

Mob Rules: Towards A Causal Model of Social Structure

Andrea Borghini and Marco J. Nathan

American Philosophical Quarterly

Abstract

This essay enriches causal models capturing the propagation of prejudice, bias, and other aggregative social mechanisms, negative or positive. These explananda include the reinforcement of economic inequality, “mob-like” behavior, peer pressure, and the establishment of social norms. The stage is set by introducing various forms of redundant causation and discussing some difficulties with mainstream preemption. Next the main proposal extends current representations of aggregative social mechanisms in two respects. First, it is more nuanced, as it identifies three distinct kinds of inferences: relevance of causes to effects, degree of redundancy of an effect, and the influence that causes exert on other causes. Second, it offers a *quantitative* – as opposed to merely qualitative – distinction between the causal contribution of preempting and preempted causes.

Keywords

Social structure; redundant causation; causal models; social norms; bias and prejudice.

1. Introduction

Human societies are complex substrata. Creatures like us interact at several different spatio-temporal levels: on an individual basis, as small-groups, and at much larger scales – neighborhood-wide, city-wide, state-wide, nationally, internationally, and even globally. As challenging as it is, capturing and explaining the mechanisms underlying these structures is crucial for understanding the world we live in.

How do we make tractable human interactions at a global scale? The short answer is by devising *causal models*, that is, abstract and idealized mathematical constructs identifying and representing the network of causal relations occurring within an individual system or population. Causal models have become an essential tool across the natural and social sciences, and have spawned into a sizeable literature. The main focus of this essay is a particular subclass of causal models. Specifically, we aim to present, assess, and enrich causal models designed to capture the propagation of prejudice, bias, and various other kinds of aggregative social mechanisms, regardless of whether their effects are treated as overall negative or positive for human societies. The explananda we have in mind include the reinforcement of economic inequality (Stiglitz 2012), the establishment of social norms (Bicchieri 2006), gender discrimination (Valian 1998, Manne 2017), and other socio-psychological phenomena which have been identified and scrutinized at least since the 19th century (Le Bon 1896).

We begin, in §2, by introducing the social dimension of prejudice. Next, we argue that current developments in the literature on redundant causation can be fruitfully employed to preliminarily explain social mechanisms (§3-§4). Finally, in §5, we sketch a model of the relevant social structures that enable aggregative social mechanisms. Our proposal adds two novel insights to the extant literature. First, it identifies three kinds of inferences: the relevance of causes to effects, the degree of redundancy of an effect, and the influence that causes exert on other causes. Second, our proposal offers a *quantitative* – as opposed to qualitative – distinction between the causal contribution of “preempting” and “preempted” causes.

2. The Social Dimension of Prejudice

The causal structure of social phenomena – in particular harm, injustice, and other potentially pernicious ones – has been the subject of much attention in philosophy, psychology, and other social sciences. Nevertheless, some readers may wonder what, exactly, there is to debate from a causal perspective. Are prejudices not simply caused by the mental states of prejudiced people and ensuing actions? It turns out that things are not as simple as they may appear at first blush. The true complexity of the situation may be illustrated via a simple toy example.

Imagine a small, relatively isolated community N , whose members – at least, most of them – stigmatize a specific, directly observable physical trait t , which is irrelevant to job performance. Instances of t include skin tone, body image, or accent. Consider a member of N , Omega, who bears trait t and, because of it, is in a position of disadvantage. One day, Omega applies for an executive position with a successful local startup. Omega is the most qualified candidate, an evident fact that is reflected in Omega’s resume. Omega is invited for a job interview. However, Beta, the director of human resources in charge of hiring on behalf of the company, has some negative prejudice against individuals with t .

¹ Consequently, Omega is turned down.

Oversimplified as this scenario may be, it should sound sadly familiar to most readers. Analogous cases of prejudice and discrimination are widespread in most real-world communities and affect various social groups, including women, minorities, immigrants, differently abled individuals, and many others.² In order to identify and, ultimately, eradicate these biases, we first need to understand them and represent them appropriately. How can we model the causal structure of commonplace scenarios, such as the one depicted above? Let’s try to capture the situation by constructing a causal model: an abstract and idealized representation that purports to depict the relevant causal interactions underlying the effect.

The plainest, most straightforward model involves a simple token causal relation, where Beta is entirely responsible, both etiologically and morally, for Omega’s unfair treatment and, thus, Beta is the sole culprit. Let’s call this model “DH” (for “direct harm”). DH provides a rather intuitive depiction of the situation. After all, Beta is *prima facie* the necessary and sufficient factor of Omega’s disadvantage, since Beta alone has the power to make the relevant decision. The assumption underlying DH is that by performing an action – such as expressing an opinion or deciding not to hire an applicant for a position – one *directly* harms another individual.³ If this

were the case, to rectify the situation, it should be enough to prevent the actions of individuals like Beta, for instance, by removing them from executive roles. Ideally, people with overt substantial prejudice, explicit or implicit, should not be put in a position to affect important aspects of other people's lives concerning, for instance, job security, health, and, more generally, their well-being.⁴

No one should dispute that DH captures *a* significant way of perpetrating harm. People can, and often do, cause harm directly, along these lines. Furthermore, making sure that overtly prejudiced individuals are not in positions of power would definitely be an important step, that, alas, we are quite far from accomplishing. But does this simple model provide an exhaustive characterization of what is going on? Is DH an adequate causal model of prejudice, in general? The answer is negative. There is a growing consensus that, intuitive as it may appear, the DH-model turns out to be overly simplistic, as it completely ignores the social, aggregative dimension of prejudice, which accounts for many of its most devastating and pervasive effects.

To pinpoint the shortcomings of DH, consider a simple variation of our example. Imagine that Omega's job interview is not conducted by Beta alone, but by Beta in collaboration with two colleagues, Gamma and Delta. After Omega leaves the room, Delta makes the following remark: "Omega's CV is only impressive due to Omega's *t*-bearing status: affirmative action these days gives *t*-bearers an unfair advantage over everybody else." Gamma nods approvingly. Beta, who is the chair of the hiring committee, makes the executive decision not to hire Omega. Here, Beta is still directly responsible for discrimination. After all, as the chair of the committee, Beta is necessary and sufficient for the unfair treatment of Omega. Yet, something else seems to be going on here. What, exactly? Simply put, direct action is not the only way of causing harm. Conversational moves, such as Delta's "*t*-phobic" remark may constitute harm by changing the conversational score and thereby enacting what is (im)permissible in a conversation thereafter (McGowan 2004, 2009).⁵

What about Gamma? Prima facie, it would appear that Gamma has no role in perpetrating Omega's mistreatment, as Gamma is not doing anything over and above tacitly assenting. Again, this is not quite accurate. In a recent paper, Ayala and Vasilyeva (2016) agree with McGowan that sexism, homophobia, and other cases of harm may be oppressive in virtue of enacting what is permissible thereafter. In addition, however, they maintain that this effect may be triggered not just by the active contribution of the speaker, but also via omissions of other interlocutors. Applying this insight to our case, Gamma's lack of intervention on Delta's *t*-phobic remark, which fails to block its contribution to the conversational score, makes Gamma also responsible, at least in part, for Omega's harm.

These simple toy examples effectively show that there is a significant difference between isolated mistreatment, which is chiefly due to the behavior of a single or restricted group of individuals, and widespread prejudice, which unfurls across a broader social group. In presenting the bias as pertaining to an entire community *N*, we implicitly acknowledge that, while Beta is responsible for the discriminating behavior, she is not *solely* responsible for it. What does this mean? While harm can be perpetrated directly, it can also be caused in a less straightforward fashion, by enacting the conversational score, or by failing to block or correct a morally impermissible presupposition.

In short, as illustrated by the contributions of McGowan and Ayala and Vasilyeva, prejudice has a pervasive social dimension.

Important as these cases are, unfortunately, they do not exhaust the ways in which prejudice and other forms of social injustice spread across societies. The goal of this article is to draw attention to a different, but no-less harmful, form of prejudice that is often neglected by extant models. To illustrate our main ideal, let's introduce another variant of our scenario. Imagine that Omega had not been interviewed by Beta but, rather, by Beta's colleagues Gamma or Delta. Can we make an informed prediction about what would have happened under these circumstances? Assuming that the bias is pervasive enough within Society N , one can reasonably conjecture that Omega would not have been offered the job anyway. This could be explained in various ways. One reason, the obvious one, is that Gamma and Delta share the same prejudice as Beta. Yet, one should also acknowledge a different kind of potential socio-dynamic underlying the mishap, which could have triggered a similar outcome. Even if Gamma and Delta are not biased towards individuals with trait t , they might be concerned about how other members of N , especially biased individuals like Beta, might perceive Omega's hiring when such decision is made public. This effect, which is a kind of *peer pressure*, makes social prejudice and discrimination a subtler and, therefore, more pernicious social mechanism. How can we introduce this further layer of complexity and represent it in our causal model?

3. Redundant Causality and Social Structures

§2 introduced a fictional case of prejudice. Oversimplified as it is, the toy example shows that a causal model focusing solely on the *actual* wrongdoing fails to do justice to the complexity of the situation. What, exactly, is wrong with models that depict Beta, Gamma, and Delta as the sole culprits of Omega's discrimination? Simply put, they are too *fragile*. Part of what makes prejudice so subtly pernicious, is that it transcends actuality. Prejudice has a *modal* aspect, in the sense that it often also affects the ways things *could* have been. To capture this modal dimension, we need to assess counterfactual claims about what would have happened to Omega under various kinds of circumstances. This brings *redundancy* into the picture.

Redundancy is a widespread feature of causation, discussed extensively by philosophers. In this section, we survey some influential treatments of redundant causality. Next, §4 applies them to our scenario. Specifically, we maintain that traditional discussions, which focus on *overdetermination* and *preemption*, only capture part of the story. §5 will then sketch a causal model of aggregative social phenomena that develops and extends these well-known concepts.

Simply put, philosophers distinguish two varieties of redundant causation: overdetermination and preemption. In cases of *overdetermination*, the occurrence of an event (or process⁶) is brought about by multiple independent causes, in a way that their aggregative power to trigger effects exceeds what it takes to bring them about. For instance, we have a case of overdetermination if each of various causes would have been sufficient to produce an effect individually. Imagine that, to ensure that a window is broken, ten pitchers simultaneously throw baseballs towards the glass

from a few yards away. Assume that all, or at least the large majority, of the balls hit the target roughly at the same time, causing the window to break. Here we have various independent events (the throws) causing the breaking of the window. Yet each of them, individually, would have been sufficient for the window to be broken.

The second kind of redundancy, *preemption*, occurs when there is a cluster of independent causal processes directed towards the same effect, where only one is brought to completion and contributes to the effect, rendering all other events or processes unnecessary. Several varieties of overdetermination and preemption have been distinguished (Hall and Paul 2013, Ch. 3). Interestingly, the variety that best serves the purposes of illustrating our view seems to pose the least problems for such accounts of causation. It comprises the so-called scenarios of “*early preemption*,” which can be illustrated through the familiar story of the backup assassin.⁷ Assassin (*a*) is hired to kill Victim (*v*). To ensure that *v* is killed, unbeknownst to *a*, Backup Assassin (*b*) is also hired to follow the scene and kill *v*, but only were *a* to fail. Imagine that *a* does not fail. As a result, *b* turns out to be redundant. Following Lewis (1973, 2000), in asymmetric cases of redundancy where there is a cause that brings about an effect and another cause that would have brought it about, had the first cause been absent or inert, call the cause that *actually* brings about the effect the “preempting” cause and the other the “preempted” alternative. In the scenario just presented, *a* is the preempting cause of *v*’s death, whereas *b* is the preempted backup. Moreover, we can classify *a* as an *early* preempting cause because *a*’s action preempts *b*’s action *before* *b* starts to bring about *v*’s death. So, *b* is not acting at all towards *v*, despite the fact that *b* is part of the causal story. At least in this sense, *b* is causally relevant.

Redundant causality has received much attention in philosophy, chiefly because of the difficulties it generates for theories of causation which treat causes as factors making a difference towards the occurrence of their effects. In the backup assassin scenario, whereas *a* causes the death of *v*, *a*’s contribution is not a *conditio sine qua non* for *v*’s death, since it does not make a difference to whether or not *v* dies. Yet, causal redundancy may also play an important role in stabilizing effects, strengthening the degree of counterfactual support, and enhancing the *robustness* of causal relations (Woodward 2010).

As noted, robustness, understood in terms of causal redundancy, has been widely discussed within the metaphysics of causation literature. Still, these insights have not yet been fully integrated in debates on aggregative social structures and mechanisms. The point we wish to stress is that, in many situations, robustness can be a blessing. This is because robustness may ensure the occurrence of an outcome that is important, actively sought, or otherwise desirable. To illustrate, additional security measures, such as extra airbags in cars or backup parachutes, can save lives. Similarly, many vital biological functions depend on the reliability of processes that are accordingly redundant. Many organisms, humans included, have two kidneys despite being perfectly capable of functioning with a single one. Yet, unfortunately, not all cases of overdetermination and preemption trigger desirable outcomes. Robustness also serves the role of perpetrating and stabilizing undesirable effects, as illustrated by the above instances of biases and prejudice. In all these scenarios, positive and negative, redundancy is key.

4. Redundant Causes and Aggregative Phenomena

Having briefly distinguished various forms of redundancy in the previous section, we are now in a position to provide a more accurate description of our scenario involving prejudice. Beta actively causes harm by treating Omega unfairly. Suppose, however, that Gamma and Delta, who are also interviewers in the same company as Beta, share Beta's prejudice against individuals bearing trait *t* or are subject to some kind of peer pressure. Hence, had Omega been interviewed by either Gamma or Delta, all things being equal, Omega would not have gotten the job anyway. How should we model the causal structure of this scenario?

As mentioned at the outset, these kinds of examples are often discussed in the literature. The standard debate revolves around the contrast between *atomistic* or *individualistic* vs. *holistic* or *structuralist* approaches. Simply put, individualists purport to reduce all social phenomena to properties of agents – intrinsic or extrinsic, depending on the variants (Epstein 2014). Structuralists, in contrast, maintain that social phenomena are inherently holistic and relational and, consequently, reducing them to properties of their components misses the mark (Garfinkel 1981, Haslanger 2016, Malinsky 2017). The scenario at hand, in and of itself, does not line up neatly with either individualism or structuralism. In order to do justice to our situation, we need to focus on *causality* (Bright et al. 2016). However, to capture it perspicuously, we shall need to borrow some key concepts from the metaphysics of causation.

How do we make sense of a situation, like our present scenario, where the presence of a cause does not make a difference to its effect? The mainstream answer, pioneered by Lewis (2000) prescribes that we supplement the initial causal intuition, where Beta – or, more precisely, Beta's prejudice – causes Omega harm, by introducing redundancy into the picture. In this specific situation Gamma and Delta, whose actions are redundant, are preempted causes; Beta, instead, is the actual preempting cause.⁸

A key trait of this mainstream approach is that it provides a common denominator for overdetermination and preemption, by regarding both as instances of *redundant causation*. *Prima facie*, this stipulation may appear metaphysically innocent. Yet, we now argue, it presupposes some substantive assumptions. Cases of overdetermination have standardly been considered genuine instances of causal relations, or parts thereof, because each of the causes is arguably both actively contributing to (part of) the outcome and is also actually making a difference to the effect (Lewis 1973). More controversially, can the same be predicated of preemption? In the preempting scenario we hypothesized, Gamma and Delta play a *sui generis* role in perpetrating the prejudice. However, such role is importantly distinct from the "direct" or "active" role of Beta. Beta is the preempting cause of Omega's unfair treatment. The mere presence of Gamma and Delta, in contrast, ensures that the outcome would have occurred even under slightly different circumstances. But, note, there is no action Gamma and Delta perform to play such a corroborating role. To paraphrase an expression introduced by Geach in the context of metaphysical analyses of properties and relations, Gamma and Delta appear to be "Cambridge causes."⁹

Are Cambridge causes genuine, bona fide causes? Or, are they merely *explanatorily* relevant, but not *causally* relevant? One could maintain that some of what we dubbed Cambridge causes are genuine causes because, although they *do* nothing, they *could* have acted in virtue of the activity of their counterparts. Hence, the possibility of Cambridge causes becoming active, “non-Cambridge” causes, is metaphysically substantiated by the objectivity of counterpart relations between Cambridge causes and their active counterparts. However, most philosophers, including yours truly, are wary of embracing a modal realism of this sort. Specifically, treating overdetermination and preemption as bona fide causes forces non-modal realists into a dilemma. On the one hand, we may regard Cambridge causes as somewhat “fictional.” Yet, by doing so, it becomes unclear how we can effectively include them in any explanation of what is going on in actuality. On the other hand, we may regard Cambridge causes as genuine causes, in which case it becomes unclear what they *do* to merit the status of cause. Either way, there are striking difficulties and counterintuitive implications.

In order to escape the trouble with Lewis’s mainstream approach discussed above, one could treat Gamma and Delta as merely *explanatory*, as opposed to truly *causal* (Strevens, forthcoming). In the preemption scenario, the argument goes, Gamma and Delta are not bona fide causes of Omega’s discrimination. Rather, they are Cambridge causes, which become relevant if we wanted to provide a comprehensive account of, say, the operations of a company’s human resource office. For instance, such variables are necessary to explain *why* Omega was not hired, or to specify the circumstances under which Omega could have been hired. Still, this does not make Gamma and Delta *causes* of Omega’s discrimination. Only what *actually* contributes to an effect, the objection runs, should be regarded as a genuine cause. Hence, agents that could have contributed, but ended up not doing so, cannot be considered causes. Call them “Cambridge causes,” if you like, but make it clear that Cambridge causes are not genuine causes at all since, by definition, they do nothing to bring about the effect.

The modal-actualist strategy just described is not devoid of controversial implications either. By taking Gamma and Delta as explanatorily relevant – but not causally relevant – to the discrimination of Omega, we must conclude that Gamma and Delta play no role whatsoever in *causing* Omega’s disadvantage. But this is a hard bullet to bite. Why is it that Gamma and Delta are blamed, at least in part, for the bias perpetrated against Omega, if they don’t cause anything? Conversely, if they are blameworthy, did they not partake in perpetrating the prejudice?

Before moving on, we should emphasize that the problem is not restricted to cases of social stigma and prejudice. The active role of Cambridge causes applies across the board to various scientific examples. For instance, Nathan (2014) discusses an example from molecular biology where concentrations of entities play an “irreducible” causal role. Nathan’s case involves a *genetic switch*, that is, the cytological structure that controls the expression of genes. Such process, he claims, cannot be causally explained by focusing solely on the individual protein that *actually* bind to operator sites. In order to provide a complete causal story, we also need to focus on the “preempted” molecules that remain idle, but would have bound to DNA had things gone slightly differently. Similar cases abound throughout the sciences, from economics to ecology, to evolution. Consider, for instance, the famous Lotka-Volterra equations, which govern the

fluctuations in the population of a predator-species and a prey-species within an ecosystem. Such dynamics can only be causally explained at the macro-level, by factoring the entire populations, not simply individual encounters between predators and prey (Garfinkel 1981). Whether aggregate entities play a bona fide causal role in these scenarios transcends our present purposes. Still, these examples show that redundant Cambridge causes are prominent across the sciences. Understanding their function is a crucial step in refining our causal models.

Let us take stock, Cambridge causes pose a dilemma. On the one hand, if we regard Cambridge causes as bona fide causes, then we must provide a plausible story about what they *do*. And this has proven to be a daunting task.¹⁰ On the other hand, if Cambridge causes are mere explanatory devices, we are left without a convincing account of why we feel compelled to saddle them with responsibility. Or, perhaps, we should not attribute responsibility, causal or moral, to individuals who do nothing in actuality. But, then, we lose the distinction between isolated action, and aggregative behavior. Either way, we are confronted by theoretical challenges.

Individualists and structuralists alike must confront this issue directly. However, neither camp is in a position to offer an obvious response. Individualists will likely contend that the contribution of redundant causes can be reduced to that of particular agents – leaving out the prominent social dimension of the causal structures under consideration. Structuralists, in contrast, will likely view the role of redundancy as an irreducible feature of social structure – leaving the underlying phenomenon unexplained. Regardless, the question remains: *how* can we capture the causal contribution of redundant causes?

5. Extending Causal Models of Aggregative Phenomena

The previous section argued that redundancy plays a substantial - albeit neglected - role in capturing the causal structure of prejudice and other kinds of aggregative social phenomena. Lewis's distinction between overdetermination and preemption is a crucial preliminary step, which allows one to distinguish the role of active vs. passive causes. However, it is only part of the story. For, clearly, there are different types of preempted causes, and they do not all play the same role. For instance, an adequate model of our toy example must distinguish between the influence that Gamma and Delta counterfactually exert on Omega from the influence that Gamma and Delta exert on Beta – the “peer pressure” that motivates Beta to perpetrate the injustice.

To understand the social aspect of prejudice via finer-grained distinctions between agents, we need to make two additions to the mainstream Lewisian approach. First, we introduce into our causal models a *quantitative*, as opposed to merely qualitative, dimension that allows us to operationalize and distribute the blame by assigning numerical values to the various parties involved, in accordance with the different causal roles they play.¹¹ Second, we further distinguish the *kinds* of roles involved in situations of redundant causation. We should make it clear that our aim is to model the structures underlying aggregative behavior and, in particular, to sort out the key joints that are linked to causal roles. Needless to say, such structures will then be realized in different forms and contexts, which may require to fine-tune our model with additional details.

To capture the basic structure of the social phenomena under scrutiny, we propose to enrich causal models of aggregative behavior with three functions. Before introducing such functions, however, some remarks on the *universe* or *domain of discourse* are in order. Intuitively, the domain of discourse is a set whose members are all the entities, processes, and events that are potentially relevant to a particular situation that we aim to represent. Precisely which features count as relevant will clearly depend on background assumptions. For instance, within a structuralist approach, the universe will include all individuals, relations, and processes that comprise the social structure under consideration (Ayala and Vasilyeva 2015; Haslanger 2015; Vasilyeva 2016). In contrast, a more individualist approach will typically work with a domain containing solely individuals (Epstein 2014).

To illustrate, in our scenario, where Omega is discriminated in the workplace by actual or potential employers, the universe must contain Omega, Beta, Gamma, and Delta, as well as the causal relations among them. By *including* additional entities, say, Pi, Rho, Sigma, and Tau, one commits to including an assessment of their role within the scenario. The commitment in question is only to *evaluating* their role, not to assigning any positive value. For any variable x in the domain, it is perfectly possible to conclude that x has no role whatsoever in producing the effect. In contrast, by *excluding* any individual or process from the model, we are thereby ruling out their contribution to the effect, as all and only members of the universe will be factored in by the model.

It should now be evident that choosing a domain is no trivial task; it is a normative act, with significant implications for the outcome. For the sake of simplicity and theoretical neutrality, we refer to the entities in the domain of discourse as “variables.”

The enriched causal model we envision rests on three functions defined over the domain of discourse. Intuitively, these functions capture the following features. (i) *The actual relevance* of a variable to the causal structure of a given *effect* in the domain of discourse. (ii) *The degree of redundancy* of a given event in the domain of discourse, that is, the influence on the effect of variables in the universe which have no actual role in bringing about the effect. And (iii) *the influence* of a given event in the domain of discourse on both redundant and non-redundant causes.¹² Let us spell out these functions in greater detail. Keep in mind that our main goal is merely to illustrate the nature of the functions themselves. Readers with different normative intuitions are welcome to change the apportionments of responsibility by manipulating the exact numerical values in the co-domain.¹³

The first step in determining the causal structure of a social event is to determine who actually brought about the effect. Thus, for instance, whose decision was it not to hire Omega? Was it Beta single-handedly? Or was there a committee constituted by Beta, Gamma, and Delta, who reached a democratic consensus? To distinguish these two scenarios, we introduce a function, “ $R(\phi, \Delta)$,” which assigns to every agent ϕ in the domain of discourse one of two possible values: either [0] or [1]. If, $R(\phi, \Delta)$ assigns to ϕ the value 1, then ϕ played an active role in bringing about Δ . In contrast, if $R(\phi, \Delta)$ assigns to ϕ the value 0, then ϕ did not play an active role in bringing about Δ .

To illustrate, assume that Beta is solely responsible for the decision (“A”) not to hire Omega. Under these circumstances, we have that $R(\text{Beta}, A) = 1$ since Beta directly sees to it that Omega gets hired in the startup. In contrast, Gamma and Delta – who, it might be worth recalling, are not actually reviewing Omega’s application – do not play a role in determining Omega’s future with the company. Consequently, we have that $R(\text{Gamma}, A) = 0$, $R(\text{Delta}, A) = 0$ and, obviously, $R(\text{Omega}, A) = 0$ since Omega cannot decide to hire herself. In contrast, suppose that Beta is part of a “democratic” (or “quasi-democratic”) hiring committee that also includes Sigma and Tau. Now, we have that $R(\text{Beta}, A) = 1$, $R(\text{Sigma}, A) = 1$, and $R(\text{Tau}, A) = 1$, for the effect would be a joint outcome of the committee’s decision.¹⁴ Note that $R(\phi, \Delta)$ merely states that ϕ plays an active role in bringing about Δ . It does not determine whether the role of ϕ is more or less significant than the role of, say ϕ^* . This is why the range of $R(\phi, \Delta)$ is either 0 or 1. It is 1 if the agent in question actually did something; it is 0 otherwise.

Once we determine who actually played a role in bringing about an effect, the following task is to introduce the factors that made a potential, preempted contribution to the effect. This is accomplished via a second function, “ $E(\phi, \Delta)$,” which specifies the potential influence that agent ϕ has on the effect Δ , that is, the capacity of ϕ to bring about Δ . The value of $E(\phi, \Delta)$ ranges over the closed interval [0-1]. The rationale for this choice should be obvious. Let’s call “sufficient” a cause that has the capacity to bring about an effect all by itself. Thus, if I am able to lift a book all by myself, then I am a sufficient cause of the lifting of the book, just like you. This can be expressed in terms of E as $E(\text{me}, \text{lifting book}) = 1$ and $E(\text{you}, \text{lifting book}) = 1$. In contrast, if it takes two of us to lift a heavier object, say, a piano, then our value for E will be lower than one. For instance, we may have that $E(\text{me}, \text{lifting piano}) = 0.5$ and $E(\text{you}, \text{lifting piano}) = 0.6$. Note that the two values of E need not add up to 1. The two of us together overdetermine the effect, in the sense that you could have lifted the piano with someone weaker than myself, as long as the combined value of E between you and the other agent is equal to or greater than 1.

Now, by adding up the value of $E(\phi, \Delta)$ for all agents ϕ in the domain of discourse, we can calculate the degree of redundancy of the event Δ . Suppose that you and I are the only agents in the domain. Then, we have that the lifting of the book has a degree of redundancy of 2, because $E(\text{me}, \text{lifting book}) = 1 + E(\text{you}, \text{lifting book}) = 1$. In contrast, the lifting of the piano has a degree of redundancy of 1.1 because $E(\text{me}, \text{lifting piano}) = 0.5$ and $E(\text{you}, \text{lifting piano}) = 0.6$. In the hiring scenario, assuming that Beta, Gamma, and Delta, all have the capacity to hire Omega at company X, and no one else does, we have that the hiring of Omega at company X has degree of redundancy 3, which is the result of $E(\text{Beta}, A) = 1$, $E(\text{Gamma}, A) = 1$ and $E(\text{Delta}, A) = 1$. In general, we can express this degree of redundancy with respect to effect A as the sum of the values $\sum E(\tau, \Delta)$, where τ is a metavariable ranging over all individual agents $x, y, z \dots$ in the domain. In short, assuming that we are dealing with finite universes, the degree of redundancy of an effect A is captured by $\sum E(\tau, \Delta)$.¹⁵ Assuming that we have ten agents in the model, none of whom would reject Omega, based on their prejudice, redundancy is 0. If 7 of them would reject Omega, the redundancy value is 7, and so on.¹⁶

$R(\phi, \Delta)$ and $E(\phi, \Delta)$, provide sharp conceptual tools that quantify the degree of redundancy, thereby introducing a clear distinction between preempting, preempted, and “collaborative” causes. Are

we done? Not yet. Framing the structure of prejudice and other aggregative social mechanisms requires one additional step. Simply put, we need to assess the role of variables (events) in the model that do not affect the outcome directly, but do so indirectly, by influencing the behavior of other actual or potential causes. These factors, which roughly correspond to peer-pressure and various other kinds of conformity to social standards, can be captured by introducing a third function $C(\chi, \phi, \Delta)$.

Consider two variants of our example. First, Beta decides not to hire Omega, despite Omega being the most qualified candidate for the position, because Beta has an overt prejudice against individuals who, like Omega, are t -bearers. As noted above, here we have that $R(\text{Beta}, A) = 1$ and $E(\text{Beta}, A) = 1$, indicating that Beta is fully responsible for the act of discrimination, although the degree of redundancy implicit in, say, $E(\text{Gamma}, A) = 1$ and $E(\text{Delta}, A) = 1$ might reveal the pervasiveness of Beta's prejudice across society.

Second, consider a scenario where Beta has no overt bias against individuals bearing trait t . However, Beta is perfectly aware that many members of her society do, in fact, have a prejudice against t -bearing individuals. In addition, Beta is concerned that she herself could become the target of discrimination, if she were not to display the same degree of prejudice against t -bearing individuals. For this reason, Beta decides not to hire Omega. In this second scenario, the values of $R(\text{Beta}, A)$ and $E(\text{Beta}, A)$ would be exactly the same as in the previous case. However, the circumstances, including Beta's degree of responsibility would arguably be different. How do we capture the significant difference between these two situations?

Simply put, the difference here boils down to the role of Gamma and Delta who, in the second variant, pressure Beta into discriminating Omega, whereas they do nothing in the original scenario. To capture this difference, we need to express how, in the former case, where Beta is prejudiced, the influence of Gamma and Delta on Beta's behavior is negligible or altogether absent. Beta would behave in the same way no matter what. In contrast, in the latter scenario, Beta is not prejudiced and would not discriminate against Omega if it were not for Gamma, Delta and the other members of society N exerting peer pressure. How do we draw this important distinction?

Let's introduce a function $C(\chi, \phi, \Delta)$, which ranges over the closed interval $[0 - 1]$ and measures the influence of a group of individuals $\chi = (z_1, \dots, z_n)$ from z_1, \dots, z_n the domain of discourse, on the proximate cause ϕ of an event Δ . The intuitive idea at play here is that $C(\chi, \phi, \Delta)$ measures whether there is a (minimal¹⁷) collection of individuals $\chi = (z_1, \dots, z_n)$ such that the presence of χ makes a difference to the fact that ϕ brings about Δ . More precisely, we have that $C(\chi, \phi, \Delta) = 0$ if and only if the presence of χ makes no difference to whether or not ϕ brings about Δ . Similarly, we have that $C(\chi, \phi, \Delta) = 1$ if the presence or absence of χ determines whether or not ϕ brings about Δ .

To illustrate, in the example at hand, let $\chi = \text{Colleagues} = (\text{Gamma}, \text{Delta})$. In the former variant, where Beta is prejudiced and the presence of Gamma and Delta makes no difference to the outcome, $C(\text{Colleagues}, \text{Beta}, A) = 0$. In contrast, in the scenario where Beta is not prejudiced and would not discriminate against Omega were it not for Gamma and Delta, we have that $C(\text{Colleagues}, \text{Beta}, A) = 1$.

Note, however, that $C(\chi, \phi, \Delta)$ can take *any* value between 0 and 1. This is important to accommodate more nuanced and realistic cases, where the effect is a combination of personal prejudice and peer pressure. For instance, were Beta to be both prejudiced and subject to peer pressure, we would have that: $R(\text{Beta}, A) = x$, where x takes as values either [0] or [1]; and $C(\text{Colleagues}, \text{Beta}, A) = y$, where y takes some value $0 < y < 1$. A value of $C(\chi, \phi, \Delta) = y$, where $0 < y < 1$ suggests that ϕ does have some degree of prejudice with regards to t-bearing individuals, but such prejudice would not be sufficient to induce discrimination in the absence of peer-pressure. Here the presence of χ is likely to influence whether or not ϕ brings about Δ , but in a non-sure-fire fashion.

Before moving on, a couple of clarifications are in order. First, in our toy example, $C(\chi, \phi, \Delta)$ is a measure of the peer pressure that Beta's colleagues Gamma and Delta exert, more or less implicitly, on Beta not to hire an employee who bears the stigmatic trait t . However, $C(\chi, \phi, \Delta)$ could also model a variety of related phenomena. Examples include well-known situations in which the presence of witnesses makes a difference to the behavior of an agent or where the presence of other helpers influences the decision to stop and help a driver in need on a highway.

Second, we should emphasize that assigning a value of $C(\chi, \phi, \Delta) = 1$ does *not* necessarily imply that ϕ is not responsible at all for her actions. What it does show is that there are different social dynamics at play, which should be considered when modeling the causal situation. Once again, this underscores the importance of introducing three distinct functions. Together, they provide a powerful tool that can be used to represent aggregative social phenomena.

Taking stock, our proposal is to enrich causal models of aggregative social phenomena by including three functions: $R(\phi, \Delta)$, $E(\phi, \Delta)$, and $C(\chi, \phi, \Delta)$ by means of which we can begin to capture, respectively, the relevance of causes to effects, the degree of redundancy of an effect, and the influence that causes exert on other causes.

One final remark. Clearly, these functions can be further enriched in various ways. For instance, one might decide to introduce a threshold in function $C(\chi, \phi, \Delta)$, which determines when causal responsibility is significant enough to praise or reprehend an agent. Or one may also want to introduce a fourth function, $P(\theta, \Delta)$, intuitively modeling the *causal production* of cause θ on effect Δ along the lines of Salmon (1998). Finally, as noted by Hall (2004), there are cases in which the function $C(\chi, \phi, \Delta)$ is greater than 0 for some of the possible collections of causes where the $E(\phi, \Delta)$ of the proximate cause is not 1. That is, there are cases in which some events act on proximate causes that do not have the ability to bring about the effect *by themselves*. These are scenarios that our account, suitably expanded, is able to capture.

6. Concluding Remarks

The goal of this essay was to set up a framework for a quantitative analysis of the distinctive roles of different types of causation. Our contribution bears on two related areas of research: the study of causal influence and the propagation of aggregative social phenomena. With respect to the first,

we suggested that, in addition to influence on effects, it is also possible to have causal influence on proximate causes, as already implicitly entailed in Dretske 1988 and Haslanger 2016. Also, and equally important, it is possible to have causal influence on causes that are not, by themselves, capable of fully bringing an effect to completion. With respect to the second area of research, our model adds one dimension of analysis for aggregative social phenomena, namely, a measure of their robustness based on redundancy.

Our analysis can be applied to several other phenomena, such as “framing effects” and other biases that are central to contemporary studies in psychology and behavioral economics (Kahneman 2011). Consider *confirmation biases* according to which people tend to favor information which confirms their current beliefs or hypotheses. To understand the social and economic effects of such distorting factors, it is not sufficient to model the cognitive and neural mechanisms that implement them in the brain. In addition, one must factor in the mechanisms that control the flow of information (serial TV commercials are an example of overdetermination in action), preempting/preempted backups in neuromarketing, and the surprising power of aggregate masses of consumers.

Finally, we should stress that our three functions do not allow one to directly derive assessments of responsibility – moral, legal, or other. We are not providing an algorithm intended to replace normative judgment. Rather, we are offering a representational tool through which one can determine and refine the relevant actors, their respective roles, and their relationships and embed them in a causal model of the social structures they comprise.

Andrea Borghini

Department of Philosophy, University of Milan, Italy

Marco J. Nathan

Department of Philosophy, University of Denver, Colorado, USA

Acknowledgments

Both authors contributed equally to this essay. We would like to express our gratitude to two anonymous reviewers for this journal, Bill Anderson, Cristina Bicchieri, Giovanni Boniolo, Larry Cahoon, Enzo Fano, Naomi Reshotko, Davide Serpico, Mika Smith, and Davide Vecchi for insightful comments on various drafts of this article. Earlier versions have been presented at the University of Lisbon, the University of Denver, the University of Parma, and the University of Florence, Pistoia campus. All audiences provided helpful feedback.

¹We shall try to minimize any substantive assumption regarding the nature of *stigma*, *bias*, and *prejudice*, important issues that transcend the scope of this essay. For present purposes, it should suffice to define “bias” as an unfair prejudice in favor or against a specific object, individual, or group. A “prejudice,” in turn, is understood as a preconceived opinion, implicit or explicit, which can potentially cause harm. Finally, by “stigma” we simply intend a trait that is – at some point in the past, present, or future – the potential object of prejudice and bias.

² Other examples of biases and stigmas concern individuals affected by diseases (Kraut 1994), AIDS/HIV (cf. United Nations Programme 2007), obesity (Puhl and Heuer 2010), and the variation of prejudices concerning body image within communities that differ in terms of ethnic composition and social status (Grogan 2008: 156-66).

³ This view presupposes an “individualistic” stance, according to which each single agent is causally and morally responsible for their beliefs, desires, and other psychological states, and for acting upon them (see Garfinkel 1981 and, for a more nuanced discussion, Epstein 2014).

⁴ Qualifiers such as “overt,” “substantial,” and “implicit or explicit” are important because there is a substantial body of psychological evidence which suggests that virtually all individuals are subject to forms of *implicit* bias (Greenwald and Banaji 1995, Banaji and Greenwald 2013, Greenwald et al. 2015). While our causal model is intended to capture all forms of influence, we should be somewhat cautious in addressing the normative implications of such claims, which are rapidly becoming an important topic of philosophical discussion (Brownstein and Saul, 2016).

⁵ Introducing morally problematic implications in a conversation can cause harm even if the presuppositions are not endorsed by any participant. As McGowan points out, this may be due to two factors. First, the normative force of the speaker’s contribution is enabled by the accommodating capacity of the conversational score, which is independent of intentions and recognition thereof. Second, from the point of view of an external observer, absence of explicit presupposition-blocking usually signals tacit endorsement by participants in the conversation.

⁶ Setting metaphysical subtleties aside, we shall not provide a clear distinction between *processes* and *events*. However, unless specified otherwise, we intend both process and events as particular (*token*), rather than universal (*type*) entities. Otherwise, as far as we can see, our account poses no other substantive assumption.

⁷ Labeling the backup-assassin scenario as a case of “early” preemption draws attention to subtle albeit significant differences with related “late” examples as, say, the case of two assassins who simultaneously try to kill a victim, one reaching the victim first and making the other redundant, but both doing some active work towards their goal.

⁸ Note that this classification depends, in part, on our choice to examine a case of preemption. Had we focused on a case of *overdetermination*, we would have had redundancy without preemption. For example, suppose that Beta, Delta, and Gamma had time to spare on the day of the interview and decided to jointly interview Omega. Since all of them bear a prejudice against trait *t*, they reject Omega for the position. In these circumstances, the presence of one of Beta, Delta, and Gamma as an interviewer would have sufficed to reject Omega’s candidacy. The simultaneous presence of all three interviewers merely overdetermined such outcome.

⁹ Humorously alluding to the metaphysical inclinations of his colleagues at Cambridge, the Oxford philosopher Peter Geach argued that a particular has a “Cambridge property” when it acquires or loses such property without any effort. For instance, a person becomes a brother in law without actually doing anything, but simply in virtue of one of his siblings getting married. Following suit, we suggest that an agent may be treated as a “Cambridge cause” when it becomes part of a causal structure without playing any active role, without performing any action, as in the case of a redundant backup.

¹⁰ As noted, a modal realist can maintain that redundant Cambridge causes do something in virtue of the activities of their counterparts. Unfortunately, only some tributary accounts of the metaphysics of modality can offer replies matching modal realism. Candidate accounts are those with a realist bent, such as dispositionalism or a theory of grounds (Borghini 2016, Chapter 7). According to these accounts, indeed, Delta and Gamma can be regarded as contributing causes to the bias against Omega because they genuinely could have exerted this bias, in virtue of a disposition to act in that way that both Delta and Gamma actually possess. Another, perhaps more ambitious line of reply may hinge on a revision of our concept of agency. Perhaps even inert entities, such as Cambridge causes, should sometimes be regarded as *doing* something. That is, a cause can be such even without acting towards its effect. In short, regarding Cambridge causes as genuine causes calls for a bold approach to either modality or causation.

¹¹ Lewis’s famous “neuron diagrams,” which illustrate graphically causal relations in terms of arrows and vertexes, arguably went in this direction (Bright et al. 2016). However, they miss some hues that are especially important to describe aggregative social phenomena, such as a numerical repartition of causal responsibility. One might want to claim that Beta, Gamma, and Delta all play a causal role in triggering the discrimination of Omega. However, in saying so, one is not committed to the claim that all three individuals are *equally* responsible.

¹² These three functions may be somewhat reminiscent of Dretske’s distinction between triggering and structuring causes (1988) and of Haslanger’s (2016) subsequent use of such notions. We should stress, however, that the present functions are not cashed out in terms of *structures*, but in terms of *degrees of redundancy*.

¹³ Equally plausible would be to try and apply causal Bayes nets to the model. In general, assigning numerical values to functions within a model is a complex task, which leaves open different options. For a comprehensive discussion

of how this can be done, a standard starting point are the volumes by Krantz et al. (1971); for a fine-tuning of different concepts involved in causal modeling, see Hitchcock (2009) and Eberhardt (2009 and 2017). We thank an anonymous reviewer for bringing this to our attention.

¹⁴ While here we are trying to model *A*, a simple case of particular causation (the event of Omega not getting hired in a scenario of preemption), it is also possible to model *general* causal relations, such as patterns of hiring decisions in companies of a certain kind across a city, state, or country. The model developed here can be generalized to conceptualize causes and effects as either *types* or as *tokens*; yet, the mode with which those models assign values to $R(\phi, \Delta)$ will likely require separate considerations.

¹⁵ Clearly, it is also possible to take a weighted measure of redundancy, by expressing the degree of redundancy in terms of the product $\prod E(\tau, \Delta)$ of all values of E , as opposed to the sum $\sum E(\tau, \Delta)$. For the sake of brevity, we shall not further explore this proposal here. We thank an anonymous reviewer for pointing this out.

¹⁶ To emphasize the difference between $R(\phi, \Delta)$ and $E(\phi, \Delta)$, we apply the distinction to some familiar cases of redundant causal influence. We begin with the infamous “backup assassin scenario” presented above, where a backup assassin (b) is hired to kill designated victim (v), but v is killed by the “regular” sniper (s). Given that, here, s is the preempting cause of the killing of v (call this event “ K ”) and b is the preempted alternative, we have that $R(s, K) = 1$ and $R(b, K) = 0$, because s actually pulls the trigger, whereas b does nothing – this is why we referred to this scenario as a case of “early preemption.” From the perspective of function $E(\phi, \Delta)$, however, we have that $E(s, K) = 1$ and $E(b, K) = 1$ because both snipers have the *ability* to fully bring the effect to completion, by themselves. In contrast, in a scenario where it is necessary for s and b to cooperate or work as a team to kill v , we have that $R(s, K) = 1$, $R(b, K) = 1$, $E(s, K) = 0.5$, and $E(b, K) = 0.5$, since both play a role in producing the effect, but neither has the ability to do it alone assuming, for the sake of simplicity, that their contribution is equal. Finally, consider a standard case of overdetermination, where agents m and n jointly lift a chair (event C) that could be easily lifted alone by either, individually. Here we have that $R(m, C) = 1$ and $R(n, C) = 1$, since both are doing the lifting, and $E(m, C) = 1$ and $E(n, C) = 1$, since both could produce the effect alone. What all of these toy examples show is that cases of preemption are distinguished from cases of overdetermination in terms of their actual contribution, captured by $R(\phi, \Delta)$. However, what characterizes all instances of preemption and overdetermination, setting them apart from non-redundant scenarios, is precisely their degree of redundancy, captured by $E(\phi, \Delta)$.

¹⁷ The qualifier that the collection be minimal purports to capture the important idea that each of the events in the collection should be necessary, that is, they should play an “indispensable, non-redundant” role, in the sense that no proper part of the collection in question should be able to trigger the same effect. Depending on the system that we are trying to model, satisfying this condition might not be trivial.

References

Ayala, Saray and Vasilyava, Nadya. 2016. “Responsibility for Silence,” Journal of Social Philosophy, vol. 47 pp. 256-272.

Bicchieri, Cristina 2006. The Grammar of Society: The Nature and Dynamics of Social Norms Cambridge: Cambridge University Press.

Banaji, Mahzarin R. and Anthony G. Greenwald. 2013. Blindspot New York: Delacorte Press.

Beebe, Helen, Christopher Hitchcock, and Peter Menzies (eds.). 2009. The Oxford Handbook of Causation Oxford: Oxford University Press.

Borghini, Andrea. 2016. An Introduction to the Metaphysics of Modality, New York: Bloomsbury Press.

Bright, Liam Kofi, Malinsky, Daniel, and Thompson, Morgan. 2016. Causally Interpreting Intersectionality Theory. Philosophy of Science vol. 83, pp. 60–81.

Brownstein, Michael. and Saul, Jennifer, (eds). 2016. Implicit Bias and Philosophy, Vol. 1 and 2 Oxford: Oxford University Press.

Dowe, Phil. 2000. Physical Causation Cambridge: Cambridge University Press.

Drestke, Fred. 1988. Explaining Behavior: Reasons in a World of Causes. Cambridge: MIT Press.

Eberhardt, Frederick. 2009. "Introduction to the Epistemology of Causation", Philosophy Compass, vol. 4, no. 6, pp. 913–925.

Eberhardt, Frederick. 2017. "Introduction to the Foundations of Causal Discovery", International Journal of Data Science and Analytics, vol. 3, no. 2, pp. 81–91.

Epstein, Brian. 2014. "What Is Individualism in Social Ontology? Ontological Individualism vs. Anchor Individualism" in Rethinking the Individualism-Holism Debate, Collin, Finn and Zahle, Julie (eds.) Dordrecht: Springer, pp. 17-38.

Garfinkel, Alan. 1981. *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven: Yale University Press.

Greenwald, Anthony G. and Mahzarin R. Banaji. 1995. "Implicit Social Cognition: Attitudes, Self-esteem, and Stereotypes", *Psychological Review*, vol. 102, no. 4, pp. 4-27.

Greenwald, Anthony G., Banaji, Mahzarin. R., and Nosek, Brian A. 2015. "Statistically Small Effects of the Implicit Association Test Can Have Societally Large Effects" Journal of Personality and Social Psychology vol. 108, pp. 553–561.

Grogan, Sarah. 2008. Body Image. Understanding Body Dissatisfaction in Men, Women, and Children New York: Routledge.

Hall, Ned. 2004. "Two Concepts of Causation." In J. Collins, N. Hall, and L.A. Paul (eds) Causation and Counterfactuals. Cambridge: MIT Press: pp. 225–276.

Hall, Ned and Paul, Laurie. 2013. Causation. A User's Guide, Oxford: Oxford University Press.

Haslanger, Sally. 2016. "What is a (Social) Structural Explanation?" Philosophical Studies vol. 173, pp. 113-130.

Hitchcock, Christopher R. 2009. "Causal Models," in Oxford Handbook of Causation, Beebe H., C.R. Hitchcock, and P. Menzies (eds.) New York: Oxford University Press, pp. 299–314.

Kahneman, Daniel. 2011. Thinking, Fast and Slow New York: Farrar, Straus and Giroux.

Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky, 1971, Foundations of Measurement Vol 1: Additive and Polynomial Representations, San Diego and London: Academic Press.

-
- Kraut, Alan M. 1994. Silent Travelers: Germs, Genes, and the "Immigrant Menace." Baltimore: Johns Hopkins University Press.
- Le Bon, Gustave. 1986. The Crowd: A Study of the Popular Mind, New York: McMillan.
- Lewis, David K. 2000. "Causation as Influence," Journal of Philosophy, vol. 97, pp. 182-197
- Lewis, David K. 1986. On the Plurality of Worlds, Oxford: Blackwell.
- Lewis, David K. 1973. "Causation." Journal of Philosophy, vol. 70, pp. 556-567.
- Malinsky, Daniel. 2017. "Intervening on Structure." Synthese vol 195, no. 5, pp. 2295-2312
- Manne, Kate. 2017. Down Girl: The Logic of Misogyny. New York: Oxford University Press
- McGowan, Mary Kate. 2004. "Conversational Excercises: Something Else We Do With our Words." Linguistics and Philosophy vol. 24, no. 1, pp. 93-111.
- McGowan, Mary Kate. 2009. "Oppressive Speech." Australasian Journal of Philosophy vol. 87, pp. 889-407.
- Nathan, Marco J. 2014. "Causation by Concentration," British Journal for the Philosophy of Science, vol. 65, pp. 191–212.
- Puhl, Rebecca M. and Chelsea A. Heuer. 2010. "Obesity Stigma: Important Considerations for Public Health." American Journal of Public Health vol. 100, no. 6: pp. 1019-1028.
- Salmon, Wesley C. 1998. Causality and Explanation New York: Oxford University Press
- Stiglitz, Joseph E. 2012. The Price of Inequality. How Today's Divided Society Endangers Our Future. New York: Norton.
- Strevens, Michael. 2008, Depth: An Account of Scientific Explanation Cambridge: Harvard University Press
- Strevens, Michael. Forthcoming. "The Causal Role of Aggregate Properties." Minnesota Studies in the Philosophy of Science.
- United Nations Programme on HIV/AIDS, 2007. Reducing HIV Stigma and Discrimination: A Critical Part of National AIDS Programmes. Geneva, Switzerland.
- Valian, Virginia. 1998. Why So Slow? The Advancement of Women. Cambridge: MIT Press.
- Woodward, James. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation," Biology and Philosophy vol. 25, no. 3, pp. 287-318.