reef resilience: links with diversity and functional roles of herbivorous fishes of the Great Barrier Reef. *Coral Reefs*, 29, 1005-1015.

Colgan, M. 1987. Coral reef recovery on Guam (Micronesia) after catastrophic predation by *Acanthaster planci*. *Ecology*, 68, 1592-1605.

Connell, J., Hughes, T., Wallace, C. 1997. A 30-year study of coral abundance, recruitment, and disturbance at several scales in space and time. *Ecological Monographs*, 67, 461-488.

Cros, A., McClanahan, T. 2003. Coral transplant damage under various management conditions in the Mombasa Marine National Park, Kenya. *Western Indian Ocean Journal Marine Science*, 2, 127-136.

Cruz, I., Waters, L., Kikuchi, R., Leao, Z., Turra, A. 2018. Marginal coral reefs show high susceptibility to phase shift. *Marine Pollution Bulletin,* 135, 551-561.

Detto, M., Molini, A., Katul, G., Stoy, P., Palmroth, S. & Baldocchi, D. 2011. Causality and persistence in ecological systems: A nonparametric spectral granger causality approach. *The American Naturalist*, 179, 000-000.

Done, T. 1983. Coral zonation, its nature and significance, in Perpectives on Coral Reefs. *Australian Institute of Marine Science*, 107-147.

Edwards, C., Friedlander, A., Green, A., Hardt, M., Sala, E., Sweatman, H. et al. 2014. Global assessment of the status of coral reef herbivorous fishes: evidence for fishing effects. *Proceedings of the Royal Society B,* 281, 20131835.

Elwert, F. 2013. Graphical causal models. *Handbook of causal analysis for social research*, 245-273.

Evans, R., Wilson, S., Fisher, R., Ryan, N., Babcock, R., Blakeway, D. et al. 2020. Early recovery dynamics of turbid coral reefs after recurring bleaching events. *Journal of Environmental Management*, 268, 110666.

Gelman, A., Jakulin, A., Grazia, P., Su, Y. 2008. A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics,* 2(4), 1360-1383.

Graham, N., Nash, K. 2013. The importance of structural complexity in coral reef ecosystems. *Coral Reefs*, 32, 315-326.

Graham, N., Jennings, S., MacNeil, A., Mouillot, D., Wilson, S. 2015. Predicting climate-driven regime shifts versus rebound potential in coral reefs. *Nature*, 518, 94-97.

Gilmour, J.P., Smith, L.D., Heyward, A.J., Baird, A.H., Pratchett, M. S. 2013. Recovery of an isolated coral reef system following severe disturbance. *Science*, 340, 69-71.

Glosh, J., Yingbo, L, Mitra, R. 2015. On the use of Cauchy prior distributions for Bayesian logistic regression. Accessed at: https://arxiv.org/abs/1507.07170

Gove, J., Williams, G., McManus, M., Clark, S., Ehses, J., Wedding, L. 2015. Coral reef benthic regimes exhibit non-linear threshold responses to natural physical drivers. *Marine Ecological Progress Series*, 522, 33-48.

Guzman, H., Cortés, J. 2007. Reef recovery 20 years after the 1982–1983 El Nino massive mortality. *Marine Biology*, 151, 401–411.

Harrison, P., Wallace, C. 1990. Reproduction, dispersal and recruitment of Scleractinian corals. In: Dubinsky Z (ed) Ecosystems of the world: coral reefs. Elsevier, Amsterdam, pp 133-207.

Huges, T. 1994. Catastrophes, phase shifts, and large-scale degradation of a Caribbean coral reef. *Science*, 265(5178), 1547-1551.

Hughes, T., Baird, A., Bellwood, D., Card, M., Connolly, S., Folke, C. et al. 2003. Climate change, human impacts, and the resilience of coral reefs. *Science*, 301, 929-933.

Hughes, T., Rodrigues, M., Bellwood, D., Ceccarelli, D., Hoegh-Guldberg, O., McCook, L. et al. 2007. Phase shifts, herbivory, and resilience of coral reefs to climate change. Current Biology, 17(4), 360-365.

Hughes, T., Graham, N., Jackson, J., Mumby, P., Steneck, R. 2010. Rising to the challenge of sustaining coral reef resilience. *Trends in Ecology and Evolution*, 25, 633-642.

Hurd, C. 2000. Water motion, marine macroalgal physiology, and production. *Journal of Phycology*, 36, 453-472.

Jennings S., Marshall S., Polunin N. 1996. Seychelles' marine protected areas: comparative structure and status of reef fish communities. *Biological Conservation*, 75, 201-209.

Johns, K., Emslie, M., Hoey, A., Osborne, K., Jonker, M. & Cheal, A. 2018. Macroalgal feedbacks and substrate properties maintain a coral reef regime shift. *Ecosphere*, 9(7), e02349.

Jouffray, J., Wedding, L., Norstrom, A., Donovan, M., Wiliams, G., Crowder, L. et al. 2019. Parsing human and biophysical drivers of coral reef regimes. *Proceedings of the Royal Society B*, 286, 20182544.

Jouffray, J., Nystrom, M., Norstrom, A., Williams, I., Wedding, L., Kittinger, J.,

Williams, G. 2015. Identifying multiple coral reef regimes and their drivers across the Hawaiian archipelago. *Philosophical Transactions of the Royal Society B*, 370, 20130268.

Koop, K., Booth, D., Broadbent, A., Brodie, J., Bucher, D., Capne, D. et al. 2001. ENCORE: the effect of nutrient enrichment on coral reefs. Synthesis of results and conclusions. *Marine Pollution Bulletin*, 42, 91-120.

Kendrick, G. & Walker, D. 1991. Dispersal distances for propagules of *Sargassum spinuligerum* (Sargassaceae, Phaeophyta) measured directly by vital staining and venturi suction sampling. *Marine Ecology Progress Series*, 79, 133-138.

Kendrick, G., Walker, D. 1995. Dispersal of propagules of *Sargassum* spp. (Sargassaceae: Phaeophyta): Observations of local patterns of dispersal and consequences for recruitment and population structure. *Journal of Experimental Marine Biology and Ecology*, 192, 273-288.

Kumagai, N., Molinos, J., Yamano, H., Takao, S., Fujii, M., Yamanaka, Y. 2018. Ocean currents and herbivory drive macroalgae-to-coral community shift under climate warming. *PNAS*, 115(36), 8990-8995.

Lapointe, B. 1997. Nutrient thresholds for bottom-up control of macroalgal blooms on coral reefs in Jamaica and southeast Florida. Limnology and *Oceanography,* 42, 1119-1131.

Larned, S., Atkinson, M. 1997. Effects of water velocity on NH4 and PO4 uptake and nutrient-limited growth in the macroalgae Dictyosphaeria cavernosa. *Marine Ecological Progress Series*, 157, 295-302.

Larsen, A., Meng, K., Kendall, B. 2019. Causal analysis in control-impact ecological studies with observational data. *Methods in Ecology and Evolution*, 10, 924-934.

Ledlie, M., Graham, N., Bythell, J., Wilson, S., Jennings, S., Polunin, N., Hardcastle, J. 2007. Phase shifts and the role of herbivory in the resilience of coral reefs. *Coral Reefs*, 26, 641-653.

Loya, Y., Sakai, K., Yamazato, K., Nakano, Y., Sambali, H., van Woesik, R. 2001. Coral bleaching: the winners and losers. *Ecology Letters*, 4, 122-131.

Madin, J., Connolly, S. 2006. Ecological consequences of major hydrodynamic disturbances on coral reefs. *Nature*, 444, 477-480.

McCook, L. 1999. Macroalagae, nutrients and phase shifts on coral reefs: scientific issues and management consequences for the Great Barrier Reef. *Coral Reefs*, 18, 357-367.

McManus, J., Rodolfo, B., Reyes, J., Cleto, L., Nanola, J. 1997. Effects of some destructive fishing methods on coral cover and potential rates of recovery. *Environmental Management*, 21, 69–78.

Mumby, P., Dahlgren, C., Harborne, A., Kappel, C., Micheli, F., Brumbaugh, D. et al. 2006. Fishing, trophic cascades, and the process of grazing on coral reefs. *Science*, 311, 98-101.

Mumby, P., Harborne, A. 2010. Marine Reserves Enhance the Recovery of Corals on Caribbean Reefs. *PLoS ONE*, 5(1), e8657.

Mumby, P., Bejarano, S., Golbuu, R., Steneck, R., Arnold, S., van Woesik, R., Friedlander, A. 2013. Empirical relationships among resilience indicators on Micronesian reefs. *Coral Reefs*, 32, 213-226.

Nakamura, T., van Woesik, R. 2001. Water-flow rates and passive diffusion partially explain differential survival of corals during the 1998 bleaching event. *Marine Ecology*

Progress Series, 212, 301-304.

Nakamura, T., Yamasaki, H., van Woesik, R. 2003. Water flow facilitates recovery from bleaching in the coral Stylophora pistillata. *Marine Ecology Progress Series*, 256, 287-291.

Pearl, J. 1988. Probabilistic reasoning in intelligent systems, San Mateo: Morgan Kaufman.

Pearl, J. 1993. Comment: Graphical models, causality, and interventions. *Statistical Science*, 8(3), 266-269.

Pearl, J. 2009. Causality: Models, Reasoning and Inference. 2nd ed. Cambridge, UK: Cambridge University Press.

Pearl, J. 2010. An introduction to causal inference. *The International Journal of Biostatistics*, 6, 1-59.

Schaffelke, B., Klumpp, D. 1997. Growth of germlings of the macroalga Sargassum baccularia (Phaeophyta) is stimulated by enhanced nutrients. *Proceedings of the 8th International 8th International Coral Reef Symposium*, 2, 1839-1842.

Sheppard, C., Spalding, M., Bradshaw, C., Wilson, S. 2002. Erosion vs recovery of coral reefs after the 1998 El Nino: Chagos reefs, Indian Ocean. *Ambio,* 31, 40-48.

Slattery, M., Lesser, M., Brazeau, D., Stokes, M., Leichter, J. 2011. Connectivity and stability of mesophotic coral reefs. *Journal of Experimental Marine Biology and Ecology*, 408(1-2), 32-41.

Soler, G., Edgar, G., Thomson, R., Kininmonth, S., Campbell, S., Dawson, T. et al. 2015. Reef fishes at all trophic levels respond positively to effective marine protected areas.

PLoS ONE, 10(10), e0140270.

Stevens, B. 2020. The ups and downs of traps: environmental impacts, entanglement, mitigation, and the future of trap fishing for crustaceans and fish. *ICES Journal of Marine Science*, fsaa135.

Strain, E., Edgar, G., Ceccarelli, D., Stuart-Smith, R., Hosack, G., Thomson, R. 2019. A global assessment of the direct and indirect benefits of marine protected areas for coral reef conservation. *Diversity and Distributions*, 25, 9-20.

Sugihara, G., May, R., Ye, H., Hsieh, C., Deyle, E., Fogarty, M., Munch, S. 2012. Detecting causality in complex ecosystems. *Science*, 338(6106), 496-500.

Tamir, R., Eyal, G., Kramer, N., Laverick, J., Loya, Y. 2019. Light environment drives the shallow-to-mesophotic coral community transition. *Ecosphere*, 10(9), e02839.

Textor, J., Zander, B., Gilthorpe, M., Liskiewicz, M. & Ellison, G. 2016. Robust causal inference using directed acyclic graphs: the R package 'dagitty'. *International Journal of Epidemiology*, 45(6), 1887-1894.

Verges, A., Vanderklift, M., Doropoulos, C. & Hyndes, G. 2011. Spatial patterns in herbivory on a coral reef are influenced by structural complexity but not by algal traits. *PLoS ONE*, 6, e17115.

Victor, S. 2008. Stability of reef framework and post settlement mortality as the structuring factor for recovery of Malakal Bay Reef, Palau, Micronesia: 25 years after a severe COTS outbreak. *Estuarine, Coastal and Shelf Science*, 77, 175-180.

Ward, S., Harrison, P. 1997. The effect of elevated nutrients on settlement of coral larvae during the ENCORE experiment, Great Barrier Reef. *Proceedings of the 8th International Coral Reef Symposium*, 1, 891-896.

Wilson, S., Robinson, J., Chong-Seng, K., Robinson, J., Graham, N. 2019. Boom and bust of keystone structure on coral reefs. *Coral Reefs*, 38, 625-635.

Wooldridge, S. 2009. Water quality and coral bleaching thresholds: Formalising the linkage for the inshore reefs of the Great Barrier Reef, Australia. *Marine Pollution Bulletin*, 58(5), 745-751.

CHAPTER 4 GLOBAL DRIVERS OF REEF FISH BIOMASS

Authors: Suchinta Arif, Matthew McLean, Aaron MacNeil

4.1 ABSTRACT

Reef fish play a key role in maintaining ecosystem health, providing key ecosystem services that are heavily relied upon for both fisheries and livelihoods. Although attempts have been made to answer how various factors effect global distribution of reef fish biomass, the methods used have not been appropriate for drawing causal conclusions. Here, we apply the Structural Causal Model (SCM) framework to answer how key socio-environmental factors influence reef fish biomass across both tropical and temperate reefs. This causal framework draws on directed acyclic graphs (DAGs) and a set of graphical rules to guide statistical adjustments required to answer specific causal queries from observational data. Our observational data comes from the Reef Life Survey (RLS), providing data from 4,357 transects and 1,844 sites worldwide. Following the application of the SCM framework, our results confirm that, across both temperate and tropical reefs, mean sea surface temperature (SST), fish diversity, marine protected areas (MPAs), and depth have positive effects on reef fish biomass, whereas degree heating weeks (DHW) and market gravity act negatively. Our results provide the first example of the SCM framework being applied on a large-scale ecological dataset and is one of the few reef studies to carry out a causal analysis on observational data. Beyond the reef-related insights gained from our analysis, our work provides a case study for applying the SCM framework across future reef studies that are reliant on largescale ecological data.

4.2 INTRODUCTION

Reef ecologists are often interested in understanding causal relationships in nature. Recent examples include aiming to quantify the impact of climate change (e.g., McClanahan et al. 2020; Hamilton et al. 2022), other anthropogenic stressors (e.g., Duprey et al. 2016; Mbaru et al. 2020), and management interventions (e.g., Mellin et al.

2016; Campbell et al. 2020) on reef ecosystems. Such questions often rely on observational data that lie out of experimental control. Global observational datasets are increasingly available through citizen science, online initiatives, and technological advances including remote-sensing, and animal-borne sensors (Sagarin and Pauchard, 2010) yet ecologists frequently lack methods to formally address causal questions.

A fundamental causal question in reef ecology remains, what are the global drivers of reef fish biomass? Reef fish play key functional roles in reef ecosystems (Bellwood et al. 2004) and are heavily relied upon for both fisheries (Grafeld et al. 2017) and ecosystem services such as recreation and tourism (Brander et al. 2007). As such, understanding the socio-environmental factors that influence reef fish biomass is important for both marine conservation and human development goals. As a starting point, previous studies have shed light on how different factors may be influencing reef fish biomass at a global scale. For example, Cinner et al. (2016) aimed to understand how 18 socioeconomic and environmental drivers impacted reef fish biomass, finding that high compliance reserves and local population growth showed a positive relationship with reef fish biomass, whereas market gravity showed a negative relationship (Cinner et al. 2016). In another global study, Duffy et al. (2016) found that temperature and biodiversity were among the strongest predictors of reef fish biomass, followed by human impacts. Several other studies at both global and localized scales have noted that reef fish biomass has a positive relationship with structural complexity, and a negative relationship with human population (e.g., Newton et al. 2007; Graham et al. 2006; Cinner et al. 2009).

Although previous observational studies have paved the way to better understand the global drivers of reef fish biomass, important limitations remain. Mainly, observational studies have not yet employed causal inference methods to understand the causal drivers of reef fish biomass. While several observational causal inference frameworks exist and have been highlighted in the ecology literature (e.g., Larsen et al. 2019; Laubach et al. 2020), these techniques have largely gone unrecognized within reef ecology. However, applying causal inference methods to observational data is essential for avoiding

statistical biases that can otherwise lead to non-causal associations between variables of interest, undermining results drawn from observational analysis.

Here we apply a recently emerging and widely applicable causal inference framework, the Structural Causal Model (SCM; Pearl 2009), to determine the global drivers of reef fish biomass in both tropical and temperate reefs. The SCM framework (Pearl 2009) combines the features of structural equation modelling (SEM; Wright 1921) and the potential outcome framework (Rubin 1974; Holland 1986) with other theories of causation to provide a framework for observational causal inference. SCM uses directed acyclic graphs (DAGs) to visualize the causal structure of a system under study and can subsequently apply a graphical rule known as the backdoor criterion to determine covariate selection required for specified causal questions (see Methods for detail). Thus far, a handful of ecological studies have employed this framework to understand varied causal relationships (Cronin and Schoolmaster Jr. 2018; Schoolmaster Jr. et al. 2020; Wilson et al. 2021). Chapter 3 also applied the SCM framework to examine how different factors affect the likelihood of coral-algal regime shifts following a climate-induced bleaching event in Seychelles and found several factors that were missed by a previously correlative study (Graham et al. 2015) using the same dataset.

Applying the SCM framework to determine the global drivers of reef fish biomass is important for two reasons. First, the application of a valid causal inference methods allows us to leverage global datasets to deepen our causal understanding of the controls of reef fish biomass, while limiting the statistical biases that can otherwise plague observational analysis. Second, the present study provides the first global ecological study that applies the SCM framework to draw causal conclusions from observational data. It can therefore be used as a template for ecologists aiming to answer causal questions from largescale ecological datasets.

4.3 METHODS

4.3.1 Observational Data

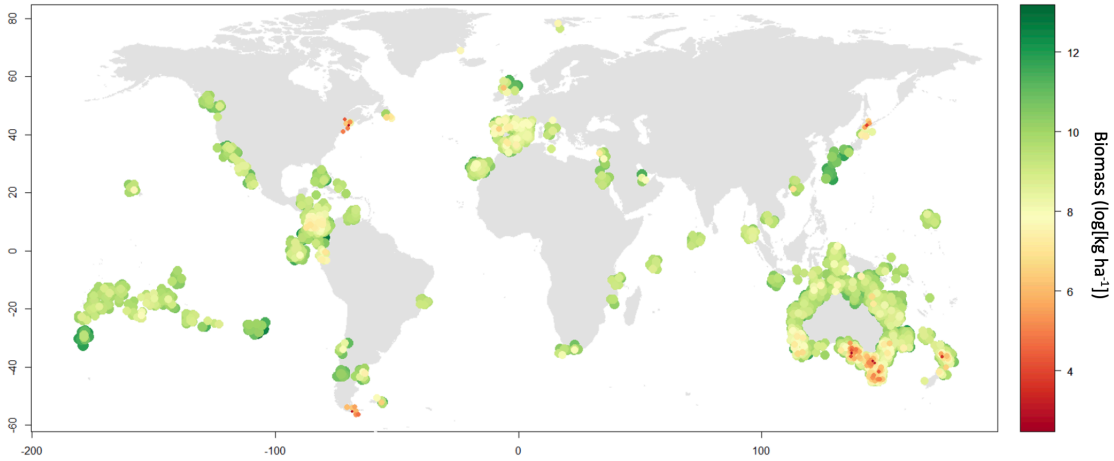

Figure 4.1 Reef fish biomass among 1,844 reef sites. Point colours vary in proportion to the average amount of fish biomass.

Observational data were obtained from the Reef Life Survey (RLS), which uses standardized visual censuses along 50-m transects, that has been proven effective for broad scale studies (Edgar and Stuart-Smith 2009). Details of fish census methods, diver training, and data quality can be found in Reef Life Survey Foundation (2019). Data were collected at 4,357 transects and 1,844 sites (mean 2.4 transects per site), across both temperate and tropical reefs from 2004 and 2013 (Figure 4.1). Fish counts per transect and size estimates were converted to biomass estimates using species-specific length-weight relationships from FishBase (www.fishbase.org). Fish diversity was obtained from the RLS data, which provided fish species richness at each transect. Depth and latitude were also recorded for each transect. At each site, management was categorized as either fished (regularly fished without effective restrictions), a no-take MPA, or a restricted MPA (active restrictions on gears or fishing effort). The Human Development Index (HDI) was obtained from the United Nations Development Program at the country level, representing a summary measure of human development encompassing: a long and healthy life, being knowledgeable, and having a decent standard of living. Human gravity represents human interactions within a reef and is determined by the population of a place divided by the squared time it takes to travel to the reef (detailed in Cinner et al. 2018); we calculated human gravity values for each transect. Data for SST mean and SST range

were obtained from the National Oceanic and Atmospheric Administration Coral Reef Watch (Lui et al. 2014). Net primary productivity data were derived from the standard Vertically Generalized Production Model (Behrenfeld et al. 1997).

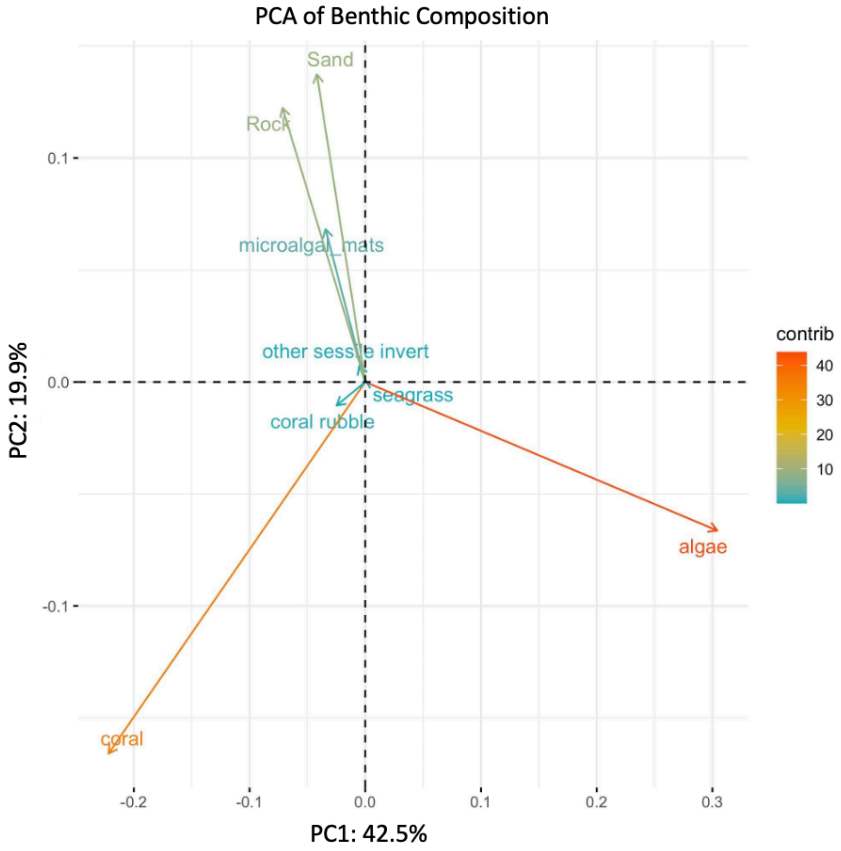

Figure 4.2 PCA of benthic composition data on tropical reefs, used to reduce the dimensionality of benthic composition. The first axes explaining 42.5% of the variation represents the amount of coral vs algae, whereas the second axes explaining 19.9% of the variation represents the amount of sand or rock vs coral or algae. These two axes are used as two separate predictor variables representing benthic composition.

Benthic composition data were derived from photo quadrat data collected at tropical reefs at the transect level (detailed in Reef Life Survey Foundation 2019). Benthic organisms were grouped into 8 categories: coral, algae, seagrass, coral rubble, rock, sand, microalgae mats, and other sessile invertebrates. For each category, percentage cover was recorded in proportions from 0-1 (e.g., 50% coral cover would equal 0.5). The data were then arc-sine transformed and a PCA was applied to reduce the dimensionality of benthic composition (Figure 4.2). The first two axes were extracted and used as predictor

variables. Because benthic data were missing for ~35% of surveys, the missing PCA axes values were imputed using a random forest approach with the R package missForest. Imputation was performed with site longitude and latitude as covariates to improve imputation accuracy. We cross validated the accuracy of imputation by randomly deleting 35% of data from sites where data existed, running the imputation, and then testing the Pearson correlation between imputed values and true values, which resulted in a value of 0.78

4.3.2 The SCM Framework

The SCM framework consists of 4 key steps (detailed in Chapter 2). First, researchers must create a DAG to represent the causal structure of their system under study (Step 1). Researchers can then refine their DAG based on DAG-data consistency (Step 2). Once a DAG is finalized, a graphical rule known as the backdoor criterion can be used to determine covariate selection required for answering specific causal queries (Step 3). Following this, a statistical model can be built to determine causal effects from observational data (Step 4). We describe these four steps below, placing an emphasis on how to deal with largescale ecological datasets.

Step 1: Creating a DAG

DAGs are used to represent a researchers' assumptions about the causal structure of a system or process under study. A DAG consists of a set of nodes (variables) that are connected to each other by edges (arrows). These arrows represent hypothesized causal relationships between variables, pointing from cause to effect, with causes preceding their effects. DAGs are also acyclic, meaning that they cannot contain bi-directional relationships (i.e., arrows need to be unidirectional) or a feedback loop where a variable causes itself (Glymour and Greenland 2008); however, see Greenland et al. (1999) for how to overcome this limitation by more finely articulating the temporal sequence of events. Importantly, DAGs must include all variables (both measured and unobserved) required to depict a system or process under study.

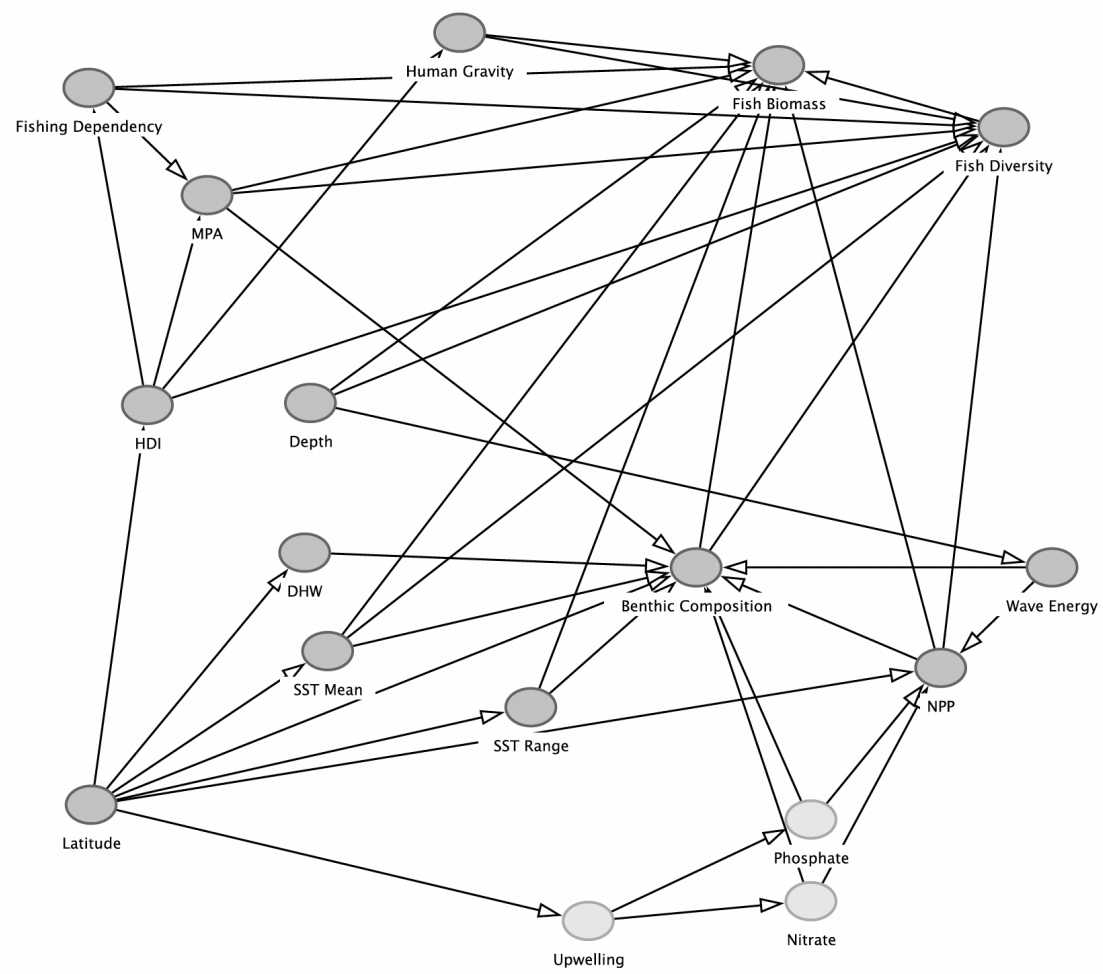

Figure 4.3 An initial DAG representing how socio-environmental drivers are hypothesized to effect reef fish biomass. Observed variables are dark grey and unobserved variables are light grey.

DAGs should be built and justified based on domain knowledge, such as past studies, expert opinion, and scientific consensus. It is critical to recognize that causality exists outside of any specific analytical method, and a DAG is one way to represent these relationships (Pearl 2009). To build our reef biomass DAG, we gathered a group of reef researchers to build an initial DAG (Figure 4.3) that included relevant environmental and socio-economic variables, their causal pathways to reef fish biomass, as well as their causal relationship with one another. We further included common cause variables, defined as variables that effect two or more variables already included in a DAG (Sprites et al. 2001; Glymour and Greenland 2008). This initial DAG was created during a week-

long workshop by the authors of the paper, as well as from reef researchers Dr. Eva Maire, Dr. James Robinson, Dr. Nicolas Loiseau, Dr. Nick Graham, Dr. Rick Stewart-Smith, and Dr. Graham Edgar.

As the aim of our research was to determine global drivers of reef fish biomass, causal relationships between variables were depicted at this scale. For example, we included latitude, and human development index (HDI) as variables in our DAG, which would not need to be included if dealing with a smaller scale study. As well, because our global dataset covered tropical and temperate shallow reefs (<50 m deep) only, we excluded causal links between depth and other variables that would exist if our dataset was inclusive of deeper reefs as well. Ultimately, the causal structure of a system may change depending on the scale and scope of the study, and researchers should be mindful of this while creating their DAG.

Step 2: Testing DAG-data Consistency

We tested our initial DAG for DAG-data consistency. Simply put, a DAG often implies many independencies (e.g., X is independent of Y) and conditional independencies (e.g., X is independent of Y, given Z) that should be consistent with the observational data, given that both the observational data and DAG are representative of the data-generating process. If DAG-data consistency is ensured, this provides support for the asserted structure of the DAG itself. On the other hand, if failed independencies exist, this can indicate potential problems. By assessing these failed independencies, researchers can update their DAG to better represent the data-generating process.

Schoolmaster et al. (2020) provide a real-world ecological example of refining a DAG based on DAG-data consistency. Here, their DAG was updated to ensure complete DAG-data consistency. However, DAG-data consistency can only be fully achieved if both the DAG and the observational dataset are perfectly representative of the data-generating process. Particularly when dealing with large-scale studies, it is important to recognize that observational datasets may not be perfectly representative of the data-generating

process (e.g., due to measurement error or coarse measurement of variables), which may also lead to failed independencies. Therefore, while failed independencies can be used to update a DAG, researchers should ensure that all updated causal links are first and foremost justified by theory, domain knowledge, and scientific ideas, and not solely to satisfy DAG-data consistency.

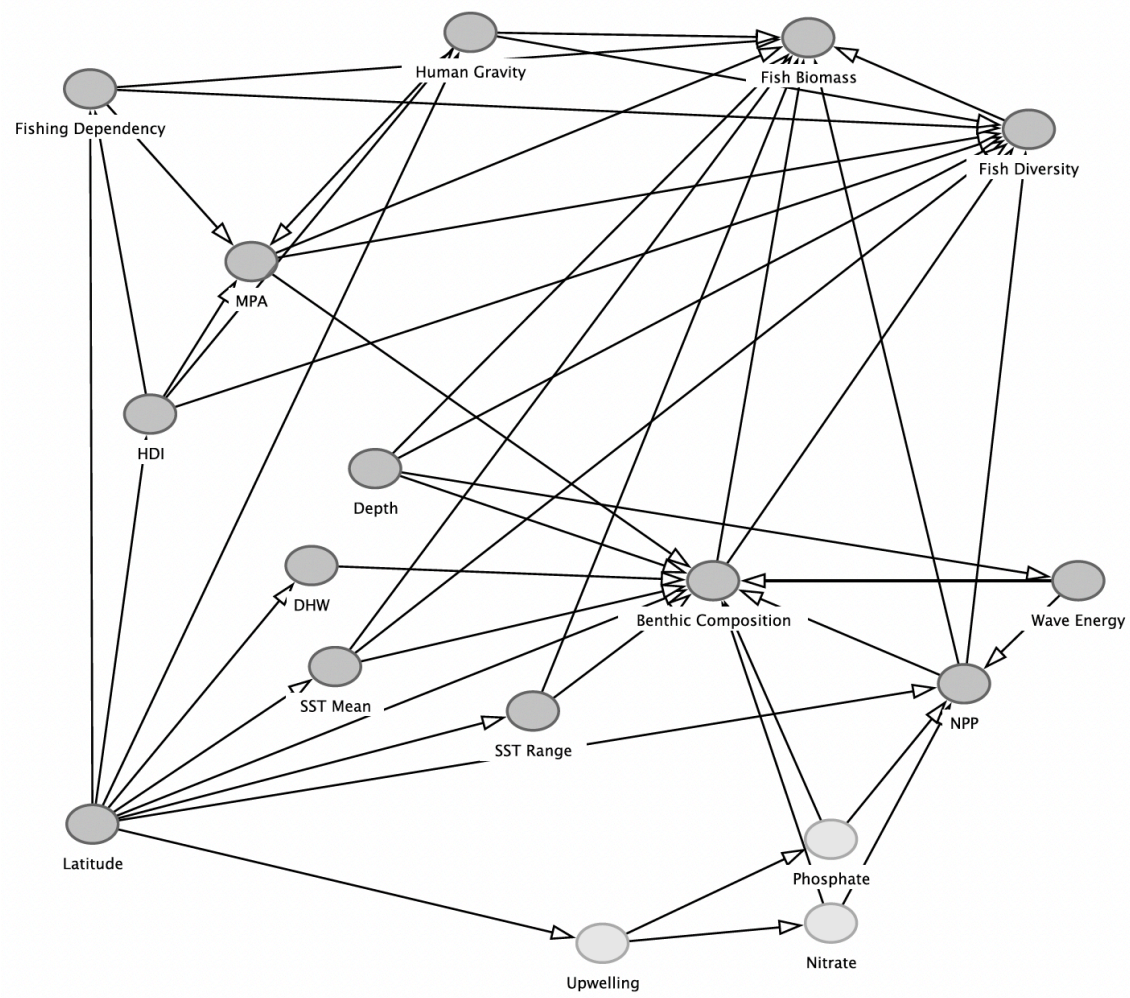

Figure 4.4 A finalized DAG representing how socio-environmental drivers are hypothesized to effect reef fish biomass. Observed variables are dark grey and unobserved variables are light grey.

To test DAG-data consistency, we used the R package 'dagitty' (Textor et al. 2016), which includes a formal test of zero correlation to signify whether independencies of a given DAG are consistent with a given dataset. As noted by Ankan et al. (2021), large

datasets can result in failed independencies between two variables even when these associations are negligible, since associations between two variables are rarely exactly zero. Therefore, we considered failed independencies when they were equal to or above an effect size of 0.2. There were 87 independencies implied by our initial DAG, of which 15 did not pass DAG-data consistency. After reviewing the overall failed independencies, we updated our DAG with four additional causal links: Latitude → Human Gravity, Latitude → Fishing Dependency, Human Gravity → MPA and Depth → Benthic Composition (Figure 4.4). This updated DAG had 6 failed independencies out of 65 total implied independencies. Although we used DAG-data consistency to update our DAG, ultimately, it was theory and domain knowledge that guided its creation. We expect some failed independencies to remain as our global observational dataset is not perfectly representative of the data-generating process. Our final DAG was approved by all researchers involved in its creation.

Step 3: Covariate Selection Based on the Backdoor Criterion

Once a DAG has been finalized, the backdoor criterion can be applied to determine the adjustment set(s) required to determine the effect of X on Y (Pearl 2009). The backdoor criterion is a set of graphical rules that instructs us to block all *backdoor paths*. Backdoor paths are sequences of nodes and arrows in a DAG with an arrow pointing into the response variable of interest; if left open, they can induce non-causal associations between variables of interest. To block a backdoor path, we can either (1) adjust for an intermediate arrow-emitting variable or (2) not adjust for a variable with two incoming arrows (i.e., a collider variable: → X ←). Once all backdoor paths for a given predictor and response variable are blocked, causal estimates can be determined from observational data without being biased by common statistical adjustments such as confounding, overcontrol or collider bias (see Chapter 2.8.2 for definitions). A detailed breakdown of the backdoor criterion is provided in Chapter 2, and the mathematical underpinnings is provided in Pearl (2009).

When dealing with large complex DAGs, however, applying the backdoor criterion become time consuming and complicated. The website www.dagitty.net (instructions within site) can be used to draw out a DAG and specify the predictor and response variable of interest, which in turn will generate the backdoor adjustment set(s). Our DAG has been saved on www.dagitty.net/mfcM6WU. We apply the backdoor criterion for each of our causal queries, creating a separate model for each of our predictor variables of interest. When more than one adjustment set is available, we selected the adjustment set with the least expected measurement error to increase the accuracy of our causal estimates.

Step 4: Causal Models

Once the backdoor criterion is applied to determine covariate selection, researchers select the best statistical model that represents the question and data at hand. Here, we have chosen a Bayesian generalized linear-mixed model, were we set reef sites nested within country as a random effect to account for the hierarchical structure of the data (MacNeil et al. 2009; Cinner et al. 2018). We created one model per predictor variable of interest, where fish biomass was the response, and additional covariates (used as controls) selected based on the application of the backdoor criterion (Table 4.1). We interpreted the standardized effect size of the predictor variable of interest for each of our causal models. We analyzed temperate and tropical reefs separately to account for the possibility of different effect sizes. All analysis was conducted using the R package 'rstanarm' using default weakly informative priors (Goodrich et al. 2022).

Table 4.1 Covariate selection for each predictor variable of interest based on the application of the backdoor criterion. For each model, reef fish biomass was set as the response, covariates were used as controls, and the coefficient of the predictor variable of interest was interpreted for their causal effects.

Predictor Variable	Control Variables
Sea Surface Temperate (SST)	Latitude

Fish Diversity	Depth, MPA, Net Primary Productivity (NPP), Fishing Dependency, Human Gravity, SST
Fishing Dependency	HDI, Latitude
MPA	HDI, Fishing Dependency, Human Gravity
Human Development Index (HDI)	Latitude
Depth	None required
Human Gravity	HDI, Latitude
Benthic Composition (coral:algae)	Depth, MPA, NPP, Latitude, SST
Benthic Composition (coral or algae:rock or sand)	Depth, MPA, NPP, Latitude, SST

4.4 Results and Discussion

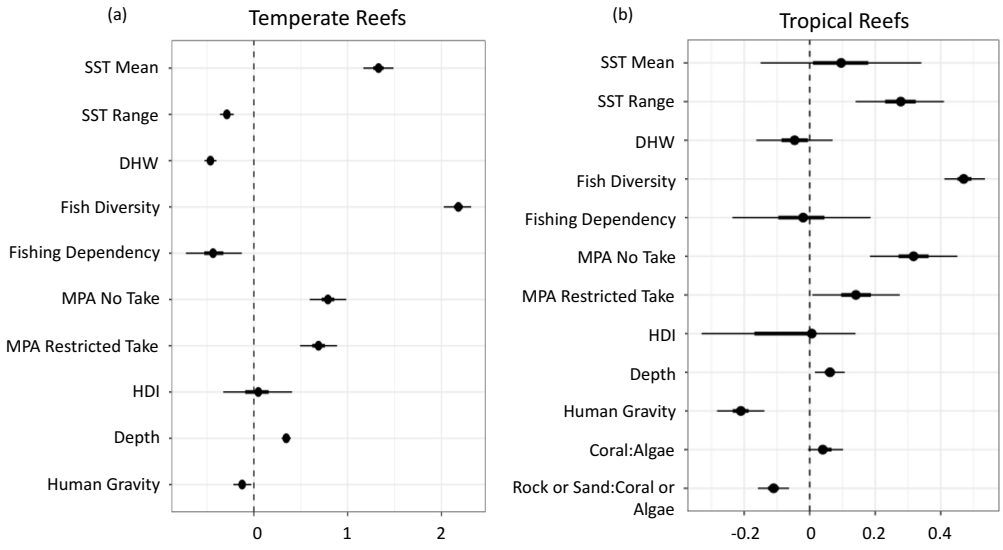

Figure 4.5 Causal effect of socio-environmental drivers on reef fish biomass on (a) temperate and (b) tropical reefs. Parameter estimates are standardized posterior mean values (dot), 95% credible intervals (thin lines), and 50% credible intervals (thick lines).

Our results confirm that temperature, represented by SST mean, is a dominant control on reef fish biomass in temperate reefs, reflecting strong temperature-dependence of biomass in cooler regions (Figure 4.5). These effects tapper off at higher temperatures found in tropical reefs, likely as reef fish reach their optimal temperature (Jobling 1997). The overall positive effect of SST mean on reef fish biomass is not surprising as temperature remains a major driver of the physiology, behaviour, abundance, and distribution of marine fishes (e.g., Brierley and Kingsford 2009; Sunday et al. 2011). As depicted by our DAG (Figure 4.5), SST mean is thought to effect reef fish biomass both directly, as well as indirectly, for example, through its effect on reef fish diversity (Duffy et al. 2016). While warmer average temperatures have a net positive effect on reef fish biomass, wider temperature fluctuations, as assessed by the effect of SST range, had a negative effect on reef fish biomass on temperate reefs but a positive effect on tropical reefs (Figure 4.5). In temperate regions, higher SST range may represent environments where the upper thermal tolerance limit of reef fish species has been surpassed, leading to an overall reduction in reef fish biomass (Munday et al. 2008; Habary et al. 2017). Previous research has found that higher temperature fluctuations can negatively impact larval dispersal and mortality (O'Connor et al. 2007; Takahashi et al. 2012), as well as reproductive success (Pankhurst and Porter 2003). However, on tropical coral reefs, sites with higher SST range may represent areas with higher historical temperature variability, which has been shown to improve physiological tolerance and performance of coral reefs under thermal stress (Safaie et al. 2018). Thus, coral reefs located in areas with higher temperature fluctuations may be more resistant to anomalous temperatures and bleaching, which in turn can have a positive effect on reef fish biomass. This is further supported by our results for thermal stress, assessed by degree heating week (DHW), which had an overall negative effect on reef fish biomass reefs, though the effect size was much weaker in tropical reefs (Figure 4.5). The lower effect size seen in tropical reefs may be because DHW is correlated with SST range; tropical coral reefs with higher historical temperature fluctuations may therefore be more protected against thermal stress assessed by DHW, leading to its reduced negative effect.

Our results also show that reef fish diversity is a major driver of reef fish biomass, having the highest effect size for tropical reefs, and second only to sea surface temperature for temperate reefs (Figure 4.5). Our results support previous global (Duffy et al. 2016; Lefcheck et al. 2021) and localized (Benkwitt et al. 2020) studies on reef ecosystems showing that biodiversity is one of the strongest predictors of reef fish biomass. It further supports a growing body of literature showing a positive effect of biodiversity on biomass and other ecosystem services (e.g., Tilman 2014). The strong positive relationship between reef fish diversity and biomass highlights the importance of conserving biodiversity across reef ecosystems to maintain biomass and other ecosystem functions and services. Reducing climate-related and other anthropogenic stressors should also be prioritized as they can also lead to biodiversity-mediated declines in reef fish biomass (Benkwitt et al. 2020).

Fishing dependency showed a negative effect on reef fish biomass on temperate reefs (Figure 4.5), which is expected as regions with higher fishing dependency likely have increased fishing pressure which in turn lead to overall reductions in reef fish biomass. Related, MPAs were shown to have a net positive impact on reef fish biomass, for both temperate and tropical regions (Figure 4.5). Our results show that general protection from fishing and other human pressures results in higher levels of reef fish biomass. This supports previous findings that fishes across trophic levels respond positively to MPAs on reefs (Soler et al. 2015). Our results are also supported by a global synthesis led by Lester et al. (2009), which show that MPAs increase fish biomass across marine ecosystems, and that MPAs have similar, if not a greater positive effect in temperate settings. Although some researchers have suggested that fished species in temperate regions are too mobile and/or long-lived to be effectively protected by MPAs (Shipp 2003; Blyth-Skyrme et al. 2006), our results show that in temperate reefs, MPAs can be highly effective.

Although MPA placement and management can be quite complex and context-dependent, our causal analysis provides a framework for assessing MPA effectiveness across future studies. We note that previous global reef studies have shown both a positive effect of

MPA (e.g., Duffy et al. 2016), as well as a lack of effect (e.g., Cinner et al. 2016), given the same dataset. These differences may have been due to statistical biases that often plague non-causal observational analysis (see Chapter 2). Here, the application of backdoor criterion allowed us to control for confounding factors assumed to impact both MPA placement and reef fish biomass, including fishing dependency, HDI, human gravity, and latitude. These controls are meant to remove bias in the socio-political processes that can accompany the selection of MPAs (e.g., Edgar et al. 2005; Ferraro et al. 2018). We encourage future studies investigating the effectiveness of MPAs to determine which variables should be controlled for, and similarly which variables should not be controlled for, through utilizing DAGs and the SCM framework.

Depth had a positive effect on reef fish biomass, particularly in temperate reefs (Figure 4.5). Although our data focused exclusively on shallow reefs (<50 m deep), our results nonetheless provide support for the deep reef refuge hypothesis (DRRH; Bongaerts et al. 2010). DRRH posits that deeper reefs are buffered from disturbances that affect shallower reefs, including fishing pressure, coastal pollution, and climate warming (Hoegh-Guldberg & Bruno 2010; Mies et al. 2020; Soares et al. 2020). As many reef fish species have a wide depth distribution (Fitzpatrick et al. 2012), they may find refugia from such stressors in deeper reefs, increasing reef fish biomass at deeper sites. Disturbance mediated coral loss from climate-induced bleaching, sedimentation, nutrient run-off, disease outbreaks, storms and other natural disturbances which may be more pronounced on shallow reefs (Slattery et al. 2011; Bridge et al. 2013) can further reduce reef fish through the loss of reef habitat (Wilson et al. 2006; Reopanichkul et al. 2009). Overall, our result highlights the potential for deeper reefs to protect reef fish species from natural and anthropogenic stressors, and we recommend that deeper reefs should continue to be incorporated into marine conservation and management strategies.

We note that previous studies that did not find an effect of depth on reef fish biomass may be due to model misspecification. For example, Cinner et al. (2016) looked at the influence of various environmental and social drivers of reef fish biomass, showing no substantial effect of depth. However, they used one statistical model with 18 covariates,

which may have biased their estimate of depth. For example, overcontrol bias occurs when an intermediate variable along a causal pathway is included as a covariate; in the pathway X → Z → Y, including Z as a covariate will remove the indirect effect of X on Y. As depth is expected to indirectly influence coral reef biomass through several mechanisms, the inclusion of 17 other covariates may have led to overcontrol bias, effectively removing the effect of depth. In contrast, given our DAG and the application of the backdoor criterion, no additional covariates were required to determine the effect of depth on reef fish biomass. This further highlights the utility of causal models for observational causal inference.

Human gravity had a strong negative effect on reef fish biomass in the tropics, and a slight negative effect in temperate reefs (Figure 4.5). Human gravity is expected to reduce reef fish biomass mostly through increased fishing pressure (Duffy et al. 2016) as well as through increased coastal development, eutrophication, and other anthropogenic stressors. The weaker effect seen in the temperate reefs may be due to the long history of human impacts endured in temperate regions; the smaller relative impact of human gravity may therefore reflect the rarity of remaining remote or pristine sites (Duffy et al. 2016).

The effect of benthic composition on tropical reefs was first assessed looking at the effect of coral to algae ratio, which had a negligible but slightly positive effect of coral on reef fish biomass (Figure 4.5). This supports recent research highlighting that both coral- and algae-dominated reefs can harbour high reef fish biomass livelihoods (Fulton et al. 2020; Robinson et al. 2019; Hamilton et al. 2022). For example, following climate-induced bleaching events, both recovering coral reefs and regime-shifted macroalgae-dominated reefs have been shown to maintain assemblages of reef fish (e.g., Robinson et al. 2019). Next, we assessed how the ratio of sand or rock to coral or algae impacted reef fish biomass, finding a negative impact of sand or rock on reef fish biomass. This is expected as both coral and algae are able provide food, shelter, and additional resources to a variety of reef fish, whereas nonliving substrates like sand or rock do not. Together, our results suggest that both coral and algae can provide a range of benefits that can increase

reef fish biomass, which in turn can sustain fisheries and livelihoods (Fulton et al. 2020; Robinson et al. 2019).

Collectively, our results can be used to prioritize reef conservation and management across both temperate and tropical reefs. Reef fish are diverse, play key functional roles in reef ecosystems, and provide food and livelihoods for over a billion people. Reef fish biomass is therefore directly related to ecosystem health, ecosystem services, and economic value (Woodhead et al. 2019). In temperate reefs, our results suggest that regions with higher average temperature, diversity, and depth, as well as lower temperature fluctuation and thermal stress and will harbour higher reef fish biomass. In tropical reefs, regions with higher average temperature, temperature fluctuation, diversity, depth, and living benthic substrate (coral or algae), as well as lower human gravity will harbour higher reef fish biomass. As our results further showed that MPAs can be effective at protecting reef fish biomass, reefs with the above characteristics should be prioritized, as they may be the most resilient against the ongoing negative impacts of climate change and other anthropogenic stressors. Understanding the drivers of reef fish biomass should therefore be used for sustaining both reef ecosystems and human development goals.

Ecologists continue to rely on observational data to better understand causal relationships in reef ecosystems, such as how key socio-environmental factors influence reef fish biomass and other ecosystem services. Although disentangling causal relationships from observational data remains a central theme in reef ecology, the methods used to answer such questions have not been appropriate for drawing valid causal conclusions. For example, most reef studies employ what is known as the ‘causal salad’ model (Bhalla 2018; McElreath 2020), whereby predictor variables of interest are placed under one statistical model and subsequently interpreted for their causal effects (e.g., Cinner et al. 2016, Darling et al. 2019). Without explicitly acknowledging how variables of interest are related to one another and the broader causal structure of a reef system, such approaches can lead to a variety of statistical biases, including (but not limited to) confounding, overcontrol, and collider bias (see Chapter 2.8.2 for definitions). As well,

several reef studies rely on predictive model selection approaches such as Akaike's Information Criterion (AIC) to select a best model, which is subsequently causally interpreted (e.g., Belwood et al. 2015; Safaie et al. 2018). However, as detailed in Arif and MacNeil (2022), predictive model selection techniques are meant for predictive inference (e.g., what model best explains a response variable, Y) and should not be conflated with causal analysis (e.g., what is the effect of X on Y).

To remedy the widespread misuse of non-causal statistical analysis, here we have highlighted the SCM framework, which is widely applicable across observational reef studies. This causal framework has several key advantages. First, it utilizes DAGs to depict the causal structure of a system under study, drawing on the domain knowledge of ecologists. Next, the application of the backdoor criterion allows researchers to create appropriate statistical models, based on their specific causal query. This step allows researchers to step away from data-driven methods such as AIC and other model selection criteria and instead allows ecological domain knowledge to be translated into effective causal models. Further, the SCM framework is transparent, as it allows researchers to communicate their causal assumptions and simultaneously allows other researchers to critically examine the conclusions of a given study. This can lead to productive discussions within reef ecologists, ultimately deepening the pace at which we understand and analyze our natural world.

CHAPTER 5 UTILIZING CAUSAL DIAGRAMS ACROSS QUASI-EXPERIMENTAL APPROACHES

A version of this work has been published: as Arif S, M MacNeil A. 2022. *Ecosphere* 13(4): e4009.

5.1 ABSTRACT

Recent developments in computer science have substantially advanced the use of observational causal inference under Pearl's structural causal model (SCM) framework. A key tool in the application of SCM is the use of casual diagrams in the form of directed acyclic graphs (DAGs), used to visualize the causal structure of a system or process under study. Here, we show how causal diagrams can be extended to ensure proper study design under quasi-experimental settings, including propensity score analysis, before after control impact (BACI) studies, regression discontinuity design (RDD), and instrumental variables (IV). Causal diagrams represent a unified approach to variable selection across methodologies and should be routinely applied in ecology research with causal implications.

5.2 INTRODUCTION

The availability and importance of observation-based research has increased in recent years due to the proliferation of both digital data and global environmental threats that cannot be manipulated experimentally (Sagarin and Pauchard, 2010). While infrequently stated, most observational studies in ecology are aimed at answering causal questions, such as ‘what is the effect of protected areas on biodiversity?’ (Gray et al. 2016). Yet the prohibition of causal language for non-experimental data promoted by Pearson and Fisher

(Pearson 1911, Fisher 1921, Glymour 2009) has constrained the use of observational data to answer fundamental causal questions in ecology. These opportunities and challenges highlight the importance of coherent methods to properly analyze observational data and attain accurate conclusions about ecological systems and processes.

Developments in observational causal inference have been spurred largely by the work of computer scientist Judea Pearl, who's structural causal model (SCM; Pearl 2009) provides a comprehensive framework that utilizes causal diagrams in the form of directed acyclic graphs (DAGs) to determine cause and effect relationships from purely observational data. DAGs explicitly state the direction of hypothesized causal associations between variables in a system and, in doing so, reveal non-causal (spurious) associations as well. By ensuring that researchers explicitly and transparently state their causal assumptions, causal diagrams invite critical reception and feedback that is typically difficult to frame.

What has gone unrecognized is that DAGs and principles from the SCM framework can also lead to effective study design across a range of quasi-experimental methods, including propensity score analysis, before after control impact (BACI) studies, regression discontinuity design (RDD), and instrumental variables (IV; Butsic et al. 2017; Larsen et al. 2019). Quasi-experimental approaches are widely used among other disciplines, and in recent years ecologists have argued for their increased use with ecological observational data (e.g., Butsic et al. 2017; Larsen et al. 2019; Wauchope et al. 2021). However, determining causal relationships from quasi-experimental methods requires proper study design and statistical analysis that benefit from explicit communication about a researchers' causal assumptions (Ferraro et al. 2019; Adams et al. 2019). Here we show how the application of DAGs combined with the core principles of SCM can be applied across quasi-experimental approaches, leading to more robust causal conclusions drawn from observational data. Using simulated data (with known causal effects), we show how the application of causal diagrams can return accurate causal estimates, as well as highlight how biases can arise when they are not considered.

5.3 STRUCTURAL CAUSAL MODEL (SCM)

SCM (Pearl 2009) provides a comprehensive theory of causation by unifying structural equation models (SEM; Wright 1921) with the potential outcome framework (Rubin 1974; Holland 1986) and other theories of causation. Although the mathematical underpinning is quite complex (Pearl 2009), one of the advantages of applying SCM is that it reduces complicated equations and mathematical theory into a graphical application of rules using DAGs to visualize and quantify causal relationships.

DAGs are causal diagrams that represent the causal structure of a system or process under study (e.g., see Grace and Irvine 2020). Specifically, nodes within a DAG represent variables, with directed arrows between nodes representing possible causal effects (e.g., X → Y shows that X affects Y). DAGs are also acyclic, meaning that they cannot contain bi-directional relationships or a feedback loop where a variable causes itself (Elwert 2014). However, they can still represent ecological systems with bi-directional relationships by more finely articulating the temporal sequence of events (Greenland et al. 1999).

DAGs are created based on researchers' domain knowledge, which can be supported by expert opinion, scientific consensus, and relevant literature (e.g., Cronin and Schoolmaster Jr. 2018; Schoolmaster et al. 2020; Grace and Irvine 2020). They must include all measured and unmeasured variables required to depict the system or process under study, as well as all common causes of any pair of variables included in the DAG (Sprites et al. 2001). For example, to determine the effect of X on Y, our DAG must include X, Y, common causes of X and Y, as well as common causes of any pair of variables that are now included in the DAG.

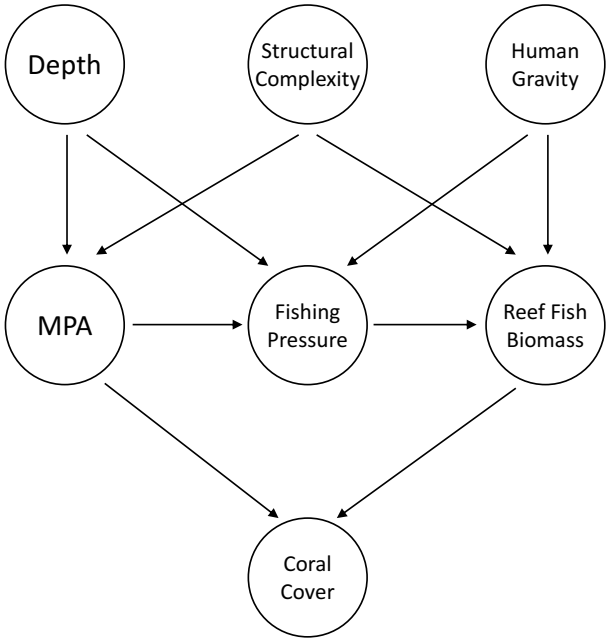

Figure 5.1 A directed acyclic graph (DAG) representing the causal structure between a marine protected area (MPA) and coral reef fish biomass.

As an ecological example, the DAG in Figure 5.1 represents the causal structure of how marine protected areas (MPAs) are expected to influence reef fish biomass for a hypothetical coral reef system, created based on past literature and expert knowledge of coral reef ecologists (Supplementary Table 5.1). We have created a simulated dataset based on the causal structure depicted in Figure 5.1, setting our *known* causal effect of MPA on reef fish biomass to **1.089** (Supplementary Material 5.8.2). We will use our DAG and associated simulated data to show how the application of SCM can lead to the accurate causal estimate of MPA on reef fish biomass.

Once a DAG has been created, it can be tested against the associated observational data to test for DAG-data consistency. Simply put, a specified DAG will have (often many) independencies (e.g., X is independent of Y) and conditional independencies between variables (e.g., X is independent of Y, given Z) that should be consistent with the associated observational data, if both the DAG and observational data are representative of the data generating process (Pearl, 2009; Textor et al. 2016; Supplementary Material 5.8.3). If *all* implied independencies are compatible with the data, it provides overall

support for a DAG. Given our DAG, there are 12 conditional independencies that are implied by our DAG (Supplementary Material 5.8.3). Using the package 'dagitty' we can test our DAG against our simulated data, which shows that all implied independencies are consistent with our simulated dataset (Supplementary Material 5.8.3). We note that if a DAG does not pass DAG-data consistency, it can be altered until DAG-data consistency is ensured, assuming that observational or simulated dataset is representative of the data-generating process depicted in a DAG (Textor et al. 2016). We also note that since several DAGs may pass DAG-data consistency, it is imperative that a finalized DAG is first and foremost justified through domain knowledge (e.g., through the literature, expert knowledge, past experiments).

Once a DAG is finalized and ensures DAG-data consistency, we can apply a graphical procedure known as the *backdoor criterion* to determine the sufficient set for adjustment required to determine the effect of X on Y (Pearl 2009), or in this case MPA on reef fish biomass. The application of backdoor criterion is based on an algebraic procedure known as do-calculas, which equates observational distributions to post-intervention distributions that would be expected under an experimental treatment (Pearl, 1995). While the application of do-calculas can make for challenging reading, based on its derived inferences rules, the backdoor criterion provides DAG-based graphical rules that can be applied to estimate causal effects from observational data.

Specifically, the backdoor criterion instructs us to block all *backdoor paths* between our predictor and response variable, X and Y. A backdoor path is a sequence of arrows and nodes connecting X and Y variables with an arrow pointing into X. If left open, these backdoor paths create bias and induce spurious correlation by providing an indirect, non-causal path along which information can flow.

The Backdoor Criterion:

The backdoor criterion (Pearl 1993, 2009) states that a set of variables, Z, is sufficient for estimating the causal effect of X on Y if variables in Z block all

backdoor paths from X to Y. To block a backdoor path between X and Y, the path must be 'd-separated' (Pearl 1988). A path between X and Y can be d-separated if either:

(i) the path contains at least one arrow-emitting variable that is in Z, or
(ii) the path contains at least one collider variable (a variable with two incoming arrows, e.g., B is a collider variable in A → B ← C) that is outside Z and has no descendant in Z

For our DAG, there are two backdoor paths between MPA and reef fish biomass that must be d-separated (i.e., blocked):

1. MPA ← Structural Complexity → Reef Fish Biomass
2. MPA ← Depth → Fishing Pressure → Reef Fish Biomass

The first backdoor path can be blocked by adjusting for structural complexity and the second backdoor path can be blocked by adjusting for depth. Therefore, to block all backdoor baths, we must adjust for both structural complexity and depth.

We note that the application of the backdoor criterion can become complicated as we move on to larger and more complex DAGs (see Supplementary Material 5.8.4). In some cases, more than one adjustment set may be available to determine the causal effect of X on Y. In these scenarios, it is best to choose the adjustment set with the lowest measurement error. Other times, the adjustment set(s) required may not be available due to the presence of unmeasured variables. To avoid this scenario, we recommend that researchers think critically and draw potential DAGs before collecting observational data. Given that application of the backdoor criterion can become difficult to apply for increasingly complex DAGs, researchers can draw their DAG on www.daggity.net (instructions within site), which will apply the backdoor criterion and generate the adjustment set(s) required to determine causal effects, given a specified DAG and causal question.

Once the backdoor criterion is applied to determine the sufficient set for adjustment, we can choose an appropriate statistical model for analysis. DAGs can be used to guide covariate selection (i.e., which variables to control for) and are not the estimation method; therefore, ecologists must select the statistical approach that best fits their data and study question. DAGs are non-parametric, meaning that they make no assumptions about the distribution of variables (e.g., normally distributed) or the functional form of causal effects (e.g., linear, nonlinear, stepwise). As such, they are compatible with a wide range of statistical analysis (e.g., hierarchical Bayesian model). Here, since our simulated data was created using linear relationships, we have applied a generalized linear model (GLM). Our GLM specifies reef fish biomass as the response variable, MPA as the predictor, and includes depth and structural complexity as controls. We interpret the coefficient of MPA as our total causal estimate on reef fish biomass. Here, our causal estimate of MPA on reef fish biomass returned an accurate estimate of 1.17 [1.08, 1.26], with the 95% confidence interval including the true causal effect of 1.089 (Supplementary Material 5.8.5).

It is important to note that covariate selection based on the backdoor criterion eliminates common statistical biases including confounding, overcontrol, and collider bias that often plague observational studies (see Chapter 2.8.2 for definitions). Confounding bias occurs when a common cause between predictor and response is not adjusted for. Given our DAG, if no adjustments are made, confounding bias would arise from depth and structural complexity, which effect both MPA and reef fish biomass. Indeed, a GLM with no adjustments gave an inaccurate estimate of 3.40 [3.32, 3.48] for the effect of MPA on reef fish biomass (Supplementary Material 5.8.6).

Lesser known, but equally important, are overcontrol and collider bias. Overcontrol bias occurs when an intermediate variable between predictor and response is adjusted for, blocking the causal association between predictor and response. Given our DAG, adjusting for fishing pressure would lead to overcontrol bias, giving an inaccurate estimate of -0.10 [-0.16, -0.04] for the MPA effect, even when depth and structural

complexity are adjusted for (Supplementary Material 5.8.6). Collider bias occurs when a variable affected by both predictor and response is adjusted for, creating a spurious association between predictor and response. Given our DAG, adjusting for coral cover would lead to collider bias, giving an inaccurate MPA estimate of -0.13 [-0.15, -0.11] (Supplementary Material 5.8.6). Collectively, the application of the backdoor criterion eliminates all three statistical biases, allowing for accurate causal estimates.

The SCM framework can be employed across a range of observational ecological studies (e.g., Cronin and Schoolmaster Jr. 2018; Schoolmaster et al. 2020; Grace and Irving 2020). Importantly, DAGs and the principles of SCM (e.g., the backdoor criterion) can be applied to ensure proper study design across other quasi-experimental approaches that have gained traction withing ecology (Butsic et al. 2017; Larsen et al. 2019), including propensity score analysis, BACI studies, RDD, and IV.

5.4 Matching Methods

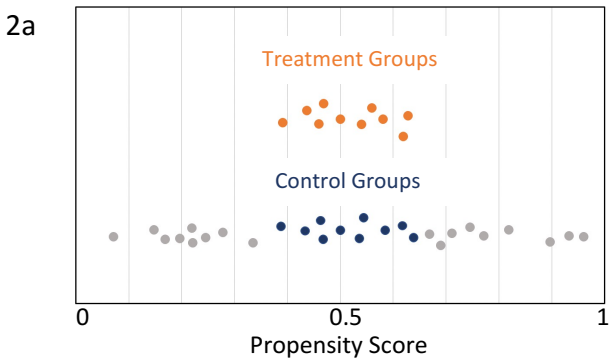

Matching Methods:

- Matching methods (e.g., propensity score analysis) selects control and treatment groups that are similar across selected covariates – based on confounding variables – to reduce confounding bias

Benefit of Causal Diagrams:

- Allows appropriate selection of variables to enter the propensity score, reducing confounding, overcontrol, and collider bias

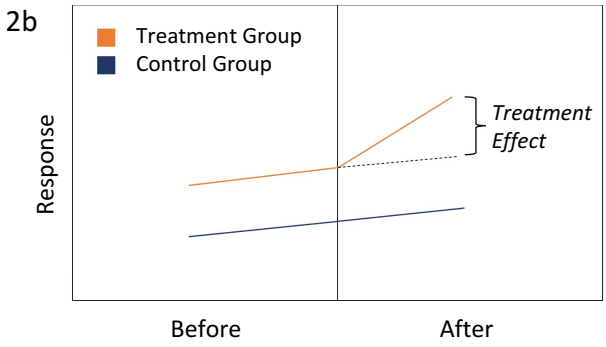

Before-after-control-impact (BACI):

- BACI measures a response both before and after an intervention for both treatment and control group(s); difference in the rate of change between treatment and control is attributed to the intervention

Benefit of Causal Diagrams:

- Clarifies whether all assumptions of the BACI approach are met; allows researchers to identify and adjust for confounding

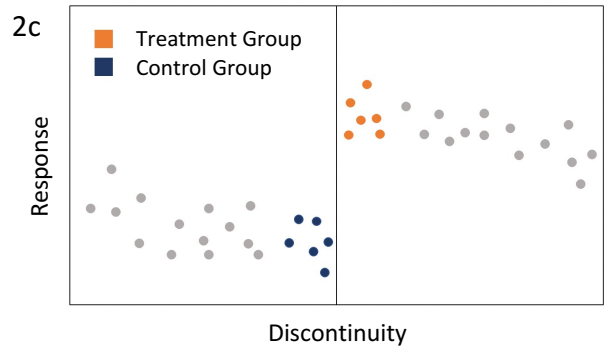

Regression Discontinuity Design (RDD):

- RDD selects treatment and control groups from either side of a discontinuity, where confounding variables are expected to be similar

Benefit of Causal Diagrams:

- Explicitly communicates a researchers' assumptions about why a chosen discontinuity is expected to eliminate bias

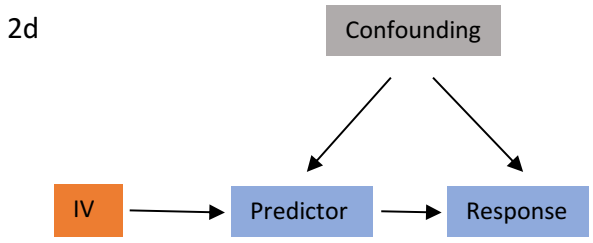

Instrumental Variable (IV):

- An instrument is used to determine the effect of a predictor on response when otherwise unfeasible (e.g., in the presence of an unmeasured confounding variable)

Benefit of Causal Diagrams:

- Clarifies whether all assumptions of the IV approach are met (e.g., exclusion criterion)

Figure 5.2 Quasi-experimental approaches: (a) propensity score analysis (b) before after control impact (BACI) (c) regression discontinuity design (RDD) (d) instrumental variable (IV).

Matching methods are often employed to remove confounding bias within observational studies (Stuart 2010). Matching methods aim to balance the distribution of covariates between treatment and control groups (Figure 2a). Covariates placed in a matching procedure should include all and only confounding variables assumed to affect both treatment assignment (e.g., MPA placement) and the outcome (e.g., reef fish biomass), thereby eliminating confounding bias (Rosenbaum and Rubin 1983). Implementing matching methods first requires the selection of a distance measure, used to define how close two units are based on selected covariates. Several distance measures are available to researchers, including propensity scores and Mahalanobis distance (Stuart 2010). Distance measures are subsequently used to match treatment and control units, which can be done through several matching methods including nearest neighbour matching, optimal matching, and exact matching (Stuart 2010).

Matching methods have been employed across a range of ecological systems to determine the causal effect of treatments, including the effect of protected areas on natural forests (Adnam et al. 2008; Herrera et al. 2019) and freshwater species (Chessman, 2013), the effect of agriculture on stream ecosystems (Pearson et al. 2016), and the impact of invasive species management on tree condition (Ramsey et al. 2019). However, although past studies state the presumed confounding variables used in their matching procedure, it is unclear how these confounding factors interact with one another within the broader causal structure of a study system. Without this knowledge, it is unclear whether there are unobserved or unmentioned variables that need to be included in the matching analysis or whether the inclusion of all selected variables may lead to other forms of bias (e.g., overcontrol or collider bias) through the accidental inclusion of non-confounding variables.

To resolve these issues, matching methods can make use of DAGs. As previously noted, after creating and finalizing a DAG, researchers can apply the backdoor criterion to determine which covariates need to be controlled for to determine the causal effect of X on Y. Under matching procedures, the set of covariates that enter the matching procedure

must satisfy the backdoor criterion (Pearl 2009). For example, if we choose propensity score matching, then given our DAG (Fig 1), to determine the effect of MPA on reef fish biomass, depth and structural complexity must enter the propensity score for MPA placement (Supplementary Material 5.8.7). To ensure covariate balance is achieved, we can employ balancing tests, which are often applied across matching procedures (Supplementary Material 5.8.7). When this is done, and our propensity score is used as a covariate adjustment, we return an accurate causal estimate for MPA of 1.17 [1.07, 1.27] (Supplementary Material 5.8.7).

Employing the backdoor criterion also eliminates other forms of bias that can occur within matching methods, including overcontrol and collider bias (Pearl 2009; Mansournia et al. 2013). For example, if all available variables entered our propensity score, this would ultimately lead to overcontrol bias (due to the inclusion of an intermediate variable, fishing pressure) and collider bias (due to the inclusion of a collider variable, coral cover), giving an inaccurate MPA causal estimate of -0.19 [-0.29, 0.11] (Supplementary Material 5.8.7). Additionally, DAGs include both measured and unmeasured variables needed to depict a complete causal structure, thereby explicitly stating any missing variables that must be considered. This is critical as the omission of unmeasured variables required in a backdoor adjustment set can lead to bias (Pearl 2009). Collectively, the application of the backdoor criterion on DAGs helps determine covariates that must and must not enter a matching procedure, while also communicating a system's assumed causal structure.

5.5 Before After Control Impact (BACI)

If observational data is available both before and after an event, BACI designs (Green 1979) can be used to assess the effect of interventions, including anthropogenic disturbances or environmental management actions. BACI works by measuring a response (e.g., reef fish biomass) both before and after an intervention (e.g., MPA placement) for both treatment and control site(s). BACI rests on the assumption that trends in the treated and control groups would be identical if the intervention did not

occur, meaning any difference in the rate of change between treatment and control is attributed to the intervention (Figure 5.2b).

BACI and its extensions (e.g., BACIPS, Stewart-Oaten et al. 1986; Progressive-Change BACIPS, Thiault et al. 2016) have been applied across various ecological studies, including to determine the effects of invasive species on invertebrates (Kadye and Booth 2012), restoration programs on biota (Bousquin and Colee 2014; Suren et al. 2011) and MPAs on coral reef fish communities (Thiault et al. 2019). Wauchope et al. (2021) further show how to analyze BACI data to determine *trend* and *immediate* change, in addition to average change, which may better capture ecological responses to interventions. However, although BACI study designs have the potential to provide valid causal inference, past reviews have noted the prevalence of improper study design, where the consideration of all relevant variables is often neglected, particularly joint consideration of both ecological and human factors (Ferraro et al. 2019; Adams et al. 2019). Here, DAGs can allow researchers to consider all relevant variables and clarify the assumptions required for appropriate BACI studies.

Let's consider a standard BACI design, which is also referred to as difference-in-difference (DiD) in some fields (Wauchope et al. 2021). Given our asserted DAG and the application of the backdoor criterion, we know that depth and structural complexity are confounding variables that must be accounted for to determine the effect of MPA on reef fish biomass. A strength of a BACI design is that it already accounts for certain confounding variables. Confounding in BACI designs occurs only if a variable (1) effects the treatment group and (2) has an effect on the outcome *trends,* which can occur when a variable has a time-varying difference between treatment groups or a time-varying effect on the outcome (Zeldow and Hatfield 2019). In our simulation, neither depth nor structural complexity have a time-varying difference between treatment groups or a time-varying effect on the outcome, so the application of a BACI analysis will return an accurate causal estimate for MPA of 1.07 [0.96, 1.20], without the need to adjust for these variables (Supplementary Material 5.8.8). Critically, when designing BACI studies, researchers must ensure that the variables in a backdoor adjustment set are accounted for,

either by design or through appropriate statistical adjustments. For example, if a bleaching event occurred after initial MPA placement, and disproportionately reduced the structural complexity across MPA sites, structural complexity would now act as a confounding variable by having a time-varying difference between treatment groups. Under these circumstances, our BACI analysis returns an inaccurate causal MPA estimate of 0.19 [0.08, 0.30; Supplementary Material 5.8.8]. However, we can return an accurate estimate of 1.06 [0.97, 1.16] by making the appropriate adjustment for structural complexity (Supplementary Material 5.8.8). We refer readers to Zeldow and Hatfield (2019), who provide instructions on how to adjust for confounding variables, when they do arise in BACI studies.

Given the need for proper study design (Ferraro et al. 2019; Adams et al. 2019), using DAGs to guide BACI studies will ultimately lead to more impactful and accurate causal estimates due to the extra care taken to understand how causal assumptions can be met. In addition, researchers can also employ placebo tests used in BACI studies to further support their causal conclusions; for example, researchers can apply a BACI analysis using only pre-treatment data, which should show a lack of causal effect (e.g., Schnabl, 2012). As such, the integration of causal diagrams with BACI can lead to additive methods for supporting causal claims, which in turn lead to more comprehensive causal conclusions.

5.6 Regression Discontinuity Design (RDD)

RDD aims to minimize the effect of confounding bias by exploiting a discontinuity in either space, time, or policy to separate observations into treatment and control groups (Imbens and Lemieux, 2008). The key assumption is that at or near this discontinuity, confounding variables are equal between treated and control groups. If the underlying confounding variables are similar before and after the change, then the treatment effect can be estimated by comparing the average difference between treated and control groups (Fig 2c). Although past review papers have highlighted the potential of RDD in ecology

(Butsic et al. 2017), it remains underutilized. We could find only one example of its use for causal inference, a conference paper studying the effect of protected areas on deforestation, population settlements, and road infrastructure that used the border of protected areas as the discontinuity with treatment and control groups being comprised of study sites from each side (Perez et al. 2017).

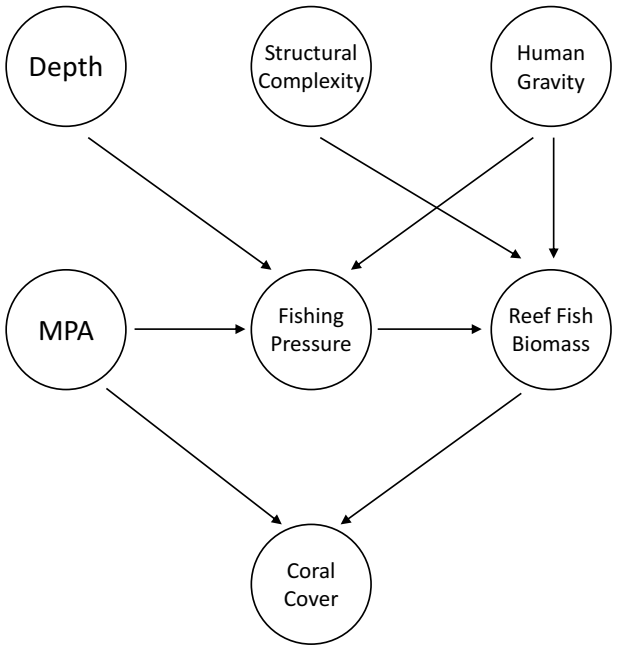

Figure 5.3 An alternative DAG representing the effect of a marine protected area (MPA) on reef fish biomass under a regression discontinuity design where only data near a discontinuity (MPA border) is considered.

Despite limited use to date, RDD provides a strong causal inference approach across ecological studies whenever there exists a sharp break between treatment groups across observational units, including protected area borders, fishing and land use zones, species ranges, and soil types (Butsic et al. 2017). Yet here again causal diagrams should be utilized to visualize how exploiting a discontinuity can break the backdoor (i.e., non-causal) paths between predictor and response. For example, to determine the effect of MPA on reef fish biomass, we can use the MPA border as our discontinuity if confounding variables (depth and structural complexity) are approximately the same on either side and fish do not readily move across the boundary. Figure 5.3 uses a DAG to

represent the causal structure between MPA and reef fish biomass within this kind of RDD design. Here, our observational data comes only from our discontinuity range, on either side of MPA border. As such, depth and structural complexity are assumed to no longer act as confounding variables, meaning the effect of MPA on reef fish biomass can be estimated without needing to adjust for additional variables (i.e., there are no backdoor paths between MPA and reef fish biomass within this discontinuity). A simulated dataset using data only from this discontinuity (following the causal structure in Figure 5.3) returned an accurate causal estimate of MPA on reef fish biomass, 1.08 [0.66, 1.49] (Supplementary Material 5.8.9).

Visualizing RDD with causal diagrams is particularly important in ecology due to the complex nature of causal connections that may exist near a chosen discontinuity. For example, it may be the case that an ecological RDD design removes some backdoor paths between two a predictor and outcome variable of interest, while still leaving a sub-set of backdoor paths open. In such a case, necessary adjustments (during the statistical analysis) can be made to ensure all backdoor paths are closed. Ultimately, DAGs allow researchers to visualize the causal structure near a discontinuity, to help ensure proper study design. Although underutilized, well thought out RDD communicated through causal diagrams can provide effective and transparent observational causal inference and should be more routinely used. Placebo tests used within RDD studies (e.g., using pre-treatment variables as placebo outcomes) can further be employed to provide additional support for causal conclusions (Eggers et al. 2021).

5.7 Instrumental Variables (IV)

The IV approach (Wright 1928; Kendall, 2015) can be used to determine the effect of X on Y in the presence of an unmeasured confounding variable, leading to confounding bias, or bi-directional relationships, which can generate simultaneity bias. For example, the DAG in Figure 5.2d shows that the effect of predictor on response cannot be determined from a simple regression analysis due to the presence of an unmeasured confounding variable. In such cases, an instrument, Z, can be used to determine the effect of X on Y if it meets three requirements (Hernan and Robins, 2006). First, Z must be

correlated with the predictor variable; the stronger the correlation, the more effective the instrument Z will be. Second, Z must not have a *direct* causal effect on the response variable and must only be associated with outcome Y through X, known as the *exclusion criterion*. Third, there must be no confounding variables that affect both Z and Y. If these three requirements are met, Z can be used as an instrument to determine the effect of X on Y through a two-stage regression (Kendall 2015).

Finding an instrument that satisfies all three criteria can be difficult practice, which may limit its use in ecological studies. However, when applicable, IV remains a powerful technique that can be used to prevent confounding and simultaneity bias across observational ecological studies. Already, several implementations of IV exist within the ecological literature: Bush and Cullen (2009) used site accessibility measures as instruments to determine the effectiveness of endangered species recovery treatments; Amin et al. (2015) used biodiversity as an instrument to determine the effect of protected areas on deforestation; and Butsic et al. (2015) used multiple instruments to determine the effect of warfare, mining, and protected areas on deforestation.

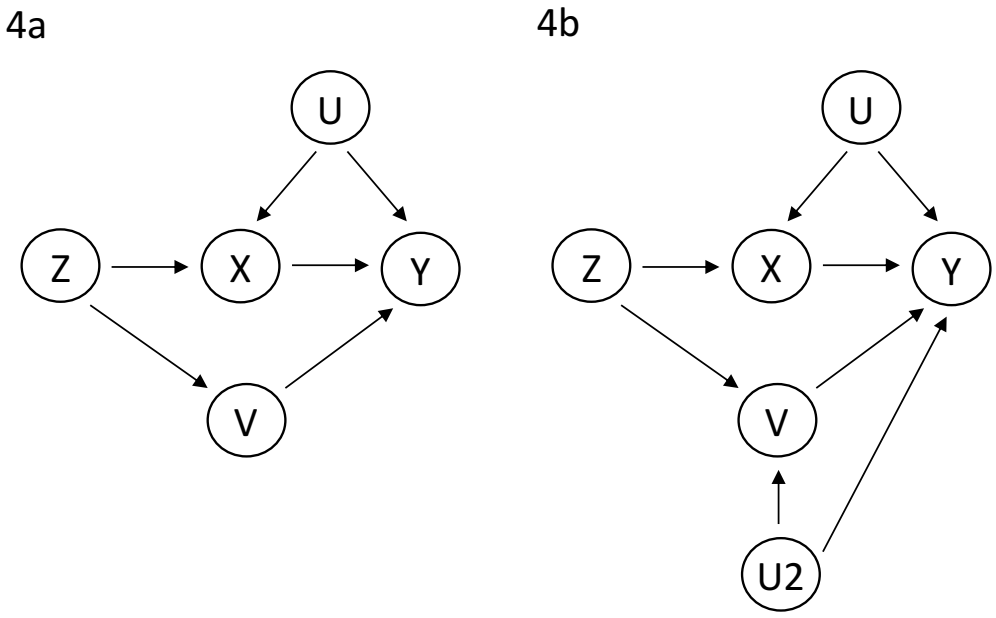

Figure 5.4 DAGs under two instrumental variable scenarios: (4a) represents a scenario where an instrument, Z, can be used to determine the effect of X on Y while (4a) represents a scenario where an instrument, Z, cannot be used to determine the effect of X on Y because it cannot meet the exclusion criterion.

When implementing IV, causal diagrams should be drawn to depict the complete causal structure of a system under study to accurately assess whether a chosen instrument meets the necessary requirements. For example, Figure 4a depicts a DAG where X and Y are confounded by an unmeasured variable U. Here, our instrument, Z, does not initially satisfy the exclusion criterion because it effects Y through another intermediate variable, V: Y ← V ← Z. To satisfy the exclusion criterion, we must block (d-separate) this pathway by adjusting for V. Once this is done, Z can be used as an instrument to determine the effect of X on Y. We can additionally test our causal assumption that our instrumental variable, Z, is sufficiently correlated with our predictor variable, X, by testing against weak instruments, which is commonly employed across IV studies (e.g., Staiger and Stock 1997; Supplementary Material 5.8.10). Following this, an IV approach on our simulated data, using Z as our instrument and adjusting for V, returns an accurate causal estimate of 1.11 [1.07, 1.15] for X on Y (known causal estimate of 1.089; Supplementary Material 5.8.10).

As another example, the DAG in Figure 5.4b also requires adjustment for V to block the additional path from Z to Y: Y ← V ← Z. However, in doing so, we open another path between Z and Y: Y ← U2 → V ← Z. V acts as a collider variable (variable with two incoming arrows) in this path, which we open when adjusting for it. To block this additional path, we must also adjust for U2. However, U2 is an unmeasured variable, and therefore cannot be adjusted. In this scenario (Figure 5.4b), Z cannot act as an instrument to measure the effect of X on Y. A two-stage regression that did not adjust for U2 returned an inaccurate estimate of 0.73 [0.69, 0.78] for X on Y (known causal estimate of 1.089; Supplementary Material 5.8.10). By utilizing causal diagrams, we explicitly communicate our assumptions about the causal structure between an instrument, predictor, and response variable, and accordingly, ensure that the assumptions required

for an IV approach are satisfied. As such, the use of causal diagrams can lead to more accurate implementation of the IV approach.

5.8 Conclusion

Although causal diagrams are underutilized within ecology, they hold tremendous potential for guiding effective causal inference across a range of observational contexts. Here, we have highlighted the utility of DAGs (within the SCM framework) across four additional quasi-experimental approaches, showing how the use of causal diagrams clarifies and unifies variable selection in non-experimental settings. Their use will also help to produce more transparent communication about causal assumptions, leading to more critical and accurate discussion about the conclusions that can be drawn from ecological research. Further, the integration of DAGs with quasi-experimental methods leads to additive methods for supporting causal claims (e.g., balancing tests for matching methods, placebo tests for BACI and RDD designs, test for weak instruments for IV approaches), which can lead to more comprehensive causal analysis. Utilizing DAGs across quasi-experimental methods can lead to more accurate and comprehensive causal analysis. The consequences of such a change are profound – from management and policy decision making, to the development of ecological laws, ecology must embrace a causal understanding of our data-rich and radically-changing natural world.

5.9 Supplementary Material

5.9.1 Coral Reef DAG

Our hypothetical DAG in Figure 5.1 represents the causal structure of how marine protected areas (MPAs) are expected to influence reef fish biomass for a hypothetical coral reef system, created based on past literature and expert knowledge of coral reef

ecologists. Supplementary Table 5.1 provides rationale for each of our directed links within our DAG.

Table 5.1 Justification for the causal links in our hypothetical coral reef DAG.

Causal Link	Rational
Depth → MPA	In our hypothetical scenario, MPAs were placed in deeper and more structurally complex reefs, and these characteristics are expected to make reefs more resilient against climatic disturbances (Graham et al. 2015).
Structural Complexity → MPA	
Depth → Fishing Pressure	Fishing pressure is less pronounced in deeper reefs (Bridge et al. 2013).
Structural Complexity → Reef fish biomass	Structural complexity increased habitat availability, which can subsequently increase reef fish biomass (Graham and Nash, 2013; Verges et al. 2011).
Human Gravity -> Fishing Pressure	Higher human gravity is associated with increased fishing pressure, Cinner et al. 2018).
Human Gravity -> Reef Fish Biomass	Higher human gravity is associated with negative human impacts in the surrounding environment (e.g., nutrient pollution), which in turn can influence reef fish biomass (Cinner et al. 2018)
MPA → Fishing Pressure → Reef fish biomass	MPAs lead to higher reef fish biomass through the reduction of fishing pressure (Edwards et al. 2014; Soler et al. 2015).

MPA → Coral Cover	MPAs protect coral cover from fishing gear and other recreational activities (McManus et al. 1997; Strain et al. 2019; Stevens 2020).
Reef Fish Biomass → Coral Cover	Herbivorous reef fish maintain coral cover through grazing on macroalgae, which subsequently enhances coral recruitment through creating space for larval settlement (McCook et al. 2001).

5.9.2 Simulating Data

We have created a simulated dataset that follows the (linear) causal structure depicted in our DAG (Fig 1) using R. The causal effect of MPA on reef fish biomass is set to **1.089**. Using simulated data (with known causal effects) will help us later determine if the application of SCM can return the accurate causal estimates.

```
# Generating simulated data set for factors influencing
reef fish biomass on a series of coral reef sites, given
our DAG in Fig 1.

# install simstudy: user friendly package for simulating
data

install.packages("simstudy")
library(simstudy)

# set seed for reproducibility

set.seed(5)

# define variables for depth, human gravity, structural
```

```
     complexity, MPA, fishing pressure, reef fish biomass
     and coral cover

def <- defData(varname = "depth", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "human_gravity", dist =
     "normal", formula = 0, variance = 1)

def <- defData(def, varname = "structural_complexity", dist
     = "normal", formula = 0, variance = 1)

def <- defData(def, varname = "MPA", dist = "binary",
     formula = "0.2 * depth + 2.8 * structural_complexity",
     link = "logit", variance = 1)

def <- defData(def, varname = "fishing_pressure", dist =
     "normal", formula = "-0.99 * MPA + -0.2 * depth + 0.3
     * human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass", dist =
     "normal", formula = "-1.1 * fishing_pressure + -0.4 *
     human_gravity + 1.65 * structural_complexity",
     variance = 1)

def <- defData(def, varname = "coral_cover", dist =
     "normal", formula = "0.5 * MPA + 2.5 *
     reef_fish_biomass", variance = 1)

# create 10000 observations

coraldata <- genData(10000,def)
```

5.9.3 Test DAG-data Consistency

After a DAG is created it can be tested against the data it represents to ensure DAG-data consistency. A specified DAG will have (often many) independencies (e.g., X is independent of Y) and conditional independencies (e.g., X is independent of Y, given Z) that must be compatible with the dataset it represents. Conditional independencies emerge from **d-separation** (dependency separation) rules:

> **d-separation** (Pearl, 1988): A set of variables, Z, is said to block (or d-separate) a path from one variable to another if either
>
> (i) the path contains at least one arrow-emitting variable that is in Z, or
> (ii) the path contains at least one collider variable (variable with two incoming arrows) that is outside Z and does not cause any variables in Z

Our DAG has 12 independencies that can be tested against our simulated data:

1. Coral Cover ⊥ Fishing Pressure | MPA, Reef Fish Biomass
2. Coral Cover ⊥ Human Gravity | MPA, Reef Fish Biomass
3. Coral Cover ⊥ Structural Complexity | MPA, Reef Fish Biomass
4. Coral Cover ⊥ Depth | Fishing Pressure, Human Gravity, MPA, Structural Complexity
5. Coral Cover ⊥ Depth | MPA, Reef Fish Biomass
6. Fishing Pressure ⊥ Structural Complexity | Depth, MPA
7. Human Gravity ⊥ Structural Complexity
8. Human Gravity ⊥ Depth
9. Human Gravity ⊥ MPA
10. Reef Fish Biomass ⊥ Depth | Fishing Pressure, Human Gravity, Structural Complexity

11. Reef Fish Biomass ⊥ MPA | Fishing Pressure, Human Gravity, Structural Complexity
12. Structural Complexity ⊥ Depth

For example, based on d-separation rules, the first conditional independency states that coral cover should be independent on fishing pressure, is MPA and reef fish biomass is controlled for. All 12 independencies must be consistent with our simulated data to pass DAG-data consistency. Here, we will use the R package 'dagitty' (Textor et al. 2016) to test DAG-data consistency between our simulated dataset and specified DAG. Dagitty uses a formal test of zero (partial) correlation for each identified independency based on d-separation rules. To mitigate problems around multiple testing (e.g., for complex DAGs with many testable implications), the p-values obtained are then corrected using the Holm-Bonferroni method. We summarize the R code required to test DAG-data consistency below and refer readers to Textor et al. 2016 for further details.

```
# Install R package dagitty
install.packages("dagitty")
library(dagitty)

# download specified DAG from dagitty.net
DAG <- downloadGraph("dagitty.net/mouvsuG")

# evaluate the d-separation implications of our DAG with
our simulated dataset
test <- localTests(DAG,coraldata)

# perform Holm-Bonferrino correction
test$p.value <- p.adjust(test$p.value)
test # should show all p values above 0.05, suggesting DAG-
data consistency
```

Our results support that our simulated dataset is consistent with all implied independencies within our DAG. This is expected as we simulated our dataset based on the (linear) causal structure of our DAG. In general, passing DAG-data consistency provides support for the causal claims within a DAG. If a DAG does not pass DAG-data consistency, it can be altered until DAG-data consistency is ensured. We note that since several DAGs may pass DAG-data consistency, it is imperative that a finalized DAG is first and foremost justified through domain knowledge (e.g., through the literature, expert knowledge, past experiments).

5.9.4 Apply the Backdoor Criterion

Once a DAG is finalized to ensure DAG-data consistency, we can apply the backdoor criterion to for model selection. Specifically, the backdoor criterion instructs us to block all *backdoor paths* between our predictor and response variable, X and Y. A backdoor path is a sequence of arrows and nodes connecting X and Y variables with an arrow pointing into both X and Y. Note that while causality follows the direction of arrows, information can flow in either direction, leading to confounding that is otherwise difficult to detect. If left open, these backdoor paths create bias and induce spurious correlation by providing an indirect, non-causal path along which information can flow.

> *The Backdoor Criterion:*
> The backdoor criterion (Pearl 2009) states that a set of variables, Z, is sufficient for estimating the causal effect of X on Y if variables in Z block all backdoor paths from X to Y. To block a backdoor path between X and Y, the path must be 'd-separated' (dependence-separated). A path between X and Y can be d-separated if either:
>
> (i) the path contains at least one arrow-emitting variable that is in Z, or
> (ii) the path contains at least one collider variable (a variable with two incoming arrows, e.g., B is a collider variable in A → B ← C) that is outside Z and has no descendant in Z

For our DAG, there are two backdoor paths between MPA and reef fish biomass that must be d-separated (i.e., blocked):

3. MPA ← Structural Complexity → Reef Fish Biomass
4. MPA ← Depth → Fishing Pressure → Reef Fish Biomass

The first backdoor path can be blocked by adjusting for structural complexity and the second backdoor path can be blocked by adjusting for depth. Therefore, to block all backdoor baths, we must adjust for both structural complexity and depth.

The application of the backdoor criterion can become complicated as we move on to larger and more complex DAGs. As an example, let's consider the DAG in Supplementary Figure 5.5.

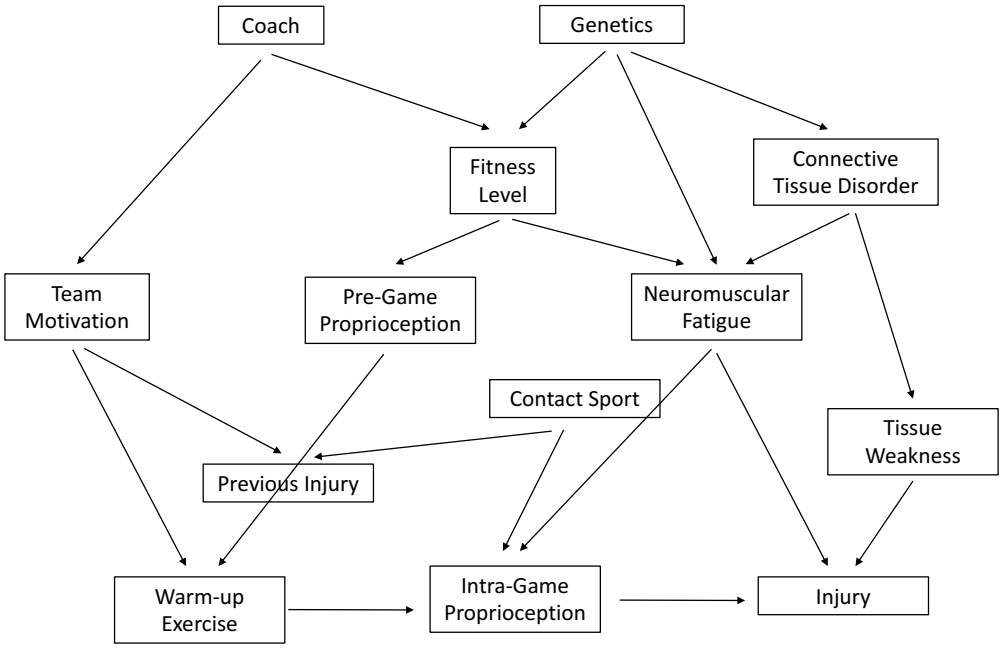

Supplementary Figure 5.5 DAG showing how different factors influence sports injury, based on Shrier and Platt (2008).

Given this DAG, if we want to look determine the effect of warm-up exercise on injury, the following backdoor paths need to be blocked:

1. Injury ← Neuromuscular Fatigue → Intra-Game Proprioception ← Contact Sport → Previous Injury ← Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable intra-game proprioception)
2. Injury ← Tissue Weakness ← Connective Tissue Disorder ← Genetics → Neuromuscular Fatigue → Intra-Game Proprioception ← Contact Sport → Previous Injury ← Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable intra-game proprioception)
3. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue ← Genetics → Fitness Level ← Coach → Team Motivation → Warm-up Exercise (already blocked because we did not adjust for either collider variables neuromuscular fatigue or fitness level)
4. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue ← Genetics → Fitness Level → Pre-game Proprioception → Warm-up Exercise (already blocked because we did not adjust for either collider variables neuromuscular fatigue)
5. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue ← Fitness Level → Pre-game Proprioception → Warm-up Exercise (already blocked because we did not adjust for either collider variables neuromuscular fatigue)
6. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue ← Fitness Level ← Coach → Team Motivation → Warm-up Exercise (already blocked because we did not adjust for either collider variables neuromuscular fatigue)
7. Injury ← Tissue Weakness ← Connective Tissue Disorder ← Genetics → Fitness Level → Neuromuscular Fatigue → Intra-Game Proprioception ← Contact Sport → Previous Injury ← Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable intra-game proprioception)

8. Injury ← Intra-Game Proprioception ← Contact Sport → Previous Injury ← Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable previous injury)
9. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue → Intra-Game Proprioception ← Contact Sport → Previous Injury ← Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable previous injury)
10. Injury ← Tissue Weakness ← Connective Tissue Disorder ← Genetics → Fitness Level ← Coach → Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable fitness level)
11. Injury ← Neuromuscular Fatigue ← Genetics → Fitness Level ← Coach → Team Motivation → Warm-up Exercise (already blocked because we did not adjust for the collider variable fitness level)
12. Injury ← Neuromuscular Fatigue ← Fitness Level → Pre-Game Proprioception → Warm-up Exercise (block by adjusting for neuromuscular fatigue, fitness level, or pre-game proprioception)
13. Injury ← Tissue Weakness ← Connective Tissue Disorder → Neuromuscular Fatigue ← Fitness Level → Pre-Game Proprioception → Warm-up Exercise (block by adjusting for tissue weakness, connective tissue disorder, neuromuscular fatigue, fitness level, or pre-game proprioception)
14. Injury ← Tissue Weakness ← Connective Tissue Disorder → Genetics → Neuromuscular Fatigue ← Fitness Level → Pre-Game Proprioception → Warm-up Exercise (block by adjusting for tissue weakness, connective tissue disorder, genetics, neuromuscular fatigue, fitness level, or pre-game proprioception)
15. Injury ← Neuromuscular Fatigue ← Fitness Level → Coach → Team Motivation → Warm-up Exercise (block by adjusting for neuromuscular fatigue, fitness level, coach, or team motivation)
16. Injury ← Tissue Weakness ← Connective Tissue Disorder ← Neuromuscular Fatigue ← Fitness Level → Coach → Team Motivation → Warm-up Exercise (block by adjusting for tissue weakness, connective tissue disorder, neuromuscular fatigue, fitness level, coach, or team motivation)

17. Injury ← Tissue Weakness ← Connective Tissue Disorder ← Genetics → Neuromuscular Fatigue ← Fitness Level → Coach → Team Motivation → Warm-up Exercise (block by adjusting for tissue weakness, connective tissue disorder, genetics, neuromuscular fatigue, fitness level, coach, or team motivation)

Our goal is to block all 17 backdoor paths with the minimal sufficient adjustment set(s) needed. In other words, we are looking for the minimal number of adjusted variables required to block all backdoor paths. We note that the first 11 pathways are currently blocked because in each of these paths, at least one collider variable was *not* adjusted for. To block the remainder of the pathways (12-17) a few options are available.

First, we see that adjusting for fitness level would block backdoor paths 10-15. However, this in turn would open backdoor paths 3, 10 and 11 (where fitness level acts as a collider variable). To re-block these paths, we could adjust for either coach, genetics, or team motivation without re-opening any additional backdoor paths. Therefore, we can satisfy the backdoor criterion by adjusting for either {fitness level and coach}, {fitness level and genetics} or {fitness level and team motivation}.

Similarly, we see that adjusting for neuromuscular fatigue would block backdoor paths 10-15. However, this in turn would open backdoor paths 4-6 (where neuromuscular fatigue acts as a collider variable). To re-block these paths, we could adjust for either tissue weakness or connective tissue disorder without re-opening any additional backdoor paths. Therefore, we can satisfy the backdoor criterion by adjusting for either {neuromuscular fatigue and tissue weakness} or {neuromuscular fatigue and connective tissue disorder}.

Another option to block backdoor paths 10-15, without opening backdoor paths 1-11 is to adjust for pre-game proprioception (blocking backdoor paths 12-14) and either coach or team motivation (blocking backdoor paths 15-17). Therefore, we can satisfy the backdoor

criterion by adjusting for either {pre-game proprioception and coach} or {pre-game proprioception and team motivation}.

Here, we have seven adjustment set(s), each requiring us to adjust for two separate variables. In scenarios with multiple minimal sufficient adjustment sets, researchers may want to choose to set with the least measurement error. As this example demonstrates, applying the backdoor criterion can become complicated with increasingly complex DAGs. To make things easier, researchers can draw their DAG on www.daggity.net (instructions within site), which will apply the backdoor criterion and generate the adjustment set(s) required to determine causal effects, given a specified DAG and causal question.

5.9.5 Choose Statistical Model

Once the backdoor criterion is applied for model selection, researchers must choose an appropriate statistical model for analysis. As DAGs are non-parametric, they make no assumptions about the distribution of variables (e.g., normal) or the functional form of effects (e.g. linear, nonlinear, stepwise), making them compatible with a wide range of statistical analysis. Since our simulated data was generated using linear relationships, here we will employ a generalized linear regression model (GLM). Our GLM specifies reef fish biomass as the response variable, MPA as the predictor, and includes depth and structural complexity as controls.

```
# run GLM following the backdoor criterion for model
selection

Model1 <- glm(reef_fish_biomass~MPA + depth +
     structural_complexity, data=coraldata)

# return estimate of MPA on reef fish biomass
```

```
summary(Model1)

# return 95% confidence interval around MPA estimate

confint(Model1)
```

The coefficient for MPA returns a causal estimate of 1.17 [1.08, 1.26] with our 95% confidence interval containing the known causal effect of 1.089.

5.9.6 Confounding, Overcontrol, and Collider Bias

When the backdoor criterion is not employed, several forms of statistical biases can arise. Below are examples of confounding, overcontrol, and collider bias.

Confounding bias occurs when a common cause between predictor and response is not adjusted for. Given our DAG, if no adjustments are made, confounding bias would arise from depth and structural complexity, which effect both MPA and reef fish biomass.

```
# Model with confounding bias
Model2 <- glm(reef_fish_biomass~MPA, data=coraldata)

# return estimate of MPA on reef fish biomass
summary(Model2)

# return 95% confidence interval around MPA estimate
confint(Model2)
```

Here, confounding bias leads to an inaccurate causal estimate of 3.40 [3.32, 3.48].

Overcontrol bias occurs when an intermediate variable between predictor and response is adjusted for, blocking the causal association between predictor and response. Given our DAG, adjusting for fishing pressure would lead to overcontrol bias

```
# Model with overcontrol bias
Model3 <- glm(reef_fish_biomass~MPA + depth +
     structural_complexity + fishing_pressure,
     data=coraldata)

# return estimate of MPA on reef fish biomass
summary(Model3)

# return 95% confidence interval around MPA estimate
confint(Model3)
```

Here, overcontrol bias leads to an inaccurate causal estimate of -0.10 [-0.16, -0.04].

Collider bias occurs when a variable affected by both predictor and response is adjusted for, creating a spurious association between predictor and response. Given our DAG, adjusting for coral cover would lead to collider bias.

```
# Model with collider bias

Model4 <- glm(reef_fish_biomass~MPA + depth +
     structural_complexity + coral_cover, data=coraldata)

# return estimate of MPA on reef fish biomass
summary(Model4)

# return 95% confidence interval around MPA estimate
confint(Model4)
```

Here, collider bias leads to an inaccurate causal estimate of -0.13 [-0.15, -0.11].

5.9.7 Matching Methods

The application of the backdoor criterion can guide proper implementation of matching methods, such as propensity score analysis. Given our DAG in Figure 5.1 and following the backdoor criterion, we will use both depth and structural complexity to determine our propensity score for MPA.

```
# Generating simulated data set for factors influencing
reef fish biomass on a series of coral reef sites, given
our DAG in Fig 1.

library(simstudy)

# set seed for reproducibility

set.seed(5)

#define variables for depth, human gravity, structural
complexity, MPA, fishing pressure, reef fish biomass and
coral cover

def <- defData(varname = "depth", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "human_gravity", dist =
     "normal", formula = 0, variance = 1)

def <- defData(def, varname = "structural_complexity", dist
     = "normal", formula = 0, variance = 1)

def <- defData(def, varname = "MPA", dist = "binary",
     formula = "0.2 * depth + 2.8 * structural_complexity",
```

```
     link = "logit", variance = 1)

def <- defData(def, varname = "fishing_pressure", dist =
     "normal", formula = "-0.99 * MPA + -0.2 * depth + 0.3
     * human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass", dist =
     "normal", formula = "-1.1 * fishing_pressure + -0.4 *
     human_gravity + 1.65 * structural_complexity",
     variance = 1)

def <- defData(def, varname = "coral_cover", dist =
     "normal", formula = "0.5 * MPA + 2.5 *
     reef_fish_biomass", variance = 1)

# create 10000 observations

coraldata <- genData(10000,def)

# Install MatchIt, a user friendly package for propensity
score analysis

install.packages("MatchIt")
library(MatchIt)

# Generate propensity score for MPA placement, based on
depth and structural complexity using a full matching
method

PS <- matchit(MPA ~ depth + structural_complexity, method =
     "full", data = coraldata, distance = "glm", link =
     "probit", discard = "both")
```

Before using our propensity score, we can check for covariate balance, the degree to which the distribution of covariates is similar across our treatment (MPA) and control (non-MPA) groups. We can do this with the summary() output in MatchIt.

```
# check for covariate balance using summary data
summary(PS)
```

There are three columns in the 'Summary of Balance for Matched Data' that we can use to assess covariate balance: standardized mean difference close to 0, variance ratios close to 1 (i.e., between 0.5-2), and low cDF statistics indicate good balance. These statistics should be considered together. Imbalance as measured by any of them may indicate a potential failure of the matching scheme to achieve distributional balance. Our results suggest that we have good balance.

Now that we have our propensity score (checked for covariate balance), we can simply run a regression analysis with our propensity score as a covariate:

```
# add propensity score to our coral reef dataset

coraldata$propensity_score <- PS$distance

# run model

Model5 <- glm(reef_fish_biomass~MPA + propensity_score,
     data=coraldata)

# return estimate of MPA on reef fish biomass

summary(Model5)

# return 95% confidence interval around MPA estimate
```

```
confint(Model5)
```

Here, we return an accurate estimate of 1.17 [1.07, 1.27].

Adding too many variables for propensity score matching can lead to biases, including overcontrol and collider bias. As an example, let's see what happens when we add all available data to calculate the propensity score for MPA:

```
# Generate propensity score for MPA placement, based on
depth, structural complexity, human gracity, fishing
pressure, coral cover

PS2 <- matchit(MPA ~ depth + structural_complexity +
     human_gravity + fishing_pressure + coral_cover, method
     = "full", data = coraldata, distance = "glm", link =
     "probit", discard = "both")

coraldata$propensity_score2 <- PS2$distance

# run model

Model6 <- glm(reef_fish_biomass~MPA + propensity_score2,
     data=coraldata)

# return estimate of MPA on reef fish biomass

summary(Model6)

# return 95% confidence interval around MPA estimate

confint(Model6)
```

Here, we return an inaccurate estimate of -0.19 [-0.29, 0.11].

5.9.8 Before After Control Impact (BACI)

Given our DAG in Figure 5.1, the backdoor criterion instructs us that depth and structural complexity (our confounding variables) must be accounted for. A main strength of BACI designs is that they account for confounding variables, as long as they are either time-invariant group attributes or time-varying variables that are group invariant. As both depth and structural complexity are time-invariant (see simulation below), they will be accounted for in our BACI design.

```
# Generating simulated before-after data set for factors
influencing reef fish biomass on a series of coral reef
sites, given our DAG in Fig 1.

library(simstudy)

# set seed for reproducibility

set.seed(5)

#define variables for depth, human gravity, structural
complexity and MPA; these variables are expected to be
similar both before and after MPA placement

def <- defData(varname = "depth", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "human_gravity", dist =
     "normal", formula = 0, variance = 1)
```

```
def <- defData(def, varname = "structural_complexity", dist
     = "normal", formula = 0, variance = 1)

def <- defData(def, varname = "MPA", dist = "binary",
     formula = "0.2 * depth + 2.8 * structural_complexity",
     link = "logit", variance = 1)

#define variables for fishing pressure, reef fish biomass
and coral cover before MPA placement

def <- defData(def, varname = "fishing_pressure_b", dist =
     "normal", formula = "-0.2 * depth + 0.3 *
     human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass_b", dist =
     "normal", formula = "-1.1 * fishing_pressure_b + -0.4
     * human_gravity + 1.65 * structural_complexity",
     variance = 1)

def <- defData(def, varname = "coral_cover_b", dist =
     "normal", formula = "2.5 * reef_fish_biomass_b",
     variance = 1)

#define variables for fishing pressure, reef fish biomass
and coral cover after MPA placement

def <- defData(def, varname = "fishing_pressure_a", dist =
     "normal", formula = "-0.99 * MPA + -0.2 * depth + 0.3
     * human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass_a", dist =
```

```
     "normal", formula = "-1.1 * fishing_pressure_a + -0.4
     * human_gravity + 1.65 * structural_complexity",
     variance = 1)

def <- defData(def, varname = "coral_cover_a", dist =
     "normal", formula = "0.5 * MPA + 2.5 *
     reef_fish_biomass_a", variance = 1)

# create 10000 observations

coraldata <- genData(10000,def)

# create a data frame specifying a time variable (0 =
before MPA placement; 1 = after MPA placement)

coraldata_before <- data.frame(depth=coraldata$depth,
     human_gravity=coraldata$human_gravity,
     structural_complexity=coraldata$structural_complexity,
     MPA=coraldata$MPA,
     fishing_pressure=coraldata$fishing_pressure_b,
     reef_fish_biomass=coraldata$reef_fish_biomass_b,
     coral_cover=coraldata$coral_cover_b)

coraldata_before$time=0

coraldata_after <- data.frame(depth=coraldata$depth,
     human_gravity=coraldata$human_gravity,
     structural_complexity=coraldata$structural_complexity,
     MPA=coraldata$MPA,
     fishing_pressure=coraldata$fishing_pressure_a,
     reef_fish_biomass=coraldata$reef_fish_biomass_a,
     coral_cover=coraldata$coral_cover_a)
```

```
coraldata_after$time=1

coraldata_DiD <- rbind(coraldata_before, coraldata_after)

# add an interaction variable between time and MPA

coraldata_DiD$interac <- coraldata_DiD$time *
coraldata_DiD$MPA

# final data frame

coraldata_DiD
```

In our dataset, time is represented by a dummy variable, where 0 represents before and 1 represents after MPA placement. We also have an interaction variable which is the product of time and MPA treatment.

To estimate the effect of MPA on reef fish biomass using before-after data, we can run a GLM with MPA, time and our interaction term as covariates:

```
# run model

Model7 <- glm(reef_fish_biomass ~ MPA + time + interac,
     data=coraldata_DiD)

# return estimate of MPA on reef fish biomass (read interac
term)

summary(Model7)
```

```
# return 95% confidence interval around MPA estimate (read
interac term)

confint(Model7)
```

Here, we see that our model returned an accurate causal estimate of MPA: 1.07 [0.96, 1.20].

As a second example, let's consider a situation where a bleaching event reduces structural complexity after MPA placement, disproportionately impacting MPA areas. Here, structural complexity is a confounding variable that is neither time- or group-invariant, invalidating the assumptions of a BACI design.

```
library(simstudy)

# set seed for reproducibility

set.seed(5)

#define variables for depth, human gravity, MPA which
remain similar both before and after MPA placement; define
structural complexity (before), fishing pressure (after),
fishing pressure (before), reef fish biomass (before), and
coral cover (before)

def <- defData(varname = "depth", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "human_gravity", dist =
     "normal", formula = 0, variance = 1)

def <- defData(def, varname = "structural_complexity_b",
```

```
     dist = "normal", formula = 0, variance = 1)

def <- defData(def, varname = "MPA", dist = "binary",
     formula = "0.2 * depth + 2.8 *
     structural_complexity_b", link = "logit", variance =
     1)

def <- defData(def, varname = "fishing_pressure_a", dist =
     "normal", formula = "-0.99 * MPA + -0.2 * depth + 0.3
     * human_gravity", variance = 1)

def <- defData(def, varname = "fishing_pressure_b", dist =
     "normal", formula = "-0.2 * depth + 0.3 *
     human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass_b", dist =
     "normal", formula = "-1.1 * fishing_pressure_b + -0.4
     * human_gravity + 1.65 * structural_complexity_b",
     variance = 1)

def <- defData(def, varname = "coral_cover_b", dist =
     "normal", formula = "2.5 * reef_fish_biomass_b",
     variance = 1)

# create 10000 observations

coraldata <- genData(10000,def)

# define structural complexity (after), which is
disproportionately reduced across MPA sites; add to pre-
existing dataset
```

```
defSC <- defCondition(condition = "MPA == 0", formula =
     "0.90 * structural_complexity_b", variance = 0.1, dist
     = "normal")

defSC <- defCondition(defSC, condition = "MPA == 1",
     formula = "0.30 * structural_complexity_b", variance =
     0.1, dist = "normal")

coraldata <- addCondition(defSC, coraldata,
     "structural_complexity_a")

# define reef fish biomass (after) and coral cover (after);
add to pre-existing dataset

coraldata$reef_fish_biomass_a <- -1.1 *
     coraldata$fishing_pressure_a + -0.4 *
     coraldata$human_gravity + 1.65 *
     coraldata$structural_complexity_a + rnorm(10000)


coraldata$coral_cover_a <- 0.5 * coraldata$MPA + 2.5 *
     coraldata$reef_fish_biomass_a + rnorm(10000)

# create a data frame specifying a time variable (0 =
before MPA placement; 1 = after MPA placement)

coraldata_before <- data.frame(depth=coraldata$depth,
     human_gravity=coraldata$human_gravity,
     structural_complexity=coraldata$structural_complexity_
     b, MPA=coraldata$MPA,
     fishing_pressure=coraldata$fishing_pressure_b,
```

```
     reef_fish_biomass=coraldata$reef_fish_biomass_b,
     coral_cover=coraldata$coral_cover_b)

coraldata_before$time=0

coraldata_after <- data.frame(depth=coraldata$depth,
     human_gravity=coraldata$human_gravity,
     structural_complexity=coraldata$structural_complexity_
     a, MPA=coraldata$MPA,
     fishing_pressure=coraldata$fishing_pressure_a,
     reef_fish_biomass=coraldata$reef_fish_biomass_a,
     coral_cover=coraldata$coral_cover_a)

coraldata_after$time=1

coraldata_DiD <- rbind(coraldata_before, coraldata_after)

# add an interaction variable between time and MPA

coraldata_DiD$interac <- coraldata_DiD$time *
coraldata_DiD$MPA

# final data frame

coraldata_DiD

# run model

Model8 <- glm(reef_fish_biomass ~ MPA + time + interac,
     data=coraldata_DiD)
```

```
# return estimate of MPA on reef fish biomass (read interac
term)

summary(Model8)

# return 95% confidence interval around MPA estimate (read
interac term)

confint(Model8)
```

This time, our interaction term returns an inaccurate estimate of 0.19 [0.08, 0.30]. Here, our BACI design does not account for structural complexity, which acts as a confounding variable. To mitigate this, we can adjust for structural complexity by adding it as a covariate. We refer readers to Zeldow and Hatfield (2019), who provide instructions on how to adjust for different types of confounding variables, when they do arise in BACI studies.

```
# run model

Model8_adjusted <- glm(reef_fish_biomass ~ MPA + time +
     interac   + structural_complexity, data=coraldata_DiD)

# return estimate of MPA on reef fish biomass (read interac
term)

summary(Model8_adjusted)

# return 95% confidence interval around MPA estimate (read
interac term)

confint(Model8_adjusted)
```

Our interaction term now returns an accurate estimate of 1.06 [0.97, 1.16].

5.9.9 Regression Discontinuity Design (RDD)

For our regression discontinuity analysis, we have MPA border as our discontinuity, with treatment and control coming from either side of this discontinuity. Our associated DAG in Figure 5.3 shows that when data from this region is considered, depth and structural complexity no longer act as confounding variables when determining the effect of MPA on reef fish biomass. Our simulation below uses a subset of data from this region, following the causal structure of our regression discontinuity DAG:

```
# Generating simulated data set for factors influencing
reef fish biomass on a series of coral reef sites, given
our DAG in Figure 5.3.

library(simstudy)

# set seed for reproducibility

set.seed(5)

#define variables for depth, human gravity, structural
complexity, MPA, fishing pressure, reef fish biomass and
coral cover

def <- defData(varname = "depth", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "human_gravity", dist =
     "normal", formula = 0, variance = 1)

def <- defData(def, varname = "structural_complexity", dist
```

```
     = "normal", formula = 0, variance = 1)

def <- defData(def, varname = "MPA", dist = "binary",
     formula = 0.5, link = "logit", variance = 1)

def <- defData(def, varname = "fishing_pressure", dist =
     "normal", formula = "-0.99 * MPA + -0.2 * depth + 0.3
     * human_gravity", variance = 1)

def <- defData(def, varname = "reef_fish_biomass", dist =
     "normal", formula = "-1.1 * fishing_pressure + -0.4 *
     human_gravity + 1.65 * structural_complexity",
     variance = 1)

def <- defData(def, varname = "coral_cover", dist =
     "normal", formula = "0.5 * MPA + 2.5 *
     reef_fish_biomass", variance = 1)

# create a subset of 500 observations

coraldata_rd <- genData(500,def)

# run model

Model9 <- glm(reef_fish_biomass~MPA, data=coraldata_rd)

# return estimate of MPA on reef fish biomass

summary(Model9)

# return 95% confidence interval around MPA estimate
```

```
confint(Model9)
```

Here, our subset of data returned an accurate estimate of 1.08 [0.66, 1.49]. We note that one limitation of the RDD design is that it can be data demanding, requiring enough observations within the discontinuity threshold.

5.9.10 Instrumental Variable (IV)

Below is the code for a simulated dataset following the causal structure of Figure 5.4a. We have set the known causal effect of X on Y to 1.089.

```
library(simstudy)

# set seed for reproducibility
set.seed(5)

#define variables for Z, V, U, X, and Y
def <- defData(varname = "Z", dist = "normal", formula = 0,
     variance = 1)

def <- defData(def, varname = "V", dist = "normal", formula
     = "0.4 * Z", variance = 1)

def <- defData(def, varname = "U", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "X", dist = "normal", formula
     = "0.9 * Z + 0.5 * U", variance = 1)

def <- defData(def, varname = "Y", dist = "normal", formula
     = "0.5 * U + 1.089 * X + 0.34 * V", variance = 1)
```

```
# create 10000 observations

coraldata_iv <- genData(10000,def)
```

Given our DAG in Figure 5.4a, to determine the effect of X on Y, we can use Z as our instrument, but must also adjust for V in order to meet the exclusion criterion. Using a two-stage least square regression, this looks like:

```
# run model

stage1 <- glm(X ~ Z + V, data = coraldata_iv)

# While running our two-stage regression, we can test
against weak instruments. A common rule of thumb is that an
F-statistic of more than 10 shows support against a weak
instrument(Staiger and Stock  1997). When dealing with one
instrument, this F-statistic is the square of the
instrument's t-statistic in the first stage.

# calculate F-statistic from first stage
summary(stage1)

# The F-statistic is 5961.9 (square of Z's t statistic
75.445), passing the weak instrument test.

stage2 <- glm(Y ~ predict(stage1) + V, data = coraldata_iv)

Model10 <- stage2

# return causal estimate of X on Y
summary(Model10)
```

```
# return 95% confidence interval around X
confint(Model10)
```

Here, we see that our results return an accurate estimate of 1.11 [1.07, 1.15].

Next, let's consider a simulated dataset following the causal structure of Figure 5.4b. Again, we set the known causal effect of X on Y to 1.089.

```
# set seed for reproducibility
set.seed(5)

# define variables for Z, U2, V, U, X, and Y

def <- defData(varname = "Z", dist = "normal", formula = 0,
     variance = 1)

def <- defData(def, varname = "U2", dist = "normal",
     formula = 0, variance = 1)

def <- defData(def, varname = "V", dist = "normal", formula
     = "0.4 * Z + 1.4 * U2", variance = 1)

def <- defData(def, varname = "U", dist = "normal", formula
     = 0, variance = 1)

def <- defData(def, varname = "X", dist = "normal", formula
     = "0.9 * Z + 0.5 * U", variance = 1)

def <- defData(def, varname = "Y", dist = "normal", formula
     = "0.5 * U + 1.089 * X + 0.34 * V + 1.6 * U2",
     variance = 1)
```

```
# create 10000 observations
coraldata_iv <- genData(10000,def)
```

Here, we can use Z as our instrument, but would also need to adjust for both V and U2 to satisfy the exclusion criterion. However, since U2 is an unobserved variable, we cannot find the true causal effect:

```
# run model

stage1 <- glm(X ~ Z + V, data = coraldata_iv)

stage2 <- glm(Y ~ predict(stage1) + V, data = coraldata_iv)

Model11 <- stage2

# return causal estimate of X on Y

summary(Model11)

# return 95% confidence interval around X

confint(Model11)
```

We have returned an inaccurate estimate of 0.73 [0.69, 0.78]. In comparison, if U2 was observed and adjusted for, we could run the following model:

```
# run model

stage1 <- glm(X ~ Z + V + U2, data = coraldata_iv)

# Again, we can check against weak instruments using the F-
statistic from our first stage regression
```

```
summary(stage1)

# The F-statistic is 5815(square of Z's t-statistic
76.257), passing the weak instrument test.

stage2 <- glm(Y ~ predict(stage1) + V + U2, data =
coraldata_iv)

Model12 <- stage2

# return causal estimate of X on Y
summary(Model12)

# return 95% confidence interval around X
confint(Model12)
```

Here, we return an accurate causal estimate of 1.08 [1.04, 1.13].

5.9.11 Literature Cited

Bridge, T., Hughes, T, Guinotte, J., Bongaerts, P. 2013. Call to protect all coral reefs. *Nature Climate Change,* 3(6), 528-530.

Cinner, J., Marie, E., Huchery, C., MacNeil, A., Graham, N., Mora, C. et al. 2018. Gravity of human impacts mediates coral reef conservation gains. *Proceedings of the National Academy of Sciences*, 115(27), E6116-E6125.

Cinner, J., Maire, E., Huchery, C., MacNeil, A., Graham, N., Mora, C. et al. 2018. Gravity of human impacts mediates coral reef conservation gains. *Proceedings of the National Academy of Sciences,* 115(27), E6116–E6125.

Graham, N., Nash, K. 2013. The importance of structural complexity in coral reef ecosystems. *Coral Reefs,* 32(2), 315-326.

Graham, N., Jennings, S., MacNeil, A., Mouillot, D., Wilson, S. 2015. Predicting climate-driven regime shifts versus rebound potential in coral reefs. *Nature,* 518 (7537), 94–97.

McCook, L., Jompa, J., Diaz-Pulido, G.. 2001. Competition between corals and algae on coral reefs: A review of evidence and mechanisms. *Coral Reefs,* 19(4), 400-417.

McManus, J., Rodolfo, B., Reyes, J., Cleto, L., Nanola, J. 1997. Effects of some destructive fishing methods on coral cover and potential rates of recovery. *Environmental Management*, 21, 69-78.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. London, England: Morgan Kaufmann.

Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Shrier, I., Platt, R. 2008. Reducing bias through directed acyclic graphs. BMC Medical Research Methodology, 30(8), 70.

Soler, G., Edgar, G., Thomson, R., Kininmonth, S., Campbell, S., Dawson, T. et al. 2015. Reef fishes at all trophic levels respond positively to effective marine protected areas. *PloS One,* 10(10), e0140270.

Staiger, D., Stock, J. 1997. Instrumental variables regression with weak instruments. *Econometrica: Journal of the Econometric Society,* 65(3), 557.

Stevens, B. 2021. The ups and downs of traps: Environmental impacts, entanglement, mitigation, and the future of trap fishing for crustaceans and fish. *ICES Journal of Marine Science,* 78(2), 584-596.

Strain, E., Edgar, G., Ceccarelli, D., Stuart-Smith, R., Hosack, G., Thomson, R. 2019. A global assessment of the direct and indirect benefits of marine protected areas for coral reef conservation. *Diversity & Distributions,* 25(1), 9-20.

Textor, J., Zander, B., Gilthorpe, M., Liskiewicz, M., Ellison, G. 2016. Robust causal inference using directed acyclic graphs: The R package 'Dagitty.' *International Journal of Epidemiology,* 45(6), 1887-1894.

Vergés, A., Vanderklift, M., Doropoulos, C., Hyndes, G. 2011. Spatial patterns in herbivory on a coral reef are influenced by structural complexity but not by algal traits. *PloS One,* 6(2), e17115.

Zeldow, B., Hatfield, L. 2021. Confounding and regression adjustment in difference-in-differences studies." *Health Services Research,* 56(5), 932-941.

.

CHAPTER 6 REDUCING BIAS IN EXPERIMENTAL ECOLOGY THROUGH GRAPHICAL CAUSAL MODELS

Authors: Suchinta Arif, Melanie Massey, Aaron MacNeil

6.1 ABSTRACT

Ecologists often rely on randomized control trails (RCTs) to prove causal relationships in nature. Many of our foundational insights of ecological phenomena can be traced back to well-designed experiments, and RCTs continue to provide valuable insights today. Although RCTs are often regarded as the "gold standard" to causal inference, it is important to recognize that they rely on a set of causal assumptions that must be justified and met by the researcher to draw valid causal conclusions. For example, common spurious associations found in observational correlative studies can also plague experimental results due to improper study design and/or statistical analysis. Here we use key ecological examples to show how biases such as confounding, overcontrol, and collider bias can occur in experimental set-ups. In tandem, we highlight how such biases can be removed through the application of the structural causal model (SCM). The SCM framework visualizes the causal structure of a system or process under study through directed acyclic graphs (DAGs) and subsequently applies a set of graphical rules to remove bias from both observational and experimental data. Here, we highlight how DAGs can be used across future ecological experimental studies to ensure proper study design and statistical analysis, leading to more accurate causal estimates drawn from experimental data.

6.2 INTRODUCTION

Experiments are a fundamental tool ecologists use to quantify causal relationships, in particular, by using randomized control trials (RCTs) that often regarded as the "gold standard" for causal inference (e.g., Hariton and Locascio 2018). Under RCTs,

researchers randomly assign units or individuals into treatment and control groups to eliminate potential confounding between treatment assignment and outcome, thereby increasing the internal validity of experiments (Rubin 1974; Holland 1986). Many of the foundational insights in biology were discovered through experiments. For instance, during early exploration of the scientific method, Francesco Redi (1626 - 1698) famously conducted his 'fly experiments' to test the theory of spontaneous generation (Gottdenker, 1979). Redi designed an experimental setup in which there were eight identical flasks containing meat; he tightly sealed four of these flasks and left four uncovered, yielding 'treatment' and 'control' groups (Gottdenker, 1979). In contrast to previously held beliefs that maggots were created within dead flesh itself, Redi's experiment revealed that only that meat which was exposed to incoming flies would eventually produce maggots, drawing the causal conclusion that, for maggots to form, "live animals must… deposit their seeds" (Gottdenker, 1979). In addition to the fundamental observation *omne vivum ex vivo* ("all life comes from life"), such experimentation would ultimately cascade into more complex tests of the causal relationships in the natural world.

Although RCTs have been invaluable in understanding numerous causal relationships in ecology, they are nonetheless susceptible to biases that can lead to erroneous causal conclusions. For example, Kimmel et al. (2021) discuss four core causal assumptions required for valid causal inference in experimental biology. These include excludability, which is the assumption that the process by which treatments are assigned has no effect on the outcome. Other causal assumptions include no interference between units, no multiple versions of treatment, and no compliance, meaning that all units receive the treatment they were assigned. Other studies have noted that RCTs can suffer from lack of generalizability, for example, because ecological treatments may not accurately represent actual ecological phenomena (e.g., Korell et al. 2019). Further, statistical approaches that have recently received criticism across observational ecological studies are also prevalent across experimental ecological studies. For example, many experimental studies employ predictive model selection techniques such as Akaike's information criterion (AIC; Akaike 1973) to select the best model for analysis (e.g., Sato et al. 2011; Cameron et al. 2013; Hunyadi et al. 2020); others place all predictor variables of interest as well as

potential confounders into one statistical model for analysis (McElreath 2020). Such approaches have been shown to be unreliable for drawing causal conclusions (McElreath 2020; Arif and MacNeil 2022). Although causal conclusions drawn from RCTs are often not questioned, biases may still arise, either through study design and/or statistical analysis. However, there is currently no unified framework that is being employed to ensure accurate causal conclusions are drawn across RCTs in ecology.

Pearl's structural causal model (SCM; Pearl 2009) is a causal inference framework that has recently been highlighted in the ecological literature as a tool for determining causal relationships from observational data (Schoolmaster et al. 2020; Laubach et al. 2020; Arif et al. 2021; Arif and MacNeil 2022). The SCM framework relies on graphical causal models, in the form of directed acyclic graphs (DAGs), to visualize hypothesized causal relationships between variables of interest, identify potential biases, and guide appropriate study design and statistical analysis required for causal inference. What has received significantly less attention is that within the SCM framework, DAGs can also be used to reduce bias across RCTs by critically visualizing the causal structure of an experimental set-up (e.g., Schoolmaster et al. 2020). Here we overview how DAGs can reduce common biases across RCTs and advocate for their widespread uptake across experimental studies.

6.3 THE SCM FRAMEWORK

The SCM framework uses DAGs to represent the causal structure of a system or process under study. DAGs consist of variables (nodes), that are connected to each other via directed arrows, pointing from cause to effect. These directed arrows communicate a causal relationship between two variables but make no assumptions about the functional form or effect size (Glymour and Greenland 2008). DAGs must also include both measured and unmeasured variables required to depict the complete causal structure of a system or process (see Cronin et al. 2018; Schoolmaster et al. 2020 for complete examples of ecological DAGs).

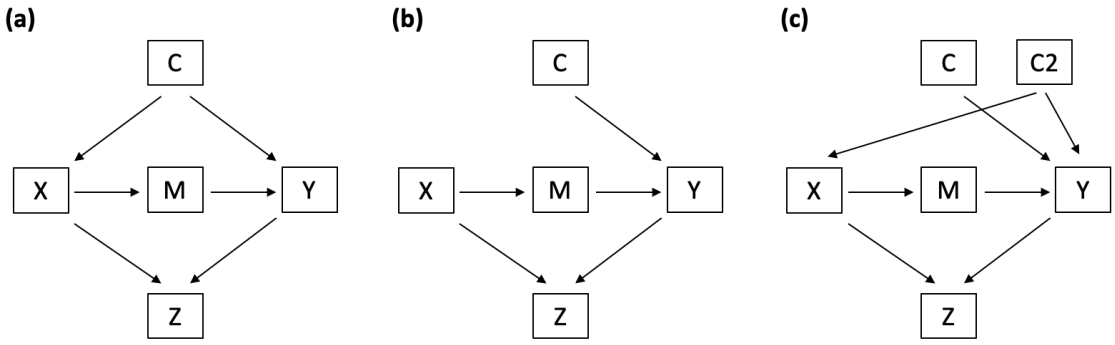

Figure 6.1 Three directed acyclic graphs (DAGs) representing (a) the causal structure in a natural setting, with a confounder (C) affecting both the variable of causal interest X and the outcome Y (b) the causal structure under a perfectly executed RCT, which breaks the association between C and X and (c) the causal structure under a RCT that introduces additional confounding from variable C2.

As an example, Figure 6.1a shows a DAG representing a natural system whereby X affects Y through mechanism M, C affects both X and Y, and both X and Y affect Z. Under the SCM framework, if we want to determine the effect of X on Y, we can apply a graphical rule known as the **backdoor criterion** to determine which variables may need to be controlled for to answer a specified causal question (Pearl 2009). In estimating the causal effect of X on Y, the backdoor criterion instructs us to block all backdoor paths between X and Y (i.e., our predictor and response variable of interest). Backdoor paths are sequences of nodes and arrows between X and Y with an arrow pointing into X; if left open, they can induce spurious (non-causal) associations between X and Y, biasing estimates. To block a backdoor path, we can either (1) control for an intermediate arrow-emitting variable or (2) not control for a variable with two incoming arrows (i.e., a collider variable, such as Z) in the pathway. To determine the effect of X on Y, given our DAG in Figure 6.1, there is one backdoor path that needs to be blocked: $X \leftarrow C \rightarrow Y$. To block this pathway, we can control for the arrow-emitting variable C. There are several ways to control for a variable, including experimental control, as well as statistical techniques including covariate adjustment, stratification, and restriction (Williams et al. 2018).

The backdoor criterion was created to remove non-causal associations that often plague observational studies, including confounding, collider, and overcontrol bias.

Confounding bias occurs when a variable that affects both the predictor and response is not controlled for. Given our DAG in Figure 6.1a, to determine the effect of X on Y, we must control for C to remove confounding bias. Here, not controlling for C would leave the backdoor path (X ← C → Y) open, leading to a non-causal association between X and Y. **Collider bias** occurs when both predictor and response affect a third common variable (or its descendant), and that variable (known as a collider variable) is controlled for. To determine the effect of X on Y, we must avoid controlling for Z. Here, controlling for Z opens a non-causal pathway (X → Z ← Y), leading to non-causal associations between X and Y (Figure 6.1a). **Overcontrol bias** occurs when an intermediate variable along a causal pathway between predictor and response is controlled for, blocking the indirect causal association between treatment and response. To determine the effect of X on Y, we must not control for M (Figure 6.1a). Here, controlling for M closes a causal pathway (X → M → Y), removing this causal association between X and Y.

A perfectly executed RCT should remove all backdoor paths between treatment and outcome through randomization. Fig 1b represents our previous DAG under a perfectly RCT where treatment X is controlled and randomized. The arrows pointing into X are removed under the assumption that only the experiment determines the value of X. Under this scenario, there are no backdoor paths that need to be blocked (because C no longer affects X), and the effect of X on Y can be estimated without bias. However, ecological experiments can often diverge from perfectly executed RCTs (Williams et al. 2018; Schoolmaster et al. 2020; Kimmel et al. 2021) and backdoor paths may be open, for example, due to additional confounding variables that arise from an imperfect treatment assignment process. For example, in Figure 6.1c, the treatment assignment process led to an additional confounding variable, C2, that affected both treatment assignment X and outcome Y. Bias can also arise from improper statistical analysis of experimental data. For example, controlling for M in Figure 6.1b and Figure 6.1c will lead to overcontrol bias, whereas controlling for Z in Figure 6.1b and Figure 6.1c will lead to collider bias. By visualizing ecological experiments through DAGs, researchers can ensure that common biases including confounding, collider, and overcontrol bias are accounted for, allowing for more accurate causal conclusions to be drawn from experiments.

Below we present DAGs representing ecological experiments to show how these biases can arise and provide solutions for how to avoid them. We further show how DAGs can be used to assess external validity, focusing on the extent to which RCTs can be generalizable to real world scenarios. The case studies and associated DAGs depicted in this paper are simplified and used for illustrative purposes. We refer readers to Chapter 2 for a comprehensive overview of creating complete DAGs for ecological research.

6.4 Confounding Bias

Ecologists are aware that confounding bias can often plague observational studies; however, with RCTs it is often assumed that the randomization process will eliminate confounding. To break any confounding between treatment assignment and outcome the excludability assumption must be met (Kimmel et al. 2021). Excludability assumes that the process by which treatments are assigned has no effect on the outcome except through its effects on variation in treatment. However, the process of treatment assignment across ecological experiments can often lead to the excludability assumption being violated, subsequently leading to confounding bias; below we present two examples:

6.4.1 Biodiversity-Ecosystem Function (BEF) Experiments

Hundreds of experiments have been carried out in the hopes of understanding the causal relationship between biodiversity metrics and various ecosystem functions (reviewed in Loreau et al. 2001). Although the drivers of ecosystem functioning are numerous and often interconnected, authors of BEF experiments rarely communicate the overall causal structure of their study system, or the causal assumptions required for valid causal inference given their experimental set-up (Schoolmaster et al. 2020). However, this is a necessary step as BEF experiments may be prone to erroneous conclusions (Schoolmaster et al 2020).

As a classic example, the Cedar Creek grassland experiments (Tilman et al. 1996) sought to determine the effect of plant species richness on productivity. In this study, each

experimental unit was a plot containing 1 - 24 species that were planted from seeds, forming a biodiversity gradient that ultimately represented the treatment. The community of species within each plot was established by randomly drawing from a candidate pool of 24 possible prairie species. A given species therefore had a 1/24 chance to be drawn into a plot with a species richness of 1, a 6/24 chance of being drawn into a plot with a species richness of 6, and so forth. It was thus assumed that the community within each plot was fully 'randomized'. Care was notably taken to ensure plots were otherwise similarly treated (*i.e.*, free of previous wild vegetation, consistent and equal weeding) throughout the experiment. After a commendable two growing seasons of experimental maintenance, Tilman et al. (1996) sampled plant biomass and concluded there was a positive causal relationship between species richness and productivity.

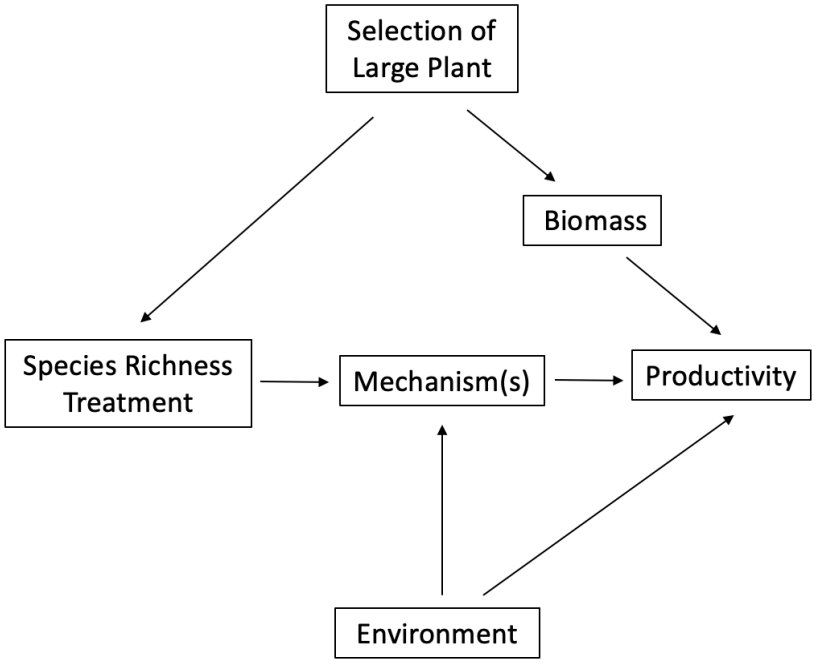

Figure 6.2 A simplified DAG representing confounding bias in a biodiversity-ecosystem function (BEF) experiment. The directed arrow from environment to species richness, which would otherwise exist in nature is removed due to the experimental treatment assignment process. However, the treatment assignment process induced an additional confounding variable, whereby selection of large plants into a treatment differentially affected high vs low species richness treatments as well as affected the productivity outcome.

Although this experiment was carefully designed, a subtle bias known as the 'selection probability effect' may have confounded the results of this study (Figure 6.2; Huston 1997). The selection probability effect occurs when there is an increasing chance of selecting a species with a specific trait as the number of sampling events increases. With respect to Tilman et al. (1997)'s study, the selection probability bias was evident as size variation existed among the 24 candidate species, and large species were more likely to be drawn into high species richness treatment plots, differentially impacting treatment assignment (Figure 6.2). Further, since plant communities are typically dominated by individuals from large species (Grime 1979), productivity data gathered from treatments with overrepresentation of large, dominant plants likely reflected effects of those dominant species, rather than species richness itself (Huston 1997). Specifically, treatments with large plants positively affected plant biomass, which in turn positively affected productivity (Figure 6.2). Therefore, the positive correlation between species richness and productivity found in this study may be due to confounding bias resulting from large plant species affecting both treatment assignment and outcome. Graphically, this is represented by a backdoor path between treatment assignment and productivity (productivity ← biomass ← selection of large plant → species richness treatment) being left open, leading to confounding bias. Although other issues with this study have also been noted (Huston 1997), this particular issue could be resolved by a study design that samples from plant species of similar height, removing the selection probability effect.

BEF experiments continue to be highly utilized to understand complex ecological relationships. At the same time, some authors have highlighted biases that may arise across BEF experiments (e.g., Huston 1997; Mora et al. 2014; Schoolmaster et al. 2020; Kimmel et al. 2021). The future uptake of DAGs within BEF studies can transparently communicate the overall causal structure of a system under study and ultimately identify and hopefully resolve any potential biases that may be at play.

6.4.2 Transgenerational Experiments

As mounting evidence suggests evolution may play a limited role in organismal responses to rapid climate change, experimental biologists have increasingly placed emphasis on phenotypic plasticity as a means of coping with climate impacts (e.g., Merila & Hendry, 2014; Seebacher et al., 2015). One form of phenotypic plasticity that is expected to contribute to organismal responses is transgenerational plasticity (TGP; sometimes termed 'anticipatory parental effects' [Marhshall & Uller, 2007]), whereby ancestral environments influence the phenotypic responses of subsequent generations non-genetically (Salinas et al., 2013; Donelson et al., 2018).

Transgenerational experiments are necessarily complex, given that an ancestral (F0) generation must be reared to sexual maturity under the desired conditions, reproduce, and then the responses of subsequent (F1, F2, etc.) generations must be recorded. Throughout the experiment, there is a risk of unexpected variables impacting the assignment of individuals into the F1 (or later) treatment group and response simultaneously.

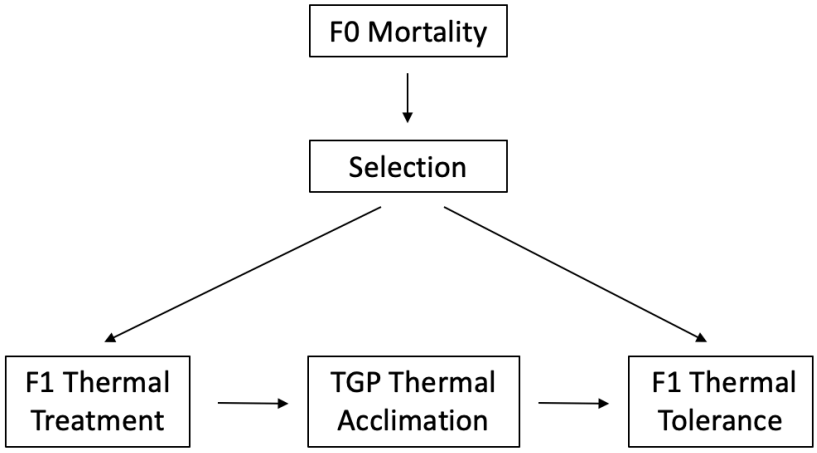

Figure 6.3 A DAG representing a transgenerational plasticity experiment, whereby differential mortality under the ancestral (F0) generation treatment leads to selection differentially affecting the subsequent (F1) treatments.

If selection exerts significant effects on both treatment assignment and response, an over- or underestimation of the strength of plasticity effects can occur. As an example, Zizzari & Ellers (2014) investigated TGP of heat tolerance in a collembolan arthropod. They

exposed F0 females to a significant heat shock, then bred them and estimated the heat tolerance of F1 offspring (Zizzari & Ellers, 2014). Although the proportion of successfully reproducing females was similar for heat shock and control treatments, a notable result of this study is that the mortality rate of heat shocked mothers was greater than double that of control mothers (19% vs. 8%). As shown in the DAG in Figure 6.3, mortality of F0 fish differentially affected selection across F1 treatment group, with potentially higher selection in the heat shock treatment. Selection may have also affected the outcome of interest, with higher selection increasing F1 thermal tolerance (Figure 6.3). In other words, the backdoor path between F1 thermal treatment and F1 thermal tolerance (F1 thermal tolerance ← selection → F1 thermal treatment) is left open, leading to confounding bias. Ultimately, F1 offspring whose mothers had greater genetic capacity to tolerate heat shock may have been overrepresented in the heat shock treatment, potentially leading to an overestimation of the strength of TGP.

In such cases as these, researchers at a minimum should be explicit in acknowledging whether their experimental treatments were subject to differential selection, and clearly rationalize how selection may have affected their conclusions (*see* Donelson et al., 2016 for an example of a clear explanation). Authors may also opt to reduce differential selection by decreasing the magnitude of treatment (e.g., reducing treatment-induced stress), or incorporate estimates of genetic effects into their statistical framework to better isolate plastic treatment effects (Merila & Hendry, 2014). Authors should be conscientious in recording treatment-dependent metadata (*e.g.,* mortality) to make informed decisions about potential confounders.

6.5 Collider Bias

Collider bias occurs when both the treatment and outcome each affect a third 'collider' variable (or its descendant), that when controlled for, leads to a non-causal association between treatment and outcome. A common way for collider bias to occur under RCTs is if both the treatment and outcome affect whether an individual or unit is included in the final analysis of a study. For example:

6.5.1 Survivorship Bias

Data analyzed from RCTs are sometimes limited to individuals or units that have survived the full term of an experiment. For example, Lusk and Del Pozo (2002) conducted an experiment to quantify the growth rates of rainforest trees under low-light and high-light environments. Seedlings from 12 Chilean rainforest tree species were grown under both low- and high-light environments, and relative growth rates (RGR) of individual plants were measured 5-6 months following the experiment. Their results showed that RGR in high-light treatment were consistently higher than low-light treatment across the 12 species.

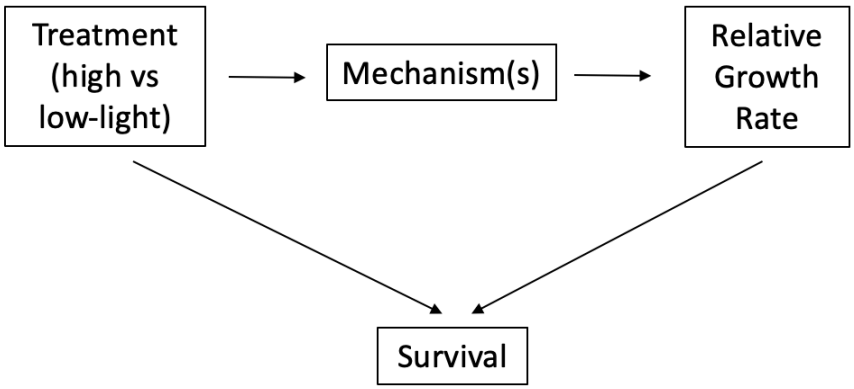

Figure 6.4 A simplified directed acyclic graph (DAG) representing collider bias in a RCT. Here, both the treatment assignment (high vs low-light conditions) and the outcome (relative growth rate) affected which plant individuals survived until the end of the experiment. Only analyzing data from plants that survived until the end of the experiment will essentially control for this collider variable. This in turn will induce a non-causal spurious correlation between treatment and outcome, leading to collider bias.

In this example, collider bias occurred since RGR was measured only for plants that survived until month 5-6 of the experiment. The study noted that mortality rates were significantly higher in low-light conditions. As well, mortality risks tend to be higher for slow-growing plants (i.e., those with lower RGR) in a population (Kobe et al. 1995). Thus, as shown by the DAG in Figure 6.4, both the treatment status (low- vs. high-light) and RGR outcome affected whether an individual plant survived long enough to be included in the final analysis. In other words, the collider variable ‘survival’ (representing

plants that survived until end of experiment) was controlled for, leading to a non-causal association between treatment and outcome (Figure 6.4). As such, low-light growth rates may be overestimated in this experiment, as only the 'winners' from low-light conditions were assessed.

Analyzing any subset of experimental data may lead to collider bias if the subset of data is affected by both the treatment and response. A second example is overviewed in Williams et al. (2018), where an RCT investigating the effect of promoting breastfeeding on child cognitive development led to collider bias when only participants who attended a post-treatment follow-up were analyzed; here, both the treatment assignment and outcome affected the likelihood of participants following up. In general, researchers should be conscientious not to control for post-treatment variables that are influenced by both the treatment and outcome. Ultimately, utilizing DAGs allow researchers to visualize whether their study design and/or data analysis may lead to collider bias.

6.6 Overcontrol Bias

Overcontrol bias occurs when an intermediate variable along a causal pathway between treatment and outcome is controlled for. Unlike confounding and collider bias, which induces non-causal associations, overcontrol bias removes indirect causal associations between treatment and outcome.

6.6.1 Intermediate Variables in Temperature Experiments

Temperature is one of the main drivers of biological functions across numerous levels of organization, influencing biotic enzyme kinetics, whole-organism physiology, population growth and distribution, and even species interactions (e.g., Wieser 1973). Given the numerous ways in which temperature can affect an outcome of interest, it is crucial to understand when variables act as mechanisms along a causal pathway. Many temperature experiments aim to reduce bias by controlling for additional variables; knowing when not to control for a variable because it is part of a causal pathway can in turn reduce overcontrol bias.

As an example, Lienart et al. (2014) conducted an experiment examining the impact of temperature and food availability on risk behaviour in fish. They collected wild juvenile *Pomacentrus chrysurus*, then randomly allocated them to one of four treatments, each representing a combination of two feeding levels and two temperature treatments. After 5 days of acclimation under experimental conditions, risk behavior was assessed.

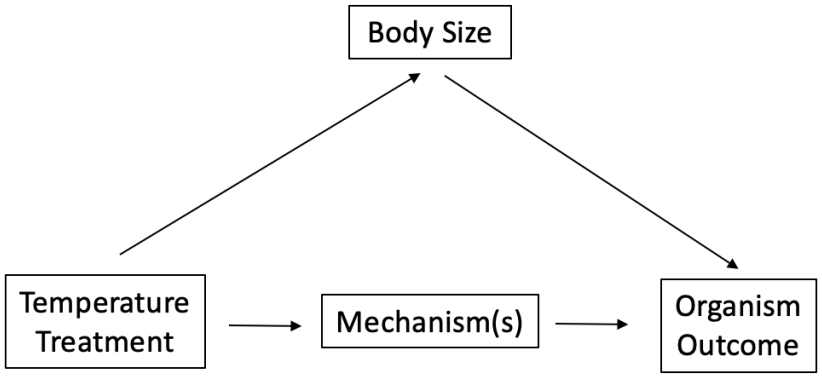

Figure 6.5 A generalized directed acyclic graph (DAG) representing the effect of temperature treatments on an organism outcome (e.g., risk behaviour). Here, body size acts as an intermediate variable between treatment and outcome. If body size is controlled for, it will lead to overcontrol bias, removing this indirect causal association between treatment and outcome.

To determine the effect of temperature and food on risk behaviour, the authors included size as a covariate in their analysis. They noted that "the manipulation of both temperature and food could have resulted in a difference in the size of the fish, which could potentially affect their antipredator response". This rationale implies that body size acts as an intermediate variable between treatment and response (Figure 6.5), and thus should not have been controlled for to determine the total effect of treatment on outcome. Although the authors ultimately found a negligible effect of size, this study nonetheless highlights the misconceptions that experimental ecologists may have about controlling mechanisms along a causal pathway. If the experiment took place over a longer timeframe, with temperature and food availability influencing risk behaviour through body size, controlling for body size would have likely led to overcontrol bias.

Variation in how researchers deal with overcontrol bias can be seen across other experimental studies that have investigated the impact of temperature of fish populations. For example, Spinks et al. (2020) conducted an RCT to investigate the parental effect of warming on reproduction and offspring quality. In their analysis, they decided whether to include post-treatment mother size (an intermediate mechanism between treatment and outcome) as a covariate based on the leave-one-out cross-validation information criterion (LOOIC), a predictive model selection technique. However, model selection techniques are meant for predictive inference (i.e., what is the best model to predict Y?) and should not be conflated with causal inference (i.e., what is the effect of X on Y?). In fact, predictive model selection techniques can often lead to overcontrol as well as other forms of bias (Arif and MacNeil 2022).

In contrast, an experimental study investigating the effects of ocean warming in marine sticklebacks noted that they "did not include egg size as a covariate as egg size is an intermediate variable that may have been affected by temperature treatments in the F0 and F1 generations" (Shama and Wegner 2013). Here, authors recognize that controlling for a mechanism should be avoided if looking for the overall effect of a treatment on outcome. Controlling for an intermediate variable is also valid if researchers are not interested in that particular causal pathway. For example, if we wanted to know the direct (vs. total) effect of temperature on outcome, then given our DAG in Figure 6.5, we would control for body size to remove the effect of this pathway.

Some researchers may include an intermediate variable as a covariate because they are also interested in their causal effect on the outcome. For example, the influence of body size may be of fundamental biological interest, even if the influence of temperature is the primary question in the study (e.g., Fuxjager et al. 2019). However, in such cases, a separate causal model should be built for each predictor variable of interest, following the backdoor criterion. Given our DAG in Figure 6.5, to determine the effect of treatment on outcome, no additional covariates need to be controlled for, as there are no backdoor paths that need to be blocked. However, to determine the effect of body size on outcome, the backdoor path outcome ← mechanism → treatment → body size can be blocked by

either controlling for ‘treatment’ or ‘mechanism’. This could be achieved, for example, by statistically adjusting for either treatment or mechanism.

A formal method for acknowledging and avoiding overcontrol bias can benefit experimental ecologists and lead to more informed experimental conclusions. DAGs allow researchers to visualize when variables may act as part of a causal pathway, subsequently allowing them to justify their exclusion or inclusion as a covariate in their analysis.

6.7 External Validity

External validity represents the degree to which results of an experiment can be generalized to subjects and situations outside of the experimental set-up (Shadish et al. 2002). RCTs can be conducted under artificial conditions and are sometimes critiqued for having low generalizability (Shadish et al. 2002). Ecologists have previously highlighted the ways in which external validity can be increased across RCTs, for example by conducting field experiments that are employed under natural settings, or replicating experiments across settings, populations, and conditions to determine whether results can be generalizable. Here, we highlight how DAGs can be used to visualize how experimental conditions may systematically differ from real-world conditions (*e.g.,* Massey & Hutchings 2020). By visualizing the mechanisms at play under experimental set-ups, researchers can more effectively communicate how causal structures may differ between experimental vs natural settings, and how this in turn may affect the causal conclusions drawn across experimental studies.

6.7.1 The Obfuscating Influence of Static Treatments

There is a growing interest in experimental biology to use treatments that better reflect natural conditions, rather than contrived or static conditions (Morash et al., 2018; Massey & Hutchings, 2020). In nature, abiotic conditions such as temperature, dissolved oxygen, salinity, moisture, and light are rarely static; instead, they vary both temporally (diurnally, seasonally, and stochastically) and spatially. Despite this, experiments often

compare static treatment conditions against one another, potentially resulting in data that lack ecological relevance, or are otherwise obfuscated by static condition-imposed pathologies.

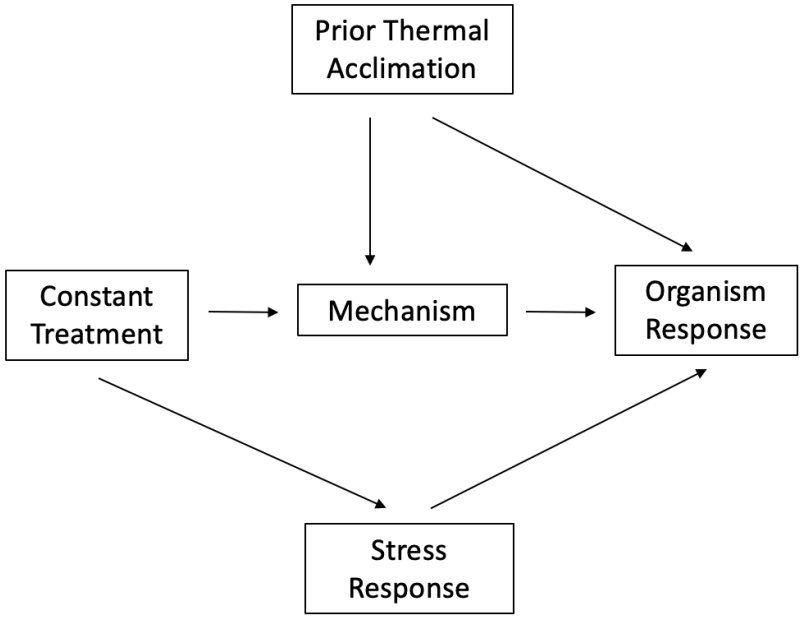

Figure 6.6 A generalized directed acyclic graph (DAG) representing how, in thermal biology, both prior thermal acclimation of an organism to static conditions and stress responses induced by static treatment conditions can influence outcomes, and thus the external validity of experiments.

In thermal biology, this problem has been explored at length, especially in the context of thermal acclimation (Angilletta, 2009). Several lines of evidence suggest that the use of constant temperatures may have serious repercussions on both individuals and even entire populations of experimental organisms. For instance, a recent study by Morgan et al. (2022) found that the laboratory-reared zebrafish, which have been kept at constant 28 °C for over 150 generations, have significantly limited capacity to plastically respond to thermal conditions when compared to lines of wild-caught zebrafish. Moreover, this acclimation capacity was limited at all levels of biological organization, from genetic, to physiological, to behavioural plasticity (Morgan et al., 2022). These findings ultimately challenge the generalizability of conclusions garnered from lab-reared zebrafish, a model organism used in at least 17 151 studies as of 2013 (Kinth et al., 2013). In this situation,

the evolutionary history of thermal adaptation to unnatural constant temperatures may affect the mechanisms and responses to treatments (Figure 6.6). Here, although the treatment effect can be quantified without bias (i.e., there are no open backdoor paths between treatment and organism response), the causal conclusions drawn will not be generalizable to what would be expected under a natural setting. The generalizability of conclusions from these experiments will thus depend on the strength of the effects of prior thermal adaptation or acclimation, and researchers should think critically about how recorded responses may differ from those in the natural system they are trying to represent.

The generality of constant temperature experiments has also widely been criticized due to the possibility of treatments imparting unintended pathologies, especially under stressful conditions (Wilson and Franklin, 2002). Natural organisms have evolved to respond to changing environments and consequently are expected to perform better when a stressor is applied at a natural time scale rather than through chronic exposure (Angilletta, 2009; Colinet et al., 2015). For example, Kingsolver et al. (2016) demonstrated that at hot constant temperatures, organismal growth becomes limited due to pathological increases in molecular coping mechanisms (*e.g.*, heat shock proteins), which reduces energy resource availability in the growth pathway. When a model based on constant temperature performance is subsequently applied to estimate growth under natural, fluctuating temperatures, growth is underestimated (Kingsolver et al., 2016; but see also: Rollinson et al., 2018). Therefore, the outcomes measured in constant temperature treatments may themselves be subject to the influence of additional and unintended mechanisms such as a stress response (Figure 6.6), and do not reflect what is expected in nature.

Although approaches that modify constant temperature models to extend their applicability have been developed (*e.g.,* controlling for Stress Response in Fig. S1; Kingsolver et al., 2016; Koussoroplis et al., 2017), many authors now advocate for the use of more ecologically relevant temperature regimes in experimental biology as a

means of generating realistic responses and conclusions (*e.g.,* Massey & Hutchings, 2020; Morash et al., 2018; Taylor et al., 2021).

6.8 Concluding Remarks

Causal diagrams in the form of DAGs are starting to gain traction across ecological observation studies but have yet to be applied in experimental ecology. Although causal conclusions drawn from RCTs are often taken at face value, ecologists are increasingly becoming aware that for causal inference to be valid, experimental approaches must be carefully designed and analyzed to avoid potential biases (Kimmel et al. 2021). By routinely utilizing DAGs, researchers can visualize and subsequently reduce confounding, collider bias and overcontrol bias across experimental studies in a coherent way. We further highlight how DAGs can be used to assess external validity of experiments by visualizing how mechanisms may differ between experimental set-ups and the natural world. By utilizing DAGs as a visual and conceptual tool, experimental ecologists can increasingly meet the causal assumptions required for valid causal inference. DAGs also allow researchers to transparently communicate their causal assumptions to others, which can facilitate more critical reception and lead to productive scientific debates that collectively deepen our understanding of ecological phenomena over time (e.g., Schoolmaster et al. 2021). Moreover, DAGs allow researchers to use their ecological domain knowledge, above all else, to build causal models, bridging the gap between ecological knowledge and statistical analysis. Ultimately, the uptake of this causal inference tool can significantly benefit experimental design, statistical analysis, and interpretation of results across experimental ecology.

CHAPTER 7 CONCLUSION

Understanding causal relationships in nature remains a central goal of ecology. It is therefore important that ecologists are equipped with an understanding of the tools required for valid causal inference. The structural causal model (SCM) framework provides a set of widely applicable tools that can be used for causal inference under both observational and experimental settings (Pearl 2009). The overarching aim of this thesis was to introduce the SCM framework to a general ecology audience as a unifying and widely applicable causal inference tool, in the hopes of increasing its uptake and relevance within the field.

7.1 A Clear Path for Observational Causal Inference in Ecology

The field of ecology has always relied on observational data (Elton 1927), with observational studies often aimed at answering causal questions in nature. For example, many of today's applied ecological problems, such as quantifying the impact of climate change and other anthropogenic stressors on ecosystems across the globe, are reliant on large-scale observational data that is not easily manipulated or controlled (Sagarin and Pauchard 2010). However, a fundamental limitation of our field is the persistent lack of causal inference methodologies applied across observational studies. The slow uptake of observational causal inference can be attributed to the combined influence of Karl Pearson and Ronald Fisher, prominent eugenicists who promoted the narrative that non-causal spurious correlations cannot be removed from observational analysis (Pearson 1911) and that randomized experiments were the only way to attain valid causal estimates (Fisher 1925), in part due to their desire to appear to be objective (Clayton 2021). This dark history created a culture whereby ecologists avoid explicitly acknowledging the causal goal of their observational research and instead use coded language that implies causality without explicitly saying so. Importantly, observational causal inference techniques are not often taught to ecologists, and developed methodologies are rarely applied within the field. Instead, common approaches used to understand causal

relationships include the ‘causal salad’ model (Bhalla 2018; McElreath 2020) as well as predictive model selection techniques (Arif and MacNeil 2022), both of which can lead to spurious non-causal associations between variables of interest. This in turn limits the credibility of causal results and conclusions drawn from observational data, reaffirming the often-repeated phrase “correlation does not imply causation” (F.A.D. 1900). The irony of this is that correlation can imply causation if valid causal inference methodologies are increasingly applied.

The current “causal revolution” being led by computer scientist Judea Pearl offers a deep potential to transform the way ecologists think about causality, particularly as it relates to observational data. Pearl’s SCM framework relies on a researcher’s domain knowledge (based scientific consensus, prior studies, expert opinion, and other informed sources), above all else, to represent the causal structure of a system or process under study, using directed acyclic graphs (DAGs; Pearl 2009). The application of graphical rules, including the backdoor and frontdoor criteria, then allows researchers to determine which variables need to be controlled for to answer causal queries from observational data (Pearl 2009). The mathematical underpinnings of the framework show that (1) valid causal conclusions can be drawn from both experimental and observational data, (2) both experimental and observational approaches require a set of causal assumptions that must be met and justified by the researcher, and importantly, (3) experimental approaches are not necessarily superior to observational approaches to causal inference (Pearl 2009). Ultimately, the validity of causal conclusions is determined by the validity of a researcher’s causal assumptions, encoded in their DAG.

A fundamental aim of this thesis was to introduce the SCM framework as a widely applicable observational causal inference tool to a general ecological audience. To this end, Chapter 2 provided a review of the SCM framework, geared towards a general ecology audience interested in observational causal inference. This chapter used simulated ecological examples to highlight the steps within the SCM framework, including how to create and finalize a DAG, as well as how to apply the backdoor and frontdoor criterion to construct causal models and subsequently draw causal conclusions

from observational data. It further defined and highlighted how biases, such as confounding, overcontrol, and collider bias often occur in observational correlative studies, and how they can be eliminated through the application of the SCM framework. Although a few studies have highlighted the utility of the SCM framework in the ecological literature, they are either presented as niche, complex and system specific (e.g., Cronin and Schoolmaster 2018; Schoolmaster et al. 2020) or are presented to a general ecological audience with missing essential steps (e.g., the review by Laubach et al. 2021 does not present the backdoor or frontdoor criterion or mention DAG-data consistency checks). Chapter 2 is the first to our knowledge to provide a complete and comprehensive guide of the SCM framework geared towards a general ecology audience, and therefore holds great potential to increase its uptake within ecological observational studies.

The few SCM ecological studies that have been published to date are predominantly reliant on either theoretical or simulated data (e.g., Cronin and Schoolmaster 2018; Schoolmaster et al. 2020; Grace et al. 2021; Wilson et al. 2021). However, it is critical for ecologists to understand how to apply the theoretical concepts of the SCM framework to real-world and often messy ecological data. To this end, Chapters 3 and 4 were the first to apply the SCM framework to answer causal questions in reef ecology, at both localized and global scales, respectively. They provide two of the first examples of how this framework can be applied to real-world ecological data and provide practical steps on how to critically examine both a DAG and available observational data to, for example, refine a DAG based on DAG-data consistency and subsequently draw causal conclusions from observational datasets. Chapter 3 and 4 have also provided novel insights into reef ecology which contrasts with conclusions drawn from correlative observational studies on the same study system. This should further encourage ecologists to use the SCM framework over non-causal correlative studies, as it can lead to more accurate and dependable causal conclusions.

Collectively, Chapters 2, 3, and 4 provide a basis for understanding how to apply the SCM framework for observational causal inference in ecology, with the aim of

encouraging widespread application of this approach across observational ecological studies.

7.2 Strengthening Quasi-experimental and Experimental Approaches

Recently, quasi-experimental approaches including matching methods, before after control impact (BACI), regression discontinuity design (RDD), and instrumental variables (IV) have been promoted in the ecological literature (Butsic et al. 2017; Larsen et al. 2019; Wauchope et al. 2021). While they provide great potential and have been employed across several ecological studies, limitations in their application remain. Each method requires that specific causal assumptions be met by a researcher to ensure accurate causal estimates; however, these assumptions are often not met due to improper study design and or statistical analysis (e.g., Mansournia et al. 2013; Ferraro et al. 2019; Adams et al. 2019). For example, matching methods such as propensity score analysis require that all and only confounding variables are used in the matching process to answer a causal question at hand (Rosenbaum and Rubin 1983). However, these is no formalized process to decide which variables act as confounders, and the addition of accidental non-confounding variables can in fact increase bias instead of reducing it (Pearl 2009; Mansournia et al. 2013). Similarly, BACI, RDD, and IV each have a set of causal assumptions that must be met before valid causal conclusions can be drawn; however, no unifying approach exists to guide in ensuring this process. To this end, Chapter 5 provided the theoretical basis of how DAGs and the principles of the SCM framework can be used to ensure proper study design and statistical analysis across quasi-experimental approaches. Here, we used theory and simulated ecological examples to demonstrate how combining quasi-experimental approaches with the SCM framework can (1) produce more transparent communication about causal assumptions (2) ensure proper study design and statistical analysis and (3) result in a more comprehensive and robust framework for causal inference.

Similar to observational and quasi-experimental approaches, randomized control trails (RCTs) also rely on a set of causal assumptions that must be communicated and justified to ensure valid causal conclusions are drawn from experimental data (e.g., Kimmel et al. 2021). Although RCTs in ecology are often regarded as the "gold standard" for causal inference, with few questioning experimental results, they too are prone to biases that lead to non-causal spurious associations between variables of interest. To this end, using core ecological examples, Chapter 6 highlighted that biases found within observational studies, including confounding, overcontrol, and collider bias can also be prevalent across experimental settings through study design and/or statistical analysis. As a solution, Chapter 6 shows that through the routine application of DAGs and the SCM framework, experimental ecologists can critically examine whether biases may be present and apply appropriate steps to remove them.

Collectively, Chapter 5 and 6 presents a significant advance in the field as it is the first time the SCM framework has been presented as a unified approach to ensure proper study design and analysis across quasi-experimental and experimental ecological studies.

7.3 Increased Uptake and Subsequent Benefits of the SCM Framework

Collectively, the core chapters within this thesis aim to increase the uptake of the SCM framework across future ecological studies. By detailing the SCM framework to a general ecology audience (Chapter 2), providing real-world ecological examples of its application (Chapter 3 and 4), and highlighting its utility across both observational (Chapter 2, 3 and 4), quasi-experimental (Chapter 5) and experimental (Chapter 6) studies, the cumulative contents of this thesis can educate a wide range of ecologists, and hopefully inspire the increased uptake and relevance of the SCM framework within ecology.

The increased uptake of the SCM framework within ecology will have a myriad of benefits. As previously mentioned, causal conclusions drawn from observational data often do not employ any causal inference methodologies, and commonly applied

approaches (e.g., the 'causal salad' model and predictive model selection techniques) can lead to a variety of non-causal biases (e.g., confounding, overcontrol, collider bias) that result in non-causal associations between variables of interest (see Chapter 2). Such biases can also arise in quasi-experimental and experimental settings if researchers do not critically examine the causal assumptions required for their chosen approach (see Chapters 5 and 6). The application of the SCM framework can significantly reduce bias across the field, leading to more accurate causal conclusions about our natural world.

Importantly, visualizing the causal structure of a system or process under study through DAGs will lead to more transparent communication about a researchers' causal assumptions and subsequent causal conclusions. This is turn can lead to more critical reception and feedback of ecological studies, leading to productive discourse that increases our collective understanding of causal relationships in nature. Already, the application of the SCM framework by Schoolmaster et al. (2020) has led to their controversial conclusion that biodiversity-ecosystem function (BEF) relationships are non-causal associations. Their DAG transparently communicated their causal assumptions, which was subsequently critiqued by Grace et al. (2021), who provided their own version of a BEF DAG, asserting that BEF correlations are indeed causal. Schoolmaster et al. (2021) followed up re-asserting their original conclusions. Although this debate has yet to be resolved, DAGs can continue to facilitate transparent and productive discourse between researchers about BEF correlations as well as other ecological phenomena.

Another key benefit of the SCM framework is that it allows researchers to bridge the gap between ecological domain knowledge and statistical models. Instead of creating statistical models based on a 'causal salad' approach, automated criteria such as AIC, or other technical approaches that rely on data-driven techniques, the SCM framework is theory-driven and relies on a researcher's domain knowledge, above all else, to answer causal queries. As noted by Pearl, "Data do not understand causes and effects; humans do (Pearl and Mackenzie 2018)." The SCM framework requires that researchers use their ecological domain knowledge to create and fine-tune DAGs, which in turn are used to

guide causal models. This should encourage ecologists intimidated by complex modelling, knowing that knowledge about how the natural world works resides with them.

7.4 Limitations of the SCM Framework

It is important to note that the SCM framework should not be the only causal inference tool ecologists are familiar with. Although we have shown that this framework can be integrated with other causal inference methodologies (including quasi-experimental approaches and RCTs) to improve causal estimates, it may not be well suited under some situations. For example, this framework is not well-suited to time-series data, which may benefit from other causal inference approaches such as convergent cross mapping, created specifically for complex ecological time-series data (Sugihara et al. 2012). The transdisciplinary field of causal inference is forever growing and consists of a variety of tools and approaches ecologists can utilize to estimate causal effects (Imbens and Rubin 2015; Morgan and Winship 2015). The best approach or set of approaches will ultimately depend on the available data and causal question at hand. It is therefore recommended that all ecologists interested in causal queries familiarize themselves with a suite of causal inference techniques and consistently update their knowledge based on ongoing development within this field.

7.5 Concluding Thoughts

Ecologists are often interested in drawing causal conclusions from data, however, developed methods for causal inference, particularly for observational data, are often not well-known amongst ecologists. This thesis centered around the SCM framework, which provides a set of widely applicable tools that can be used to draw causal conclusions across observational, quasi-experimental and experimental settings. The SCM framework uses DAGs to visualize the causal structure of a system or process under study, allowing researchers to explicitly communicate their causal assumptions to their audience. Once a DAG has been built that is sufficient to characterize a system or process under study, graphical rules including the backdoor or frontdoor criterion can be employed to guide

the construction of causal models. Doing so can improve causal conclusions drawn across a wide range ecological contexts and will ultimately increase the depth and pace of ecological research.

REFERENCES

Adams, V., Setterfield, S., Douglas, M., Kennard, M., Ferdinands, K. 2015. Measuring benefits of protected area management: Trends across realms and research gaps for freshwater systems." *Philosophical Transactions of the Royal Society B: Biological Sciences,* 370(1681), 20140274.

Adams, V., Barnes, M., Pressey, R. 2019. Shortfalls in conservation evidence: Moving from ecological effects of interventions to policy evaluation. *One Earth,* 1(1), 62-75.

Akaike, H. 1973. Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & B. F. Csaki (Eds.), Second International Symposium on Information Theory, (pp. 267–281). Academiai Kiado: Budapest.

Alkhamis, M., Fountain-Jones, N., Aguilar-Vega, C., Sanchez-Vizcaino, J. 2021. Environment, vector, or host? Using machine learning to untangle the mechanisms driving arbovirus outbreaks. *Ecological Applications*, 31(7), e02407.

Amin, A., Choumert, J., Combes, J., Motel, P., Kere, E., Olinga, J., Schwartz, S. 2015. A spatial econometric approach to spillover effects between protected areas and deforestation in the Brazilian amazon. Accessed from: https://halshs.archives-ouvertes.fr/halshs-00960476v2.

Andam, K., Ferraro, P., Pfaff, A., Sanchez-Azofeifa, G., Robalino, J. 2008. Measuring the effectiveness of protected area networks in reducing deforestation. *Proceedings of the National Academy of Sciences,* 105(42), 16089-16094.

Anderson, B. 2006. Response of tropical marine macroalgae to thermal stress. M.Sc. Thesis, Florida Atlantic University, Boca Raton, FL.

Angilletta Jr, M.J. *Thermal adaptation: a theoretical and empirical synthesis* (Oxford, 2009; online edn, Oxford Academic, 1 May 2009).

Ankan, A.,Wortel, I., Textor, J. 2021. Testing graphical causal models using the R package 'Dagitty.' *Current Protocols,* 1(2), e45.

Arif, S., MacNeil, A. 2022. Predictive models aren't for causal inference. *Ecology Letters,* 25(8), 1741-1745.

Baum, J., Worm, B. 2009. Cascading top-down effects of changing oceanic predator abundances. *Journal of Animal Ecology,* 78(4), 699-714.

Baumrind, D. 1983. Specious causal attributions in the social sciences: The reformulated stepping-stone theory of heroin use as exemplar. *Journal of Personality and Social Psychology,* 45, 1289-1298.

Behrenfeld, M., Falkowski, P. 1997. Photosynthetic rates derived from satellite-based chlorophyll concentration. *Limnology and Oceanography,* 42, 1-20.

Bellwood, D., Hughes, T., Folke, C., Nyström, M. 2004. Confronting the coral reef crisis. *Nature,* 429(6994), 827-833.

Bellwood, D., Hughes, T., Connolly, S., Tanner, J. 2005. Environmental and geometric constraints on Indo-Pacific coral reef biodiversity. *Ecology Letters,* 8(6), 643-651.

Benkwitt, C., Wilson, S., Graham, N. 2020. Biodiversity increases ecosystem functions despite multiple stressors on coral reefs. *Nature Ecology & Evolution*, 4, 919-926.

Bhala, J. 2018. The book of why: How a 'causal revolution' is shaking up science. Accessed at: https://bigthink.com/errors-we-live-by/judea-pearls-the-book-of-why-brings-news-of-a-new-science-of-causes

Bielby, J., Mace, G., Bininda-Emonds, O., Cardillo, M., Gittleman, J., Jones, K. et al. 2007. The fast-slow continuum in mammalian life history: An empirical reevaluation. *The American Naturalist,* 169(6), 748-757.

Blyth-Skyrme, R., Kaiser, M., Hiddink, J., Edwards-Jones, G., Hart, P. 2006. Conservation benefits of temperate marine protected areas: variation among fish species. *Conservation Biology*, 20, 811–820.

Bollen, K., Pearl, J. 2013. Eight myths about causality and structural equation models. In S.L.. Morgam (Ed.), Handbook of Causal Analysis for Social Research, Chapter 15, 301-328, Springer.

Bongaerts, P., Ridgway, T., Sampayo, E. M., Hoegh-Guldberg, O. 2010. Assessing the "deep reef refugia" hypothesis: Focus on Caribbean reefs. *Coral Reefs*, 29, 309-327.

Bousquin, S., Colee, J. 2014. Interim responses of littoral river channel vegetation to reestablished flow after phase I of the Kissimmee river restoration project. *Restoration Ecology,* 22(3), 388-396.

Brander, L, Van Beukering, P., Cesar, H. 2007. The recreational value of coral reefs: A meta-analysis. *Ecological Economics*, 63(1), 209–218.

Bridge, T., Hughes, T., Guinotte, J., Bongaerts, P. 2013. Call to protect all coral reefs. *Nature Climate Change*, 3, 528-530.

Brierley, A., Kingsford, M. 2009. Impacts of climate change on marine organisms and ecosystems. *Current Biology*, 19(14), R606-R614.

Brown, J., Gillooly, J., Allen, A., Savage, V., West, G. 2004. Toward a metabolic theory of ecology. *Ecology,* 85(7), 1771-1789.

Bruno, J., Peters, L., Harvell, C., Hettinger, A. 2003. Nutrient enrichment can increase the severity of two Caribbean coral disease. *Ecology Letters*, 6, 1056-1061.

Busch, J., Cullen, R. 2009. Effectiveness and cost-effectiveness of yellow-eyed penguin recovery. *Ecological Economics,* 68(3), 762-776.

Butsic, V., Baumann, M., Shortland, A., Walker, S., Kuemmerle, T. 2015. Conservation and conflict in the Democratic Republic of Congo: The impacts of warfare, mining, and protected areas on deforestation. *Biological Conservation,* 191, 266-273.

Butsic, V., Lewis, D., Radeloff, V., Baumann, M., Kuemmerle, T. 2017. Quasi-experimental methods enable stronger inferences from observational data in ecology. *Basic and Applied Ecology,* 19, 1-10.

Cameron, S., Wynn, M., Wilson, R. 2013. Sex-specific trade-offs and compensatory mechanisms: bite force and sprint speed pose conflicting demands on the design of geckos (Hemidactylus frenatus). *Journal of Experimental Biology*, 216(20), 3781-3789.

Campbell, S., Darling, E., Pardede, S., Ahmadia, G., Mangubhai, S., Maire, E. 2020. Fishing restrictions and remoteness deliver conservation outcomes for Indonedia's coral reef fisheries. *Conservation Letters*, 13(2), e12698.

Capdevila, P., Linares, C., Aspillaga, E., Riera, J., Hereu, B. 2018. Effective dispersal and density-dependence in mesophotic macroalgal forests: Insights from the Mediterranean species *Cystoseira zosteroides*. *PLoS ONE*, 13(1), e0191346.

Carmel, Y., Kent, R., Bar-Massada, A., Blank, L., Liberzon, J., Nezer, O. et al. 2013. Trends in ecological research during the last three decades – A systematic review. *PLoS ONE*, 8(4), e59813.

Chessman, B. 2013. Do protected areas benefit freshwater species? A broad-scale assessment for fish in Australia's murray-darling basin. *Journal of Applied Ecology,* 50(4), 969-976.

Cinner, J., McClanaham, T., Daw, T., Graham, N., Maina, J., Wilson, S., Hughes, T. 2009. Linking in social and ecological systems to sustain coral reef fisheries. *Current Biology,* 19, 206-212.

Cinner, J., Huchery, C., MacNeil, A., Graham, N., McClanahan, T., Maina, J. et al. 2016. Bright spots among the world's coral reefs. *Nature,* 535, 416-419.

Cinner, J., Maire, E., Huchery, C., Mouillot, D. 2018. Gravity of human impacts mediates coral reef conservation gains. *Biological Sciences*, 115(27), E6116-E6125.

Cinner, J., Zamborian-Mason, J., Gurney, G., Graham, N., MacNeil, A., Hoey, A., et al. 2020. Meeting fisheries, ecosystem function, and biodiversity goals in a human-dominated world. *Science*, 368, 307-311.

Choy, S., O'Leary, R., Mengersen, K. 2009. Elicitation by design in ecology: Using expert opinion to inform priors for Bayesian statistical models. *Ecology,* 90(1), 265-277.

Clayton, A., 2021. Bernoulli's Fallacy: Statistical Illogic and the Crisis of Modern Science. Columbia University Press.

Colgan, M. 1987. Coral reef recovery on Guam (Micronesia) after catastrophic predation by *Acanthaster planci*. *Ecology*, 68, 1592–1605.

Colinet, H., Sinclair, B. J., Vernon, P., & Renault, D. 2015. Insects in fluctuating thermal environments. *Annual Review of Entomology*, 60, 123.

Côté, I., Darling, E. 2010. Rethinking ecosystem resilience in the face of climate change. *PLoS Biolology,* 8(7), e1000438.

Cronin, J., Schoolmaster Jr., D. 2018. A causal partition of trait correlations: using graphical models to derive statistical models from theoretical language. *Ecosphere,* 9(9), e02422.

Cruz, I., Waters, L., Kikuchi, R., Leao, Z., Turra, A. 2018. Marginal coral reefs show high susceptibility to phase shift. *Marine Pollution Bulletin, 135,* 551-561.

Cunningham, Scott. 2021. *Causal Inference: The Mixtape*. *Causal Inference*. Yale University Press.

Dajka, J., Wilson, S., Robinson, J., Chong-Seng, K., Harris, A., Graham, N. 2019. Uncovering drivers of juvenile coral density following mass bleaching. *Coral Reefs,* 38, 637-649.

Darling, E., Alvarez-Filip, L., Oliver, T., McClanahan, T., Côté, I. 2012. Evaluating life-history strategies of reef corals from species traits. *Ecology Letters*, 15, 1378-1386.

Darling, E., McClanahan, T., Maina, J., Gurney, G., Graham, N., Januchowski-Hartley, F. et al. 2019. Social–environmental drivers inform strategic management of coral reefs in the Anthropocene. *Nature Ecology & Evolution,* 3, 1341-1350.

de Leeuw, J. (1985). Reviews. *Psychometrika*, 50, 371–375.

DeCarlo, T., Gajdzik, L., Ellis, J., Coker, D., Roberts, M., Hammerman, N. et al. 2020. Nutrient-supplying ocean currents modulate coral bleaching susceptibility. *Science Advances*, 6(34), eabc5493.

Drescher, M., Perera, A., Johnson, C., Buse, L., Drew, C., Burgman, M. 2013. Toward rigorous use of expert knowledge in ecological research. *Ecosphere,* 4(7), 1-26.

Donelson, J. M., Salinas, S., Munday, P. L., Shama, L. N. 2018. Transgenerational plasticity and climate change experiments: where do we go from here? *Global Change Biology,* 24(1), 13-34.

Duffy, J., Lefcheck, J., Stuart-Smith, R., Navarrete, S., Edgar, G. 2016. Biodiversity enhances reef fish biomass and resistance to climate change. *Proceedings of the National Academy of the Sciences*, 113(22), 6230-6235.

Duncan, O. 1975. *Introduction to structural equation models*. New York: Academic Press.

Duprey, N., Yasuhara, M., Baker, D. 2016. Reefs of tomorrow: eutrophication reduces coral biodiversity in an urbanized seascape. *Global Change Biology,* 22(11), 3550-3565.

Edgar, G., Bustamante, R., Farina, J., Calvopina, M., Martinez, C., Toral-Granda, M. 2004. Bias in evaluating the effects of marine protected areas: The importance of baseline data for the Galapagos Marine Reserve. *Environmental Conservation*, 31(3), 212-218.

Edgar, G. J., Stuart-Smith, R. D. 2009. Ecological effects of marine protected areas on rocky reef communities: a continental-scale analysis. *Marine Ecology Progress Series*, 388, 51-62.

Eggers, A., Tuñón, G., Dafoe, A. 2021. Placebo tests for causal inference. Accessed at: https://pelg.ucsd.edu/Eggers_2021.pdf.

Elmhirst, T., Connolly, S., Hughes, T. 2009. Connectivity, regime shifts and the resilience of coral reefs. *Coral Reefs,* 28, 949–957.

Elton, C.S. 1927. Animal Ecology. University of Chicago Press, Chicago.

Elwert, F. 2014. "Graphical Causal Models." In *Handbook of Causal Analysis for Social Research*. 245-273. Dordrecht: Springer.

Ethier, D., Nudds, T. 2017. Complexity of factors affecting bobolink population dynamics communicated with directed acyclic graphs: Complexity of factors affecting bobolink. *Wildlife Society Bulletin,* 41(1), 4-16.

F. A. D. 1900. "The Grammar of Science." *Nature* 62 (1594): 49–50.

Ferraro, P., Sanchirico, J., Smith, M. 2018. Causal inference in coupled human and natural systems. *Proceedings of the National Academy of the Sciences*, 116(12), 5311-5318.

Fisher, R. 1925. Statistical Methods for Research Workers. Edinburgh: Oliver and Boyd.

Fisher, R. 1935. The Design of Experiments. Edinburgh: Oliver and Boyd.

Fitzpatrick, B.M., Harvey, E.S., Heyward, A.J., Twiggs, E.J., Colquhoun, J. 2012. Habitat specialization in tropical continental shelf demersal fish assemblages. *PLoS One*, 7(6), e39634

Freedman, D. A. 2004. Graphical models for causation, and the identification problem. *Evaluation Review*, 28, 267–293.

Froese, R., Pauly, D. 2011. FishBase. (http://www.fishbase.org)

Fulton, C., Berkstrom, C., Wilson, S., Abesamis, R., Bradley, M., Akerlund, C., et al. 2020. Macroalgal meadow habitats support fish and fisheries in diverse tropical seascapes. *Fish and Fisheries*, 21(4), 700-717.
Fuxjager, L., Wanzenbock, S., Ringler, E., Wegner, K., Ahnelt, H., Shama, L. 2019. Within-generation and transgenerational plasticity of mate choice in oceanic stickleback under climate change. *Philosophical Transactions of the Royal Society B*, 374(1768), 20180183.

Gabry, J., Goodrich, B. 2017. Rstanarm: Bayesian applied regression modeling via Stan. R package version 2.15.3. Accessed at: http://mc-stan.org/rstanarm/.

Galles, D., Pearl, J. 1998. An axiomatic characterization of causal counterfactuals. *Foundations of Science,* 3(1), 151-182.

Gause, G. 1932. Experimental studies on the struggle for existence: 1. Mixed population of two species of yeast. *Journal of Experimental Biology*, 9, 389–402.

Gelman, A., Hill, J. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models.* Cambridge: Cambridge University Press.

Gilmour, J., Smith, L. D., Heyward, A., Baird, A., Pratchett, M. 2013. Recovery of an isolated coral reef system following severe disturbance. *Science,* 340, 69-71.

Glymour, M., Greenland, S. 2008. Causal Diagrams, in *Modern Epidemiology*. Lippincott Williams & Wilkins, Philadelphia.

Glymour, C. 2009. "Causation and Statistical Inference." CIn *The Oxford Handbook of Causation*. Oxford: Oxford University Press.

Goetze, J., Langlois, T., Egli, D., Harvey, E. 2011. Evidence of artisanal fishing impacts and depth refuge in assemblages of Fijian reef fish. *Coral Reefs,* 30, 507–517.

Goldthorpe, J. 2001. Causation, statistics, sociology. European Sociological Review, 17, 1–20.

Goreau, T., McClanahan, T., Hayes, R., Strong, A. 2000. Conservation of coral reefs after the 1998 bleaching event. *Conservation Biology*, 14, 5-15.

Goodrich, B., Gabry, J., Ali, I., Brilleman, S. 2022. "rstanarm: Bayesian applied regression modeling via Stan." R package version 2.21.3, https://mc-stan.org/rstanarm/.

Gottdenker, P. 1979. Francesco Redi and the fly experiments. *Bulletin of the History of Medicine*, 53(4), 575-592.

Gove, J., Williams, G., McManus, M., Clark, S., Ehses, J., Wedding, L. 2015. Coral reef benthic regimes exhibit non-linear threshold responses to natural physical drivers. *Marine Ecology Progress Series*, 522, 33-48.

Grace, J., Anderson, T., Olff, H., Scheiner, S. 2010. On the specification of structural equation models for ecological systems. *Ecological Monographs*, 80(1), 67-87.

Grace, J., and Irvine, K. 2020. Scientist's guide to developing explanatory statistical models using causal analysis principles. *Ecology,* 101(4), e02962.

Grace, J., Loreau, M., Schmid, B. 2022. A graphical causal model for resolving species identity effects and biodiversity–ecosystem function correlations: Comment. *Ecology,* 103(2), e03378.

Grafeld, S., Oleson, K. L. L., Teneva, L., Kittinger, J. N. 2017. Follow that fish: Uncovering the hidden blue economy in coral reef fisheries. *PLoS One*, 12(8), e0182104.

Graham, N., Wilson, S., Jennings, S., Polunin, N., Bijoux, J. Robinson, J. 2006. Dynamic fragility of oceanic coral reef ecosystems. *Proceedings of the National Academy of Sciences*, 103, 8425-8429.

Graham, N., Nash, K. 2013. The importance of structural complexity in coral reef ecosystems. *Coral Reefs*, 32, 315-326.

Graham, N., Jennings, S., MacNeil, A., Mouillot, D., Wilson, S. 2015. Predicting climate-driven regime shifts versus rebound potential in coral reefs. *Nature*, 518, 94-97.

Graham, N., Robinson, J., Smith, S., Govinden, R., Gendron, G., Wilson, S. 2020. Changing role of coral reef marine reserves in a warming climate. *Nature Communications*, 11, 2000.

Graham, N., Jennings, S., MacNeil, A., Mouillot, D., Wilson, S. 2015. Predicting climate-driven regime shifts versus rebound potential in coral reefs. *Nature,* 518(7537), 94–97.

Granger, C. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica,* 37, 424–438.

Gray, C., Hill, S., Newbold, T., Hudson, L., Borger, L., Contu, S., et al. 2016. Local biodiversity is higher inside than outside terrestrial protected areas worldwide. *Nature Communications*, 7, 12306.

Green, R. 1979. *Sampling Design and Statistical Methods for Environmental Biologists*. Nashville, TN: John Wiley & Sons.

Greenland, S., Pearl, J., Robins, J. 1999. Causal diagrams for epidemiologic research. *Epidemiology,* 10(1), 37–48.

Griffith, G., Morris, T., Tudball, M., Herbert, A., Mancano, G., Pike, L., et al. 2020. Collider bias undermines our understanding of COVID-19 disease risk and severity. *Nature Communications,* 11, 5749.

Grubbs, R., Carlson, J., Romine, J., Curtis, T., McElroy, W., McCandless, C. et al. 2016. Critical assessment and ramifications of a purported marine trophic cascade. *Scientific Reports,* 6(1), 20970.

Guttman, L. 1977. What is not what in statistics. *Journal of the Royal Statistical Society: Series D (The Statistician),* 26, 81–107.

Guy, C., Ratcliffe, J., Mideo, N. 2020. The influence of bat ecology on viral diversity and reservoir status. *Ecology and Evolution*, 10(12), 5748-5758.

Guzman, H., Cortés, J. 2007. Reef recovery 20 years after the 1982–1983 El Nino massive mortality. *Marine Biology*, 151, 401–411.

Habary, A., Johansen, J., Nay, T., Steffensen, J., Rummer, J. 2017. Adapt, move or die–how will tropical coral reef fishes cope with ocean warming? *Global Change Biology,* 23, 566-577.

Hamilton, M., Robinson, J., Benkwitt, C., Wilson, S., MacNeil, A., Ebrahim, A., Graham, N. 2022. Climate impacts alter fisheries productivity and turnover on coral reefs. *Coral Reefs*, 41, 921-935.

Hariton, E. Locascio, J. 2018. Randomised control trails – the gold standard for effectiveness research. *BJOG*, 125(13), 1716.

Hayden, B., Palomares, M., Smith, B., Poelen, J. 2019. Biological and environmental drivers of trophic ecology in marine fishes - a global perspective. *Scientific Reports,* 9, 11415.

Hempson, T., Graham, N., MacNeil, A., Hoey, A., Wilson, S. 2017. Ecosystem regime shifts disrupt trophic structure. *Ecological Applications*, 28(1), 191-200.

Hernan, M., Robins, J. 2006.Instruments for Causal Inference - An Epidemiologists Dream? *Epidemiology,* 17(4), 360-372.

Hernan, M. 2018. The C-Word: Scientific euphemisms do not improve causal inference from observational data. *AJPH*, 108(5), 616-619.

Herrera, D., Pfaff, A., Robalino, J. 2019. Impacts of protected areas vary with the level of government: Comparing avoided deforestation across agencies in the Brazillian Amazon. *Proceedings of the National Academy of Sciences,* 116(30), 14916-14925.

Hershberger, S. 2003. The growth of structural equation modelling: 1994-2001. *Structural Equation Modeling*, 10(1), 35-46.

Hoegh-Guldberg, O., Bruno, J. 2010. The impact of climate change on the world's marine ecosystems. Science, 328(5985), 1523-1528.

Holland, P. 1986. Statistics and causal inference. *Journal of the American Statistical Association,* 81(396), 945–960.

Hughes, T. 1994. Catastrophes, phase shifts, and large-scale degradation of a Caribbean coral reef. *Science,* 265(5178), 1547-1551.

Hughes, T., Rodrigues, M., Bellwood, D., Ceccarelli, D., Hoegh-Guldberg, O. et al. 2007. Phase shifts, herbivory, and resilience of coral reefs to climate change. *Current Biology*, 17(4), 360–365.

Hughes, T., Anderson, K., Connolly, S., Heron, S., Kerry, J., Lough, J. et al. 2018. Spatial and temporal patterns of mass bleaching of corals in the Anthropocence. *Science*, 359(6371), 80-83.

Hughes, T., Graham, N., Jackson, J., Mumby, P., Steneck, R. 2010. Rising to the challenge of sustaining coral reef resilience. *Trends in Ecology & Evolution*, 25, 633-642.

Hume, D. 1739. A Treatise of Human Nature. London. Reprint, Prometheus Books, 1992.

Hume, D. 1748. An Enquiry Concerning Human Understanding. London. Reprint, Dover Publications, 2004.

Hurd, C. 2000. Water motion, marine macroalgal physiology, and production. *Journal of Phycology*, 36, 453-472.

Huston, M. 1997. Hidden treatments in ecological experiments: re-evaluating the ecosystem function of biodiversity. *Oecologia*, 110, 449–460.

Imbens, G., Lemieux, T. 2008. Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics,* 142(2), 615-635.

Imbens, G., Rubin, D. *Causal inference in statistics, social, and biomedical sciences.* Cambridge University Press; 2015.

Imbens, G. 2020. Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature,* 58(4), 1129-1179.

Jiang, T., Smith, M., Street, A., Seegulam, V., Sampson, L., Murray, E. et al. 2021. A comorbid mental disorder paradox: Using causal diagrams to understand associations between posttraumatic stress disorder and suicide. *Psychological Trauma*, 13(7), 725-729.

Jobling, M. 1997. Temperature and growth: modulation of growth rate via temperature change. In: Wood CM, McDonald DG (eds) Global warming: implications for freshwater and marine fish. Cambridge University Press, Cambridge, pp 225–253

Johns, K., Osborne, K., Logan, M. 2014. Contrasting rates of coral recovery and reassembly in coral communities on the Great Barrier Reef. *Coral Reefs*, 33, 553-563.

Joppa, L., Pfaff, A. 2009. "High and Far: Biases in the Location of Protected Areas." *PLoS ONE,* 4(12), e8273.

Johnson, P., Rohr, J., Hoverman, J., Kellermanns, E., Bowerman, J., Lunde, K. 2012. Living fast and dying of infection: Host life history drives interspecific variation in infection and disease risk. *Ecology Letters,* 15(3), 235-242.

Kadye, W., Booth, A. 2012. Detecting impacts of invasive non-native sharptooth catfish, *Clarias gariepinus*, within invaded and non-invaded rivers. *Biodiversity and Conservation,* 21(8), 1997-2015.

Kendall, B. 2015. A statistical symphony: Instrumental variables reveal causality and control measurement error. Published as chapter 7 of G.A. Fox, S. Negrete-Yankelevich, and V.J. Sosa, eds. *Ecological Statistics: Contemporary Theory and Application* (Oxford University Press).

Kimmel, K., Dee, L., Avolio, M., Ferraro, P. 2021. Causal assumptions and causal inference in ecological experiments. *Trends in Ecology and Evolution*, 36(12), 1141-1152.

Kingsolver, J., Woods, H. 2016. Beyond thermal performance curves: modeling time-dependent effects of thermal stress on ectotherm growth rates. *The American Naturalist*, 187(3), 283-294.

Kinth, P., Mahesh, G., Panwar, Y. 2013. Mapping of zebrafish research: a global outlook. *Zebrafish*, 10(4), 510-517.

Koop, K., Booth, D., Broadbent, A., Brodie, J., Bucher, D., Capone, D. et al. 2001. ENCORE: the effect of nutrient enrichment on coral reefs. Synthesis of results and conclusions. *Marine Pollution Bulletin*, 42, 91-120.

Korell, L. et al. 2019. We need more realistic climate change experiments for understanding ecosystems of the future. *Global Change Biology,* 26, 325-327.

Koussoroplis, A. M., Pincebourde, S., Wacker, A. 2017. Understanding and predicting physiological performance of organisms in fluctuating and multifactorial environments. *Ecological Monographs*, 87(2), 178-197.

Kubicek, A., Breckling, B., Hoegh-Guldberg, O., Reuter, H. 2019. Climate change drives trait-shifts in coral reef communities. *Scientific Reports*, 9, 3721.

Kubicek, A., Muhando, C., Reuter, H. 2012. Simulations of long-term community dynamics in coral reefs – How perterbations shape trajectories. *PloS Computational Biology*, 8(11), e1002791

Kuhnert, Petra M., Tara G. Martin, and Shane P. Griffiths. 2010. A guide to eliciting and using expert knowledge in Bayesian ecological models. *Ecology Letters,* 13(7), 900–914.

Kumagai, N., Molinos, J., Yamano, H., Takao, S., Fujii, M., Yamanaka, Y. 2018. Ocean currents and herbivory drive macroalgae-to-coral community shift under climate warming. *Proceedings of the National Academy of the Sciences,* 115(36), 8990-8995.

Larsen, A., Meng, K., and Kendall, B. 2019. Causal analysis in control–impact ecological studies with observational data. *Methods in Ecology and Evolution,* 10(7), 924-934.

Laubach, Z., Murray, E., Hoke, K., Safran, R., Perng, W. 2021. A biologist's guide to model selection and causal inference. *Proceedings of the Royal Society B: Biological Sciences,* 288(1943), 20202815.

Lefcheck, J., Edgar, G., Stuart-Smith, R., Bates, A., Waldock, C., Brandl, S. et al. 2021. Species richness and identity both determine the biomass of global reef fish communities. *Nature Communication*, 12, 6875.

Lester, S., Halpern, B., Grorud-Colvert, K., Lubchenco, J., Ruttenberg, B., Gaines, S. et al. 2009. Biological effects within no-take marine reserves: a global synthesis. *Marine Ecological Progress Series*, 384, 33-46.

Lienart, G., Mitchell, M., Ferrari, M., McCormick, M. 2014. Temperature and food availability affect risk assessment in an ectotherm. *Animal Behaviour*, 89, 199-204.

Linares, C., Pratchett, M., Coker, D. 2011. Recolonisation of *Acropora hyacinthus* following climate-induced coral bleaching on the Great Barrier Reef. *Marine Ecology Progress Series*, 438, 97–104.

Lindfield, S. J., Harvey, E. S., Halford, A. R., McIlwain, J. L. 2016. Mesophotic depths as refuge areas for fishery-targeted species on coral reefs. *Coral Reefs*, 35, 125–137.

Loreau, M., S. Naeem, P. Inchausti, J. Bengtsson, J. P. Grime, A. Hector, D. U. Hooper, et al. 2001. Biodiversity and ecosystem functioning: Current knowledge and future challenges. *Science,* 294(5543), 804–808.

Loya, Y., Sakai, K., Yamazato, K., Nakano, Y., Sambali, H., van Woesik, R. 2001. Coral bleaching: the winners and losers. *Ecology Letters*, 4, 122-131.

Lu, X., Chen, X., Zhao, X., Lv, D., Zhang, Y. 2021. Assessing the impact of land surface temperature on urban net primary productivity increment based on geographically weighted regression model. *Scientific Reports*, 11, 22282.

Lui, G., Heron, S., Eakin, C., Muller-Karger, F., Vega-Rodriguez, M., Guild, L. et al. 2014. Reef-scale thermal stress monitoring of coral ecosystems: New 5-km global products from NOAA coral reef Watch. *Remote Sens*, 6(1), 11579-11606.

Lusk, C., Del Pozo, A. 2002. Survival and growth of seedlings of 12 Chilean rainforest trees in two light environments: Gas exchange and biomass distribution correlates. *Austral Ecology*, 27, 173-182.

Madin, J., Connolly, S. 2006. Ecological consequences of major hydrodynamic disturbances on coral reefs. *Nature*, 444, 477-480.

MacDonald, R. 2004. The specific analysis of structural equation models. *Multivariate Behavioral Research,* 39, 687-713.

Mansournia, M., Hernán, M., Greenland, S. 2013. Matched designs and causal diagrams. *International Journal of Epidemiology,* 42(3), 860-869.

Marcot, B., Steventon, J., Sutherland, G., McCann, R. 2006. Guidelines for developing and updating bayesian belief networks applied to ecological modeling and conservation." *Canadian Journal of Forest Research,* 36(12), 3063-3074.

Marshall, D., Uller, T. 2007. When is a maternal effect adaptive? *Oikos,* 116(12), 1957-1963.

Marshall, K., Anderson, K., Brown, N., Dytnerski, J., Flynn, K., Bernhardt, J. et al. 2021. Whole-organism responses to constant temperatures do not predict responses to variable temperatures in the ecosystem engineer *Mytilus trossulus*. *Proceedings of the Royal Society B,* 288(1947), 20202968.

Martin, T., Burgman, M., Fidler, F., Kuhnert, P., Low-Choy, S., McBride, M, Mengersen, K. 2012. Eliciting expert knowledge in conservation science. *Conservation Biology: The Journal of the Society for Conservation Biology,* 26(1), 29-38.

Mascia, M., Claus, C., Naidoo, R. 2010. Impacts of marine protected areas on fishing communities. *Conservation Biology,* 24, 1424–29.

Massey, M. D., Hutchings, J. 2020. Thermal variability during ectotherm incubation: a review and synthesis. *Journal of Experimental Zoology Part A*, 335(1), 59-71.

May, R. 1977. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269, 471–477.

Mbaru, E., Graham, N., McClanahan, T., Cinner, J. 2020. Functional traits illuminate the selective impacts of different fishing gears on coral reefs. *Journal of Applied Ecology*, 57(2), 241-252.

McClanahan, T., Maina, J., Moothien-Pillay, R., Baker, A. 2005. Effects of geography, taxa, water flow, and temperature variation on coral bleaching intensity in Mauritius. *Marine Ecology Progress Series*, 298, 131-42.

McClanahan, T. 2008. Response of the coral reef benthos and herbivory to fishery closure management and the 1998 ENSO disturbance. *Oecologia*, 155, 169–177.

McClanahan, T., Maina, J., Darling, E., Guillaume, M., Muthiga, N., Dagata, S. et al. 2020. Large geographic variability in the resistance of corals to thermal stress. Global *Ecology and Biogeography*, 29(12), 2229-2247.

McCook, L. 1999. Macroalgae, nutrients and phase shifts on coral reefs: scientific issues and management consequences for the Great Barrier Reef. *Coral Reefs*, 18, 357-367.

McCook, L., Jompa, J., Diaz-Pulido, G. 2001. Competition between corals and algae on coral reefs: a review of evidence and mechanisms. *Coral Reefs*, 19, 400-417.

McElreath, R. 2020. *Statistical rethinking: A Bayesian course with examples in R and Stan*. CRC Press, Boca Raton, FL.

McNay, R., Marcot, B., Brumovsky, V., Ellis, R. 2006. A Bayesian approach to evaluating habitat for woodland caribou in north-central British Columbia. *Canadian Journal of Forest Research,* 36, 3117-3133.

Mellin, C., MacNeil, A., Cheal, A., Emslie, M., Caley, J. 2016. Marine protected areas increase resilience among coral reef communities. *Ecology Letters*, 19(6), 629-637.

Merilä, J., Hendry, A. 2014. Climate change, adaptation, and phenotypic plasticity: the problem and the evidence. Evolutionary Applications, 7(1), 1-14.

Merton, R. 1968. The Matthew Effect in science. *Science*, 159(3810), 56-63.

Mies, M., Francini-Filho, R., Zilberberg, C., Garrido, A., Longo, G., Laurentino, E., Güth, A., Sumida, P., Banha, T. 2020. South Atlantic coral reefs are major global warming Refugia and less susceptible to bleaching. *Frontiers in Marine Science,* 7, 1-13.

Millard, J., Outhwaite, C.L., Kinnersley, R., Freeman, R., Gregory, R., Adedoja, O. *et al*. 2021. Global effects of land-use intensity on local pollinator biodiversity. *Nature Communications*, 12, 2902.

Mittelbach, G. 2012. Community ecology. Sinauer Associates, Sunderland, Maryland, USA.

Mora, C., Danovaro, R., Loreau, M. 2014. Alternative hypotheses to explain why biodiversity-ecosystem functioning relationships are concave-up in some natural ecosystems but concave-down in manipulative experiments. *Scientific Reports*, 4, 5427.

Morash, A., Neufeld, C., MacCormack, T., Currie, S. 2018. The importance of incorporating natural thermal variation when evaluating physiological performance in wild species. *Journal of Experimental Biology,* 221(14), jeb164673.

Morgan, R., Finnøen, M. H., Jensen, H., Pélabon, C., Jutfelt, F. 2020. Low potential for evolutionary rescue from climate change in a tropical fish. *Proceedings of the National Academy of Sciences*, 117(52), 33365-33372.

Morgan, Stephen L., and Christopher Winship. 2014. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. 2nd ed. Analytical Methods for Social Research. Cambridge: Cambridge University Press.

Mumby, P., Dahlgren, C., Harborne, A., Kappel, C., Micheli, F. et al. 2006. Fishing, trophic cascades, and the process of grazing on coral reefs. *Science*, 311, 98-101.

Mumy, P., Hastings, A. 2008. The impact of ecosystem connectivity of coral reef resilience. *Journal of Applied Ecology*, 45, 854-862.

Mumby, P., Harborne, A. 2010. Marine reserves enhance the recovery of corals on Caribbean reefs. *PLoS ONE*, 5, e8657.

Munday, P., Jones, G., Pratchett, M., Williams, A. 2008. Climate change and the future for coral reef fishes. *Fish and Fisheries*, 9(3), 261-285.

Muthen, B. 1984. A general structural equation model with dichotomous, ordered categorical, and continuous latent variable indicators. *Psychometrika,* 49, 115-132.

Myers, Ransom A., Julia K. Baum, Travis D. Shepherd, Sean P. Powers, and Charles H. Peterson. 2007. Cascading effects of the loss of apex predatory sharks from a coastal ocean. *Science,* 315(5820), 1846-1850.

Nakamura, T., Van Woesik, R. 2001. Water-flow rates and passive diffusion partially explain differential survival of corals during the 1998 bleaching event. *Marine Ecology Progress Series*, 212, 301-304.

Nakamura, T., Yamasaki, H., Van Woesik, R. 2003. Water flow facilitates recovery from bleaching in the coral *Stylophora pistillata*. *Marine Ecology Progress Series*, 287-291.

Nash, K., Graham, N., Jennings, S., Wilson, S., Bellwood, D. 2015. Herbivore cross-scale redundancy supports response diversity and promotes coral reef resilience. *Journal of Applied Ecology*, 53, 646-655.

Newton, K., Cote, I.M., Pilling, G.M., Jennings, S., Dulvy, N.K. 2007. Current and future sustainability of island coral reef fisheries. *Current Biology*, 17, 655-658.

O'Connor, M., Bruno, J., Gaines, S., Halpern, B., Lester, S., Kinlan, B., Weiss, J. 2007. Temperature control of larval dispersal and the implications for marine ecology, evolution, and conservation. *Proceedings of the National Academy of Sciences*, 104, 1266-1271.

Paine, R. Food Web Complexity and Species Diversity. *The American Naturalist*, 910 (1966), 65–75.

Pankhurst, N.W., Porter, M.J.R. 2003 Cold and dark or warm and light: variations on the theme of environmental control of reproduction. *Fish Physiology and Biochemistry,* 28, 385-389.

Pearce, N., Lawlor, D. 2016. Causal Inference—so much more than statistics. *International Journal of Epidemiology,* 45(6), 1895-1903.

Pearl, J. 1993. Graphical Models, Causality and Intervention. *Statistical Science,* 8(3), 266-269.

Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika,* 82(4), 669-688.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Pearl, J. ([2000] 2009). Causality: Models, Reasoning and Inference. 2nd ed. Cambridge, UK: Cambridge University Press.

Pearl, J. 2012. The causal foundations of structural equation modeling. Technical Report.

Pearl, J. 2014. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. London, England: Morgan Kaufmann.

Pearl, J., Mackenzie, D. 2018. The book of why: The new science of cause and effect. Harlow, England: Penguin Books.

Pearson, C., Ormerod, S., Symondson, W., Vaughan, I. 2016. Resolving large-scale pressures on species and ecosystems: Propensity modelling identifies agricultural effects on streams. *The Journal of Applied Ecology,* 53(2), 408–17.

Pearson, K. 1911. Grammar of Science (3rd ed). London: A. and C. Black.

Pereira, P., Macedo, C., Nunes, J., Marangoni, L., Bianchini, A. 2018. Effects of depth on reef fish communities: Insights of a "deep refuge hypothesis" from Southwestern Atlantic reefs. *PLoS One*, 13(9), e0203072.

Perez, J., Higuera, I., Bonilla-Mejia, L. 2017. Marginal effects of overlapping protected areas: A regression discontinuity approach. 64th Annual North American Meetings of the Regional Science Association International. Vancouver, British Columbia, Canada.

Pierce, C.S., Jastrow, J. 1885. On small differences in sensation. *Memoires of the National Academy of Sciences*, 3, 73-83.

Polunin N., Roberts C. 1993. Greater biomass and value of target coral-reef fishes in two small Caribbean marine reserves. *Marine Ecology Progress Series*, 100, 167-176.

Pratchett, M., McWilliam, M., Riegl, B. 2020. Contrasting shifts in coral assemblages with increasing disturbances. *Coral Reefs*, 1-11.

Ramsey, D., Forsyth, D., Wright, E., McKay, M., Westbrooke, I. 2019. Using propensity scores for causal inference in ecology: Options, considerations, and a case study. *Methods in Ecology and Evolution,* 10(3), 320-331.

Reef Life Survey Foundation, 2019. Standardised Survey Procedures for Monitoring Rocky & Coral Reef Ecological Communities. Reef Life Survey Foundation, Hobart, Australia available at. https://reeflifesurvey.com/wp-content/uploads/2019/02/ NEW-Methods-Manual_150815.pdf.

Rebolledo, A. P., Sgrò, C. M., Monro, K. 2021. Thermal performance curves are shaped by prior thermal environment in early life. *Frontiers in Physiology*, 12.

Reopanichkul, P., Schlacher, T., Carter, R., Worachananant, S. 2009. Sewage impacts coral reefs at multiple levels of ecological organization. *Marine Pollution Bulletin*, 58(9), 1356-1362.

Robinson, J., Wilson, S., Jennings, S. & Graham, N. 2019a. Thermal stress induces persistently altered coral reef fish assemblages. *Global Change Biology*, 25, 8, 2739-2750.

Robinson, J., Wilson, S., Robinson, J., Gerry, C., Lucas, J. & Assan, C. 2019b. Productive instability of coral reef fisheries after climate-driven regime shifts. *Nature Ecology and Evolution*, 3(2), 183-190.

Robinson, J., Wilson, S. & Graham, N. 2019c Abiotic and biotic controls on coral recovery 16 years after mass bleaching. *Coral Reefs*, 38(6), 1255-65.

Rollinson, N., Holt, S. M., Massey, M. D., Holt, R. C., Nancekivell, E. G., Brooks, R. J. 2018. A new method of estimating thermal performance of embryonic development rate yields accurate prediction of embryonic age in wild reptile nests. *Journal of Thermal Biology*, 74, 187-194.

Rogers, A., Blanchard, J., Mumby, P. 2014. Vulnerability of coral reef fisheries to a loss of structural complexity. *Current Biology*, 24, 1000-1005.

Rohrer, Julia M. 2018. Thinking clearly about correlations and causation: Graphical causal models for observational data. *Advances in Methods and Practices in Psychological Science,* 1(1), 27-42.

Rosenbaum, P. *Observational Studies.* 2. New York City: Springer; 2002.

Rosenbaum, P., Rubin, D. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41-55.

Ross, G., Doropoulos, C., Zupan, M., Rogers, A., Steneck, R., Golbuu, Y. et al. 2015. Phase shift facilitation following cyclone disturbance on coral reefs. *Oecologia,* 178, 1193-1203.

Rubin, D. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688-701.

Rubin, D. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association,* 100, 322-331.

Runge, J., Bathiany, S., Bollt, E., Camps-Valls, G., Coumou, D., Deyle, E. et al. 2019. Inferring causation from time series in Earth system sciences. *Nature Communications,* 10, 2553.

Salinas, S., Brown, S. C., Mangel, M., Munch, S. B. 2013. Non-genetic inheritance and changing environments. *Non-Genetic Inheritance*, 1.

Safaie, A., Silbiger, N., McClanahan, T., Pawlak, G., Barshis, D., Hench, J. 2018. High frequency temperature variability reduces the risk of coral bleaching. *Nature Communications*, 9, 1671.

Sagarin, R. Pauchard, A. 2010. Observational approaches in ecology open new ground in a changing world. *Frontiers in Ecology and the Environment,* 8(7), 379-386.

Saji, N., Goswami, B., Vinayachandran, P., Yamagata, T. 1999. A dipole mode in the tropical Indian Ocean. *Nature*, 401, 360-363.

Sam, K., Koane, B., Jeppy, S., Novotny, V. 2014. Effect of forest fragmentation on bird species richness in Papua New Guinea. *Journal of Field Ornithology,* 85(2), 152-167.

Sato, K., Shiomi, K., Marshall, G., Kooyman, G., Ponganis, P. 2011. Stroke rates and diving air volumes of emperor penguins: implications for dive performance. *Journal of Experimental Biology,* 214(17), 2854-2863.

Schaffelke, B. Klumpp, D. 1997 Growth of germlings of the macroalga Sargassum baccularia (Phaeophyta) is stimulated by enhanced nutrients. *Proceedings of the 8th International Coral Reef Symppsium*, 2, 1839-1842.

Scheffer, M., Carpenter, S., Foley, J., Folke, J., Walker, B. 2001. Catastrophic shifts in ecosystems. *Nature*, 413, 591-596.

Schielzeth, H. 2010. Simple means to improve the interpretability of regression coefficients. *Methods in Ecology and Evolution*, 1(2), 103-113.

Schnabl, P. 2012. The international transmission of bank liquidity shocks: Evidence from an emerging market. *The Journal of Finance,* 67(3), 897-932.

Schoolmaster Jr., D., Zirbel, C., Cronin, J. 2020. A graphical causal model for resolving species identity effects and biodiversity–ecosystem function correlations. *Ecology*, 101(8), e03070.

Schoolmaster Jr., D., Zirbel, C., Cronin, J. 2022. A graphical causal model for resolving species identity effects and biodiversity–ecosystem function correlations: Reply. *Ecology,* 103(2), e03593.

Seebacher, F., White, C., Franklin, C. 2015. Physiological plasticity increases resilience of ectothermic animals to climate change. *Nature Climate Change*, 5(1), 61-66.

Selig, E., Bruno, J. 2010. A global analysis of the effectiveness of marine protected areas in preventing coral loss. *PloS ONE*, 5(2), e9278.

Shama, L., Wegner, K. 2014. Grandparental effects in marine sticklebacks: transgenerational plasticity across multiple generations. *Journal of Evolutionary Biology*, 27(11), 2297-2307.

Sheppard, C., Spalding, M., Bradshaw, C., Wilson, S. 2002. Erosion vs recovery of coral reefs after 1998 El Nino: Chagos reefs, Indian Ocean. *Ambio*, 31, 40-48.

Shipp, R. 2003. A perspective on marine reserves as a fishery management tool. *Fisheries,* 28, 10-21.

Shipley, B. 2000. A new inferential test for path models based on directed acyclic graphs. *Structural Equation Modeling: A Multidisciplinary Journal,* 7(2), 206-218.

Shipley, B. 2016. Cause and correlation in biology: A user's guide to path analysis, structural equations and causal inference with R. Second edition. Cambridge University Press, Cambridge, UK.

Shrier, I. 2009. Letter to the editor. *Statistics in Medicine,* 27, 2740-2741.

Simberloff, D. 1976. Experimental zoogeography of islands: Effects of island size. *Ecology*, 57(4), 629-648.

Sjölander, A. 2009. Letter to the Editor: Propensity Scores and M-Structures. *Statistics in Medicine,* 28, 1416-1420.

Slattery, M., Lesser, M., Brazeau, D., Stokes, M., Leichter, J. 2011. Connectivity and stability of mesophotic coral reefs. *Journal of Experimental Marine Biology and Ecology*, 408(1-2), 32-41.

Soares, M., Tavares, T., Carneiro, P. 2019. Mesophotic ecosystems: Distribution, impacts and conservation in the South Atlantic. *Diversity and Distributions*, 25(2), ddi.12846.

Sobel, M. 2008. Identification of causal parameters in randomized studies with mediating variables. *Journal of Educational and Behavioral Statistics*, 33, 230-251.

Soler, G., Edgar, G., Thomson, R., Kininmonth, S., Campbell, S., Dawson, T. et al. 2015. Reef fishes at all trophic levels respond positively to effective marine protected areas. *PLoS ONE*, 10(10), e0140270.

Spinks, R., Bonzi, L., Ravasi, T., Munday, P., Donelson, J. 2020. Sex- and time-specific parental effects of warming on reproduction and offspring quality in a coral reef fish. *Evolutionary Application*, 14(4), 1145-1158.

Spirtes, P., Glymour, C., Scheines, R. 2001. *Causation, Prediction, and Search*. 2nd ed. Adaptive Computation and Machine Learning Series. Cambridge, MA, USA: A Bradford Book.

Staiger, D., Stock, J. 1997. Instrumental variables regression with weak instruments. *Econometrica: Journal of the Econometric Society,* 65(3), 557.

Steneck, R., Mumby, P., MacDonald, C., Rasher, D., Stoyle, G. 2018. Attenuating effects of ecosystem management on coral reefs. Scientific Advamces, 4:eaao5493.

Stewart-Oaten, A., Murdoch, W., Parker, K. 1986. Environmental impact assessment: 'Pseudoreplication' in time? *Ecology,* 67(4), 929-940.

Stolberg, M. 2006. Inventing the randomized double-blind trail: the Nuremberg salt test of 1835. *Journal of the Royal Society of Medicine,* 99(12), 642-643.

Stuart, E. 2010. Matching methods for causal inference: A review and look forward. *Statistical Science*, 25(1), 1-25.

Sugihara, G. May, R., Ye, H., Hsieh, C., Deyle, E., Fogarty, M. et al. 2012. Detecting Causality in Complex Ecosystems. *Science*, 338(6106), 496-500.

Sunday, J.M., Bates, A.E., Dulvy, N.K. 2011. Global analysis of thermal tolerance and latitude in ectotherms. *Proceedings of the Royal Society B: Biological Sciences*, 278, 1823-1830.

Suren, A, Lambert, P., Sorrell, B. 2011. The impact of hydrological restoration on benthic aquatic invertebrate communities in a New Zealand wetland. *Restoration Ecology,* 19(6), 747-757.

Takahashi, M., McCormick, M., Munday, P., Jones, G. 2012. Influence of seasonal and latitudinal temperature variation on early life-history traits of a coral reef fish. *Marine and Freshwater Research,* 63, 856-864.

Tamir, R., Eyal, G., Kramer, N., Laverick, J., Loya, Y. 2019. Light environment drives the shallow-to-mesophotic coral community transition. *Ecosphere*, 10(9), e02839.

Taylor, E, Diele-Viegas, L., Gangloff, E., Hall, J., Halpern, B., Massey, M.D. et al. 2021. The thermal ecology and physiology of reptiles and amphibians: A user's guide. *Journal of Experimental Zoology Part A: Ecological and Integrative Physiology,* 335(1), 13-44.

Textor, J., Zander, B., Gilthorpe, M., Liskiewicz, M., Ellison, G. 2016. Robust causal inference using directed acyclic graphs: the R package 'dagitty'. *International Journal of Epidemiology*, 1887-1894.

Thiault, L., Kernaléguen, L., Osenberg, C., Claudet, J. 2017. Progressive-change BACIPS: A flexible approach for environmental impact assessment. *Methods in Ecology and Evolution,* 8(3), 288–96.

Thiault, L., Kernaléguen, L., Osenberg, C., Lison de Loma, T., Chancerelle, Y., Siu, G., Claudet, J. 2019. Ecological evaluation of a marine protected area network: A progressive-change BACIPS approach. *Ecosphere,* 10(2), e02576.

Thistlethwaite, D., Campbell, D. 1960. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology,* 51(6), 309-317.

Tilman, D., Wedin, D., Knops, J. 1996. Productivity and sustainability influenced by biodiversity in grassland ecosystems. *Nature,* 379(6567), 718-720.

Tilman, D., Isbell, F., Cowles, J. 2014. Biodiversity and ecosystem functioning. *Annual Review of Ecology, Evolution & Systematics,* 45, 471-493.

Tyler, E., Speight, M., Henderson, P., Manica, A. 2009. Evidence for a depth refuge effect in artisanal coral reef fisheries. *Biological Conservation*, 142, 652-667.

Van Hooidonk, R. Maynard, J., Tamelander, J., Gove, J., Ahmadia, G., Raymundo, L. et al. 2016. Local-scale projections of coral reef futures and implications of the Paris Agreement. *Scientific Reports*, 6, 39666.

Victor, S. 2008 Stability of reef framework and post settlement mortality as the structuring factor for recovery of Malakal Bay Reef, Palau, Micronesia: 25 years after a severe COTS outbreak. *Estuarine Coastal Shelf Science,* 77, 175-180.

Wieser, W. 1973. Effects of temperature on ectothermic organisms. Springer Berlin Heidelberg.

Williams, I., Polunin, N., Hendrick, V. 2001. Limits to grazing by herbivorous fishes and the impact of low coral cover on macroalgal abundance on a coral reef in Belize. *Marine Ecology Progress Series*, 222, 187-196.

Williams, T., Bach, C., Matthiesen, N., Henriksen, T., Gagliardi, L. 2018. Directed acyclic graphs: a tool for causal studies in paediatrics. *Pediatric Research*, 84, 487-493.

Wilson, S., Graham, N., Pratchett, M., Jones, G., Polunin, N. 2006. Multiple disturbances and the global degradation of coral reefs: are reef fishes at risk or resilient? *Global Change Biology*, 12(11), 2220-2234.

Wilson, S., Graham, N., Polunin, N. 2007. Appraisal of visual assessments of habitat complexity and benthic composition on coral reefs. *Marine Biology*, 151, 1069-1076.

Wilson, S., Graham, N., Fisher, R., Robinson, J., Nash, K., Chong-Seng, K., Polunin, N., Aumeeruddy, R., Quatre, R. 2012. Effect of macroalgal expansion and marine protected areas on coral recovery following a climatic disturbance. *Conservation Biology*, 26(6), 995-1004.

Wilson, S., Robinson, J., Chong-Seng, K., Robinson, J., Graham, N. 2019. Boom and bust of keystone structure on coral reefs. *Coral Reefs*, 38, 625-635.

Wilson, S., Nudds, T., de Vries, A. 2021. A causal modelling approach to informing woodland caribou conservation policy from observational studies. *Biological Conservation*, 254, 109370.

Woodhead, A., Hicks, C., Norstrom, A., Williams, G., Graham, N. 2019. Coral reef ecosystem services in the Anthropocene. Functional Ecology, 33(6): 1023-1034.

Wright, I., Reich, P., Westoby, M., Ackerly, D., Baruch, Z., Bongers, F. et al. 2004. The Worldwide Leaf Economics Spectrum. *Nature,* 428(6985), 821-827.

Wright, S. 1920. The relative importance of heredity and environment in determining the piebald pattern of guinea-pigs. *Proceedings of the National Academy of Sciences*, 6, 320-332.

Wright, S. 1921. Correlation and causation. *Journal of Agricultural Research,* 20, 557-585.

Wright, P. *The Tariff on Animal and Vegetable Oils.* New York, NY: The MacMillan Company; 1928.

Zeldow, B, Hatfield, L. 2021. "Confounding and Regression Adjustment in Difference-in-Differences Studies." *Health Services Research,* 56(5), 932-941.

Zizzari, Z., Ellers, J. 2014. Rapid shift in thermal resistance between generations through maternal heat exposure. *Oikos,* 123(11), 1365-1370.