4

IDEAL OBSERVATIONS

Information and Causation in Biological Practice

OLIVER M. LEAN

1. INTRODUCTION

From the way biologists tend to talk, it seems like biology involves two very different types of phenomena simultaneously. DNA is used to produce proteins by complex molecular interactions, for example, yet in doing so it is also said to "code for" those proteins. When a neuron in the brain depolarizes, it's said to be passing along a "message." When a gazelle leaps high into the air, it is "signaling" to others around it. In all these cases, everyone agrees that there is a physical process unfolding according to causal principles of some kind. Yet at the same time, talk of coding and signals and messages suggests that these processes also involve *information*. References to information and related terms are common in biological sciences covering a range of scales, from the molecular level of genetics to ethology (the study of animal behavior) and neuroscience. *Information-talk*, as I will call it, is a common and widely accepted feature of the language of biologists.

What is less clear, however, is how to interpret this fact. Information-talk has a number of features that appear at first glance to be puzzling. For example, biological information seems to be intimately related to yet importantly different from the physical stuff of biological phenomena: A physical thing can be said to carry different pieces of information, and the same information can appear in different physical forms. Hence, information can be "transmitted" from one physical thing to another, potentially very different, thing. On its face, these are observations with *metaphysical* significance: What is this "information" biologists seem so comfortable talking about, and how does it fit into the wider scientific picture of the biological world? In particular, how does it relate to the relatively uncontroversial *causal* or *physical* aspects of biology? Is information-talk just a convenient shorthand for what are in fact garden-variety physical phenomena (Sarkar 2000; 2005)? Or does information play a deeper and more essential role that cannot be accounted for in purely physical terms (Barbieri 2016)?

Why does this matter? The idea of biological systems trading in information lies at the heart of various debates—some intrinsic to science and its methods, some with wider social significance. For some, the role of information is deeply important to what distinguishes living things from nonliving matter (Stotz 2019), perhaps even a *biosignature* that may help us recognize life elsewhere in the universe (Walker et al. 2018). For others, talk of information is a historical relic that confuses more than it clarifies and so should be dispensed with. Some ethologists, for instance, have argued that talk of information in animal signals is misleading because focusing on an intangible "content" to signals diverts attention from important physical aspects of signal design. Instead, they argue, we should understand animal communication as the attempt by senders to *manipulate* receivers (Rendall et al. 2009; Owren et al. 2010). Some even argue that information-talk should be resisted on political grounds: the idea that genes carry information, some believe, descends from outmoded and harmful ideas about the primacy of genetic factors over others in development—that certain traits are "genetically encoded" and hence immutable (Francis 2003).

What we have, then, is a question with a metaphysical flavor and that concerns the biological world and the study of it. In this chapter, I offer a way to address this question of how to interpret information-talk in biology. Importantly, my discussion goes further than an analysis of the logical or linguistic relationships between information and other concepts intrinsic to biological theory; it is not just an exercise in *metaphysics of science*, in the sense discussed in the introduction to this volume. Instead, it has implications for how the use of that concept drives the success of biological practices—specifically, practices of investigating and explaining how biological things solve what I will call *coordination problems*. Because of this, it qualifies as a work of *scientific metaphysics* in the sense outlined in the introduction to this volume and represented throughout other chapters. That is, it has implications for not just what biological theory is like but what the biological world is like.

I take the view (alongside other contributors to this volume) that such questions are best done with an explicit and well-defined purpose to motivate and constrain our theorizing. To that end, my particular aim here is to make sense of biological information in a way that suits what I henceforth call *naturalized epistemology.* The "naturalism" I refer to here is one that allows scientific developments to place constraints on other areas of our thinking. With that in mind, naturalized *epistemology* is the project of situating epistemic concepts—including *belief, knowledge,* and *meaning*—within a scientific picture of the world, potentially reshaping those concepts in order to make them fit. If we want to continue to understand ourselves as knowing subjects with beliefs about the world that can be true or false, a commitment to naturalism compels us to reckon seriously with science's picture of humans as physically embodied organisms evolved to solve physical, biological, and social problems.

Many take the concept of *information* to play a vital role in fitting these two pictures together. On the one hand, information as a concept has its origin in mentalistic talk: it invokes epistemic ideas like instruction or evidence. On the other hand, as seen previously, information appears to be commonplace in respectable scientific discourse about biological phenomena. Because information appears to stand astride these two realms, it's thought that we may be able to narrow the conceptual chasm between them by tracing the origins of "full-blown" mental phenomena to biological processes that are more generally informational (Dretske 1981; Millikan 2004; Skyrms 2010; Sterner 2014; Garson and Papineau 2019).

To achieve this, however, one must first lay the aforementioned metaphysical groundwork: we need to define exactly what is meant when we say that biological phenomena are trading in "information." To do the work of naturalizing mental or epistemic notions, this account of biological information should satisfy at least two closely related criteria.

First, the account must imply that biological systems are "really" informational in a literal or substantive sense (Collier 2008). In other words, information should not be merely something we project onto what are in fact plainly physical phenomena. To use a classic example, the rings of a tree are often said to carry information about its age, but this is apparently just information *for us,* the observers. Instead, what we need is a sense in which something is information *for the organism* itself, in some sense independently of us. Without this, we lack a basis for recognizing epistemic phenomena in the natural world because "information" remains confined to the minds of the biologists.

Secondly, we need a clear sense of how information so defined relates to the *causal* aspects of biological phenomena. Everyone agrees that causality is vitally important to our understanding of how biology works. The key question is what talk of "information" adds to that picture—what explanatory role it can play—that ordinary causal/physical language cannot do by itself. Some have addressed this challenge by effectively *reducing* the latter to the former—that is, explicating information *in causal terms.* Holly Andersen (2017), for instance, expresses information as a measure of pattern over a causal structure. Elsewhere, recent and widely discussed work by Griffiths and colleagues has analyzed biological information simply as a measure of *causal specificity*—of the extent to which a cause yields fine-grained or precise control over its effect (Griffiths et al. 2015; Stotz and Griffiths 2017; Stotz 2019; see section 3.3). Yet for the particular purposes of reconceiving *epistemic* phenomena in biological terms, this reductive approach seems at best incomplete: What does causal specificity have to do with knowledge and related concepts? It's proven difficult to answer that question in a way that naturalism would deem acceptable.

Taken together, these interwoven criteria establish the puzzle that an adequate characterization of information in biology should answer: What we need is a way of understanding biological information that is both (1) metaphysically integrated into the established scientific picture of biology and (2) playing a substantive explanatory role within that picture.

While there is still much debate about exactly how to solve this puzzle, there has emerged a broad (though far from universal) agreement about what the notion of biological information is broadly about; namely, that it has something to do with *use.* This general idea is represented in a wide variety of works in a wide variety of fields. For example, it is represented in the teleosemantic approach to representational content of Millikan (2004; 2013), Neander (1995; 2017), and others, which (very roughly) ties the content of a representation to the biological function of that representation's consumer. In short, something means what it has the function of meaning within a system evolved to solve adaptive problems. Variations on this general theme of information are also present in Dretske (1988), Maynard Smith (2000a; 2000b), Collier (2008), Skyrms (2010), Shea (2007; 2013), Bergstrom and Rosvall (2009), Anderson and Rosenberg (2008), Seyfarth et al. (2010), Robinson and Southgate (2010), Kight et al. (2013), and Lean (2014), to name just a few. It is also strongly represented in the ecological approach to cognitive science (Gibson 1966; 1979; Chemero 2003; Baggs and Chemero 2018), which treats the environment as containing information that is perceived by the organism as affordances for achieving its own biological aims. Similar appeal to use and purpose is at work in the context of cognitive science: Bechtel (2009), for one, argues that the proper conception of "information" is the one offered by control theory, which is concerned with the design of systems that exert control over factors affecting those systems by sensing and responding to changes in those factors. This covers an extremely diverse range of works, whose similarities and differences are too many and too diverse to discuss here. The key point at present is that they broadly agree on at least one thing: namely, they treat "information" as something that is being exploited by an agent (or something like an agent) to achieve some functional purpose.

I will call this general point of agreement the *use-consensus*, and this consensus will be the central focus of this chapter. As is sometimes but not always acknowledged, this is inherited directly or indirectly from Peirce's theory of signs and might be seen as an extension of Peirce's account of meaning to the living world in general (Short 2007). One way these views differ is in how exactly they understand the idea of function that fixes the meaning of a representation or even whether the "information" in question is best thought of as meaningful at all. I use the term *use-consensus* as an attempt to remain as neutral as possible about these issues—to capture what is shared between these approaches and to abstract from their differences. Rather than proposing yet another variation of this view, I offer a *reconstruction* of the use-consensus from the perspective of practice as an operational principle that organizes certain kinds of scientific investigation.

A key feature of my account is that it is closely modeled on Woodward's (2003; 2010) interventionist analysis of causation, which similarly aims to understand causal reasoning as a cognitive tool designed to serve various scientific aims. Ereshefsky and Reydon (2023) see Woodward's framework as similar in spirit to their grounded functionality account (GFA) of natural kinds: in their terms, reasoning about causes serves various scientific functions in ways that are grounded in the world. With this in mind, one might think of this chapter as something like a grounded functionality account of biological information—one that is explicitly compared and contrasted with Woodward's account of causal reasoning.

A benefit of this method of analysis, I hope to show, is that it brings to the foreground answers to the preceding conditions for a substantive account of biological information: First, it lays in sharp relief how this information is related to yet importantly distinct from the causal features of those phenomena. Specifically, it finds that causation and information need not be ontologically distinct in any strong sense; rather, they correspond to different ways of thinking about a given system for different purposes. Second, it shows when and why information can be said to be "literally" or substantively at work in biological phenomena. The answer it gives to this question embodies a form of naturalism in the spirit of Price (2011): since scientific observers are themselves living things, information "for us" is simply a special case of information "for the organism." In a sense, then, all information is biological information.

The chapter proceeds as follows. Section 2 outlines the interventionist account of causation as developed by Woodward, highlighting key features that I will borrow and adapt in order to explicate informational reasoning in similar terms. Section 3 then develops this framework for understanding informational reasoning: First, I characterize it in general terms that include non-biological cases of informational thinking using a simple case of scientific measurement. There I introduce the concept of an *ideal observation*, which makes sense of the informational aspects of a system and how these differ from both its causal and statistical features. Then I show how the same framework can be used to ground a substantive view of *biological* information. Section 4 concludes.

2. CAUSAL REASONING: CONTROL AND INTERVENTION

In this section, I outline key features of the interventionist account of causation, particularly the version defended by Woodward (2003; 2010; 2014)—henceforth *ICW*. As we will see, this framework is founded on explicitly pragmatic ideas: it aims to understand causality by asking what purposes the use of the concept serves in scientific practice rather than aiming to ground causality in a particular scientific or metaphysical theory. This focus on the function of causal reasoning, rather than on some intrinsic nature of causes as such, is embodied in the way it understands the commitments made by causal claims. It is this method of analysis that I aim to replicate in my analysis of the concept of information in the next section.

The first and most important feature of ICW is that it takes causality to be intimately tied to the notion of *control*: humans are not passive observers of our reality; we *act* in the world to effect changes, and our idea of causality guides reasoning about the differences our actions can make. Jenann Ismael (2017) expresses this motivation for interventionism in terms of affordances:[1] "To the embedded agent who doesn't just observe, but also intervenes in, his environment, the world is chock-full of opportunities and affordances. The terms in which he represents the world will be designed to disclose them. Causal relations are the generic form of these opportunities and affordances" (117).

The idea that causal reasoning reveals affordances for control is built into the way ICW analyzes claims about causal relationships. Consider three variables connected by edges, representing a physical system (Figure 4.1). These variables form a *directed acyclic graph* or DAG:

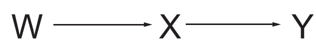

Figure 4.1. A simple causal chain represented by a directed acyclic graph (DAG).

In the preceding graph, the edges connecting *W* and *X* and *X* and *Y* represent the idea that *W* causes *X* and *X* causes *Y*. (The language of familial relationships is useful here: *W* is a "parent" of *X*, and *X* and *Y* are "descendants" of *W*.) Interventionists in general hold that no amount of probabilistic or statistical information about these variables can express the idea that they are *causally* related: there is no way to express that *X* is a cause of *Y*, for instance, in terms of conditional dependencies between the variables. Instead, to say that *X* is a *cause* of *Y* is to say that *Y* would change if *X* were manipulated from outside the system—that *changing* X *would change* Y.

Hence, causality on this view is essentially tied to the idea of *intervention*. Pearl (2009), whose work informs ICW, characterizes interventions algebraically: an intervention on *X* "breaks" its upstream dependencies on its parents and turns it into an independent variable. In contrast, a peculiar feature of ICW in particular is that it characterizes interventions as themselves causal processes of a very particular kind, which Woodward calls *ideal interventions*: An ideal intervention on *X* with respect to *Y* is "a causal process that changes the value of *X* in an appropriately exogenous way, so that if a change in the value of *Y* occurs, it occurs only in virtue of the change in the value of *X* and not through some other causal route" (Woodward 2003, 94). Treating interventions as themselves causal processes makes it possible to explicitly represent them, as shown in Figure 4.2.

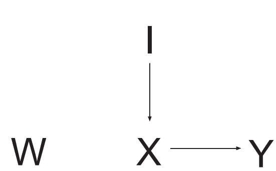

Figure 4.2. An intervention (represented by *I*) on *X* in the causal chain from Figure 4.1. In ICW, interventions are considered causal processes themselves, hence representable as such in DAGs.

Hence, for Woodward, to claim that *X* is a cause of *Y* is to claim that an ideal intervention on *X* changes *Y*, or at least the probability distribution over *Y*. Since the intervention detaches *X* from its causal dependency on its parent *W*, as seen in Figure 4.2, any remaining correlation between *X* and *Y* must be due to a causal relationship between them. This view of the meaning of causal claims aims to account for how scientists actually test and refine their causal models of the world (see Woodward 2014 for more detail).

As well as distinguishing causes from non-causes, this framework allows us to differentiate *between* causes of some effect along a number of different dimensions (Woodward 2010). For example, one cause variable is more *specific* than another if it exerts more fine-grained control over the value of the effect variable; that is, if it has a greater number of values that each specify a particular value of the effect. As mentioned previously, it is this *causal specificity* that Stotz and Griffiths (2017) claim underwrites the notion of informational relationships in biology—an issue to which I'll return in section 3.3.

ICW has met with several criticisms, the answers to which will be useful for understanding what follows. One class of objections centers around the view that it is overly anthropomorphic and subject-focused. We generally want to understand causality as something *objective* that actually governs nature independently of us and our beliefs about it, it is argued. Given that, it may seem that defining causality in terms of interventions implies that there are no causes in the natural world without agents who act on it. However, there is a way to understand how causal claims on ICW can be understood to be objective—that is, meaningful and true in a way that is, in an important sense, independent of the subject(s) making the claims. Since I will be adapting this aspect of ICW in my complementary analysis of information, I will take some time to address it.

The objectivity of causal claims in ICW comes from the fact that the interventions they posit are *idealized*. To idealize in scientific theorizing is to introduce assumptions that are false or unrealistic as a useful theoretical device. With that in mind, what false assumptions are made with respect to ideal interventions? Firstly, the ideal interventions central to ICW are essentially *counterfactual*: they posit differences in one and the same event if the value of a variable *were* different. This is an idealization because, by definition, it is not possible to realize and compare a set of mutually exclusive counterfactuals; instead, the best we can do is *approximate* this counterfactual experiment by repeating a process under controlled conditions in which only the (putative) cause is varied. Because they are counterfactual, then, ideal interventions are hypothetical (Woodward 2003, 40).

Secondly, ideal interventions are ideal in the sense that they are assumed to have an ideal degree of surgical precision: they are both causally and statistically unrelated to any other variable in the system. This is an idealizing assumption because it is always at least logically possible that there is a confounding factor that we have not controlled for or that our intervention had some other unintended effect. The best we can do, then, is to *approximate* these ideal interventions by controlling as much as possible. This account fits well with various aspects of scientific practice with respect to testing causal claims: For example, I've argued elsewhere (Lean 2020) that to maximize *binding specificity* in drug design is effectively to approximate an ideal intervention on the drug target.

It is these idealized features of interventions that render the associated causal claims objective in a specific but important sense. In particular, it makes sense of the idea that one thing can cause another whether or not we actually perform the kind of interventions necessary to test this claim. We can, for example, coherently talk about causal relationships in the deep past or between distant celestial bodies despite the fact that we can't go back in time and lack the power to divert the course of planets: when we discuss these causal relationships, we are discussing *hypothetical* interventions on the causes, not actual ones. For example, to say that a meteorite caused the extinction of the dinosaurs means that if, say, the meteorite had been diverted to miss Earth, the mass extinction would not have occurred. This has the important semantic consequence that the truth of causal claims is not relativized to the actual abilities and actions of agents: they are about manipulability *in principle*, not necessarily in practice.

In addition to making causal claims independent of our abilities, this also makes them independent of our beliefs: even if our beliefs about the causal structure of the world are based on well-executed controlled experiments, those beliefs can still be wrong. This is because the experiments we actually perform to test causal claims can only be approximations of the counterfactual experiments that those causal claims are about, and so actual experimental results do not logically entail the truth or falsity of causal claims. Of course, this disconnect between the *meaning* of causality and the empirical observations we take as *evidence* of causal connections is central to Hume's discussion about the concept. Nevertheless, causality is central to our reasoning about the world despite these worries (as Hume observed), and ICW provides a rich account of this central role as it manifests in scientific practice.

Broadly, then, the effect of idealizing the interventions posited in causal claims is that it removes from those interventions any explicit reference to agency or purposeful action. It is "heuristically useful," Woodward says, to think of an ideal intervention as the action of an agent; however, it is possible to characterize ideal interventions without endowing them explicitly with agency; instead, we can simply stipulate causal and statistical conditions. If an action is a relationship between an agent and the system on which they act, successful control depends on both. In idealizing the role of the agent in that relationship—that is, granting all the necessary statistical and causal features of the intervener—the absence of an effect can therefore be attributed to a lack of causal power in the target variable, not to any failure to intervene properly on it. This feature, in my view, strikes a careful balance between acknowledging the essential role of agency and purpose in causal reasoning while maintaining a sense of contact with an external reality. This, I take it, is why Ereshefsky and Reydon (2023) connect ICW to their grounded functionality account of natural kinds: While causal reasoning—indeed, *all* reasoning—is tied to the purposes for which we as agents engage in it, in doing so we are tapping into features of the world on which we depend for those purposes to be satisfied. (See Bausman 2023 for an exploration of this "tapping into" notion in terms of adaptation.)

To summarize, ICW begins by supposing that causal reasoning is designed to serve certain purposes and then develops its account in virtue of what purposes it serves and how. First, it holds that we distinguish causal from noncausal relationships because only the former are affordances for control and that causal models are designed to reveal these affordances in a general-purpose format. Second, causal information is expressed—can only be expressed—as information about how a system would change under various interventions from outside it. Third, the interventions contained in the meaning of causal claims are idealized in various ways in order to lend objectivity to our causal models; we aim to make claims that are true or false independently of our beliefs about a system or our ability to intervene on it. These hypothetical interventions are impossible to perform; nevertheless, in forming and testing causal hypotheses, we aim where possible to *approximate* those ideal interventions through controlled experiment.

3. INFORMATIONAL REASONING: COORDINATION AND OBSERVATION

The previous section sketches the key features of Woodward's interventionist framework for causality that I will adapt in developing a general account of information of the kind represented by the use-consensus. A key lesson of ICW, as I understand it, is that the concept of causality is *pre-theoretic*: Calling a relationship causal per se does not commit to some specific theory, from physics or somewhere else, of what grounds or underwrites that relationship. Instead, what we mean when we make such a claim is that it is a potential "lever" for bringing about changes. Of course, much of scientific inquiry involves figuring out what kind of levers they are—for example, by developing mechanistic models for how that phenomenon is realized. However, its being causal *qua causal* is independent of these specific details. What's more, focusing in on these details does not reduce away the "intervention" aspect: Firstly, because intervention remains essential for distinguishing the relevant from the irrelevant properties in those models with respect to the phenomenon in question. Secondly, because abstraction is a vital aspect of explanatory generalization even when the more concrete details are available (Levy and Bechtel 2013).

In short, I take causality in ICW to be more than just a "thin concept" (Cartwright 2005)—that is, more than a placeholder for a range of richer, more informative, context-dependent "thick" concepts. Instead, I take causality in the interventionist sense to be an indispensable abstraction: it is the feature that those relations all share despite their differences—a feature that is vitally important to agents who interact purposefully with the world around them.

This idea will also be an overarching motivation for the complementary account of *information* that I develop here. There is already precedent in the literature for the analogous approach to "information"—that it need not have a predefined theoretical underpinning in order to be useful for empirical inquiry. This view is held by Beckett Sterner (2014), for example: following Bergstrom and Rosvall (2009), Sterner argues that it would be antithetical to biological inquiry to define "information" in a way that establishes its extension based on a priori theoretical principles about semantics. Instead, the concept of information serves simply as a diagnostic tool: it is understood in a way that guides and constrains how we gather and organize empirical cases. Given this, the only conceptual assumptions we should bring into our empirical inquiry are those that tell us how to recognize informational processes when we see them. Crucially, those processes once recognized can and should reciprocally update our theoretical notions about what biological information amounts to.

In a similar vein, Kelle Dhein (2020) details a fascinating case study illustrating how and why ethologists ascribe semantic content to behavior. His case focuses on the study of navigation by ants of the genus *Cataglyphis*, who are surprisingly adept at finding the shortest possible route back to their nest despite long and meandering outward journeys. A research program built around these insects has been highly successful at explaining the ants' navigation ability: as it turns out, ants achieve this through a process of "path integration"—by storing a vector representing the distance and direction from the nest that is constantly updated during its journey. It is important to note, however, that the scientists did not begin this process with a particular theory about semantic information and how it manifests in neurophysiology. Instead, the first step is simply the identification of a surprisingly reliable match between an animal's behavior and its circumstances—in this case, the ant's consistent choice of the shortest journey home. Given a surprising observation of this kind, the search for an explanation is guided by the notion that the animal must possess some means of reliably choosing a successful behavior. The search for an explanation for this phenomenon is understood as the search for an adequate information channel exploited by the ant, whatever precise form that channel turns out to take.

I will return to these works in due course. For now, the key point is that, as Sterner and Dhein both illustrate, the use of "information" in these biological contexts does not depend on some precise scientific or philosophical definition of the concept. Instead, it is an operational principle that guides inquiry into a certain kind of biological phenomenon. The phenomenon in question is what I will henceforth call *coordination*—a concept I will elaborate later. For now, I intend this term to capture, as generally as I can muster, any kind of fortuitous or adaptive "match" between a functional entity or process and its circumstances—one that is surprising enough to invite investigation into how that match comes about. "Information" refers to whatever turns out to play a particular role in that process of coordination. A significant feature of the account I develop is that it elaborates on this idea in terms adapted from Woodward's analysis of causal reasoning: informational reasoning pertains to coordination *in the same way that causal reasoning pertains to control.*

One implication of this is that biological information (in the sense shared by the use-consensus) is not fundamentally different from the everyday human uses of "information" that we tend to find less controversial; rather, they are all instances of a general type of reasoning that I characterize here. In other words, information "for us" (human observers) is simply a special case of information "for the organism." To drive this point, I will take a detour from explicitly biological examples and develop this account of informational reasoning, and its relationship with causal reasoning, in a human context. Following that, I will return in section 3.3 to the case of information in biological practice to show how it exemplifies that same form of reasoning.

3.1 Ideal Observations: The Case of Measurement

Consider a familiar physical system of a simple set of mechanical scales, shown in Figure 4.3. This system consists of a variable *load* representing different weights that might be applied to the scale and a variable *readout* representing different values shown by the dial.

Figure 4.3. A DAG representation of a generic scale. By design, the load placed on the scale causally affects the value displayed on the readout.

To simplify, assume that the possible weights take integer values (e.g., *load* = [0 kg, 1 kg, 2 kg, . . .]) and that the output values are similarly discrete (e.g., *readout* = ["0 kg," "1 kg," "2 kg," . . .]). As one would hope, there is a causal relationship between these two variables: under some range of normal background conditions and within a given range of variation, changing the weight on the scale moves the dial on the readout. This causal relationship is what the edge between the variables represents.

Now, consider the obvious point that this particular physical setup is of a special kind: it is a *measuring instrument.* Hence, we can ask the following question: Is what we've said about this system so far sufficient to account for what makes it a measuring instrument? The answer seems to be "no": in general, for some *Y* to count as a *measurement* of some *X,* it isn't sufficient that *X* causes *Y.* The reason is stated simply by Peter Kosso (1992) (in the context of "observation"): "An observation, *as it is to contribute to knowledge by motivating and testing claims about the world,* must be an epistemic event and not just a physical event. Not only must the object causally affect the viewer, it must produce an *informative effect* of some sort" (26, emphasis added).

In other words, while measurements (and observations) may certainly involve a causal relationship between the input and the output, measurement is more than that: it is essentially about producing *knowledge* about the measured object. A similar sentiment is expressed by van Fraassen (2008): the outcome of a measurement is a *representation,* with all the intentionality that entails and which causal relationships, qua causal, are lacking. This is one reason I begin my analysis of information with a human example: in biological cases, it may be possible in principle to understand biological processes without intentional terms, opting instead for more physicalist language of causal role functions and so on. But this, I think, is a symptom of the fact that in those cases we are viewing both the object and its perceiver from the third person, from which it's easier for us to see both as just physical systems that are causally related in some way. It is harder to justify when we and our measurements *are* one of those systems. Embedded in the here and now of contemporary scientific practice, we cannot avoid taking the results of our measurements—and things derived from them such as data models—to be *about* something (van Fraassen 2008; Finkelstein 1994; Cropley 1998a; 1998b).

What must we add, then, to make sense of this measurement setup *as a measurement setup*? As suggested by Kosso and van Fraassen, we must clarify the notion of a measurement not merely as a causal process but as a conduit for *information*—of the outcome having an "informative effect." With that in mind, an obvious starting point would be to point out that (under certain assumptions) the load's causal influence on the readout implies that the two will share *mutual information* according to the mathematical theory of communication, or *MTC* (Shannon 1948; 1949). This is the case whenever the values of two variables are correlated; in that case, it's said that one variable "reduces uncertainty" about the other, in the sense that knowing the value of one variable allows you to make a more reliable guess at the value of the other.

Yet there is reason to doubt that this fully captures the informative effect we're looking for. To see why, suppose that the instrument is poorly *calibrated*, that it consistently overestimates the load by, say, 1 kilogram. Crucially, the miscalibration is invisible to the MTC measure of mutual information: as long as every value of *load* corresponds to exactly one value of *readout,* mutual information is at a maximum and there appears to be nothing wrong. Yet calibration is a necessary part of good measurement procedures (Soler et al. 2013), and so neither causality per se nor the mutual information that it underwrites is sufficient as a means of evaluating the scale's adequacy as a conduit for information.

Or so it may seem. There is a conflicting intuition: Surely as long as we *knew* that the readout overestimates the load, we could use the instrument without issue simply by subtracting 1 kg from the indicated value after the fact. To point this out is to say that, *in principle*, there is nothing wrong with the scale itself: the information has made it to the readout intact, though we may lack the ability or background knowledge to properly acquire it. In this sense, the problem is instead with the *interpretation* of its output by a user. Yet the fact remains that not all bangs and scrapes are lessons; that is, merely being causally in contact with the world is insufficient for being *informed* about it. To be informative, our measuring instruments must be set up so that proper interpretation is possible.

To summarize, we seem torn between two ways of looking at this situation: On the one hand, there is "objectively" nothing wrong with the measurement setup. On the other hand, this objective state of affairs is importantly beside the point: measuring setups must take into account an actual subject or user and make it possible for that subject to acquire knowledge. If it doesn't, the setup is inadequate.

As I will show, we can resolve these competing intuitions by considering information as essentially tied to its use. Again, this in itself is far from new. My contribution is to show that this use-consensus about information can be neatly expressed by adapting the conceptual tools of ICW. Recall that ICW's conceptual device of ideal interventions resolves a similar tension about causality: namely, that it is essentially tied to the notion of intervention from without yet at the same time an "objective" property of systems independently of any actual interventions. I propose that information of the sort discussed here can be understood in an analogous way: as a concept, it is essentially tied to the agential notion of *observation,* yet in an important sense it is objectively "there" to be observed.

Following ICW as closely as possible, we can express this idea using directed graphs, albeit with some embellishments. Consider again our miscalibrated scale represented by *load* → *readout*. We've established the sense in which their relationship is causal. Our aim now is to analyze the claim that the readout carries *information* about the load in the sense of being a source of knowledge to a user—the as-yet-undefined relation in Figure 4.4.

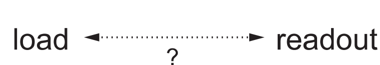

Figure 4.4. The scale is designed so that the readout carries information about the load. What exactly defines this informational relationship?

As I've argued, this sense is not fully captured by the mutual information measure of MTC, since that measure is blind to issues such as calibration that are essential to the notion of a good measurement setup. So if that will not do, how should we interpret the relation instead? Recall that, despite being miscalibrated, there is a sense in which the scale is adequate as a measuring instrument. The intuition is that, *in principle,* it's possible to learn about the load from the readout, and in that sense there is "objectively" information to be acquired about the load whether or not anyone manages to acquire it. To turn this from an intuition into a theory, we need to consider more explicitly the agent that can be said to be "learning" about the load. We can represent this as a variable *A* ("agent"), as in Figure 4.5.

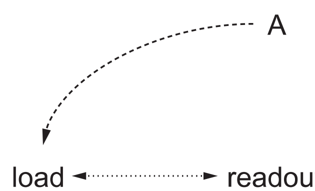

Figure 4.5. An agent *A*'s aim is to learn about—coordinate with—the readout. This coordination relationship is represented by the dashed arrow.

The relationship between *A* and *load* noted previously is what I refer to as *coordination.* We can now clarify this further: *A* refers either to an agent or something used by an agent—for example, a system of representations—whichever is the most useful way of depicting the situation in a given scientific context. We can think of coordination as pertaining to a function $f: A \rightarrow$ *load* that states which value of *A* is appropriate for which value of *load.* One could interpret this as a semantic correspondence relationship in the classical sense—as *load* giving the truth conditions of every value of *A*. However, my aim here is to also include other types of relation that are not (or not obviously) semantic. One non-semantic relation I wish to include is *adaptation:* An adaptive phenotype is not "about" its environment, yet it is *coordinated* with it in the sense I use here. For the sake of characterizing informational reasoning in the abstract, the details of particular instances of this relationship are irrelevant; in any case, the idea that coordination is designed to capture is a match of some kind between the value of *A* and the value of *load.*

A helpful way to understand this idea of coordination is that it relates to ICW's idea of control by inverting its direction of fit. While control is changing the world to suit the agent (or some extension of agency), coordination is fitting the agent to the world. Measurement clearly fits into the latter class: while we may test and calibrate an instrument by manipulating the input to determine its relation to the output, its use *as a measuring instrument* ultimately means pointing the setup at an as-yet-unknown target and hopefully getting a correct value.

With that in mind, suppose that the goal in this case is a coordination between *A* and *load,* and the question is how that might be achieved. The ideal case, of course, would be to have the value of *A* be directly determined by *load.* Yet, as is often the case, we lack such direct access to the object of our interest. In this case, of course, the hope is that the value of *readout* provides this information—that one can instead form true beliefs about the load by conditioning those beliefs on what the readout says. Here we can begin to make this claim both more precise and more general: Whether *readout* contains information about *load,* I propose, is in effect a question of whether *readout* can be exploited to produce a coordination between *A* and *load.*

As we've seen from the previous issue of calibration, there is an apparent tension between two ways of interpreting this question; specifically, about whether to understand this exploitability in relation to actual users, or in a more in-principle sense that is independent of actual users' contingent limitations. Importantly, there is an analogous tension at work in the interventionist notion of causation, and both can be resolved in the same way: recall that causal claims aim to be objective in the sense of not being relative to the contingent knowledge or abilities of actual agents, which is why they are implicitly about *ideal* interventions. I claim that the same thing applies to informational reasoning: In reasoning about the information that is "objectively" in a system, we are in effect reasoning about *ideal observations.* An ideal observation is shown in Figure 4.6.

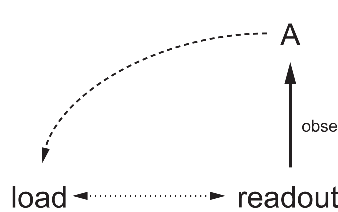

Figure 4.6. By *observation* of the readout, the agent *A* is able to coordinate with the load. This observation is a causal relationship of a particular kind between the readout and the agent.

As the name suggests, ideal observations are information's counterpart to ICW's ideal interventions. Though they share the property of being causal relations, the direction is reversed: while interventions determine the value of their target, observation variables are *determined by* the target of the observation. In other words, while interventions turn their target from a dependent variable in the system to an independent one, observation turns *A* from an independent to a dependent variable. With this in mind, claims about information are about whether and to what extent such ideal observations improve *A*'s chances of coordination; that is, whether hooking *A* up to *readout,* causally speaking, improves the chances of satisfying the coordination function relative to some prior distribution of *A*'s values.

I will avoid committing to certain particulars in how to interpret this question; for example, which interpretation of probability is at work may depend on the particular context to which this general framework is applied. However, there are certain features of ideal observations that should be specified to allow them to do the work required of them; that is, to render the claims about information that they involve "objective" in an important sense.

First, and most obviously, the observation variable must be ideally causally sensitive to its target variable; that is, it should adopt exactly one value for each value of the target. (This allows that more than one value of the observer may be "correct" for the same value of *load*; this is discussed further in section 3.2, but not the reverse.) This is analogous to the condition for ideal interventions that they fully determine the value of their target.

Second, relative to the system in question, the observer variable *A* must be causally independent of every other variable except for the observed one. This requirement is analogous to the condition of ideal interventions that they *uniquely* determine their target and is required for analogous reasons: With interventions, this is to guarantee that any resulting change in the effect is exclusively due to the intervention on the cause. Similarly, ideal observations should guarantee that any resulting coordination with the variable of interest—with *load* in the present case—is exclusively due to the observation of *readout* and for no other reason.[2] For example, suppose we wanted to put palmistry to the test. We may accept that a palm reader can find things out about their subject through a combination of conversation, body-language, "cold reading," and so on. But what we really want to know is whether they can learn anything significant about the person specifically by attending to particular features of their hand. Any controlled test of this claim will have to "block" those other potential sources of information, ensuring that anything the reader can learn about their subject (or at least any *improvement* in their knowledge) is for the claimed reason alone.

Finally, as well as simply being causally sensitive to its target and only to its target, the details of this causal relationship must be of the "right" kind; that is, it must have the right observation function *f: readout* → *A*. The purpose of this condition is to account for cases of miscalibration as discussed previously: an actual observer may be fully sensitive to the readout, yet assume by default that when the readout says "2 kg," the load is 2 kilograms. An *ideal* observer, however, is by definition one that makes the necessary correction. This captures a notion of error that will be developed more later: the difference between the information the readout *appears* to carry, from an actual user's subjective viewpoint, and what it objectively carries. The idealization simply posits an observer with the necessary background knowledge, in other words, the right observation function.

The overarching purpose of the concept of an ideal observation is to capture the idea of information being an "objective" property of the world: on this framework, to ask whether some *X* carries information about some *Y* is to ask whether it is possible for some agent *with some ideal set of capacities* to coordinate with *Y* by conditioning its state or behavior on *X*. Here lies a critical point: very often, it will simply not be possible for *X* to be used to coordinate with *Y*, no matter what imagined capacities an observer of *X* has. For example, if the scale were broken and the load did not affect the readout at all, there would be no possible way to learn about one by observing the other: changing this fact would mean changing the system itself, not the observer of it. Hence, tying claims about information to idealized observations does not make them trivially true: sometimes information simply doesn't exist to be exploited.

However, that there is information in the world to be exploited in this sense does not mean that actual agents (for example, living things) are able to exploit this information. Nevertheless, ideal observation remains an important concept because it establishes a sense in which there "really is" information to be acquired, independently of whether there are in fact any observers with the right properties to acquire it. Of course, determining what is possible for an ideal observer requires a third-person viewpoint from which we can "see" the whole system—that is, from which we can see both the observed variable and the variable of interest independently. From there, we can determine what kind of relationship (causal, geometric, and so on) exists between the two variables and what capacities an agent would need to have in order to tap into this relationship for some purpose.

Overall, this framework aims to express the sense in which informational reasoning is tied to agency, in just the same way that ICW does. Both causal and informational reasoning are designed to suit the purposes of the agents engaging in that reasoning, hence the presence of interveners and observers in the analysis of the corresponding claims. However, in *idealizing* these interventions and observations, they end up losing their agent-like properties: Intervention and observation variables are simply causal variables with certain particular features. Nevertheless, both causal reasoning and informational reasoning remain intimately tied to agency in that they are designed to reveal affordances—affordances for control in the case of causal reasoning and for coordination in the case of informational reasoning. Yet, as Ismael (2017) argues, causal models reveal these affordances in "generic form": they are not explicitly purposeful because they aim to be neutral about what they might be used for. The same, I claim, applies to information: it is possible to model the ways in which a system *might* conceivably be exploited for coordination without committing to particular purposes for doing so, hence the existence of non-semantic, or *purpose-neutral*, measures of information such as MTC.

A graphical summary of the difference between causation and information is shown in Figure 4.7. This analysis depicts informational reasoning as intimately related to but importantly distinct from causal reasoning: In reasoning about the causal structure of a system, we are reasoning about how it would change if we acted on it in various ways. In contrast, to reason about the information in a system is to reason about what we might learn about one part by observing another. In short, they are two different ways of thinking about a system for different purposes. Given this difference in the basic function of the two reasoning types, each has its own hypothetical, idealized interaction with that system within the types of claim being made. In this sense, then, causation and information are *mutually irreducible*.

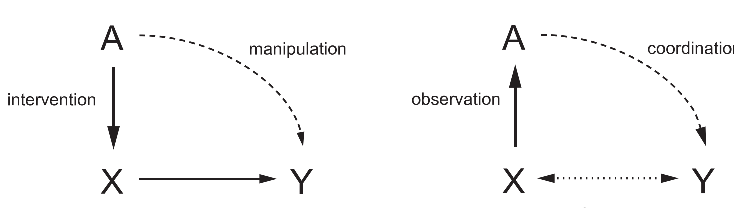

Figure 4.7. The generic form of causal and informational relationships: If *X* causes *Y*, an ideal agent can manipulate *Y* by intervening on *X*. If *X* carries information about *Y*, an ideal agent can coordinate with *Y* by observing *X*.

3.2 What Is the System For—Coordination or Control?

In section 3.3, I'll relate the framework for informational thinking I've developed here to the notion of information in biological practice. Before that, however, I'll offer some further illustration of why informational thinking is in an important sense distinct from causal thinking; that is, why information does not *reduce* to causation in any meaningful sense. This is evidenced by the fact that causal and informational reasoning *evaluate* one and the same physical system in distinctly different ways: a given system may be perfectly adequate as a tool for control and yet *in*adequate as a tool for transmitting information, or vice versa. For one, as is well known, two things don't have to be causally related for one to carry information about the other; correlations without a direct causal link from one to the other are informational nonetheless; for example, they can be related by a common cause. Yet even when the relation in question *is* causal, as with the preceding case of the scale, evaluating that system in terms of the information flowing through that causal process requires a different set of conceptual tools.

To illustrate, consider the causal relationship represented in Figure 4.8. The figure represents not just the causal relationship between *X* and *Y* but which *value* of *Y* results from each possible intervention on *X:* an intervention setting x_1 produces y_1, while x_2 and x_3 both bring about y_2.

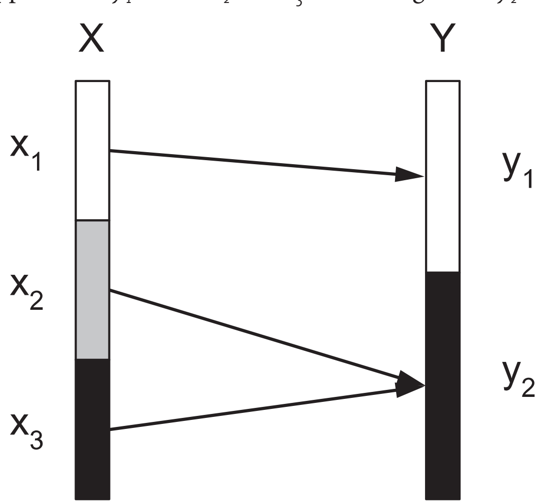

Figure 4.8. A causal relationship between *X* and Y, showing the possible values of each and how they relate. Here, *X* taking value x_1 causes Y to take value y_1, while both x_2 and x_3 cause y_2.

First, consider this system from the causal perspective. That is, let us think about this system in terms of its adequacy as a means of *controlling* Y. If this is the aim, then the ideal case would be one in which for any value of Y we might want to bring about, there is an intervention on X that can reliably do so. This is indeed the case in Figure 4.8: there is no value of Y that cannot reliably be brought about by some intervention on X. Hence, provided one *could* intervene on X—remember, whether we actually can doesn't bear on the system's objective features—one could fully control Y. (In fact, if we had such a means at our disposal, we may even come to think of ourselves as setting the value of Y directly!) We can say that x_2 and x_3 are causally *redundant,* and causal redundancy does not threaten control.

Now consider this same system's viability as a means of *transmitting information*—that is, consider the extent to which an *observer* of the effect Y can potentially be informed about (coordinate with) the value of the upstream cause. Again, it is useful to consider the ideal case: a system is an ideal conduit for information when any observed value of Y corresponds to exactly one value of X. (Analogously to what was shown previously, this may lead us to talk as though we have observed X directly.) Importantly, from this perspective the preceding case is *not* optimal: while observing y_1 guarantees that $X = x_1$, observing y_2 only tells the observer—*even an ideal one*—that $X = x_2$ *or* x_3. To use Dretske's (1981) term, y_2 *equivocates* between x_2 and x_3. The system in Figure 4.8 is therefore adequate as a means of manipulation but inadequate as a conduit for information: information is *lost* between X and Y.

When we consider the reverse case in Figure 4.9, the reverse will be true: In this case, setting x_1 uniquely produces y_1, but setting x_2 can produce either y_2 or y_3. This situation is suboptimal for purposes of control: if we want y_2, there is no way to reliably bring it about through intervention on X. As a conduit for information, however, nothing is wrong: an ideal observer of Y could always coordinate with X perfectly. In doing so, we might interpret the ideal observer as treating y_2 and y_3 as semantically equivalent—or, more generally, by adopting the same state in response to both.

The view I've developed here embodies a kind of dualistic perspective on the relationship between information and the causal properties of a system, though one that should not ruffle a naturalist's feathers. On the one hand, we needn't think of information as a kind of substance that is separate from the physical features of a system: talk of information "in" or "flow-

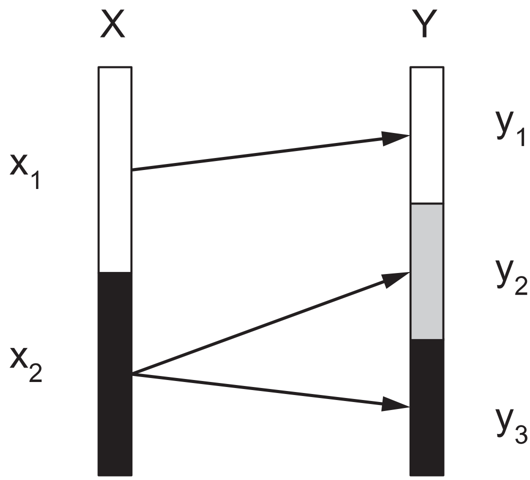

Figure 4.9. In contrast to Figure 4.8, setting X to x_2 can result in Y taking values y_2 or y_3. This situation is also reversed in terms of its adequacy as a conduit for information or control.

ing through" a system can be understood simply as a way in which that system can be used. On the other hand, different kinds of reasoning about a system suit different kinds of use to which that system can be put. So in some sense at least, the two forms of reasoning are designed to access different properties or aspects of that system.

3.3 Coordination Functions in Biology

As I've argued so far, causation and information can be thought of as alternative ways of thinking about a system that each suit different purposes. This makes my account of information and its relationship to causation explicitly *functional*: which perspective is most appropriate to take on a phenomenon depends on what we aim to achieve by doing so. So far I've developed this account in a nonbiological context involving humans and their artifacts. I turn now to how this framework can be used to understand at least one notion of information at work in biological practice. As discussed, a sticking point in the discussion of biological information is in whether we can attribute to information a literal or substantive place in the biological processes we study—whether it is ever information "for the organism," not just for us. Since biology in general is deeply concerned with the purposes or functions of the systems it studies, biological information lends itself well to being understood in functional terms.

As we've seen, the use-consensus broadly takes biological information to be something that is exploited by a biological system in carrying out some function. Sterner and Dhein both put an explicitly pragmatic spin on this idea: Scientists find it useful to *attribute* information to biological phenomena satisfying a certain set of criteria, and doing so serves certain investigative purposes including (but not limited to) guiding investigation that elucidates the mechanisms for those phenomena. In turn, the "information" at work in these phenomena is whatever turns out to play a particular kind of role within those mechanisms. In that sense, information in these practices is not defined by a theory—whether biological, physical, or mathematical. Rather, it is an operational concept that leaves open how particular phenomena are realized and that also permits abstraction and comparison with functionally similar but differently realized phenomena in other species.

I've used the term *coordination* as a label for the kinds of biological phenomena for which this sort of informational reasoning is appropriate. In line with Sterner and Dhein, I hold that (what I call) coordination is presumed by or a precondition of discussions of information in these biological contexts. Dhein's case study of insect navigation is one example: There, investigation begins with the understanding of the coordination the ant is achieving. But biology is replete with others: For instance, the idea of *positional information* (Wolpert 1969) was introduced in the context of explaining how individual cells in a developing embryo differentiate into the "correct" cell type given its position in the mass. Elsewhere, Bechtel's (2009) discussion about the application of *control theory* to cognitive science comes from the idea that brains are functional systems whose sensory capacities serve to detect and respond to environmental circumstances that matter to the organism. What these cases and others have in common is that there is a range of possible states—cell types, behaviors, and so on—that are each appropriate under different circumstances relative to some goal or function. The thing to be explained is how the "correct" choice is reliably made despite the many ways to be *un*coordinated.

However, the reasoning process needn't necessarily *begin* by specifying the coordination relation in question: Sometimes, we begin with a behavior and then ask what state of affairs that behavior is supposed to match with; that is, what its informational function is (Lean 2014). For example, the debate about stotting in gazelles amounts to a discussion of what coordination function it serves: Is it used by other gazelles to indicate danger or by the predator as a sign that the gazelle is too nimble to bother chasing? This is the question whose answer establishes the informational content of stotting as a signal or indeed whether it is a signal at all. So while the study of insect navigation went from an agreed function to the discovery of mechanisms for it, in the case of stotting the "how" was established before the "why." Yet whatever the order, the notion of coordination is a necessary background for the reasoning process. Given some coordination—whether established or simply posited—*information* is whatever makes that coordination possible.

Here we can see why information should not be tied to some particular theory such as MTC: in the case of insect navigation, for example, ants and bees are able to record information about their direction of travel by sensing polarized light. The fact that light polarization can be taken to indicate direction of travel is not readily comprehensible as a statistical relationship between the two, but as something like a *geometric* relation between the light and the ant's orientation.[3] This example illustrates that it is not just statistical regularities in the world that can support reliable inferences and actions; this can also be done by phenomena as diverse as landmarks, one-off causal events, spatial relations, and so on.[4] Yet whatever it is that underwrites the coordination, *what makes that feature of the world "information" is that it is so exploited or that it can be exploited by an ideally endowed observer.* In other words, information can be thought of as an abstract property covering a range of relations in the world that can be exploited to serve coordination functions. Again, this is analogous to my previous claim about causation—that it is an abstract concept linking disparate phenomena by their common capacity to underwrite control.

The third key concept in the triad—observation—is the relationship between the property bearing this informational relationship with the variable of interest on the one hand and the *observer* on the other. My use of "observer" need not imply a strong claim about rich intentionality; we can think of it simply as the thing that is causally responsive to the information carrier, and whose response to that carrier can be seen as bearing a functional match to the circumstances (see Lean 2014). This could apply to part of a mechanism within an organism, such as an *adapter* that links an input to an output for functional purposes (Lean 2019).

In any case, the relationship I call observation is a key component of practices that investigate coordination phenomena in biology. In fact, the *absence* of such a relationship can potentially refute hypotheses about the mechanisms that explain cases of coordination. For example, even if there exists a relationship between light polarization and direction of travel in an insect, this would be irrelevant to the explanation if insects lacked the ability to detect polarized light. That they in fact possess this ability required empirical demonstration of the organism's sensitivity to that feature (Wehner 1997; Rossel and Wehner 1984a; 1984b; Wehner and Müller 2006, all cited by Dhein 2020). Yet by appeal to an *ideal* observer, we can distinguish the information that is objectively "out there" from the subset of that information that is actually playing a role in biological functions. The latter, I suggest, is information in the literal or substantive sense for biological purposes.

Finally, in reconstructing the use-consensus's notion of information in the same way in which ICW understands causation, we have at our disposal a clear sense of how the two are related. This outcome prompts a subtle but important shift in how we think of biological information relative to those who consider it in purely causal terms. In particular, recall the claim by Stotz and Griffiths (2017) that informational relationships in biology are simply causal relationships with high specificity. With the preceding framework in mind, we can interpret the relation between information and causal specificity somewhat differently: rather than saying that informational relationships *are* specific causal relationships, we can instead ask why causal specificity might be useful in systems that are used for informational purposes. Woodward (2010) argues that specific causal relationships are useful in situations when we want fine-grained *control* over some effect. Evidently, this is reasoning in the causal mode—the one that considers the effect of actions. Alternatively, we might instead consider specific causal relationships in the informational mode, in terms of observations and the coordination that they afford. From that perspective, specificity is valuable because it allows fine-grained *discernment* between possible states of the variable of interest, and hence the ability to adjust oneself to those states in a fine-grained way. Similar translations into the informational perspective may be applied to Woodward's other causal properties of stability and proportionality.

In summary, I argued in section 3.1 that while ICW takes causation to be affordances for control, informational reasoning can be understood as reasoning about the affordances a system offers for coordination—about what can be learned, in principle, about one part by observing another. Coordination with one's circumstances through observation—through conditional response—is a basic necessity for agents alongside that of manipulating the world through action. This view lends itself to making sense of when biological systems are trading in information in a literal or substantive sense: this is the case when an organism, or some functional biological system in general, is known or believed to involve some coordination function—to be successfully conditioning its behavior on its circumstances. Relative to that coordination relation, information refers to the relation being exploited by that system to achieve a successful match. In short, biological information exists when something like agency, or at least function, is being attributed to the system. This, I argue, distills in the most general terms the notion of information shared by the use-consensus.

4. CONCLUSION

The methodology by which this chapter relates causation and information aims to embody the lessons of what is sometimes called the "practice turn" (Soler et al. 2014): that is, it aims to consider scientific reasoning as essentially practical, as something we *do* in order to achieve goals in a certain context. This approach is also taken by ICW: it is based on the notion that concepts like "cause" serve certain purposes and can only be fully understood by considering those purposes. In turn, then, the very *origin* of the notion of a cause—the reason we have such a notion in the first place—is because we are actively engaged with the world we're trying to understand. Causality, as we understand it, is intimately connected to *action:* causes are things that can in principle be acted on to influence the world.

I intend my analysis of information in this chapter to be not just analogous but complementary with that view: just as causal reasoning ties to action, informational reasoning ties to *perception:* Information is the world's influence *on us,* insofar as it guides or constrains our inferences and decisions. Just like ICW, however, the aim of objectivity in scientific inquiry leads us to consider information as something that is "out there," in a sense, independently of those inferences and decisions. To meet this need, claims about information can be understood as claims about what a system offers to an *idealized* observer—one that is free of the contingencies and limitations of actual agents. This, I've argued, is complementary to ICW's use of idealized interventions, which serve to express causality as an objective feature independent of actual interveners.

From a naturalist point of view, this conclusion about the complementarity of causal and informational reasoning is unsurprising: Perception and action are the two means by which agents (qua agents) interact with the world around them, and so it makes sense for each to be associated with a distinct type of reasoning engaged in by those agents. Yet those types of reasoning, while in a sense distinct, are intimately related because perception and action are ultimately inseparable: The two jointly constitute an iterative feedback loop, and hence each can only be understood in connection with the other (Dewey 1896; Hurley 2001). Nevertheless it is possible and often useful to decompose that overall process and to view it either from the causal or informational perspective; that is, to view a system either as a target of action or as a conduit for the information that guides action. The fact that biological sciences often treat their objects in at least quasi-agential terms—as doing things for reasons—underwrites what is distinctly biological about biological information.

NOTES

1. Ismael's use of "affordance" is evidently quite different from its use in ecological psychology—a notion I cannot discuss at length here. It should also be noted that Ismael does not subscribe to Woodward's particular interpretation of interventionist causation.

2. However, we can allow a statistical relationship with other variables pre-observation. For example, when one variable screens off another, it may be relevant to point out that the screened-off variable offers no information to some observer over and above what has already been acquired. This may make information claims relative to some prior state of an observer, but it needn't make them subjective.

3. Of course, this evolved system means that light polarization then constrains ants' direction of travel with high probability, which *can* be understood in causal and statistical terms. However, the geometric relationship holds independently of any evolution in the ant's sensory capacities and is the *reason* why those sensory capacities evolved in the first place. Hence, there are two different relations at work here even if the relata are superficially similar.

4. For this reason in particular, the framework I develop should not be taken to lean heavily on the use of directed graphs, which are primarily designed to represent probabilistic relationships. Other means of graphically representing this relationship may be appropriate for other types of coordination phenomena.

REFERENCES

Andersen, H. 2017. "Patterns, Information, and Causation." *Journal of Philosophy* 114, no. 11: 592–622.

Anderson, M. L., and G. Rosenberg. 2008. "Content and Action: The Guidance Theory of Representation." *The Journal of Mind and Behavior* 29: 55–86.

Baggs, E., and A. Chemero. 2018. "Radical Embodiment in Two Directions." *Synthese*. Berlin: Springer.

Barbieri, M. 2016. "What Is Information?" *Philosophical Transactions of the Royal Society A* 374: 20150060.

Bausman, W. C. 2023. "How to Infer Metaphysics from Scientific Practice as a Biologist Might." In *From Biological Practice to Scientific Metaphysics,* edited by W. C. Bausman, J. K. Baxter, and O. M. Lean. Minneapolis: University of Minnesota Press.

Bechtel, W. 2009. "Constructing a Philosophy of Science of Cognitive Science." *Topics in Cognitive Science* 1: 548–69.

Bergstrom, C. T., and M. Rosvall. 2009. "The Transmission Sense of Information." *Biology and Philosophy* 26, no. 2: 159–76. https://doi.org/10.1007/s10539-009-9180-z.

Cartwright, N. 2005. "Causation: One Word, Many Things." *Philosophy of Science* 71, no. 5: 805–20.

Chemero, Anthony. 2003. "Information for Perception and Information Processing." *Minds and Machines* 13: 577–88.

Collier, J. 2008. "Information in Biological Systems." In *Handbook of the Philosophy of Science, Vol. 8: Philosophy of Information*, edited by P. Adriaans and J. van Benthem, 763–87. Amsterdam: Elsevier Science.

Cropley, D. H. 1998a. "Towards Formulating a Semiotic Theory of Measurement Information—Part 1: Fundamental Concepts and Measurement Theory." *Measurement* 24: 237–48.

Cropley, D. H. 1998b. "Towards Formulating a Semiotic Theory of Measurement Information—Part 2: Semiotics and Related Concepts." *Measurement* 24: 249–62.

Dewey, J. 1896. "The Reflex Arc Concept in Psychology." *Psychological Review* 3, no. 4: 357–70.

Dhein, K. 2020. "What Makes Neurophysiology Meaningful? Semantic Content in Insect Navigation Research." *Biology and Philosophy* 35, no. 52.

Dretske, F. I. 1981. *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press.

Dretske, F. I. 1988. *Explaining Behavior: Reasons in a World of Causes.* Cambridge, Mass.: MIT Press.

Ereshefsky, M., and T. A. C. Reydon. 2023. "The Grounded Functionality Account of Natural Kinds." In *From Biological Practice to Scientific Metaphysics,* edited by W. C. Bausman, J. K. Baxter, and O. M. Lean. Minneapolis: University of Minnesota Press.

Finkelstein, L. 1994. "Measurement and Instrumentation Science: An Analytical Review." *Measurement* 14, no. 1: 3–14.

Francis, R. 2003. *Why Men Won't Ask for Directions: The Seductions of Sociobiology.* Princeton, N.J.: Princeton University Press.

Garson, J., and D. Papineau. 2019. "Teleosemantics, Selection and Novel Contents." *Biology and Philosophy* 34, no. 36.

Gibson, J. J. 1966. *The Senses Considered as Perceptual Systems*. Boston: Houghton-Mifflin.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin.

Griffiths, P. E., et al. 2015. "Measuring Causal Specificity." *Philosophy of Science* 82: 529–55.

Hurley, S. 2001. "Perception and Action: Alternative Views." *Synthese* 129: 3–40.

Ismael, J. 2017. "An Empiricist's Guide to Objective Modality." In *Metaphysics and the Philosophy of Science: New Essays*, edited by M. H. Slater and Z. Yudell. Oxford: Oxford University Press.

Kight, C. R., et al. 2013. "Communication as Information Use: Insights from Statistical Decision Theory." In *Animal Communication Theory: Information and Influence*, edited by U. E. Stegmann. Cambridge: Cambridge University Press.

Kosso, P. 1992. "Observation of the Past." *History and Theory* 31, no. 1: 21–36.

Lean, O. M. 2014. "Getting the Most Out of Shannon Information." *Biology and Philosophy* 29, no. 3: 395–413.

Lean, O. M. 2019. "Chemical Arbitrariness and the Causal Role of Molecular Adapters." *Studies in History and Philosophy of Biological and Biomedical Sciences*. https://doi.org/10.1016/j.shpsc.2019.101180.

Lean, O. M. 2020. "Binding Specificity and Causal Selection in Drug Design." *Philosophy of Science* 87, no. 1: 70–90.

Levy, A., and W. Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80, no. 2: 241–61.

Maynard Smith, J. 2000a. "The Concept of Information in Biology." *Philosophy of Science* 67, no. 2: 177–94.

Maynard Smith, J. 2000b. "Reply to Commentaries." *Philosophy of Science* 67, no. 2: 214–18.

Millikan, R. G. 2004. *Varieties of Meaning*. Cambridge, Mass.: MIT Press.

Millikan, R G. 2013. "Natural Information, Intentional Signs and Animal Communication." In *Animal Communication Theory: Information and Influence*, edited by U. E. Stegmann. Cambridge: Cambridge University Press.

Neander, K. 1995. "Misrepresenting and Malfunctioning." *Philosophical Studies* 79: 109–41.

Neander, K. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, Mass.: MIT Press.

Owren, M. J., D. Rendall, and M. J. Ryan. 2010. "Redefining Animal Signaling: Influence versus Information in Communication." *Biology and Philosophy* 25, no. 5: 755–80.

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

Price, H. 2011. *Naturalism without Mirrors*. Oxford: Oxford University Press.

Rendall, D., M. J. Owren, and M. J. Ryan. 2009. "What Do Animal Signals Mean?" *Animal Behaviour* 78, no. 2: 233–40.

Robinson, A., and C. Southgate. 2010. "A General Definition of Interpretation and Its Application to Origin of Life Research." *Biology and Philosophy* 25, no. 2: 163–81.

Rossel, Samuel, and R. Wehner. 1984a. "Celestial Orientation in Bees: The Use of Spectral Cues." *Journal of Comparative Physiology A* 155, no. 5: 605–13.

Rossel, Samuel, and R. Wehner. 1984b. "How Bees Analyse the Polarization Patterns in the Sky." *Journal of Comparative Physiology A* 154, no. 5: 607–15.

Sarkar, S. 2000. "Information in Genetics and Developmental Biology: Comments on Maynard Smith." *Philosophy of Science* 67, no. 2: 208–13.

Sarkar, S. 2005. "How Genes Encode Information for Phenotypic Traits." In *Molecular Models of Life: Philosophical Papers on Molecular Biology*, 261–83. Cambridge, Mass: MIT Press.

Seyfarth, R. M., et al. 2010. "The Central Importance of Information in Studies of Animal Communication." *Animal Behaviour. Elsevier Ltd* 80, no. 1: 3–8.

Shannon, C. E. 1948. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27, no. 4: 623–56.

Shannon, C. E. 1949. "Communication in the Presence of Noise." *Proceedings of the IRE*, pp. 10–21.

Shea, N. 2007. "Representation in the Genome and in Other Inheritance Systems." *Biology and Philosophy* 22: 313–31. https://doi.org/10.1007/s10539-006-9046-6.

Shea, N. 2013. "Inherited Representations Are Read in Development." *The British Journal for the Philosophy of Science* 64, no. 1: 1–31. https://doi.org/10.1093/bjps/axr050.

Short, T. L. 2007. *Peirce's Theory of Signs*. Cambridge: Cambridge University Press.

Skyrms, B. 2010. *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press.

Soler, L., et al. 2013. "Calibration: A Conceptual Framework Applied to Scientific Practices Which Investigate Natural Phenomena by Means of Standardized Instruments." *Journal for General Philosophy of Science* 44: 263–317.

Soler, L., S. Zwart, M. Lynch, and V. Israel-Jost. 2014. *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. New York: Routledge.

Sterner, B. 2014. "The Practical Value of Biological Information for Research." *Philosophy of Science* 81, no. 2: 175–94.

Stotz, K. 2019. "Biological Information in Developmental and Evolutionary Systems." In *Evolutionary Causation: Biological and Philosophical Reflections*, edited by T. Uller and K. N. Laland, 323–44. Cambridge, Mass.: MIT Press.

Stotz, K., and P. E. Griffiths. 2017. "Biological Information, Causality and Specificity: An Intimate Relationship." In *From Matter to Life: Information and Causality*, edited by S. I. Walker, P. Davies, and G. Ellis, 366–90. Cambridge: Cambridge University Press.

van Fraassen, B. C. 2008. *Scientific Representation*. Oxford: Oxford University Press.

Walker, S. I., et al. 2018. "Exoplanet Biosignatures: Future Directions." *Astrobiology* 18, no. 6: 779–824.

Wehner, Rüdiger. 1997. "The Ant's Celestial Compass System: Spectral and Polarization Channels." In *Orientation and Communication in Arthropods*, edited by Miriam Lehrer, 145–85. Basel: Birkhäuser.

Wehner, Rüdiger, and M. Müller. 2006. "The Significance of Direct Sunlight and Polarized Skylight in the Ant's Celestial System of Navigation." *Proceedings of the National Academy of Sciences* 103, no. 33: 12575–79.

Wolpert, L. 1969. "Positional Information and the Spatial Pattern of Cellular Differentiation." *Journal of Theoretical Biology* 25, no. 1: 1–47.

Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Woodward, J. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy* 25, no. 3: 287–318.

Woodward, J. 2014. "A Functional Account of Causation; or, a Defense of the Legitimacy of Causal Thinking by Reference to the Only Standard That Matters—Usefulness (as Opposed to Metaphysics or Agreement with Intuitive Judgment)." *Philosophy of Science* 81, no. 5: 691–713.