Causation by Concentration*

Marco J. Nathan†

British Journal for the Philosophy of Science (2012)
doi:10.1093/bjps/axr056

Abstract

This essay is concerned with concentrations of entities, which play an important—albeit often overlooked—role in scientific explanation. First, I discuss an example from molecular biology to show that concentrations can play an irreducible causal role. Second, I provide a preliminary philosophical analysis of this causal role, suggesting some implications for extant theories of causation. I conclude by introducing the concept of *causation by concentration*, a form of statistical causation whose widespread presence throughout the sciences has been unduly neglected and which deserves to be studied in more depth.

*I would like to express my gratitude to Thomas Blanchard, John Collins, Laura Franklin-Hall, Anubav Vasudevan, Ken Waters, Vicki Weafer, Brad Weslake, and, especially, Philip Kitcher for constructive comments on various versions of this essay. I am also grateful to anonymous reviewers for some valuable suggestions, and to Giulia Cugnasca for helping with the illustrations.

†Contact information: Department of Philosophy, University of Denver, 257 Sturm Hall, 2000 E. Asbury Ave., Denver, CO, 80208. Email: marco.nathan@du.edu

1 Introduction

Statements about concentrations of entities within a domain are widespread in scientific explanations. Sometimes concentrations figure *qua* explananda, as when a chemist wants to understand the presence of a particular number of calcium ions within a volume of solution. Other times, concentrations are invoked as explanantia. For instance, the reliability of the technique of titration in analytic chemistry can be explained by reference to the amount of a reactant in a solution. Despite the frequent explanatory appeal to concentrations within scientific practice, philosophers have generally overlooked their significance—though Waters' ([1998]) discussion of biological generalizations that individuate distributions of entities constitutes a notable exception.

This essay investigates the role of concentration in scientific explanation. More specifically, I have three goals. The first is to show that concentrations can have an irreducible causal role in the triggering of events. Second, I employ Lewis' concept of *preemption* to provide a preliminary analysis of explanations involving concentrations, suggesting some implications for extant theories of causation and causal explanation. Finally, I define and introduce *causation by concentration*: a widespread form of causation that has been unduly neglected and deserves to be studied independently of other types of causes and in more depth. I conclude by mentioning some further applications and refinements.

The claim that concentrations have causal powers might sound not particularly surprising or, worse, trivial. After all, concentration is just a measure of relative quantity, and we are all familiar with events where relative quantity is causally relevant to the outcome of a process; for example, when some bricks piled at the center of a roof cause the roof to collapse. However, I identify a specific way in which concentrations can bring about effects, a probabilistic causal relation that, differently from the case of the bricks, cannot be straightforwardly reduced to the additive contribution of its components. This particular form of causation, omnipresent in science, is introduced by discussing, at length, an example from molecular biology: the functioning of a genetic switch. The focus on biological detail is motivated by the significance of the case study. Genetic

switches constitute the foundation of the operon model of gene regulation and are thus a paradigm of an extremely important style of molecular-developmental explanation, one that is central to present (and, most likely, future) developmental biology. Hence, a general theory of causation that cannot account for these cases is defective in important respects. Nonetheless, scientists and scientifically inclined philosophers might still wonder what is new. Anyone who is familiar with analytic and physical chemistry is likely to have encountered examples analogous to the case presented here. In this respect, the value of the present discussion does not lie in the exposition of a familiar phenomenon. Rather, my contribution consists in motivating the claim that the causal role of concentrations in these scientific processes is indeed irreducible, and thus ought to be covered by general accounts of causation. This point is appreciated best when the example is described in some detail.

Before moving on, some brief remarks about terminology. Throughout the essay, I refer to the quantity of x within a specified volume ('there are n molecules of x per unit volume inside a cell') as the *concentration* of x. When talking about the ratio of substance x to substance y, I use the expression *relative concentration*, while I employ the term *distribution* to refer to the allocation of entities according to some spatial parameter (the position of molecules of x within the cell).[1] It should be noted that while many distributions are compatible with a single concentration, a difference in concentration, in a given volume, implies a corresponding difference in distribution. In standard philosophical jargon, concentrations *supervene* on distributions.

2 Solving Lillie's Paradox: Lysogenic induction in phage λ

Phenotypic differences among organisms and among parts of an organism can be explained by appealing to *differentiation*, the generation of cellular diversity. For a long time, the understanding of differentiation was complicated by a major puzzle, known

[1] To be sure, focusing solely on quantities and ratios of chemical substances represents a simplification, since other processes and entities—such as the presence of ions and dissociating molecules—are also important. However, for present purposes, we can safely restrict our attention to the number of molecules per unit volume and set other biochemical details aside.

as the 'Developmental Paradox' (Amundson [2005]) or 'Lillie's Paradox' (Burian [2005]) after the embryologist Frank Rattray Lillie. Given that (almost) every somatic cell contains the same genetic material, it cannot be the DNA alone that is responsible for the determination of the fate of a cell. If the same information is encoded in the genome of the precursors of both muscle cells and neural cells, then why do the former develop into myocytes and not into neurons? The answer to Lillie's paradox resides in the fact that not every gene is simultaneously expressed in every cell. At any given time, most genes are silenced and the network of active genes is different in every type of cell. The silencing of genes is not a permanent state: depending on signals that can be either internal (e.g. epigenetic) or external (e.g. environmental) to the cell, the same gene can be selectively transcribed at different times. Much of what we know about differential gene expression stems from research on viruses that infect bacteria ('bacteriophages' or 'phages,' for short). Hence, phages are a good place to begin investigating gene regulation. By focusing on the development of phage λ, we shall illustrate and explain the functioning of a *genetic switch*, the molecular structure that controls the expression of genes at different times and in different places.[2] I suggest that, besides their unquestionable scientific value, genetic switches also have great philosophical interest, for they instantiate features of causation that are often neglected or overlooked.

Phage λ is an obligate parasite of the bacterium *Escherichia coli*. Obligate parasites are organisms that cannot live independently of their host: in order to survive and reproduce, phage λ must introduce its genetic material within the genome of a bacterium. The virus attaches by its tail to the surface of an *E. coli*, drills a hole through the cell wall, and injects its chromosome into the bacterium, leaving its empty coat behind. Once infected, the bacterium faces two possible fates. Under normal circumstances, the virus undergoes a process of extensive replication, called *lytic growth*, and, within

[2]We should note that, strictly speaking, the development of phage λ does not instantiate Lillie's paradox because infected bacteria do not develop into different *kinds* of cells, and thus do not undergo *differentiation*. However, many of the relatively simple mechanisms underlying differential pathways in viruses are analogous to the (generally more complex) molecular processes that occur in higher organisms, and thus can be used to study and explain the basic processes of cellular differentiation in eukaryotes as well.

some 45 minutes, about a hundred new progeny phages are lysed[3] by the bacterium and are ready to infect other *E. coli*. Alternatively, under certain conditions, the λ chromosome inserts itself into the host chromosome and, as the bacterium grows and divides, the phage is passively replicated and passed along to its bacterial progeny. As long as no perturbation occurs, this passive state, commonly referred to as *lysogeny*, is maintained indefinitely.[4] However, if the cell is 'disturbed' by an external agent such as ultraviolet (UV) light, the phage goes through a process called *lysogenic induction* and enters the lytic cycle. Our target here is an explanation of the molecular process underlying lysogenic induction: the event that leads a passive virus to initiate (and maintain) an active process of replication terminating with the lysis of the bacterial cell.

Here is a simplified reconstruction.[5] DNA sites where protein binds to enhance or inhibit the transcription of genes are called 'operators.' Whether the phage is passively replicated (lysogenic growth) or extensively replicated (lytic growth) depends on the kind of molecules that bind to specific operators on the λ chromosome. When a protein called 'clear 1' ('cI,' for short) is bound to these sites, then *cI*—the gene that keeps the phage in a lysogenic state—is transcribed, while gene *cro* is silenced. As a result, the phage is passively replicated (figure 1a). In contrast, when a protein named 'control of repressor operator' (Cro) binds to the operators, gene *cI* is silenced and gene *cro* is transcribed, initiating a cascade of other cellular events that induce lytic growth (figure

[3]In cytology the term 'lysis' refers to the death of a cell by breaking of the cellular membrane, which causes the contents to spill out.

[4]The fate of the phage on initial infection is determined by the stability of a transcription activator protein called Clear 2 (cII). When stable, cII reaches high concentration and leads to lysogeny, whereas the phage enters the lytic cycle when the concentration of cII is lowered by degradation. Factors that favor lysogeny include low temperature, cell starvation, and high multiplicity of infection. The differential fate of phage λ has an intuitive explanation: when a phage infects a bacterium in good health, it lyses, because of high protease activity, due to sufficient nutrients. In contrast, in low protease activity 'starving' cells the phage will lysogenize, waiting for a more favorable environment to reproduce. However, recent research suggests that whether an infected bacterium initially lyses or lysogenizes might also depend on physical differences between individual cells (St-Pierre & Endy [2008]).

[5]For a complete explanation, see (Ptashne [2004]).

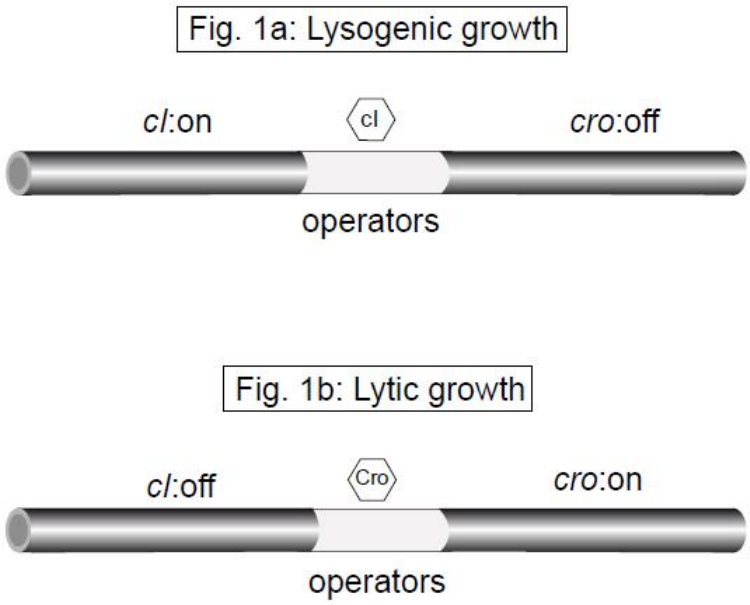

Figure 1: Differential gene expression underlying lysogenic and lytic growth. Transcription of *cI* and repression of *cro* leads to passive lysogenic replication (fig. 1a). Transcription of *cro* and repression of *cI* triggers active lytic replication (fig. 1b).

1b).[6] In short, lysis and lysogeny are a consequence of differential gene expression, which depends on interactions between protein and nucleotides.

The last aspect left to be explained is the effect of external perturbation. As said, when the phage is undergoing lysogenic growth, *cI* is turned on and *cro* is turned off. UV light flips the switch,[7] activating *cro* and silencing *cI*, thereby triggering lysogenic induction. How does this process work? In order to bind to the operators, cI molecules need to 'dimerize' into roughly a dumbbell-shaped structure. Irradiation by UV light

[6]Throughout the paper, I follow the terminological convention of italicizing the gene name and not italicizing the protein. Thus '*cro*' and '*cI*' name genes, whereas 'Cro' and 'cI' name proteins, sequences of amino acids coded by the gene.

[7]In general, the effect is triggered by any mutagen—i.e. physical or chemical agent that modifies the genetic material of an organism increasing the frequency of mutations above the natural threshold—or other source of stress. However, since considerations raised concerning UV light can be straightforwardly extended to other factors, I will refer to UV radiation as *the* efficient cause of lysogenic induction.

damages the host DNA, changing the behavior of RecA, a bacterial protein whose normal function is to catalyze recombination between DNA molecules (whence its name). When DNA is damaged, RecA becomes a highly specific enzyme that randomly cleaves cI dimers in the *E. coli* cell, irreparably turning them into monomers. In the monomeric form, the affinity of cI molecules with operators drops drastically so that it becomes very difficult for them to bind. Cro molecules, which do not dimerize and are left untouched by the action of RecA, are thus free to attach to the operators and transcribe *cro*, the gene that initiates the lytic cycle.

We can summarize the above explanation by isolating three causal relations:

1. The binding of cI to the operators causes and maintains passive lysogenic growth.
2. The binding of Cro to the operators causes and maintains active lytic growth.
3. UV radiation causes the switch to flip by replacing cI molecules with Cro molecules at the operator sites.

Let us focus on these causal processes. (1) and (2) are instances of standard interactions between molecular gears and are not particularly challenging from a philosophical perspective. In contrast, the causal relation described by (3) is more interesting. UV radiation triggers the random cleavage of cI dimers (*via* the action of RecA), which replaces cI molecules with Cro at the operators, inducing lytic growth. But how can the *random* cleavage of cI dimers have a causal effect on the genetic switch, without actually interacting with the nucleotides?[8] The key to answering this question lies in the role played by the relative concentration of proteins within the cell. In the next two sections I elaborate the function of cI and Cro concentrations in infected *E. coli*, arguing that they play a genuine and irreducible causal role. The rest of the essay provides a philosophical analysis of this sort of causal relation.

[8]The indirect interaction between RecA and DNA resembles ordinary cases of 'causation at a distance,' such as familiar (and problematic) cases of double prevention (Collins [2000]; Hall [2004]). However, the present example is further complicated by the fact that the random cleavage process does not affect all dimers, and thus not every cI dimer is prevented from binding with the switch. Our problem here is to understand the mechanisms that ground causation at a distance in these probabilistic cases.

3 Repressor concentration and the tuning of the switch

Recall from above that phage λ can reproduce in one or the other of two modes (slow-acting lysogeny or accelerated lytic growth) and that the mode is determined by a mechanism of differential gene expression, that is, a genetic switch. When cI is found at the operators the phage grows as a lysogen, whereas it lyses when Cro binds. The problem is how the random cleavage of cI dimers in the bacterial cell (by RecA) can bring about a change of molecules at the operators, inverting the expression of genes.

The first thing to note is the instability of the binding process. When a molecule binds to DNA, it does not remain there indefinitely: proteins fall off and reattach all the time. In spite of this, *cI* and *cro* are not continuously turned on and off, depending on whether cI or Cro is found at the operators. The switch is, at the same time, a finely tuned gear and an extremely stable system: it is surprisingly sensitive to changes that take the concentration of repressor above or below a certain threshold, but it is insensitive to the occasional binding of the 'wrong' protein at the operators. This is because, provided that the relative concentration of cI:Cro is appropriate, the operators are very likely to be filled by the right molecules at any given instant and the phage remains in the selected state virtually indefinitely.

All of this is a consequence of the switch's sophisticated structure. The O_R operator complex, which lies between *cI* and *cro*, contains three 17 base-pair individual operators called O_{R1}, O_{R2}, and O_{R3} (figure 2). Each operator serves a different function and has a specific affinity for binding with proteins, meaning that for any fixed amount of substance circulating in the cell, the probability that each site is bound is not identical (table 1).

Table 1: Affinity patterns governing the binding of repressor at O_R operators.

Molecule type	**Affinity pattern**
cI (ignoring cooperative binding)	$O_{R1} > O_{R2} = O_{R3}$
cI (factoring cooperative binding)	$O_{R1} = O_{R2} > O_{R3}$
Cro (no cooperative binding)	$O_{R3} > O_{R2} = O_{R1}$

Figure 2: The O_R operator complex on λ DNA.

Lysogenic growth requires simultaneous binding of cI at both O_{R1} and O_{R2}. O_{R1}, the operator closest to the *cro* gene and with the highest cI-affinity, exerts negative regulation: it prevents the transcription of *cro*, the gene that brings about the lysogenic induction. However, O_{R1} exerts no positive regulation, since it is too far from *cI* to activate it. For there to be positive enhancing, cI must bind to O_{R2} as well. In spite of the significant difference in affinity, it virtually never happens that cI binds to O_{R1} but not to O_{R2} because the cI dimers bind *cooperatively*. Molecules at O_{R1} interact with O_{R2}, facilitating the binding of cI to the lower affinity site O_{R2} as well, so that both operators are bound simultaneously at a concentration of cI that would be sufficient to bind only O_{R1} if the two sites were binding independently (non-cooperatively).[9] Finally, consider the operator closest to *cI*: O_{R3} has a lower cI-affinity and does not bind cooperatively, meaning that the concentration of cI will normally not be high enough for dimers to regularly bind to this site. However, when the concentration increases above a certain threshold—for example, if cell division is temporarily inhibited—then cI dimers start attaching also to O_{R3}, with the effect of repressing the transcription of *cI* and further synthesis of cI. In short, O_{R3} acts like a sink: it has no positive control over gene expression, only the negative capacity to inhibit the production of the cI protein by silencing *cI*.

[9]Binding of repressor to O_{R2} but not to O_{R1} would lead to the simultaneous transcription of both genes, with disastrous consequences for the cell. However, since the affinity of O_{R2} is tenfold lower than O_{R1}, under normal circumstances (i.e. in non-experimental settings), the binding of dimers at O_{R2} but not O_{R1} is extremely unlikely.

The binding pattern of Cro is similar, with two significant differences. First, no cooperative binding occurs. Second, while transcription of *cI* requires both positive and negative regulation (repression of *cro* and activation of *cI*), transcription of *cro* requires no positive regulation. The presence of Cro at a single operator (O_{R3}), blocking *cI* transcription, is thus sufficient for *cro* to be transcribed, even in the absence of protein at O_{R2} and O_{R1}. Cro binds also to the lower affinity sites O_{R2} and O_{R1} only when Cro concentration becomes too high, to repress the expression of *cro* and downregulate its own transcription.[10] (See table 2 for an overview.)

Table 2: An overview of the binding of repressor at the O_R operator complex.

Molecule type	**Operator(s) bound**	**Effect**
cI	O_{R1}	cI binds to O_{R2} through cooperative binding
cI	O_{R1}, O_{R2}	*cI* transcribed, *cro* silenced: lysogenic growth
cI	O_{R1}, O_{R2}, O_{R3}	*cI* and *cro* silenced, lysogeny maintained
Cro	O_{R3}	*cro* transcribed, *cI* silenced: lytic growth
Cro	O_{R3}, O_{R2}, O_{R1}	*cro* and *cI* silenced, lysis maintained

With all of this in mind, we can turn back to the problem posed at the beginning of this section: how does repressor concentration stably regulate the switch, given that protein molecules continuously fall on and off from DNA? The explanation, in the case of cI, runs as follows (the explanation of Cro is analogous). Suppose that a cI dimer detaches from O_{R1}. If there is enough cI in the cell, there is a high probability that a nearby cI dimer (the same one or another) will take its place for, as said, O_{R1} has high affinity with cI. The same thing happens when a cI molecule detaches from O_{R2} because, in spite of the difference in affinity, the two sites bind cooperatively. Due to its low affinity, O_{R3} is usually left empty. However, if the cI concentration goes above a certain threshold, molecules will bind also to this site, inhibiting the production of cI (without promoting Cro) until the concentration goes back to normal levels.

[10] In spite of the fact that Cro does not bind cooperatively, O_{R2} and O_{R1} are always bound at the same time because, as shown in table 1, they have the same Cro-affinity.

The fine-tuning of the switch to concentration makes it immune to perturbation. The presence of isolated Cro molecules in the cell, which might wander in the proximity of the operators and attach to them, has negligible effects on the fate of the phage. This is because, when the relative concentration of cI:Cro is high enough, it is very unlikely that enough Cro molecules will bind to the operators at the same time: as soon as that single Cro molecule drops off, its place will be taken by cI. And the right concentrations are maintained by both positive self-regulation and negative control that inhibits the production of molecules if the concentration becomes too high.[11]

All of this shows that considering the structure and dispositions of the molecules interacting with the operators is not sufficient to explain the functioning of the switch. What ensures the stability of gene expression is the relative concentration of cI:Cro in the whole *E. coli* cell, which determines the probability that, at any time, enough molecules of the same kind are in the neighborhood of the operators, ready to bind.

4 Concentration and causality

Having established the important role of concentration in providing stability to cellular processes, we can now turn back to the question of the relationship between concentration and causality. In presenting the explanation of lysogenic induction, I suggested that the change in repressor concentration, triggered by UV radiation, figures among the causes of the flipping of the switch. However, this claim could be challenged. In particular, one could acknowledge that concentrations play an important role within the process, while doubting that the relation between repressor concentration and gene expression is really causal. In other words, the objection runs, talking about concentrations is a useful

[11]It should be apparent by now that, for the switch to function properly, the relative concentration of cI:Cro must be under constant control. This regulation is possible in virtue of the fact that *cI* and *cro* encode precisely cI and Cro, the proteins that regulate the transcription of those same genes. Transcription of *cI* ensures both that enough cI is being produced in the cell and that Cro is not synthesized. Vice versa, transcription of *cro* brings it about that Cro is synthesized while the production of cI is blocked. *Ergo*, the regulation of the quantity of protein produced is really a matter of positive feedback and feedforward, i.e. self-regulation of the genes involved.

generalization that captures concisely what is going on in the cell. But, strictly speaking, there is nothing that is 'caused' by concentrations: all causal interactions occur at the level of individual protein molecules binding, or failing to bind, to operators on λ DNA.[12] Against this view, I argue that, independently of the particular view of causation we adopt, the relation between changes in concentration and differential gene expression is genuinely and irreducibly causal (in a sense to be made clear below).

Consider the following question: how many molecules are required to operate the switch? To maintain lysogenic growth, *cro* must be silenced while *cI* is transcribed. This requires two cI dimers (one at O_{R1} and one at O_{R2}), and that means four monomers of protein total. Even less protein is required for lytic growth—repression of *cI* and transcription of *cro*—since a single Cro monomer (at O_{R3}) is sufficient. This 'snapshot'[13] (correctly) describes the switch as being operated by four molecules, at most.[14] However, something important is missing from the snapshot. In support of this claim, I will show that the production of the molecules binding the operators is necessary but not sufficient for the switch to function properly; what is missing from the explanans is precisely an account of the role of repressor concentration.

The problem with the snapshot is twofold, concerning both the activation and the

[12]For instance, advocates of *process* theories of causation might be inclined to reject the claim that concentrations are genuine causes because no mark transmission (Salmon [1984]) or exchange of conserved quantities (Dowe [2000]) is involved in the relation between the cause and the effect.

[13]By a 'snapshot' I simply mean a description of all the physical interactions between molecules, relevant to a given explanation, occurring at time t.

[14]Strictly speaking, this is an oversimplification because, while four molecules are indeed sufficient to regulate the O_R complex, the differential fate of the bacterium is also governed by another operator complex, called O_L, which is located 2.4 kb away on the λ chromosome and interacts with O_R by folding DNA in a 'hairpin' loop. O_L contains three individual operators, called O_{L1}, O_{L2}, and O_{L3}, which, like the O_R operators, bind to cI or Cro to increase the repression of genes. As a result, the regulation of gene expression in the phage requires a few more repressor molecules to control O_L as well as O_R. However, given that the O_R-O_L interactions are extremely complicated, not fully understood, and, more importantly, considerations regarding the regulation of O_R can be straightforwardly extended to O_L without affecting the main philosophical argument, we shall set the O_L complex aside and focus solely on O_R.

stability of gene regulation. Consider activation first. Suppose that only two cI dimers are transcribed in an infected *E. Coli* cell. In principle, this would be enough to trigger lysogenic growth, for it is certainly possible that both dimers find themselves in exactly the right position. However, the probability of such an event occurring is extremely low. This is because proteins essentially need to reach the switch by diffusion in order to have a chance to interact with the nucleotides. And the probability that all molecules are at the right place at the right time is governed by the relative concentration of repressor substance found within the bacterial cell or (the reductionist might say) of the trajectories of all individual molecules after they have been synthesized. Furthermore—and this leads us to the second issue—even if, against all odds, both dimers did find themselves in the right position, still the system would not operate efficiently, for it would lack the necessary stability. As said, protein constantly falls off from operators, and when this happens, there must be a sufficient number of potential replacement pieces in the neighborhood to maintain the switch in a particular direction across time. In conclusion, even though a small number of molecules are, in principle, sufficient to govern the switch, a lot more are required to ensure that all pieces fall in the right place and for the process to remain stable. What controls gene expression is not whether any specific group of dimers binds or fails to bind to the operators, but the relative concentration of cI:Cro.

In addition to being descriptively accurate, this explanation is also backed up by patterns of counterfactual dependence. Consider the following statements:

1. If the cI (Cro) molecules that are now bound at the operator sites were not there, lytic (lysogenic) growth would be activated and maintained.

2. If the relative concentration of cI to Cro were below (above) a certain threshold, lytic (lysogenic) growth would be activated and maintained.

(1) is very likely to be false, for even if these dimers were elsewhere or did not exist, still the *cro* gene would not be transcribed (or the state would not be maintained), provided that the relative concentration of cI:Cro satisfies a certain threshold. In contrast, even

if it is physically possible that (2) comes out false, it has an overwhelming probability of being true: when the concentration satisfies the threshold value, molecules at the operators are replaced, and stably so.

In short, the patterns of counterfactual dependence regarding both stasis and change of the switch are captured at the level of repressor concentration, for the presence or absence of particular molecules (or sufficiently small subsets of molecules) does not make a difference.[15] I should make it very clear that the bottom line is not to dismiss the explanatory role played by the snapshot, which provides a useful description of the interaction, at an instant, between molecules and nucleotides. The point is rather that this description is only partial: there is an important causal component of the explanation that is not captured by a snapshot and can only be provided by a broader description that takes into account the contribution of concentration.

What does all of this show? At the beginning of this section, I claimed that concentrations can have 'genuine causal powers' that are 'irreducible' to interactions between molecules and operators. To clarify the relevant sense of reduction, contrast the effect of repressor concentration on the genetic switch with the example, mentioned at the outset, of the bricks causing the roof to collapse. The crucial difference between the two cases can be appreciated by introducing the notion of *net causal influence*, the sum of all individual causal components of an effect (Hitchcock [2001]).[16] The concentration of bricks that causes the roof to collapse is identical to the net sum of all the bricks that are on the roof. Adding up the component effects—the forces exerted by every single

[15]The reason for this should be obvious by now: RecA acts at the level of concentration. By randomly cleaving dimers, it reduces the quantity of repressor within the whole cell, drastically lowering the probability that once a dimer naturally falls off from DNA there will be a replacement repressor in the neighborhood ready to take its place. But no specific dimer has to be cleaved for the flipping to succeed.

[16]Hitchcock notes that the relation between a cause C and an effect E can be extremely complex, with many—possibly independent—routes, and distinguishes between *net effects* and *component effects* along a causal route. Each causal pathway constitutes a component effect. The *net causal influence* of C on E is the sum of all the component effects, after we have balanced the positive and negative effects of all routes, factoring out those that cancel each other and considering non-additive interactions. An analogous distinction is found in (Woodward [2003]) in terms of 'total' vs. 'contributing' causes.

brick—we obtain the net causal influence, which corresponds to the force exerted by the concentration of bricks. In contrast, the causal influence exerted by repressor concentration is *not* identical to the net sum of all component interactions between molecules and nucleotides (the additive sum of all the snapshots) because, in addition to all the molecules that actually interact with the nucleotides, the repressor concentration comprises also of many proteins that might never physically interact with the operators. Nonetheless, these potential causes play a counterfactual supporting role in virtue of their disposition to activate the switch, if bound. The presence of redundant molecules, which play a potential 'backup' for proteins that actually bind to the operators, constitutes the crucial causal feature of concentrations that makes them irreducible to actual mechanical interactions.[17]

Before moving on, let us respond to one final objection. The claim that concentrations play a genuine causal role could be resisted by insisting that only mechanical interactions, such as those between molecules and nucleotides, constitute acceptable causal relations. However, such stipulation comes at a cost, for it completely severs the connection between causality and counterfactual dependence. Now, surely, simple counterfactual dependence is not a necessary condition for an event to figure among the causes of another, as shown by ordinary examples of redundant causation, in which an effect does not depend upon any particular cause because of the presence of other—actual or potential—causes.[18] Nevertheless, there is a widespread assumption within the philo-

[17]Here it is important to distinguish between a weaker and a stronger sense in which the effect triggered by a concentration could be 'reduced' to interactions between its constituents. On the weaker reading, to say that a concentration is reducible to the effects of its components is to say that the *totality* of constituent-level causal relations in the system (at a time) is sufficient to ground the concentration-level dependence relations. While this claim raises important and controversial issues—such as whether all causal interactions take place at the fundamental, as opposed to a 'higher,' level—the argument developed here is intended to remain agnostic with respect to it. Our primary concern is with a different, and stronger sense of reduction: the present argument rejects the possibility of reducing the causal effect of concentrations to a subset of the fundamental relations, namely the actual interactions between molecules and operators.

[18]Essentially, this is what happens in the case of protein-nucleotide interactions: each binding of a particular repressor molecule at an operator is a cause of gene expression,

sophical community that, whatever causes turn out to be, the fact that the probability of an effect E directly depends upon the occurrence of an event C is sufficient ground for C to be a cause of E. In the example at hand, patterns of counterfactual dependence are captured by events described at the level of concentration (a decrease in cI concentration would make lysogenic induction vastly more probable), and cannot be reduced to actual molecule-nucleotide interactions. Hence, unless we are willing to concede the possibility of an event that does not counterfactually depend upon *any* of its causes, we must conclude that concentrations are causes of the flipping of the switch.[19] If this is the case, then statements such as 'the fall in the concentration of repressor within the cell caused the switch to flip' should be taken at face value, and not as shorthand generalizations for causal interactions occurring at a lower level.

5 *Preemption* in concentrations: Analysis and implications

The upshot of the above discussion is that there are two different kinds of causal processes regulating gene expression within an infected bacterium: a mechanical one, instantiated

but the effect does not counterfactually depend upon the occurrence of these causes.

[19]In spite of its popularity, the claim that counterfactual dependence is a sufficient criterion for causation has been questioned, on the grounds that it contrasts with other allegedly uncontroversial theses about causes, such as locality, intrinsicness, and transitivity (Hall [2004]). Even granting that counterfactual dependence fails as a general criterion for causation (an important issue that is too large to be discussed in full in the present essay), the situation in the case of the switch is different from alleged counterexamples to the sufficiency claim, where event A prevents event B from preventing C from causing D. In such cases of double prevention, event E does counterfactually depend upon C; what is at question is whether A's double preventing action should be included among the causes of E, since without A, C would not have occurred. In contrast, if we set aside the causal role of concentrations in the genetic switch, we end up with an event (lysogenic induction) that does not depend, even in part, upon any of its direct causes, for even considering the totality of repressor-nucleotide interactions will not make the effect robust. In short, the problem here is not just that it is possible to have causation without counterfactual dependence (a well-known fact) or, perhaps, counterfactual dependence without causation. The point is that the causal role of concentrations is necessary to make lysogenic induction counterfactually depend upon its *direct* causes.

by protein binding to operators, and a statistical one, instantiated by the relative quantity of molecules satisfying a concentration threshold.[20] But how can the distinction between these two forms of causation be cashed out?

Following David Lewis ([1986], [2000]), in asymmetric cases of redundant causation where there is a cause that brings about an effect and another cause that would have brought about the effect had the first cause been absent, let us call the cause that actually brings about the effect the *preempting* cause, and the other the *preempted* alternative, or backup. As remarked, concentrations give us the right counterfactuals because the occurrence of the effect (lysogenic induction or maintenance of lysogeny) does not depend on any particular dimer (or small subset of dimers) actually binding to the operators, as long as there are enough backup molecules that would take their place had those molecules not been there. Hence, the difference between the two kinds of causal claims can be explained as follows: the snapshots merely provide the preempting causes, but this is not enough. To capture the full causal story we need to specify also the preempted causes that make the process stable. In short, the role of concentration is to provide the causal backup, the redundant causes that would have brought about the effect had the circumstances been slightly different. Throughout the rest of this essay, our goal will be to investigate the causal role of concentrations through the concept of preemption and to discuss its implications for extant theories of causation.

The counterfactual-supporting role of concentration challenges any theory of causation and causal explanation that purports to reduce all causal events to actual interactions between their components.[21] In contrast, basic counterfactual accounts such as (Lewis [1973*a*], [1973*b*], [2000]) or (Woodward [2003]) will recognize concentration as a

[20] To be sure, nothing we said here prevents us from following (Hall [2004]) and treat production and counterfactual dependence as *two different concepts of causation*, insofar as we admit the change in concentration as a cause of the flipping of the switch.

[21] As noted, process theories of causation (Salmon [1977], [1984], [1998]; Dowe [2000]) fail to capture the explanatory relevance of concentrations in causal relations (cf. Hitchcock [1996]). Also mechanistic theories of causation (Bechtel & Abrahamsen [2005]; Glennan [2002]; Machamer, Darden & Craver [2000]) seem committed to an 'actualism' of this sort and thus cannot straightforwardly admit concentrations as irreducible causes within their framework.

cause of lysogenic induction. However, as I will argue, both these influential theories lack a specific concept that discriminates between the causal-types instantiated by concentrations and other kinds of causes. To emphasize, causes involving concentrations are *not* a counterexample to Lewis or Woodward's account, nor do they necessitate a new theory of causation. Nonetheless, they show that counterfactual frameworks should be more fine-grained and need to be supplemented in order to pinpoint particular causal relations.

Let us focus on manipulability theory first.[22] Within an interventionist framework, we can make sense of the causal role of concentration by saying, for instance, that most interventions that set the concentration of cI dimers below the relevant threshold will result in lysis, or such interventions will result in lysis with high probability.[23] Note, however, that the basic manipulability framework does not discriminate between actual causes and backups, and thereby lacks a precise concept that distinguishes between manipulations that affect actual causes or net causal influences, and manipulations that affect potential (or redundant) causes. Supplementing the framework with *actual-* and *potential difference makers* (Waters [2007]) is a step in the right direction but, by itself, will not suffice. Such distinction discriminates between causes that make a difference vs. causes that do not make a difference (in a given context); however, it does not pinpoint the case at hand, where both protein-nucleotide interactions and concentrations are actual, not merely potential difference makers. Here potential difference makers are part of a concentration, which is an actual cause of the flipping of the switch.

Moving on to a different framework, Lewis ([2000]) explicitly recognizes that cases of preemption constitute a serious problem for counterfactual theories of causation, and

[22]Since such account is well known among philosophers, I will not summarize it here in full. For our purposes, it should suffice to recall that Woodward provides a non-reductive theory of causation, based on the idea that causal relationships occur between properties, and are governed by patterns of counterfactual dependencies concerning manipulability: what would happen if the value of certain properties were to be manipulated.

[23]This statistical qualification is necessary because some interventions will set the trajectories of all, or almost all, of the dimers so that they will not bind with the operators. Under these unnatural distributions—which are physically possible, albeit extremely unlikely—lysis will occur despite the fact that the concentration is below the threshold.

proposes to account for them by amending his original ([1973*a*]) account and interpreting causation in terms of influence.[24] It is questionable whether this modification successfully solves all difficulties, but let us set these problems aside. For the purposes at hand, I want to focus on a more general problem with extant counterfactual accounts of causation—whether in the original or the amended 'causation as influence' form—namely that they fail to distinguish between *kinds* of preemption. In particular, while some redundant causes depend solely, or primarily, on concentrations of entities, others do not. To illustrate, consider the following scenarios.

1. To make sure that a person is killed, an assassin is hired to shoot a victim and, unbeknownst to the killer, a deadly poison is also administered, in the eventuality that the killer fails. The assassin successfully performs her job; the poison turns out to be redundant.

2. To make sure that a person is killed, a mobster hires a number of assassins who roam the city in search of the designated victim. One of the assassins finds her and kills her; all other killers turn out to be redundant.

Both examples involve preemption. We have an actual preempting cause and redundant preempted causes that would have brought about the effect had the former cause failed. However there is an important difference: in the second scenario, but not in the first, the number (concentration) of killers plays an important role in explaining why the victim was guaranteed (or almost guaranteed) to die. Of course, the concentration of potential causes is not the only important factor. Mental states of the assassins, their shooting ability, etc. are also relevant to the explanation of the murder. Likewise, the chemical structure of repressor molecules, their ability to bind with the operators, and

[24]'Where C and E are distinct actual events, let us say that C *influences* E if and only if there is a substantial range $C_1, C_2...$ of different not-too-distant alterations of C (including the actual alteration of C) and there is a range $E_1, E_2...$ of alterations of E, at least some of which differ, such that if C_1 had occurred, E_1 would have occurred, and if C_2 had occurred, E_2 would have occurred, and so on. Thus we have a pattern of dependence of how, when, and whether upon how, when, and whether.' (Lewis [2000], p. 190)

so forth, are also important in genetic switches. However, whereas in the poison case the concentration of backup causes plays no significant role in the explanation of the necessity (or quasi-necessity) of the effect, this factor is important in the other examples. If the city is large enough, too small a number of assassins might not be sufficient to account for the inevitability of the murder.[25] Note that analogous distinctions between kinds of preemption are also found in actual science. For instance, in molecular biology, the epigenetic mechanism of DNA methylation, that prevents protein from binding to genes, may be viewed as a form of preemption, albeit preemption of a different sort from the one discussed here, where concentration plays a significant role.

In sum, even though concentrations count as genuine causes in both counterfactual and manipulability frameworks, both accounts lack a specific concept that discriminates between different kinds of redundant causation. In the next section, I will introduce a new causal concept to overcome this problem. However, a legitimate worry should be addressed first. I suggested that concentrations play a specific causal role that deserves to be classified and studied independently of other types of causation. But what is the explanatory pay-off that offsets the cost of proliferating types of causes and explanations?

As noted by (Waters [2007]) and (Woodward [2010]), two independent projects belonging to the philosophy of causation should be kept distinct. The first consists in providing criteria that discriminate causal from non-causal relations. The second project—that has received much less attention—is to distinguish among different *types* of causal relationships. This essay is not intended as a contribution to the former project (except insofar as it jeopardizes actualist theories of causation). Rather, it is supposed to draw the attention to a specific form of causal relation that plays an important, albeit overlooked, role in scientific explanation. Among several arguments that show the benefit of specifying a variety of causal types, one is linked to *causal pluralism* (cf. Cartwright [2007]). If causal relations do not constitute a single monolithic block, but come in a

[25]Of course, the former scenario could be slightly modified in a way that makes concentrations relevant as well. For instance, instead of a single deadly poison, the victim could be administered a large number of less-than-lethal poisons, each of which has a small chance of actually killing the victim.

diverse array, then a context- or topic-dependent analysis is the only way to specify the structure of this dappled world of causes. Another—distinct but related—reason for adopting a wide range a causal types, stems from the precision that we gain in analyzing and explaining scientific phenomena. The isolation of specific causes provides us with the conceptual tools to analyze and compare scientific claims across various fields, in spite of remarkable methodological differences. In this respect, the frequent appeal to concentrations in science renders the development of a concept that underlies their causal role extremely useful. This is the task that I set out to fulfil in the following, concluding section.

6 *Causation by Concentration*: General definition, refinements, and further applications

In the first part of the essay, we presented an explanation in which concentrations are an irreducible causal factor. Next, we employed the concept of preemption to elucidate this causal factor and distinguish it from net causal influence. Two observations follow. First, the biological example shows how causal redundancy, a notion traditionally discussed by philosophers in the context of everyday (or imaginary) settings and ordinary language—backup assassins, stones breaking windows, spells trumping one another, etc.—has fruitful applications in science as well. Second, as noted, basic theories of causation based on patterns of counterfactual dependence (e.g. causation as influence, manipulability theory) lack a specific concept that discriminates between various kinds of preemption.

In order to make the causal role of concentration more precise, let us focus on its general characteristics. First, the cause is *multiply-realizable*, since concentrations supervene on and can be instantiated by several particular distributions. Second, the cause-effect relation is *probabilistic*[26] as opposed to deterministic, for it is physically

[26]To say that a causal relation is 'probabilistic' is simply to say that cause and effect are correlated by statistical—as opposed to deterministic—laws. Probabilistic analyses of causation were originally advanced by (Reichenbach [1956]; Good [1961]; Suppes [1970])

possible (though very unlikely) that the repressor concentration in the cell is well above the required threshold, yet the phage does not lyse. This could occur, for instance, if all cI dimers are located near the boundaries of the cell, too distant to have a causal effect on the switch. Third, the causal relation is *redundant*, leaving space for both preempting and preempted causes. In addition, the redundancy must depend, at least in part, on the quantity of potential causes as opposed, for instance, to differences in qualitative states. I suggest that we call *causation by concentration* (CC) a causal relation that is multiply-realizable, probabilistic, and 'quantitatively' redundant, since the effect is triggered by changes in the concentration of actual or potential causes. When all three conditions are met, the effect is irreducible, in the sense discussed above, to actual causal interactions. Supplementing extant accounts of causation with the concept of CC allows us to distinguish between cases of preemption that, like the genetic switch, depend on concentrations, and other cases (such as the standard backup assassin scenario or DNA methylation) where preemption is independent of such feature.[27]

This general characterization of CC was inspired by a particular example: lysogenic induction in phage λ. However, concentrations figure as a causal factor in the explanation of a variety of other phenomena, both within and outside molecular biology, suggesting that CC can be widely and fruitfully applied in other areas of science as well. Extending the concept to new fields necessitates subtler distinctions and additional refinements.

While the simple genetic switch discussed here involves a single relative concentration, in more complex scenarios several substances are simultaneously at play, generating a hierarchy of redundant causes. For example, in the anterior-posterior segmentation of certain thoracic areas of *Drosophila*, the high concentration of nanos protein brings

and, more recently, in amended form, by (Salmon [1984]; Eells [1991]).

[27]To be sure, the three criteria are intended to provide a general characterization of causation by concentration, not a precise definition. To fix the right set of phenomena with sufficient precision, it is necessary to further analyze the relation between concentrations and distributions, especially with respect to multiple-realizability and probabilistic considerations. For example, one important issue concerns how distributions affect concentrations and the probability that an entity will come into contact and bind with operators. Questions of this kind, worthy of further careful analysis, transcend the scope of this paper.

about an effect that would be triggered anyway by the low concentration of bicoid and, vice versa, high concentration of bicoid can causally preempt the action of nanos. In these cases, it is useful to distinguish between *preempting* and *preempted concentrations*.

Another important difference between instances of CC lies in the presence of *thresholds*. In phage λ, when the relative concentration of cI:Cro goes above a certain (remarkably precise) value, lysogenic induction becomes very likely to occur, while it will probably not occur when the concentration is below the threshold. Compare this with the following examples: rates of chemical reactions; cross effects that appear in accordance with the Onsager reciprocal relations in irreversible thermodynamics; oscillating color change in Belousov-Zhabotinsky reactions; and stationary Turing patterns that arise from the diffusion of substances. Like the genetic switch, these cases exhibit CC because the effect depends on and is triggered by complex patterns involving concentrations and concentration gradients. There is, however, an important difference. While the relative concentration of cI:Cro has a threshold, these other examples exhibit no single tipping point where a certain effect becomes likely to occur; differences in concentrations generate differences in patterns. The conflation of these types of CC can be avoided by referring to the former cases as *causation by concentration with threshold*, or *CCt*.

A detailed examination of CC in physics, chemistry, and biology deserves an independent discussion that transcends the scope of this essay. However, let me clarify one last issue. In thermodynamics and other areas of the physical sciences, scale-invariant properties, such as pressure or temperature (average kinetic energy of the constituent molecules), are usually referred to as *intensive* properties. In contrast, additive properties like heat, that are directly proportional to the system size or the amount of material contained in the system, are known as *extensive* properties. If we compare the intensive property of gas pressure with the extensive property of heat, where all kinetic interactions (as well as rotational, vibrational, and other forms of energy) need to be taken into account, it is tempting to classify the exhibition of intensive properties as a form of CC. Indeed pressure is a function of concentration and it is the concentration of gas that brings about the pressure's effects. Yet, the redundancy occurring in the case of

many scale-invariant properties, where average kinetic energy is measured, is different in important respects from the causal type exhibited by the genetic switch, where the binding of some molecules on DNA may preempt the binding of others. To illustrate, when we measure pressure at a surface, only the interactions between molecules and walls of the container matter. In this case, the collision of some molecules preempts other interactions, and the effect is triggered by CC. In contrast, when we consider the uniform pressure of a gas at equilibrium (important, for instance, for plotting phase changes on phase diagrams and considering the effect of shock waves), what happens at the walls is not the only thing that matters; no genuine preemption occurs. To avoid confusion, I suggest that general cases of *causation by intensive properties* be kept distinct from instances of CC (lysogenic induction, pressure at a surface), where concentrations in the proper sense of the term are at issue and preemption occurs.

Finally, let me note some possible applications of CC to the 'higher' sciences. Consider the sort of functional explanation we often find in economics and evolutionary biology. In a sense, we can explain the profit of a local store, or the coevolution of rabbits and foxes in the same territory, by summing up the store-customer/rabbits-foxes interactions. Focusing on the former example, suppose that at time t, the store sells item x to customer a; at time t' the store sells item y to customer b, and so forth. If the sum of all the sales is greater than the overall costs, then we have an explanation of why the store made profit. Notice, however, that this itemized description of all the monetary transactions is not fully explanatory of the success of the economic strategy for—just like the snapshots of the switch—it is not stable. As a contingent matter of fact, it was customers a, b, and c who purchased the items. But if the strategy is robust, even if a, b, or c (or all of them) had not bought the items, other customers would have made the purchase, provided that the actual distribution of customers satisfies an appropriate concentration. In short, it is not the actual transactions between store and customers that explains the success and robustness of a simple economic strategy, but rather its efficacy against the background of the concentration of potential customers.[28]

[28]My considerations regarding counterfactual robustness in economics bear important

All of this is to say that the success of simple economic and evolutionary strategies can be explained in terms of *CCt*: the causal relation is multiply-realizable, statistical, and quantitatively redundant, with a threshold for break even points and rates of survival.

In conclusion, examples involving backup assassins, pressure, customers, rabbits and foxes, and so forth, have been discussed at length and for a long time. I have attempted to show that, in spite of relevant differences, they are all instances of CC, and that it is important to recognize them as such. As a form of causation instantiated in many areas of the natural and social sciences, CC deserves to be analyzed more in detail and further refinements need to be introduced to extend its range of application. The present essay is intended as a contribution to this project.

analogies to Pettit's ([1996], [2000]) remarks concerning 'virtual selection' and 'black boxes' in functional explanation in the social sciences. There are also important similarities with Jackson and Pettit's ([1990]) notion of a 'program explanation', with the proviso that, on my view, concentrations are both causally relevant *and* causally efficacious, in their sense.

References

Amundson, R. [2005]: *The Changing Role of the Embryo in Evolutionary Thought. Roots of Evo-Devo*, Cambridge: Cambridge University Press.

Bechtel, W. & Abrahamsen, A. [2005]: 'Explanation: A mechanist alternative', *Studies in the History and Philosophy of Biology and Biomedical Sciences (Special Issue: 'Mechanisms in Biology')*, **36**, pp. 421–41.

Burian, R. M. [2005], 'Lillie's paradox—or some hazards of cellular geography', in R. M. Burian, 2005, *The Epistemology of Development, Evolution, and Genetics. Selected Essays*, Cambridge: Cambridge University Press, pp. 183–209.

Cartwright, N. [2007]: *Hunting causes and using them: approaches in philosophy and economics*, Cambridge: Cambridge University Press.

Collins, J. [2000]: 'Preemptive preemption', *The Journal of Philosophy*, **97**, pp. 223-34.

Dowe, P. [2000]: *Physical Causation*, Cambridge: Cambridge University Press.

Eells, E. [1991]: *Probabilistic causality*, Cambridge: Cambridge University Press,.

Glennan, S. [2002]: 'Rethinking mechanistic explanation', *Philosophy of Science (Supplement)*, **69**, pp. S342–53.

Good, I. [1961]: 'A causal calculus i-ii', *British Journal for the Philosophy of Science*, **11**, pp. 305–18.

Hall, N. [2004]: 'Two concepts of causation', *in* J. Collins, N. Hall & L. Paul (eds), 2004, *Causation and Counterfactuals*, Bradford, MIT Press, pp. 225–76.

Hitchcock, C. R. [1996]: 'The mechanist and the snail', *Philosophical Studies*, **84**, pp. 91–105.

Hitchcock, C. R. [2001]: 'A tale of two effects', *The Philosophical Review*, **110**(3), pp. 361–96.

Jackson, F. & Pettit, P. [1990]: 'Program explanation: A general perspective', *Analysis*, **50**(2), pp. 107–17.

Lewis, D. K. [1973*a*]: 'Causation', *The Journal of Philosophy*, **70**, pp. 556–67.

Lewis, D. K. [1973*b*]: *Counterfactuals*, Oxford: Blackwell.

Lewis, D. K. [1986]: 'Postscript e to "Causation"', in *Philosophical Papers Vol. 2*, Oxford University Press, pp. 193–212.

Lewis, D. K. [2000]: 'Causation as influence', *The Journal of Philosophy*, **97**, pp. 182–97.

Machamer, P., Darden, L. & Craver, C. [2000]: 'Thinking about mechanisms', *Philosophy of Science*, **67**, pp. 1–15.

Pettit, P. [1996]: 'Functional explanation and virtual selection', *British Journal for the Philosophy of Science*, **47**, pp. 291–302.

Pettit, P. [2000]: 'Rational choice, functional selection and empty black boxes', *Journal of Economic Methodology*, **7**(1), pp. 33–57.

Ptashne, M. [2004]: *A Genetic Switch. Phage λ Revised*, 3rd edn, Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

Reichenbach, H. [1956]: *The Direction of Time*, Berkeley and Los Angeles, CA: The University of California Press.

Salmon, W. [1984]: *Scientific explanation and the causal structure of the world*, Princeton, NJ: Princeton University Press.

Salmon, W. C. [1977]: 'The "at-at" theory of causal influence', *Philosophy of Science*, **44**, pp. 215–24.

Salmon, W. C. [1998]: *Causality and Explanation*, New York: Oxford University Press.

St-Pierre, F. & Endy, D. [2008]: 'Determination of cell fate selection during phage lambda infection', *Proceedings of the Natural Academy of Science USA*, **105**(52), pp. 20705–10.

Suppes, P. [1970]: *A Probabilistic Theory of Causality*, Amsterdam: North Holland.

Waters, C. K. [1998]: 'Causal regularities in the biological world of contingent distributions', *Biology and Philosophy*, **13**, pp. 5–36.

Waters, C. K. [2007]: 'Causes that make a difference', *The Journal of Philosophy*, **104**(11) pp. 551–79.

Woodward, J. [2003]: *Making things happen. A theory of causal explanation*, New York: Oxford University Press.

Woodward, J. [2010]: 'Causation in biology: stability, specificity, and the choice of levels of explanation', *Biology and Philosophy*, **25**, pp. 287–318.