

CAUSATION AND COUNTERFACTUAL DEPENDENCE IN
ROBUST BIOLOGICAL SYSTEMS¹

ABSTRACT

In many biological experiments, due to gene-redundancy or distributed backup mechanisms, there are no visible effects on the functionality of the organism when a gene is knocked out or down. In such cases there is apparently no counterfactual dependence between the gene and the phenotype in question, although intuitively the gene is causally relevant. Due to relativity of causal relations to causal models, we suggest that such cases can be handled by changing the resolution of the causal model that represents the system. By *decreasing* the resolution of our causal model, counterfactual dependencies can be established at a higher level of abstraction. By *increasing* the resolution, stepwise causal dependencies of the right kind can serve as a sufficient condition for causal relevance. Finally, we discuss how introducing a temporal dimension in causal models can account for causation in cases of non-modular systems dynamics.

1. INTRODUCTION

Counterfactual dependence accounts of causation have several problems accounting for causation in complex biological systems (Mitchell 2009, Strand and Oftedal 2009). Often perturbations on such systems do not have any clear-cut phenotypic effects, and consequently there is no direct counterfactual dependence between the cause candidate intervened on and the effect considered. For example, many gene knockouts and knockdowns have no detectable effect on relevant functionality, even though the genes in question are considered causally relevant in non-perturbed systems (Shastry 1994, Wagner 2005).

Two different mechanisms give rise to such stability: (1) gene redundancy; the workings of backup-genes explain the lack of counterfactual dependence between the effect and the preempting cause, and (2) distributed robustness; the system readjusts functional dependencies among other parts of the system rather than invoking backup genes. The latter cases challenge not only the necessity of coun-

1 Thanks to Henrik Forssell, Sara Green, Carsten Hansen, Heine Holmen, Veli-Pekka Parkkinen, the audience at the ESF Philosophy of Systems Biology Workshop, Aarhus University, and the participants at the colloquium for analytic philosophy, Aarhus University, for helpful comments and suggestions.

terfactual dependence for causation, but also our thinking about truth conditions for the relevant counterfactuals. We suggest that such cases can be handled by a counterfactual dependence account of causation by changing the resolution of the causal model.

There is a related problem concerning non-modularity of complex systems. Modularity is by some seen as a requirement for making adequate causal inferences (Woodward 2003). The intuitive idea is that different mechanisms composing a system are separable and in principle independently disruptable (Hausman and Woodward 1999). However, research indicates that compensatory changes in response to disruptions in biological systems can change functional relations between relevant variables and thereby violate modularity.

In the following we first introduce the core elements of a counterfactual dependence based philosophical analysis of causation. Then we present gene redundancy and distributed robustness using biological examples and argue that changing representational resolution helps understand causal dependence in these cases. Finally, we discuss how introducing a temporal dimension in causal models gives a grip on non-modular systems dynamics.

2. SKETCH OF AN ANALYSIS OF CAUSATION

Causes typically make a difference to their effects, and many philosophers argue that this idea should be at the core of the philosophical analysis of causation (e.g. Lewis 1973, 2004, Woodward 2003, Menzies 2004). We agree and suggest the following general definition of causation, where X and Y are variables and M is a causal model, i.e. a set of variables and functional relations between them:

Causal Relevance: X is a cause of Y relative to M if and only if there is a change of X that would result in a change of Y when we hold some subset of variables (allowing this set to be empty) in M fixed at some values.

This definition is in line with other well discussed difference-making accounts of causation viewing causal relata as variables (e.g. Menzies and Woodward). Such views capture the idea that causal relations are exploitable for purposes of manipulation and control.

The requirements of *some* subset of variables being held fixed at *some* values are chosen with care. The main idea is that this definition states causal relevance in the broadest sense, and that different explications of the relevant subset of variables and their relevant values give different kinds of causal relevance. Letting the subset be empty, for example, gives the notion of a total cause:

Total Cause: X is a total cause of Y relative to M if and only if there is a change of X that would result in a change of Y.

Also the notion of a direct cause comes out as a special case:

Direct cause: X is a direct cause of Y relative to M if and only if there is a change of X that would result in a change of Y when we hold all other variables in M fixed at some values.

Causal paths are understood as chains of relations of direct causation. Using this idea, we can cash out the notion of a contributing cause:

Contributing Cause: X is a contributing cause of Y relative to M if and only if there is a change of X that would result in a change of Y when we hold all variables in M not on the relevant path from X to Y fixed at some values.

‘Relevant path’ is a placeholder for the chain of direct causal relations between X and Y over which we are checking for mediated causal relevance among X and Y. The existence of a mediating chain of direct causal relationships is not itself sufficient for causal relevance due to counterexamples to transitivity. Furthermore, actual causation can be specified in terms of all variables not on the relevant path between X and Y being held fixed at their actual values.

These distinctions, which mirror Woodward’s 2003 distinctions, are not exhaustive. We add two additional notions here, tentatively called *Restricted Causal Relevance* and *Dynamic Causal Relevance* (Section 5). The idea behind restricted causal relevance is to capture causal understanding often implicit in actual scientific practice, where variables not tested for causal relevance are held fixed at assumed normal or expected values. This is *restricted* causal relevance because it requires counterfactual dependence under a limited range of values of the variables held fixed.

Restricted Causal Relevance: X is a restricted cause of Y relative to M if and only if there is a change of X that would result in a change of Y when we hold all variables in M not on the relevant path from X to Y fixed at their normal values.

There will be a variety of different notions of restricted causal relevance. The one stated here is analogous to contributing cause, and should be sufficient to illustrate the core idea.

Counterexamples to the claim that counterfactual dependence is necessary for causation feature a redundancy of cause candidates: preemption (a cause preempts a backup cause), overdetermination (two individually sufficient causes), and trumping (a cause trumps another cause candidate). Moreover, distributed robust-

ness found in biological systems presents yet another counterexample, one that has not received sufficient attention in the philosophical literature, but has some interesting and perhaps surprising philosophical consequences.

It is important to be aware that preemption cases are only problematic for the more demanding varieties of causal relevance. For a standard case of preemption to arise in the first place, both cause candidates C_1 and C_2 must be causally relevant to E in the broad sense of causal relevance (see Section 4). When philosophers ask which of C_1 and C_2 are actual or restricted causes of E problems occur, because asymmetry intuitively should arise for these more demanding notions. It is only if C_2 is causally relevant in the broad sense that it can be a preempted backup cause in actual or normal circumstances where it is not an actual or restricted cause itself. Allowing for a proper description of type-level preemption cases is the main role of the notion of restricted causal relevance in this paper.

3. GENETIC REDUNDANCY AND DISTRIBUTED ROBUSTNESS

Gene knockout and knockdown experiments investigate the functioning of genes by effectively deleting or silencing specific genes (e.g. Xie et al. 2005). Hypothesis-driven experiments of this sort often involve causal reasoning of the form that if the procedures make a difference to a particular phenotypic trait, then the gene in question is causally relevant for that trait. However, due to system robustness, very often gene perturbations do not have any apparent effect on the functionality of the system at hand (Shastri 1994, Wagner 2005).

Robustness is a ubiquitous property of living systems and allows systems to maintain their biological functioning despite perturbations (Kitano 2004, 826). Mutational robustness can be described as functional stability against genetic perturbations (Strand and Oftedal 2009), and two types are recognized in the literature; genetic redundancy and distributed robustness (Wagner 2005). Genetic redundancy involves multiple copies of a gene or genes with similar functionality (so-called duplicate genes) that can take the role of the perturbed gene. Distributed robustness is more complex and involves organizational changes of multiple causal pathways in such a way that the system manages to compensate for the genetic disturbance (Hanada et al. 2011).

Genetic redundancy was investigated in Kuznicki et al. (2000), where duplicate genes were found to contribute to robustness in the nematode *C. elegans*. GLH proteins (GermLine RNA Helicases) are constitutive components of the nematode P granules. These granules are distinctive bodies in the germ cells found to have roles in the specification and differentiation of germ line cells. The genes associated with the proteins GLH-1 and GLH-4 belong to the multi-gene GLH family in *C. elegans*, and the GLHs are considered important in the development of egg cells (oogenesis). Still, no effect on oogenesis could be detected either from

a RNAi knockdown of the gene associated with GLH-1 or with GLH-4. However, the combinatorial knockdown of both the GLH-1 and GLH-4 genes resulted in 97% sterility due to lack of egg cells and defective sperm. The results indicate that GLH-1 and GLH-4 are duplicate genes that can compensate for each other when one or the other is lacking. We return to this example in Section 4.

Distributed robustness was investigated in Edwards and Palsson 2000a and 2000b where chemical reactions in *E. coli* were perturbed. Rather than knocking out or knocking down genes, chemical reactions in the process of glycolysis (the metabolic process of converting glucose into pyrovalate and thereby produce the energy rich compounds ATP and NADPH) were blocked one by one to find whether any of the reactions were essential to cell growth. Only seven of the 48 reactions were found to be essential, and of the 41 remaining, 32 reduced cell growth by less than 5%, and only nine reduced cell growth with more than 5%. For example, the blocking of the enzyme G6PD (glucose-6-phosphate dehydrogenase) resulted in growth at almost normal levels. However, the elimination of this reaction had major systemic consequences (Wagner 2005). Instead of producing two-thirds of the cell's NADPH (a coenzyme needed in lipid and nucleic acid synthesis) by the pentose phosphate pathway, more NADH was produced through a different path, the tricarboxylic acid cycle, and this NADH was then transformed into NADPH via a highly increased flux through what is called the transhydrogenase reaction. In other words, practically all the NADPH needed for upholding normal cell growth was still produced, but through different pathways. We return to this example in Section 5.

4. DEALING WITH REDUNDANCY

When a duplicate gene takes the role of a silenced gene, there is typically no phenotypic change that indicates causal relevance of the silenced gene². Consider a standard case of late preemption (Figure 1).




Figure 1

² Certain fine-grained redescrptions of the phenotypic effect may still reveal dependencies. This could be cashed out in terms of a change of causal model or different restrictions on the values of the variables.

According to our definition of type level causal relevance, both cause candidates are causally relevant, but only C_1 is actual cause. Type level preemption cases can be formulated in terms of restricted causal relevance. C_1 and C_2 may both be broadly causally relevant, while both should not be restricted causes. Suppose that normal values of C_1 and C_2 are that both are present. Then we should be able to say that in normal cases, C_1 is a restricted cause, while C_2 is not, even though C_2 would be a cause in abnormal knockout cases where C_1 is not present.

One might think that this asymmetry could be accounted for in terms of specificity (Woodward 2010) or fine-tuned influence (Lewis 2004, 92). The idea is that there is an asymmetry between the preempting cause and the preempted backup, because the relation between the preempting cause and the effect is such that one can make minor changes to the cause that are followed by minor changes in the effect, while there is no analogue for the preempted backup. Intervening to slightly alter the preempted backup will not change the effect at all. This strategy is promising, and can cover several cases, but it does not work in full generality. In particular, it does not handle cases of threshold causation, where fine-tuning of the preempting cause either changes the effect from occurring to not occurring, or makes no difference at all.

On our definitions, like on Woodward's (2003) and Menzies' (2004), causal relations obtain relative to a causal model. On this background, we see how causal dependencies can be masked and/or revealed by changing the resolution of the causal model, for example by invoking a more coarse-grained model. Consider a simplified gene-redundancy scenario (Figure 2). Binary variables v_1 , v_2 , and v_3 represent the presence or absence of three functionally similar genes. Consider another representation involving only one binary variable, v_4 , that takes the value present when at least one of v_1 , v_2 , or v_3 take the value present, and the value absent when all of v_1 , v_2 , and v_3 take the value absent.




Figure 2

Interventions changing the value of v_4 will directly affect the effect variable in this setup, and the counterfactual dependence of the effect on v_4 is straightforward. Applied to the example of redundancy in the previous section, this corresponds to

considering the disjunction of the two mutually compensating genes GLH-1 and GLH-4 as the relevant variable in relation to germ cell formation and not the presence of the individual genes. Whether higher-level representations like these have, or should have, realist interpretations in terms of e.g. modally robust entities is a further question (see Strand and Oftedal 2009). Alternatively, one may introduce a fine-grained effect-variable in order to reveal fine-tunable causal influence.

5. DEALING WITH DISTRIBUTED ROBUSTNESS

In systems exhibiting distributed robustness, organization and causal paths can be rewired under perturbations. Such systems can retain their biological functions by changing their causal structure compared to the structure they would have had in the absence of that perturbation. Systems with such behavior can be non-modular in the sense that the intervention on one causal factor, for example a gene, changes causal relations between other factors in the system.³ *Prima facie*, the perturbed gene is a cause in the normal case even if there is a distributed back-up mechanism at play in the perturbed case. The causal analysis should account for the gene being causally linked to the relevant phenotypic trait in the normal case, even in the absence of the right kind of direct counterfactual dependence.

We consider two options. One is to decrease the representational resolution by abstracting away from details, and thereby in effect treating systems with distributed robustness as modules that are not internally modular. Interventions on such systems will be radical; wipe out the whole module. The other is to increase the representational resolution, in the sense that one zooms in on the relevant gene and the causal paths leading from that gene to the effect in question. This is done by introducing causal intermediaries and tracking stepwise causal dependencies. If one could establish stepwise counterfactual dependence, one could for example take the ancestral relation of counterfactual dependence, and thus establish the causal status among distant nodes that are not related directly by counterfactual dependence.⁴ We elaborate on the second option in the following.

First, consider a relatively simple abstract case of distributed robustness (Figure 3).

3 E.g. Woodward (2003) and Cartwright (2001) discuss modularity.

4 The idea in David Lewis' original account is that whenever you have a chain c_1, c_2, \dots, c_n and each c_m and c_{m+1} are related by counterfactual dependence, any two distant elements in that chain is causally related by definition. Taking the ancestral was crucial to secure transitivity for Lewis, but transitivity is questioned in the contemporary debate (e.g. Woodward 2003, Paul 2004).




Figure 3

Imagine a knockout on v_1 that changes the functional dependencies between some of the other parts. The point is not exactly what changes are being induced, rather that such changes indeed occur. In the redundancy case, no such changes occur, it is simply a backup gene performing the function of the knocked out gene.

Now, consider the path v_1, v_2, v_5 , to e . If this is a causal path, there will be causal interventions on v_1 that change the value of v_2 , causal interventions on v_2 that change the value of v_5 , and causal interventions on v_5 that change the value of e . However, the changes in v_2 brought about by changes of v_1 may not be such that they induce changes in v_5 . Rather, it may be that the interventions on v_2 that do change v_5 cannot be induced by intervening on v_1 . In such a case, there will be no direct counterfactual dependence, but there will still be a path between v_1 and e . The question is whether v_1 qualify as a cause of e ? If we straightforwardly take causation to be the ancestral of counterfactual dependence, v_1 will qualify as a cause of e . However, this is too weak and deems some non-causal relations causal.

On the other hand, one might think that causal relevance between v_1 and e is mediated via a causal path if and only if there is a causal intervention on v_1 that changes the value of e (Woodward 2003 requires this). This, however, is too strong. In a case of distributed robustness, changes of v_1 that brings about certain changes in v_2 might also trigger distributed backup mechanisms that affect whether v_2 and/or v_5 can bring about changes in e . If the system is non-modular, it will be impossible to control for such backups by holding other variables fixed. We need to find some middle ground between taking the ancestral which is too weak, and requiring direct counterfactual dependence which is too strong.

Here is a tentative account. It is sufficient for causal relevance that there are changes of v_1 that result in changes of v_2 , changes within the range of changes that can be brought about on v_2 by changing v_1 that result in a range of changes in v_5 , and finally the same for v_5 and e . If there is such a series of ranges of changes, then

v_1 is a cause of e even if there are no changes of v_1 that would result directly in changes of e . Let's label this the relevance requirement.

An example is the metabolic reactions in *E. coli* presented previously. The distributed robustness of these reactions makes sure that practically the same amount of the energy rich compound NADPH is produced even though the pentose phosphate pathway, which normally is considered the main source of NADPH, is blocked by knocking out the enzyme G6PD (Edwards and Palsson 2000a, 2000b). As shown in the figures below, NADPH is mainly produced through the pentose phosphate pathway when no chemical reactions are blocked. When G6PD is knocked out, however, NADPH production goes through different pathways. The tricarboxylic acid cycle produces NADH at elevated levels and this NADH is transformed into NADPH through the transhydrogenase reaction.

The following illustrations are adapted from Edwards and Palsson (2000b). Additional nodes are introduced to represent the key chemical reactions DH (dehydrogenation) and DC (decarboxylation). Black arrows represent the main causal pathways (high flux). Grey arrows represent minor causal pathways (low flux). Dashed arrows represent no-flux pathways. Figure 4 shows a causal representation of glucose metabolism in *E. coli* under normal circumstances. NADPH is mainly produced through the pentose phosphate pathway (in black).




Figure 4

Figure 5 shows a causal representation of glucose metabolism in *E. coli* when the enzyme G6PD is knocked out. No NADPH is produced through the pentose phosphate pathways (now in grey). Rather, there is an increased activity in different pathways ultimately leading through the tricarboxylic acid cycle and the transhydrogenase reactions (details omitted from the figure).

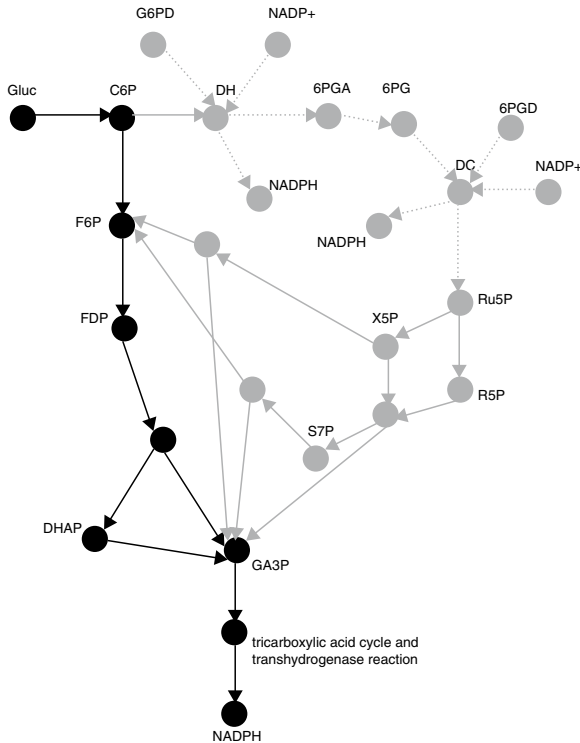


Figure 5

Even though there appears to be no counterfactual dependence between the enzyme G6PD that catalyzes parts of the pentose phosphate pathway and the total NADPH production, G6PD should still be considered causally relevant for the production of NADPH. Translating our tentative account of how causal relevance can be established into this example gives the following: It is sufficient for causal relevance that there are changes on G6PD that result in a range of changes of DH, changes within that range of changes on DH that result in a range of changes in 6PGA and NADPH, the same for 6PGA and 6PG, for 6PG and DC, and for DC and NADPH. If there is such a series of ranges of changes, then G6PD is a cause of NADPH even if there are no changes on 6PGA that would result in changes in NADPH production directly.

This account demonstrates a need to represent the effects of perturbations on the dynamic evolution of systems. Ignoring the dynamic dimension, and assuming that the systems remain fixed under perturbations, is deeply problematic when confronted with more complex system behavior. Intervening at a certain point in the dynamic evolution of a system may change the upcoming development, and when intervening at a later point we have the choice of intervening on a system that has not been disturbed, or on a *different system*, namely the one that was disturbed at an earlier point. These will be represented by different system trajectories along the dynamic dimension.

Such systems are not modular since we cannot intervene on causal paths of interest without changing other aspects of the systems. According to some philosophers (e.g. Woodward 2003), systems that are sufficiently non-modular are not causal. We think this response is too hasty. Moreover, eschewing this as a problem for dependence accounts of causation still leaves the problem of how we should understand counterfactual claims about non-modular functionally robust systems. We will therefore proceed treating it as a question about causation, trusting that it has philosophical value even if one should choose to label it otherwise.

Consider the following notion of causal relevance:

Dynamic Causal Relevance: *X* is a dynamic cause of *Y* relative to *M* if and only if there is a possible change of *X* that would result in a change of *Y* when we hold all variables in *M* fixed at some values at the time of intervention on *X*.

This notion allows for systemic changes over time due to earlier perturbations. We can represent the system in *n*-dimensional space, where *n* is the number of variables describing the system. The changes in the dynamic dimension can include changes of the functional relations among different parts of the system. In effect some system trajectories in this dimension will represent different systems than other trajectories. It might happen, for example in cases of distributed robustness, that the post-intervention system changes not only the values of the variables, i.e. its state, but also the functional dependencies among the variables. For such cases, we need an account that tells which counterfactual scenarios are relevant for evaluating counterfactual claims about the non-perturbed system.

Prima facie, there are two options. First, one may consider the perturbed system at a later point in its dynamic evolution, but this can mask causal relations since the system may have changed, and backups may have been triggered as a result of the perturbation. Second, one may intervene on a non-disturbed system identical to the system of interest up to the time of interest. Which choice we make can affect what relations come out as causal.

Since this is a question of when we can infer mediated causal relationships we need to get clearer on the general question of causal transitivity. The simplest general form of the standard counterexamples to transitivity requires a setup like the following:

Variables	v_1	v_2	e
Possible values	{a, b}	{1, 2, 3}	{@, \$}
Possible changes	(a, b)	(1, 2)	(@, \$)
		(1, 3)	(\$, @)
		(2, 1)	
		(2, 3)	
		(3, 1)	
		(3, 2)	

v_1 and e are binary variables, while v_2 can take three different values. Possible changes are represented by ordered pairs of values of the variables. In general, for a n-ary variable there will be $n(n-1)$ possible changes of values. The counterexamples arise when the changes that can be brought about in v_2 by intervening on v_1 do not overlap with the changes in v_2 that will result in changes on e .

Woodward gives an example where a dog bites his right hand. At the type level this can be represented by a binary variable taking the values {dog bites, dog does not bite}. The bite causes him to push a button with his left hand rather than with his right hand. This intermediate cause can be represented by the triadic variable {pushes with right hand, pushes with left hand, does not push}. The pushing of the button causes a bomb to go off, represented by the binary variable {bomb explodes, bomb does not explode}. Appeal to causal intuition tells us that the bite causes the pushing, the pushing causes the explosion, but the bite does not cause the explosion. When there are no changes of v_1 that result in changes in e , v_1 is not a cause of e even though there is a chain of direct causal dependencies connecting v_1 and e . Woodward’s way of dealing with the counterexamples is to deny that the ancestral of direct causal dependence is sufficient for causal relevance. To get a sufficient condition he requires interventions on v_1 that change the value of e when all variables not on the path from v_1 to e are fixed at suitable values. This latter requirement, however, is too strong.

What is needed to block the counterexamples to transitivity is a requirement of relevance and not of direct dependence. The changes brought about in v_2 by changing v_1 must be such that inducing some of *those changes* in v_2 results in changes of e . This relevance requirement accounts for the problem cases of transitivity more surgically. In particular, it leaves open the possibility that the existence of *the right kind* of causal chain is sufficient for causal relevance, even in the absence of direct dependence. In light of our earlier discussion of non-modular systems exhibiting distributed backup mechanisms, we can understand how such cases may arise. A variable can be causally relevant for an effect further downstream a certain causal path P, even if changes of that variable trigger distributed backup mechanisms that counterbalance or nullify the effect of further changes that would have been brought about along P.

In cases of distributed robustness there may be backup mechanisms that mask causal relations by ruling out counterfactual dependence between the cause and effect. We have suggested that such cases can be handled by establishing mediated causal relations that are not grounded directly by counterfactual dependence. This requires a chain of mediating causal relations of the right kind, given by the relevance requirement. In dynamic cases with distributed robustness, how should we think about the relevance requirement and about the truth conditions for the counterfactual dependencies?

Our tentative suggestion is that relevant counterfactuals should be evaluated by looking at systems that are similar to the systems of interest *at the time changes are induced*. When inducing multiple changes at different times, the counterfactual scenarios involve systems that are similar to the system of interest up to the point of the relevant change. Even if it is a variable upstream that we are interested in checking the causal relevance of, we should let the counterfactual target system evolve like the normal system up to the point of changes in downstream variables. In this way we avoid that distributed backups potentially triggered by earlier changes mask the mediated causal relationships we want to reveal. The way to think about truth-conditions for causal counterfactuals about systems exhibiting distributed robustness and non-modular behavior is to compare a normal system with different counterfactual systems subject to the same dynamic evolution as the normal system up to the time of changes of the mediating variables.

This is a tentative definition of causal relevance, in the broad sense, for systems changing their dynamic evolution as a result of perturbations. It is designed to be a special case of the general philosophical analysis of causation that we started out with. There will also be dynamic analogues to restricted and actual causation, by restricting the relevant values to normal values and to actual values respectively. Developing a full-fledged philosophical account along these lines is a task for future work, but we hope to have made a convincing case for the philosophical interest of representing the dynamics of causal systems.

6. CONCLUDING REMARKS

We have used biological examples of gene-redundancy and distributed robustness to suggest some extensions and revisions of the philosophical understanding of causation. The focus has been on cases of causation where there are no direct variable-on-variable counterfactual dependencies, and we have suggested that changing the resolution of the causal representation is a natural move in such cases. This can be done by increasing or decreasing the resolution of the causal model. Either way you go, causal claims face the tribunal of experience in concert. The relativization to a model puts the focus of causal investigation where it should

be; namely on generating good causal models, rather than establishing singular causal claims in isolation.

REFERENCES

- Cartwright, N., 2001, "Modularity: It Can, and Generally Does, Fail", in: M.C. Galavotti, P. Suppes, and D. Constantini (Eds.), *Stochastic Causality. CSLI Lecture Notes*. Stanford, CA: CSLI Publications, pp. 65-84.
- Edwards, J. S. and Palsson B. O., 2000a, "Robustness Analysis of the *Escherichia coli* Metabolic Network", in: *Biotech Progress* 16, pp. 927-939.
- Edwards, J. S. and Palsson B. O., 2000b, "The *Escherichia coli* MG1655 *in silico* Metabolic Genotype: Its Definition, Characteristics and Capabilities", in: *Proceedings of the National Academy of Science of the United States of America* 97, pp. 5528-5533.
- Hanada, K., Sawada Y., Kuromori T., Klausnitzer R., Saito K., Toyoda T., Shinozaki K., Li W.-H., and Hirai M. Y., 2011, "Functional Compensation of Primary and Secondary Metabolites by Duplicate Genes in *Arabidopsis thaliana*", in: *Molecular Biology and Evolution* 28, pp. 377-382.
- Hausman, D. and Woodward J., 1999, "Independence, Invariance and the Causal Markov Condition", in: *British Journal of the Philosophy of Science* 50, pp. 521-583.
- Kitano, H., 2004, "Biological Robustness", in: *Nature Reviews Genetics* 5, pp. 826-837.
- Kuznicki K. A., Smith P.A., Leung-Chiu W. M. A., Estevez A. O., Scott H. C. and Bennett K. L., 2000, "Combinatorial RNA Interference Indicates GLH-4 Can Compensate for GLH-1; These Two P Granule Components are Critical for Fertility in *C. elegans*", in: *Development* 127, pp. 2907-2916.
- Lewis, D., 1973, "Causation", in: *Journal of Philosophy* 70, pp. 556-567.
- Lewis, D., 2004, "Causation as Influence", in: J. Collins, N. Hall and L.A. Paul (Eds.), *Causation and Counterfactuals*. Cambridge (Mass.): The MIT Press, pp. 75-117.
- Menzies, P., 2004, "Difference Making in Context", in: J. Collins, N. Hall and L.A. Paul (Eds.), *Causation and Counterfactuals*. Cambridge (Mass.): The MIT Press, pp. 139-180.
- Mitchell, S., 2009, *Unsimple Truths. Science, Complexity and Policy*. Chicago: The University of Chicago Press.

- Paul, L. A., 2004, "Aspect Causation", in: J. Collins, N. Hall and L.A. Paul (Eds.), *Causation and Counterfactuals*. Cambridge (Mass.): The MIT Press, pp. 205-224.
- Shastry, B. S., 1994, "More to Learn from Gene Knockouts", in: *Molecular and Cellular Biochemistry* 136, pp. 171-182.
- Strand, A. and Oftedal G., 2009, "Functional Stability and Systems Level Causation", in: *Philosophy of Science* 76, pp. 809-820.
- Wagner, A., 2005, "Distributed Robustness versus Redundancy as Causes of Mutational Robustness", in: *BioEssays* 27, pp. 176-188
- Woodward, J., 2003, *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Woodward, J., 2010, "Causation in Biology", in: *Biology and Philosophy* 25, pp. 287-318.
- Xie, J., Awad K. S., Guo Q., 2005, "RNAi Knockdown of Par-4 Inhibits Neurosynaptic Degeneration in ALS-linked Mice", in: *Journal of Neurochemistry* 92, pp. 59-71.

Anders Strand

Department of Philosophy, Classics, History of Art and Ideas
University of Oslo
Box 1020 Blindern
0315, Oslo
Norway
anders.strand@ifikk.uio.no

Gry Oftedal

Department of Philosophy, Classics, History of Art and Ideas
University of Oslo
Box 1020 Blindern
0315, Oslo
Norway
gry.oftedal@ifikk.uio.no