

CONSTITUTION AND CAUSATION IN MECHANISMS

Constitución y causalidad en los mecanismos

MAX KISTLER ^{a, b}

<https://orcid.org/0000-0001-6613-5774>

mkistler@univ-paris1.fr

^a Université Paris 1 Panthéon-Sorbonne, Paris, France.

^b Institut universitaire de France, Paris, France.

Abstract

Craver (2007) has argued that constitutive relevance can be discovered by mutual manipulability, based on interventions (Woodward, 2003). However, the requirements on interventions make mutual manipulability of mechanisms and their constituents impossible (Baumgartner & Gebharter, 2016). Two proposals for providing empirical criteria for constitutive relevance are examined. Both provide only part of the information needed for constructing models of mechanisms that represent both causal links and constitutive relations.

Baumgartner, Casini, and Krickel's (2020) account is not adequate for mechanisms that contain 1) causal chains of activities and/or 2) activities working in parallel. Furthermore, their requirement to detect two or more activities simultaneously does not fit standard scientific practice. Craver, Glennan, and Povich's (2021) thesis that constitutive relevance can be reduced to "causal betweenness", between the input and the output condition of a mechanism, is not adequate for multi-level mechanisms. Models that do not contain any representation of constitution relations might represent the activities of some hypothetical fundamental level. However, both the existence of such a fundamental level and the application of causation to that level are problematic. Craver, Glennan, and Povich's account of the construction of models for mechanisms leads to the paradoxical result that there are no levels and thus no multi-level mechanisms, but only causal chains of fundamental level activities.

The key to understanding how models of multi-level mechanisms can be constructed on the basis of empirical information is that 1) the relevant experiments directly provide only information about causal relations (contrary to what Craver 2007 and Baumgartner, Casini, and Krickel 2020 claim), but that (contrary to what Craver, Glennan, and Povich 2021 claim) this information about causal relations can bear on variables at different levels. A multi-level model is built in two steps. 1) First, partial purely causal models are built for each hypothetical constituent variable Φ_i , on the basis of top-down and bottom-up experiments that modify or measure Φ_i in a level-specific way, 2) second, those partial models are merged in a comprehensive model containing both causal and constitution relations between variables, on the basis of information about each variable and spatio-temporal constraints.

Key words: Metaphysics of Science; Constitution; Constitutive Relevance; Causation; Mechanism; Intervention; Surgical; Interventionism; Level; Higher-Level; Multi-Level; Fundamental; Eliminativism; Nihilism; Mereological.

Resumen

Craver (2007) ha argumentado que la relevancia constitutiva puede descubrirse mediante la manipulabilidad mutua, basada en intervenciones (Woodward, 2003). Sin embargo, los requisitos sobre las intervenciones hacen imposible la manipulabilidad mutua de los mecanismos y sus constituyentes (Baumgartner & Gebharter, 2016). Se examinan dos propuestas para proporcionar criterios empíricos de relevancia constitutiva. Ambas proporcionan solo parte de la información necesaria para construir modelos de mecanismos que representen tanto los vínculos causales como las relaciones constitutivas.

La explicación de Baumgartner, Casini y Krickel (2020) no es adecuada para mecanismos que contienen 1) cadenas causales de actividades y/o 2) actividades que funcionan en paralelo. Además, su requisito de detectar dos o más actividades simultáneamente no se ajusta a la práctica científica estándar. La tesis de Craver, Glennan y Povich (2021) según la cual la relevancia constitutiva puede reducirse a la «causalidad entre», entre la condición de entrada y la de salida de un mecanismo, no es adecuada para los mecanismos multinivel. Los modelos que no contienen ninguna representación de las relaciones constitutivas podrían representar las actividades de algún nivel fundamental hipotético. Sin embargo, tanto la existencia de dicho nivel fundamental como la aplicación de la causalidad a ese nivel son problemáticas.

La explicación de Craver, Glennan y Povich de la construcción de modelos de mecanismos conduce al resultado paradójico de que no hay niveles y, por tanto, no hay mecanismos multinivel, sino solo cadenas causales de actividades de nivel fundamental.

La clave para entender cómo se pueden construir modelos de mecanismos multinivel a partir de información empírica es que 1) los experimentos relevantes proporcionan directamente solo información sobre relaciones causales (contrariamente a lo que afirman Craver, 2007 y Baumgartner, Casini & Krickel, 2020), pero que (contrariamente a lo que afirman Craver, Glennan & Povich, 2021) esta información sobre relaciones causales puede afectar a variables de distintos niveles. Un modelo multinivel se construye en dos pasos. 1) En primer lugar, se construyen modelos parciales puramente causales para cada hipotética variable constitutiva Φ_i , sobre la base de experimentos descendentes y ascendentes que modifican o miden Φ_i de manera específica para cada nivel, 2) en segundo lugar, esos modelos parciales se fusionan en un modelo global que contiene tanto relaciones causales como de constitución entre variables, sobre la base de información sobre cada variable y restricciones espacio-temporales.

Palabras clave: Metafísica de la ciencia; Constitución; Relevancia constitutiva; Causalidad; Mecanismo; Intervención; Quirúrgica; Intervencionismo; Nivel; Nivel Superior; Multinivel; Fundamental; Eliminativismo; Nihilismo; Mereológico.

1. Introduction

One of the aims of the life sciences, and in particular of cognitive neuroscience, is to understand the mechanisms that govern the behavior of animals and humans. Some mechanisms, such as those that structure the behaviour of the roundworm nematode, *c. elegans*, are now well understood (White et al., 1986, Gray et al., 2005, Goodman & Sengupta, 2019). One of these mechanisms makes the worm back up when its head is touched. What does it take to explain how this mechanism works, in other words, to explain why the worm backs up when its head is touched? Let us take an organism S of species *c. elegans*, and call Ψ the activity that corresponds to the mechanism of backing up¹. Let us take for granted that, in order to understand and explain Ψ it is necessary and sufficient to discover how the activities of its constitutive parts are structured, i.e., in which causal and non-causal relations they stand to each other, and to the input Ψ_{in} and output Ψ_{out} of Ψ ². Much effort has been spent on trying to provide an analysis of the empirical criteria scientists use for judging that activity Φ_i of part X_i of organism S is *constitutive* of S 's mechanism Ψ . Craver (2007) has argued that the relation of constitution between the mechanism of Ψ -ing, characteristic of some system S , and the activities of Φ_i -ing of S 's parts X_i can be experimentally established by mutual manipulation (MM), where "manipulation" is understood according to Woodward's (2003) analysis of causation³.

¹ Following a widely used convention, upper case letters of the Latin alphabet, such as X , Y , W , are used for variables that represent types of properties or types of objects, whereas Greek letters, such as Ψ , represent types of processes or activities. Furthermore, I follow the convention of using S as a variable representing a whole mechanism, and $X1$, $X2$, etc., as variables representing specific parts of that mechanism.

² Here is one characterization of the concept of mechanism as it will be used in this article. "A mechanism [...] is a set of entities and activities organized such that they exhibit the phenomenon to be explained." (Craver, 2007, p. 5). I provisionally accept the metaphysical framework introduced by Machamer, Darden, and Craver (2000), which takes "entities" and "activities" to be the only fundamental ontological categories needed to analyze the structure of mechanisms. However, I will eventually suggest to complement this framework with the category of standing conditions, i.e., conditions that do not change.

³ I assume in what follows that it is equivalent to say " X 's Φ -ing is constitutive of S 's Ψ -ing" and to say " X 's Φ -ing is constitutively relevant to X 's Ψ -ing". Sometimes, "constitutive relevance" is used to designate an epistemic concept characterizing explanations, whereas "constitution" is used to designate a metaphysical relation (which exists if the explanation is true), but I will neglect this distinction here. Moreover, I follow Craver's usage of calling the activity of the mechanism under investigation interchangeably either " Ψ " or " Ψ -ing".

(MM) “ X ’s Φ -ing is constitutively relevant to S ’s Ψ -ing if the two [i.e., X and S ; M.K.] are related as part to whole and the relata are mutually manipulable. There should be some ideal intervention on Φ under which Ψ changes, and there should be some ideal intervention on Ψ under which Φ changes” (Craver 2007, p. 154).

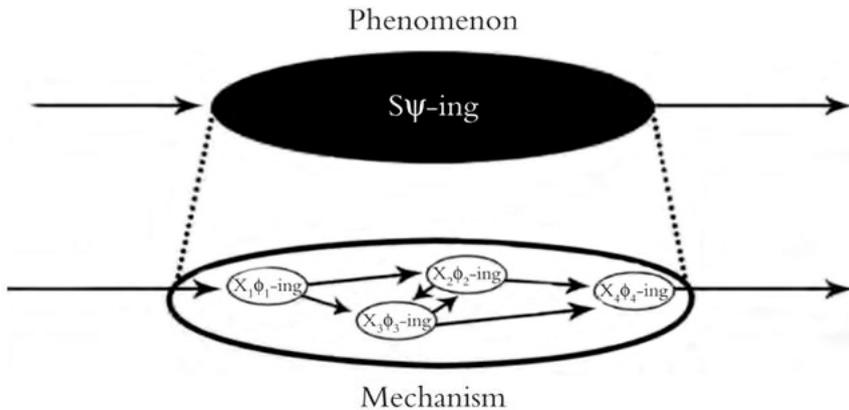


Fig. 1. Constitution, from Craver (2007, p. 7 and p. 121)

To use the variables of fig. 1, let S be a system, or object, such as a worm of species *c. elegans*, that has a mechanism of Ψ -ing, such as backing up. X_1, X_2 , etc. are parts of S . In order to establish that, e.g., the activity of Φ_3 -ing of part X_3 of S is *constitutive* of the Ψ -ing of S , it is necessary and sufficient to perform both the relevant top-down and bottom-up experiments (Craver, 2007, p. 144 ff.). In a bottom-up experiment, one intervenes on a part of the mechanism, either stimulating or inhibiting⁴ its activity, to observe modifications of the activity of the whole mechanism; in a top-down experiment, one intervenes on the whole mechanism by stimulating its activity and observes modifications in the activities of hypothetical constituents.

The existence of such experiments is not controversial, nor is their relevance for the explanation of mechanisms in general and for the discovery of constitution relations in particular, but their interpretation is. Romero (2015) and Baumgartner and Gebharder (2016)⁵ have shown that a

⁴ These two possibilities characterize two sorts of bottom-up experiments, which are, respectively, excitatory or inhibitory.

⁵ See also Baumgartner and Casini (2017), Baumgartner (2018).

top-down experiment cannot be interpreted as an experiment in which the experimenter *intervenes* on a variable at the level of the whole mechanism, such as S 's Ψ -ing, with respect to an activity of one of its constituent's parts, such as X_3 's Φ_3 -ing. Such an intervention could not possibly satisfy Woodward's (2003) requirements on interventions⁶. I will use the *mereological* notion of level, according to which complex objects belong to a higher level than their parts, and the properties and activities of a whole mechanism lie at a higher level than the properties and activities of its parts⁷.

Without going into details, the main problem is that experiments that manipulate higher-level variables are necessarily "fat-handed" in the sense that any variable I that causes Ψ must be a common cause of: 1) Ψ , 2) at least one variable at the level of Ψ 's constituents Φ_i , and 3) of at least one variable at each lower level: at the level of Φ_i 's constituents, at the level of the constituents of Φ_i 's constituents etc. This is because Ψ supervenes on its constituents Φ_i , so that there can be no modification of the value of Ψ without simultaneous modification of the value of at least one variable at each level below the level of Ψ . For this reason, even if the variable I represents an experimental modification of Ψ , it cannot be an intervention variable with respect to Ψ in Woodward's (2003) sense. The point is often expressed by saying that interventions must be "surgical", which means that an intervention variable I with respect to variable X must be such that modifying its value modifies only the value of X but allows keeping other variables fixed. However, if X is modified the variables in X 's supervenience base cannot be held fixed, and if Ψ represents the performance of a mechanism, Ψ 's constituents Φ_i belong to Ψ 's supervenience base.

Several strategies have been devised for finding an alternative conceptual analysis of top-down and bottom-up experiments. The task is to spell out a conceptually coherent experimental criterion for constitution, i.e., a sufficient condition for one variable Φ_i being a constituent of variable

⁶ In Baumgartner and Gebharter's words, "there cannot possibly exist (M)-(IV)-defined intervention variables for macro variables w.r.t. [with respect to] their micro supervenience bases such that the latter could be changed by intervening on the former" (Baumgartner & Gebharter, 2016, p. 743) (where (M) and (IV) refer to Woodward's (2003, p. 59) definitions of direct cause and of an intervention variable). Baumgartner and Gebharter (2016) only question the conceivability of top-down interventions but not the possibility of bottom-up interventions, within the framework of Woodward's (2003) conception of interventions. "Bottom-up interventions on a mechanism may well exist" (Baumgartner & Gebharter, 2016, p. 743, note 8).

⁷ On levels, see Craver (2007, chap. 5; 2015), Eronen (2013, 2015, 2021), Kistler (2021). Craver (2007) calls *mereological* levels "levels of composition" (Craver, 2007, p. 184), in opposition to levels corresponding to scientific reductions, which he calls "levels of science" (Craver, 2007, p. 172).

Ψ , where all variables are both observable and at least in principle experimentally manipulable.

2. Attempts of Analyzing Constitution in Terms of Modified Notions of Interventions

2.1. Horizontally surgical interventions

Woodward (2015) has introduced a new notion of intervention appropriate for research on systems that can only be appropriately modeled with variables at different levels⁸. This notion is defined by conditions that relax Woodward's (2003) original constraints on interventions, replacing conditions (I.3) and (I.4) by (I.3*) and (I.4*).

Woodward (2003) defines intervention variables I by four conditions⁹.

I1. I causes X .

I2. I acts as a switch for all the other variables that cause X . That is, certain values of I are such that when I attains those values, X ceases to depend on the values of other variables that cause X and instead depends only on the value taken by I .

I3. Any directed path from I to Y goes through X . That is, I does not directly cause Y and is not a cause of any causes of Y that are distinct from X except, of course, for those causes of Y , if any, that are built into the $I - X - Y$ connection itself; that is, except for (a) any causes of Y that are effects of X (i.e., variables that are causally between X and Y) and (b) any causes of Y that are between I and X and have no effect on Y independently of X .

I4. I is (statistically) independent of any variable Z that causes Y and that is on a directed path that does not go through X . (Woodward, 2003, p. 98)

If clause (I3) of the original definition (IV) of an intervention variable is modified in the following way (Woodward, 2015, pp. 333-334)¹⁰,

⁸ For similar proposals, see Shapiro and Sober (2007), Shapiro (2010), Raatikainen (2010).

⁹ In the next few paragraphs, where I report Woodward's definitions of intervention variables, and in particular clauses (I3), (I4), (I3*) and (I4*), I stick to the variables I , X , Y , Z , which are used in Woodward's (2003) and (2015).

¹⁰ "...when (non-causal) supervenience relationships are present, the characterization (IV) should be interpreted in such a way that in condition (I3) a directed path counts as 'going from I to Y through X ' even if I also changes (as it must) the supervenience base SB(X) of X , as well as the value of X . Similarly, the reference in (I4) to 'any variable Z '

interventions on variables X w.r.t. (i.e., with respect to) Y are possible even in models that also contain variables $SB(X)$ on which X supervenes. (I3*) differs from (I3) by the expressions in italics.

(I3*). Any directed path from I to Y goes through X (*even if I also changes (as it must) the supervenience basis $SB(X)$ of X , as well as the value of X*). That is, I does not directly cause Y and is not a cause of any causes of Y that are distinct from X except, of course, for those causes of Y , if any, that are built into the I - X - Y connection itself *and except the supervenience basis $SB(X)$ of X* ; that is, except for (a) any causes of Y that are effects of X (i.e., variables that are causally between X and Y) and (b) any causes of Y that are between I and X and have no effect on Y independently of X ...

Similarly, (I4*) differs from (I4) by the expression in italics.

(I4*). I is (statistically) independent of any variable Z (*different from the supervenience basis $SB(X)$ of X*) that causes Y and that is on a directed path that does not go through X .

According to (I.3*) and (I.4*), it is not required that the variables in the supervenience base of X be held fixed in order to intervene on X w.r.t. to Y . Applied to the analysis of mechanisms, there can be interventions on S 's Ψ -ing in the sense of "intervention" specified by (I.3*) and (I.4*), even if the model also contains variables representing parts and their activities at lower levels, i.e., at the level of constituents of the mechanism of S 's Ψ -ing, and at the level of the constituents of these constituents etc. A manipulation of Ψ can count as an intervention in this sense even though it is "fat-handed", i.e., non-surgical. Such a fat-handed intervention causes both, and simultaneously, a modification at the level of the whole mechanism, i.e., a modification of S 's Ψ -ing (more precisely it causes a modification of the input condition Ψ_{in} of S 's Ψ -ing), and a modification, at the next lower level in the mechanistic hierarchy, of the activity Φ_i^k of some constituent X_k of Ψ_{in} ¹¹.

It remains controversial whether this relaxed notion of intervention is appropriate for accounting for higher-level causation, and in particular for top-down causation. One problem is that Woodward's (2015) account seems to have the consequence of mistakenly attributing causal efficacy to variables representing activities of whole mechanisms, which is really due to activities of their constituents. In Blanchard's terms, Woodward's account results in "over-inclusion: it grants to composites causal abilities that belong to their parts only" (Blanchard, 2023, p. 2685). We need not

should be interpreted as 'any variable Z other than those in the supervenience base $SB(X)$ of X ' (Woodward, 2015, pp. 333-334).

¹¹ For the rest of this section, I adopt Baumgartner, Casini and Krickel's (2020) convention, in which Φ_i^k represents the i -th activity of part X_k of system S .

enter the debate on this issue here. The question relevant for our present enquiry is whether such a fat-handed intervention on both Ψ and Φ_i^k can reveal (or provide empirical evidence for) the fact that one of its effects, Φ_i^k , is constitutive of its other effect, Ψ .

According to Baumgartner, Casini and Krickel (2020) (henceforth, “BCK”), there is a special sort of fat-handed intervention that provides a sufficient condition for constitution: “horizontally surgical” interventions (HSI). Such interventions are modeled by variables satisfying the following condition (H).

- (H) $I_{\Phi_i^k}$ is a horizontally surgical intervention variable on a part¹² Φ_i^k of Ψ w.r.t. Ψ iff:
- (i) $I_{\Phi_i^k}$ is a cause of Φ_i^k ;
 - (ii) if $I_{\Phi_i^k}$ causes changes in both Φ_i^k and Ψ , these changes occur simultaneously;
 - (iii) $I_{\Phi_i^k}$ is a direct cause of at most one behavior on every lower level than Ψ 's. (BCK, p. 425-426)

According to BCK, if Φ_i^k is a spatiotemporal part of Ψ , then the existence of a horizontally surgical intervention (HSI) on Φ_i^k w.r.t. to Ψ is a sufficient condition for Φ_i^k to be a constituent of Ψ . In other words, (hFAT) provides a sufficient condition for constitution.

- (hFAT) Φ_i^k constitutes Ψ iff
- the instances of Φ_i^k are spatiotemporal parts of instances of Ψ ;
 - there exists a (possible) horizontally surgical intervention $I_{\Phi_i^k} = i$ on Φ_i^k w.r.t. Ψ that causes changes in both Φ_i^k and Ψ (BCK, p. 428).

BCK emphasize that this thesis “is not subject to MM’s conceptual flaws” (BCK, p. 429), where “MM” designates Craver’s (2007) condition of mutual manipulability quoted above. It is not the case, as it is for MM, that the existence of instances of the definiens of (H) “is excluded on a priori grounds” (BCK, p. 429). Recall that given Woodward’s (2003) original definition of interventions, there cannot be interventions on Ψ w.r.t. to Φ_i^k , which would correspond to the “top-down” part of Craver’s (2007) mutual manipulability¹³.

¹² In BCK’s (pp. 420-421) terminology, Φ_i^k can represent both an activity that is a constituent of mechanism Ψ and a part of system S , whereas I distinguish between the part X_k and its activity Φ_i^k . This should not create any confusion.

¹³ “[T]here cannot possibly exist an intervention variable as defined by M and IV for

Interventions that satisfy (H) are conceptually possible. However, BCK's condition (hFAT) for constitution suffers from the following problems. In mechanisms with more than one constituent, and in particular in biological systems, condition (iii) of the definition (H) of horizontally surgical interventions is in general not satisfied. Before I argue for this thesis, note that BCK's model uses variables that are not explicitly specific for time, i.e., time-indexed (Spohn, 2006). This raises no problem if, at a given level, a mechanism is purely serial, i.e., if it contains no parallel branches. In that case, BCK's model can represent the fact that the activities corresponding to the different constituents at that level are ordered in time. In the convention I have adopted following BCK, the upper index k represents the entity, whereas the lower index represents the activity, so that Φ_i^k represents entity k 's activity i , where indices k and i are ordered as the natural numbers. If there are at level Φ no parallel (and in particular, no redundant) activities, and if every relevant entity performs only one activity, we may adopt the numbering convention $i=k$, so that the activities that are constituents of the mechanism appear as ordered in a simple causal chain, in which Φ_1^1 is the first link, Φ_2^2 the second, etc. In that case, we may simplify the representation and drop one of the indices, and represent the chain of activities of the parts Φ of mechanism S , by Φ_1, Φ_2, \dots . If some entities Φ_k perform more than one activity, but if these activities are still ordered in a unique causal chain, the numbering with a unique index corresponds to a loss of information about which entity performs activity i . This will make no difference to my argument. However, it is crucial for my argument that there are in general several parts some of whose activities occur in parallel, i.e., occur at the same time. If such activities are ordered in time, so that the lower index runs from 1 to n , there are in this case, for some temporal step i , two or more parts $\Phi_i^m, \Phi_i^n, \dots$, with $n \neq m$, whose activities occur in parallel. We will get to the problem this raises in a moment.

Let Ψ represent the activity of the whole macroscopic mechanism that is the target for the research on its constituents. Ψ is extended over time. An intervention on Ψ must modify its initial temporal stage, t_1 . Let Φ_1 represent the activity of a part of mechanism Ψ that is constitutive of Ψ and occurs in this initial phase t_1 . If Ψ represents *c. elegans'* mechanism of backing up, Φ_1 may represent the activity of registering a touch on the head by a sensory neuron, such as the ALM neuron (Goodman & Sengupta, 2019).

any macro-variable Ψ with respect to any micro-variable Φ_i , such that Φ_i changes when Ψ is wiggled" (Baumgartner & Gebharder, 2016, p. 743).

BCK's second and third conditions on a HSI are:

- (ii) if $I_{\Phi_i^k}$ causes changes in both Φ_i^k and Ψ , these changes occur simultaneously
- (iii) $I_{\Phi_i^k}$ is a direct cause of at most one behavior on every lower level than Ψ 's (BCK, p. 426).

Here are two obstacles to satisfying (iii).

1) Let us suppose that the activity Ψ of the whole mechanism is realized, at the first level Φ below the level of Ψ , by a series of n activities, $\Phi_1, \Phi_2, \dots, \Phi_n$. They are all simultaneous with Ψ in the sense that their times of occurrence overlap the duration of Ψ .

$I_{\Phi_i^k}$ can fulfil the requirement (iii) of being a *direct* cause only for the *first* link in the series Φ_i . If $I_{\Phi_i^k}$ causes Φ_1 directly, it can cause all other Φ_i , for $i \neq 1$, only indirectly, by first directly causing Φ_1 . Thus, hFAT is inadequate for all constituents, Φ_2, \dots, Φ_n , whose activity occurs later than that of the very first, Φ_1 . This also holds at all lower levels, i.e., at the level of Φ_1 's constituents, etc.¹⁴ (hFAT) is inadequate for mechanisms that contain, at one level at least, more than one activity.

The existence of mechanisms in which the series of activities at level Φ , $\Phi_1, \Phi_2, \dots, \Phi_n$, contains only one member (i.e., the degenerate case where the series consists only of Φ_1) is conceptually possible but biologically unrealistic. In *c. elegans*, the mechanism for touch perception is constituted by a series of neurons, through which the signal for touch detection propagates: CEP, ASH, ALM, PDE, PVD, PLM.

2) There is a second problem, which would arise even if (H) were employed as a criterion for the first constituent, Φ_1 , in a temporally extended chain of activities, or for analyzing the improbable mechanism in which the time span of the first activity at level Φ coincides with the time span of the activity of the whole mechanism Ψ .

Even for the first activity, at any level, the conditions for constitution in hFAT can only be satisfied if that first activity is not accompanied by parallel activities, which may or not be redundant. At the initial temporal part of Ψ , there will in general be two or more activities $\Phi_i^m, \Phi_i^n, \dots$, with

¹⁴ Krickel makes a similar point against earlier versions of accounts of constitution in terms of "fat-handed" interventions. "Fat-handedness approaches represent the phenomenon by only one variable. Thereby, they cannot account for the fact that mechanisms are temporally extended and involve changes over time." (Krickel, 2018, p. 66). Krickel (2018, p. 66) suggests that the condition hFAT for constitution in BCK is only appropriate for the constituents of the very first instant at which the mechanism begins to operate.

$n \neq m$, at the level of parts X_m and X_n of S , which evolve in parallel. The possibility of a mechanism without activities (at the level of parts) running in parallel (some of which could be redundant) is conceptually possible but biologically unrealistic. In *c. elegans*' mechanism of retracting, the relevant parts are neurons; their relevant activities are firings. In this mechanism, as is typical for biological mechanisms, several constitutive activities run in parallel: Several neurons detecting harsh touch react in parallel: ASH, IL1, IL2, OLQ, ADE, CEP, and others (Goodman & Sengupta, 2019, p. 29).

Both problems 1) and 2) arise again at each level below the level of Φ . Let χ_i^k be variables at the next lower level below Φ , i.e., variables that represent activities of molecular parts of individual neurons.

Let us use hFAT at this level, so that it indicates the conditions for activity χ_i^k being a constituent of activity Φ , where Φ is one particular constituent of Ψ , and where activity χ_i^k represents the i -th activity of the k -th relevant part of X , which performs Φ . Recall that, applied to the issue of whether activity χ_i^k is a constituent of activity Φ , the second and third conditions on a HSI are:

- (ii) if $I_{\chi_i^k}$ causes changes in both χ_i^k and Φ , these changes occur simultaneously.
- (iii) $I_{\chi_i^k}$ is a direct cause of at most one behavior on every lower level than Φ 's.

Let us suppose, for the sake of the argument, that there is an organism with a mechanism for touch sensation that is even simpler than that of *c. elegans*, in that touch is detected by exactly one sensory neuron. No parallel activities, no chain of activities at the level of the activities Φ of neurons. For such a biologically unrealistic, maximally simple mechanism, problem 2) arises at the level of the constituents of that unique sensory neuron. Mechanisms at the level of neurons are typically mediated by proteins inserted in the neuron's membrane, some of them functioning as ionic channels that change their configuration in function of various input conditions. One type of channel that is involved in the neurons registering touch in *c. elegans* is the egl-19 voltage gated calcium channel (Suzuki et al., 2003, Goodman & Sengupta, 2019, p. 31). Let us represent the activity that corresponds to a change in conformation of one such channel by χ_i^k . Problem 2) arises at this point because in each neuron, there are many individual channel molecules of any particular type whose changes in configuration occur in parallel. If the change χ_i^k in conformation of an individual channel protein molecule is a constituent of the mechanism Φ , by which a given individual sensory neuron detects touch (Φ itself being the first constituent

of mechanism Ψ), then (iii) is not satisfied in real neurons because the relevant intervention variable does not “change at most one behaviour” at the level of ionic channels.

There is a third reason for which hFAT is not practically adequate, i.e., does not provide an analysis of experiments that are used in neuroscience in order to discover constituents of mechanisms. As I have noted above, the variables in hFAT are not time-indexed. The fact that both Φ_i^k and Ψ are directly caused by $I_{\Phi_i^k}$, suggests that Φ_i^k and Ψ are intended to represent simultaneous activities, immediately following their common cause, i.e., the change in the value of $I_{\Phi_i^k}$ ¹⁵, although the fact that Φ_i^k and Ψ are simultaneous is not explicitly represented in BCK’s analysis. However, real top-down and bottom-up experiments are causal, in the sense that the experimenter first modifies one of the two variables Φ_i^k and Ψ at time t_1 , then measures or observes the value of a second variable a little later, at t_2 , and sometimes, as we will see in a moment, a further variable at t_3 . Such experiments, both bottom-up and top-down, track causal influence, rather than the simultaneous occurrence of two effects of an intervention acting as their common cause.

The importance of this fact and of the corresponding time delay between the cause and effect variables, will become clearer in the light of a second proposal of analyzing the experimental justification of constitution relations in mechanisms.

2.2. Identification of constitutive relevance with “causal betweenness”

Harinen (2018), Prychitko (2021) and Craver, Glennan and Povich (2021) (henceforth, CGP) have elaborated another proposal intended to model the experimental justification of constitution relations. As we have seen, “mutual manipulability” of the activity of a whole mechanism and of the simultaneous activity of one of its constituents is impossible if it is construed with Woodward’s (2003) concept of intervention. CGP have elaborated an analysis of bottom-up and top-down experiments in neuroscience that does not require mutual manipulability; it does not require any problematic interventions on an activity of a whole w.r.t. a simultaneous activity of one of its parts, or vice versa, as in Craver’s (2007) analysis, and it does not require detection of simultaneous effects of common causes, as in BCK’s analysis. According to CGP’s analysis, bottom-

¹⁵ The times of occurrence of all constituent activities Φ_i^k overlap the duration of the activity Ψ of the whole mechanism S of which Φ_i^k are constituents, but as I have argued above, the conditions of hFAT could only be satisfied by activities in the very first temporal segment of the process of mechanism Ψ .

up experiments do not involve interventions on variables representing the activity of a part with respect to a variable representing the simultaneous activity of the whole mechanism, and top-down experiments do not involve interventions on variables representing the activity of the whole mechanism with respect to variables representing the simultaneous activity one of its parts.

According to CGP, such experiments contribute to the construction of causal models that do not directly contain any representation of constitution but nevertheless license inferences to constitution relations. In the relevant models, the constitutive activities Φ_i of the relevant parts X_i of mechanism S 's activity of Ψ -ing are represented as causal intermediaries between the variable Ψ_{in} , representing the input to mechanism Ψ , and the variable Ψ_{out} , representing the output of mechanism Ψ .

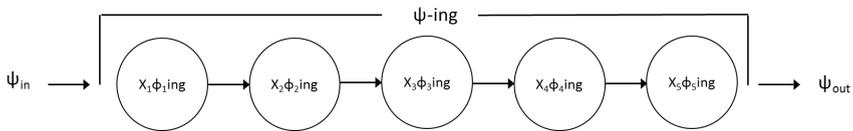


Fig. 2 from (CGP, p. 10), similar to fig. 1 in Prychitko (2021, p. 2)

CGP argue that interlevel experiments, such as bottom-up stimulation, bottom-up inhibition and top-down stimulation, are used to construct models that have the structure sketched in fig. 2. More precisely, CGP propose the following “matched interlevel experiments condition” (MIE, CGP, p. 8822) for being a constituent of a mechanism. (MIE) contains four conditions concerning the relevant interlevel experiments. These four conditions are “jointly sufficient” to establish that “an entity X and its activity Φ are constitutively relevant to a mechanism that Ψ s” (CGP, p. 8822).

(Bottom-up inhibiting experiments) “(CR1i) If an experiment initiates conditions Ψ_{in} while a bottom-up intervention, I , prevents or inhibits X 's Φ -ing, alterations to or preventions of Ψ 's terminal conditions, Ψ_{out} , are detected.”

(Bottom-up excitatory experiments) “(CR1e) If a bottom-up intervention, I , stimulates X 's Φ -ing, Ψ 's terminal conditions, Ψ_{out} are detected.”

(Top-down experiments) “(CR2*) If a top-down experiment initiates conditions Ψ_{in} and detects Ψ 's terminal conditions, Ψ_{out} , X 's Φ -ing is also detected.”

“(Matching) The activities Φ_i activated or inhibited in bottom-up experiments (CR1i and CR1e) must be of the same kind as, and occur

within quantitatively overlapping ranges with, the activities Φ_i detected in top-down experiments (CR2*)." (CGP, p. 8822)

The model built on the basis of MIE contains the input (Ψ_{in}) and output (Ψ_{out}) of the whole mechanism, as well as the constituents X_i 's Φ_i -ing, together with the causal relations linking them. However, the model 1) neither contains any variable representing the activity of the mechanism as a whole nor 2) any representation of (non-causal) constitution relations¹⁶. In other words, the model is flat and purely causal. "Flat" means that it represents only one level, and "purely causal" means that it does not contain relations other than causal relations, and in particular no non-causal constitution relations.

CGP claim that such a model contains information about constitution relations without explicitly representing them. The model explicitly represents only causal but no constitutive relations, which are non-causal. One way of expressing this idea is to say that their model accomplishes a *reduction* of non-causal constitution relations in terms of purely causal structure, where "reduction" means "conceptual analysis": claims about constitution can be analyzed in a purely causal vocabulary.

According to CGP, a mechanism Ψ is just a single causal chain that is intermediate between Ψ_{in} and Ψ_{out} . The meaning of the expression "is a constituent of the mechanism Ψ -ing" is shown by analysis to be the same as the meaning of the expression "is a node in the causal chain connecting Ψ_{in} to Ψ_{out} ". In their words, "constitutive relevance is causal betweenness" (CGP, p. 8807)¹⁷, and "constitutive relevance ... amounts to a kind of causal mediation" (CGP, p. 8821).

CGP sometimes express this idea in terms of truthmaking. A purely causal model that consists of a single causal chain without any explicit representation of constitution nevertheless contains information about constitution because the causal structure represented in the model is the "truth-maker" of the non-causal relation of constitution. "The constitutive relevance relation MIE detects is a three-place relation of causal betweenness, the ontological truthmaker for claims of constitutive relevance" (CGP, p. 8825).

¹⁶ Craver's (2007) model of a mechanism sketched in fig. 1 contains both: Ψ -ing is a variable within the model that represents an activity of the whole mechanism, and the dotted lines represent the synchronous non-causal constitution relation between the activities of the parts and the activity of the whole.

¹⁷ Similarly, Prychitko writes: "Demonstrating that a part of S is caused by S_{in} and causes S_{out} establishes that it is a component of S 's Ψ -ing. [...] Showing that something lies on this causal chain would establish that the part is constitutively relevant to S 's Ψ -ing" (2021, p. 1838).

Let me say a word about the assumption, which I will take for granted in what follows, that parts and wholes cannot be causally related, and that therefore, the activities of the constituents of mechanisms cannot be causes of the activity of the mechanism as a whole.

This thesis, which is accepted by many participants in the debate on the applicability of interventionism on the discovery of constitution, which is a relation between (activities of) parts and (activities of) wholes, has recently been challenged, in two ways.

1. A first proposal consists in construing constitution to be a diachronic relation (Leuridan & Lodewyckx, 2020, Kiverstein & Kirchhoff, 2021). If constitution is diachronic, instead of being synchronic, constitution relations can be explored by the methods appropriate for causal relations. For lack of space, I will not examine this proposal here.

2. An even more direct strategy for making constitution directly accessible to methods of discovering causes, which doesn't call into question constitution's being synchronic, consists in taking constitution to be a form of causation (Leuridan, 2012, Wilson, 2018, Friend, 2019, Kiverstein & Kirchhoff, 2021).

It is not trivial that constituents cannot be causes of their wholes. That they cannot might seem to be a consequence of the thesis that cause and effect must be distinct, in the sense that they must not overlap, i.e., must not share any part. For *C* to be a cause of *E*, "*C* and *E* must be distinct events — and distinct not only in the sense of nonidentity but also in the sense of nonoverlap and nonimplication. It won't do to say that my speaking this sentence causes my speaking this sentence; or that my speaking the whole of it causes my speaking the first half of it; or that my speaking causes my speaking it loudly, or vice versa." (Lewis, 2004, p. 78)¹⁸ However, it is not obvious that parts of complex objects are parts of those objects *in the mereological sense*. In fact, I think they are not¹⁹: The atoms contained in a molecule are not mereological parts of those molecules. This can be seen from the fact that the existence of the atoms, by themselves, is not sufficient to guarantee the existence of a molecule, whereas it is sufficient for the existence of a *mereological* whole. Without a bond, the atoms don't constitute any molecule, although they do constitute a mereological whole.

¹⁸ Lewis (2004) says that there cannot be causal relations between terms that are related as parts and whole in the mereological sense. Like many others, CGP take it that it follows directly that there cannot be causal relations between part and whole in the mechanistic sense. "The Φ -ings that are constitutively relevant to a mechanism are parts of, and hence at a lower level than, the Ψ -ing they constitute, and thus cannot cause or be caused by the Ψ -ing" (CGP, p. 8813).

¹⁹ "The mereological composition of particulars is a myth" (Mellor, 2012, p. 402).

In other words, the notions of part and whole that are relevant for the analysis of mechanisms are *not* the mereological notions of part and whole used in Lewis' criterion of distinctness. Therefore, the acceptance of Lewis' requirement that causes and effects must be distinct or non-overlapping does not entail that parts of a mechanism cannot be causes of their mechanism.

I will not try to adjudicate here whether "parts can cause their wholes" (Friend, 2019) or whether there can be synchronic causation. Fortunately, we can leave these issues open. If there is synchronic causation, in addition to diachronic causation, the following discussion can be rephrased in terms of two sorts of causation: diachronic and synchronic²⁰. In that vocabulary, the question how knowledge about constitution can be extracted from knowledge about causal relations, would become the question how knowledge about synchronic causation can be extracted from knowledge about diachronic causation.

For our present purposes, I will stick to the traditional terminology, in which constitution is a non-causal synchronic dependence relation, whereas causation is always diachronic.

The fact that, as we have seen above, interventionism cannot be used in a straightforward way to construct models that contain constitution relations, is simply a consequence of the fact that constitution relations are not causal. This fact entails that bottom-up experiments cannot be interpreted in terms of Woodwardian (2003) interventions on a part with respect to a corresponding whole, and that top-down experiments cannot be interpreted in terms of such interventions on the whole with respect to one of its parts. According to CGP, none of these two sorts of intervention is needed. What bottom-up and top-down experiments establish is that a given component of S 's Ψ -ing, X_i 's Φ_i -ing, is causally intermediate between Ψ_{in} , the input condition of the mechanism, and Ψ_{out} , its output condition.

However, CGP's solution is only partial. One argument for this claim is that the mechanism itself is not explicitly represented in their model. True, figure 2 contains the expression " Ψ -ing", which represents the mechanism's activity of Ψ -ing. However, " Ψ -ing" is not a variable in the model, i.e., it is not the term of any relation explicitly represented in the model, causal or not.

Here are two interpretations of the meaning of the expression " Ψ -ing" in fig. 2, none of them satisfactory.

1. The whole mechanism and its activity of Ψ -ing appear in figure 2,

²⁰ In Karen Bennett's (2017) terminology, both causation and constitution, or both "diachronic and synchronic causation" are species of a more general relation of "building".

but not as part of the model. The symbol “ Ψ -ing” in fig. 2 names *part of the model*, rather than *the system represented* by the model. In other words, “ Ψ -ing” is a second-order variable, rather than a first-order variable such as Φ_i . Therefore, the model does not represent any relation between Φ_i and Ψ . Experiments do not justify the existence of constitution relations, in addition to causal relations, and thus do not contribute to construct models of constitution relations. Constitution is a relation between representations, not between activities. It characterizes relations between parts of models, not relations between entities represented by such models.

2. The mechanism is defined by its input and output conditions, as whatever mediates between input and output. The causal chain between Ψ_{in} and Ψ_{out} is identical with the mechanism, and each variable in the chain is a constituent of the mechanism.

According to the first interpretation, constitution is not a relation between activities, which is tantamount to the eliminativist thesis that there are no constitution relations. The second interpretation is incomplete, for at least two reasons, only one of which is acknowledged by CGP.

1. One limitation (acknowledged by CGP) of this interpretation of CGP’s conception of what it means to establish that some activity is a constituent of a whole mechanism is that it cannot account for any structures that are more complex than a simple chain. In many if not most mechanisms there are, at every level, parallel and often redundant activities. I have already mentioned parallel neural pathways in *c. elegans*’ mechanism of retracting²¹, and parallel changes in the configuration of ionic channels in the mechanism of the activation of any individual neuron. In what follows I will concentrate on the following second point.
2. Mechanisms are typically structured into many levels. To account for two levels, in addition to the level of the whole mechanism Ψ , it is necessary to account for constitution relations among constituent parts (of Ψ) lying at different levels. However, this doesn’t seem to be possible in a model containing only causal relations within one single level.

Let me explain. Take a mechanism that has more than two levels. Level n corresponds to system S ’s mechanism of Ψ -ing. Level $n-1$ contains activities Φ_i of S ’s parts X_i that are constitutive of S ’s Ψ -ing; level $n-2$ contains activities χ_j -ing of X_i ’s parts Z_j that are constitutive of X_i ’s Φ_i -ing,

²¹ Goodman and Sengupta (2019).

and so on, until eventually a fundamental level is reached that contains elementary activities that are not mechanistically decomposable. According to CGP's analysis, the n-level mechanism

$$\Psi_{\text{in}} - \Psi - \Psi_{\text{out}}$$

is equivalent to a chain of activities at level n-1:

$$\Psi_{\text{in}} - X_1\text{'s } \Phi_1\text{-ing} - X_2\text{'s } \Phi_2\text{-ing} - X_3\text{'s } \Phi_3\text{-ing} - \dots - \Psi_{\text{out}}$$

In general, each constituent Φ_i of mechanism Ψ is itself a mechanism with its own input and output conditions. Let us spell out the mechanistic analysis of the first two activities at level n-1, Φ_1 and Φ_2 , which shows them to be equivalent to chains of activities at level n-2.

$$\Phi_{1\text{in}} - X_1\text{'s } \Phi_1\text{-ing} - \Phi_{1\text{out}} \text{ is equivalent to}$$

$\Phi_{1\text{in}} - Z_1\text{'s } \chi_1\text{-ing} - Z_2\text{'s } \chi_2\text{-ing} - \dots - \Phi_{1\text{out}}$, where Z_i are parts of X_1 , and χ_i their activities;

$$X_2\text{'s } \Phi_2\text{-ing is equivalent to}$$

$\Phi_{2\text{in}} - W_1\text{'s } \theta_1\text{-ing} - W_2\text{'s } \theta_2\text{-ing} - \dots - \Phi_{2\text{out}}$, where W_i are parts of X_2 , and θ_i their activities.

Putting together the analyses of all the mechanisms X_i 's Φ_i -ing at level n-1, the mechanism of S 's Ψ -ing, at level n, can be analyzed as a flat (or "one-level") causal chain of activities at level n-2:

S 's Ψ -ing is equivalent to

$$\Psi_{\text{in}} - Z_1\text{'s } \chi_1\text{-ing} - Z_2\text{'s } \chi_2\text{-ing} - \dots - W_1\text{'s } \theta_1\text{-ing} - W_2\text{'s } \theta_2\text{-ing} - \dots - \Psi_{\text{out}}$$

If the activities χ_i at level n-2 are not elementary, they can themselves be analyzed in an analogous manner, until eventually a bottom level 0 is reached, which contains variables representing simple activities that cannot be analyzed any more in terms of mechanisms.

The result of the complete analysis of a multi-level mechanism is a unique chain linking the input Ψ_{in} to the output Ψ_{out} of the whole mechanism, through a chain of variables representing fundamental activities belonging to a mechanistically fundamental level.

If this analysis were correct it would show that there really are no mechanisms. Once the analysis is achieved, the mechanism is shown to consist of a flat series of activities all belonging to a single fundamental level. One may doubt whether it makes sense to speak of a level in an analysis that contains only one level, but let us call that unique level, "level 0".

The only way to recover the mechanistic levels that seemed to exist before that analysis has been achieved, seems to correspond to our first interpretation above, in the following sense.

One may add a series of tags to the chain at level 0, beginning with tags corresponding to a (nominal) level 1. In figure 3, Z_1 's χ_1 -ing, ..., W_j 's

²² $\Phi_{1\text{in}}$ is just $\Psi_{1\text{in}}$, $\Phi_{1\text{out}}$ is the input condition $\Phi_{2\text{out}}$ of the following activity Φ_2 .

θ_j -ing, ... represent fundamental variables at level 0. Each level-1 tag, illustrated by tags Φ_1 and Φ_2 in fig. 3, marks a part of the fundamental chain (belonging to level 0) that corresponds to an activity of level 1; then one may add a series of level-2-tags, each of which marks a series of level-1 tags that together correspond to an activity at level 2, and so on, until one reaches the tag Ψ that groups together the whole chain.

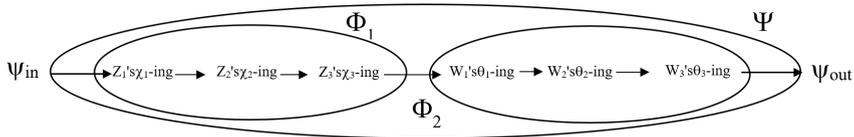


Fig. 3. The chain of fundamental activities, represented by variables Z_i 's χ_i -ing, ..., W_j 's θ_j -ing, ..., is conceptually structured at two levels. Small ellipses, tagged Φ_1 and Φ_2 , represent different chunks of that chain; the large ellipse, tagged Ψ , represents the whole chain.

Such tags are, as in interpretation 1 above, comments on the mechanistic model, which do not represent parts of the mechanism, but parts of the (flat) model of the mechanism. The tags, and the levels to which they correspond, are the result of a conceptual operation of abstraction, or, to use Norton's (2003) term, "chunking".

If these tags are just comments on the model, not parts of the model, adding them or not doesn't correspond to any empirical difference. There is only a difference in representation, but no difference in empirical content, between

1) a causal chain without any mechanism, and without any constitution,

2) a mechanism with two levels, as in fig. 1, where a mechanism corresponds to (is identical with) the whole chain (which includes all the variables between the input and the output variables), and

3) mechanisms with more than one level, such as the LMTP mechanism²³, where different parts (or "chunks") of the bottom-level chain correspond to nodes in chains of higher-level variables.

According to this interpretation of GCP's model, claims concerning constituents above the fundamental level have no empirical meaning. The multi-layer structure of a complex mechanism is not empirically but only conceptually constrained. The only reason for introducing, by convention,

²³ Craver (2002, 2007).

tags to represent chunks of the causal chain at level 0 is cognitive: a multi-layered structure may be cognitively easier to process. The choice of introducing a tag, and a layer corresponding to several tags, is merely a matter of representation. Different representations are possible, which are all empirically equivalent. They are compared and assessed, not by empirical criteria, but according to their usefulness given the context of explanation in which the model is to be used. Given that tags are not variables, they are not part of the model and they do not represent real activities that play causal roles, and in particular they cannot be intervened on and measured²⁴.

CGP's analysis leads to a nominalist conception of mechanisms. All real causation happens at the bottom level, i.e., only activities at the fundamental level are real and causally efficacious. The whole mechanism and the activities constituting it at all levels except the fundamental level, are just names or "tags", chosen by convention to make the representation of the fundamental-level causal chain cognitively easier to handle.

This nominalist conception of mechanisms is eliminativist for activities of all levels except the fundamental. Such mechanistic nihilism is ontologically parsimonious but it accounts neither for scientific practice nor for scientific models. Mechanistic nihilism is analogous to compositional nihilism, which takes only fundamental components to exist, but no

²⁴ At best, relations between such higher-level tags might correspond to what Kim (1984) calls "epiphenomenal causation". According to Kim, "modern theoretical science treats macrocausation as reducible epiphenomenal causation" (1984, p. 96). The "epiphenomenal causation" between an instance of macroscopic property F and an instance of macroscopic property G can be reduced to the fact that there are microscopic properties $m(F)$ and $m(G)$, which are the respective supervenience bases of F and G , such that " $m(F)$ and $m(G)$ are appropriately causally connected" (p. 97). Epiphenomenal causation at the macroscopic level is only apparent but it is always grounded on microscopic causation that is real and not only apparent. Kim borrows Salmon's example of "two successive spots of light on the wall" (p. 93). The pseudo-process of the succession of spots on a wall mimics a causal process and is "apt to be mistaken for such" (p. 93) although there is no "process involving a real causal chain" (p. 93) between the spots of light. Salmon (1984, pp. 141-142) illustrates his concept of a pseudo-process with a spotlight in uniform rotation placed in the center of a hollow cylinder, projecting a light ray directed towards the lateral wall of the cylinder. The series of events that consists in the spot of light appearing on the wall of the cylinder and apparently moving across that wall is Salmon's paradigmatic example of a pseudo-process. It satisfies the two conditions defining that concept: first, the series of events is not a causal process, in the sense that the events constituting it are not related among each other as causes and effects, and second, it is a world line exhibiting structural uniformity so that it has the illusory appearance of a causal process. (See Kistler 2006, pp. 59-61). However, all epiphenomenal causal relations between macroproperties F and G "are reducible to more fundamental causal relations" (p. 94) between the microproperties $m(F)$ and $m(G)$ that are supervenience bases of F and G .

composite entity²⁵. Apart from the paradoxical result that there are no mechanisms, no levels, and no activities except at the fundamental level, mechanistic nihilism raises two major problems:

1. There may be no bottom level (Schaffer, 2003)
2. If the bottom level is (present-day) fundamental physics, there are reasons to doubt that the concept of causation can be used there²⁶.

We have reached the surprising result that the mechanistic philosophers Craver, Glennan and Povich defend a view that entails that there really are no mechanisms, but only causal chains of fundamental activities, whereas mechanisms are just names.

Even for a mechanism that has only one level below the whole, and whose parts are arranged in a simple chain (such as the mechanism represented in fig. 2), CGP's claim that "constitutive relevance is causal betweenness" (p. 8807) has no empirical content. If the mechanism were simply identical with the chain of variables (or conditions) between input and output, mechanisms would not be empirically discovered but conceptually constructed.

Such a construction would require 4 steps:

- 1) In a first step the input and output conditions Ψ_{in} and Ψ_{out} are chosen according to practical or explanatory interests.
- 2) In a second step²⁷, a causal chain of variables $\chi_1, \chi_2, \dots, \theta_1, \dots, \theta_n$, is discovered linking Ψ_{in} to Ψ_{out} .
- 3) In a third step, which is purely conceptual, the model-builder decides to group together the whole chain. This grouping, which is represented in fig. 3 by the large ellipse, receives the tag Ψ corresponding to the whole mechanism.
- 4) In a series of further steps, the model-builder introduces intermediate "chunks" of the chain. Fig. 3 represents two such groupings by the two small ellipses, tagged Φ_1 and Φ_2 .

²⁵ See Korman (2020), Kistler (2022). Merricks argues that "there are no inanimate macroscopic objects such as ... rocks or stars" (Merricks, 2001, p. vii) because they would be "causally redundant" (Merricks, 2001, p. viii). What seems to be caused by rocks, such as events of windows scattering, is in fact caused by the rocks' "constituent atoms, acting in concert" (Merricks, 2001, p. 56). Merricks (2001, chap. 4) defends eliminativism with respect to ordinary objects, but not radical nihilism with respect to all composite objects: persons and some other composites must be recognized on account of their nonredundant causal powers.

²⁶ "Fundamental physics is not a hospitable context for causation" (Woodward, 2009, p. 257). See also Norton (2003), Field (2003), Lange (2009), Woodward (2014, p. 702), Blanchard (2016).

²⁷ Harinen (2018) sketches this step in his fig. 3, Prychitko (2021) in her fig. 2, CGP in their fig. 2.

Mechanistic nihilism matches neither scientific practice nor scientific theory. The nihilist thesis that whole mechanisms do not exist, but that only their fundamental components exist, doesn't match scientific practice because biologists, cognitive scientists and neuroscientists directly manipulate, observe and measure not only those fundamental parts, but also larger parts and whole mechanisms. In the course of the exploration of the mechanism of backing up, scientists stimulate worms, and not only specific fundamental parts of worms, and they observe not only fundamental parts of worms, but also parts at intermediate levels as well as whole worms. The fact that non-fundamental parts of worms and whole worms can be manipulated is a strong reason for taking them to be real, and therefore also an argument against mechanistic nihilism. Ian Hacking has illustrated his case for entity realism, i.e., the thesis that the entities appearing in scientific theories can be interpreted as really existing, and not just as useful instruments or fictions, with an experiment in which positrons are sprayed on a niobium ball. Hacking's argument is that if something can be manipulated then it is real, so that, speaking of those niobium balls, "if you can spray them then they are real" (Hacking, 1983, p. 23). Furthermore, the fact that biologists and cognitive neuroscientists construct models of mechanisms, and not only models of fundamental components of such mechanisms, shows that mechanistic nihilism doesn't match scientific theory either.

Thus, the task remains of providing an analysis that justifies the existence and causal interactions of activities at higher levels, in which constitution plays the role of a non-causal relation between activities at different levels. The task is not straightforward: Constitution cannot be directly empirically tested, as is possible with causal influence between variables (Woodward, 2015, Baumgartner & Gebharder, 2016).

3. Integrating Causal Chains at Different Levels in a Multi-Level Model

Here is a suggestion of how the task might be accomplished²⁸. Multi-level mechanisms can be causally explored at different levels. The key to solving the problem is to distinguish 1) the project of constructing a model of a mechanism from 2) the project of analyzing the structure of experiments that provide empirical reasons for justifying parts of the model. "Vertical",

²⁸ What follows is a raw sketch of an analysis of the strategy for building multi-level models on the basis of experiments revealing causal relations between variables at specific levels. Much more work will be required to elaborate the sketch.

i.e. non-causal, constitution relations must be part of an adequate model of a mechanism, but they cannot, and need not, be part of what the relevant bottom-up and top-down experiments can directly establish.

Let me suggest a way in which the analysis of such experiments by Harinen (2018), Prichytko (2021) and CGP can be modified so as to construe them as part of the experimental strategy of discovering constitution in multi-level mechanisms.

The crucial idea is that interventions on, and measurements of, variables employ techniques that are *specific* to these variables. The specificity of these experimental methods includes in particular their being tuned to the level to which a variable belongs. I propose to call an *intervention* I on variable X w.r.t. variable Y *level-specific* if and only if I belongs to the same level as X . It is important that variable Y may lie at a lower level than the level of X , as in top-down experiments (TDE), or at a higher level than the level of X , as in bottom-up experiments (BUE). *Measurements* are always level-specific. A measurement of X consists in observing the value of Y (the indication of a measurement instrument), where Y is at the same level as X and where the causal influence of X on Y is specific, in the sense that each value of X is mapped onto a different value of Y ²⁹.

The following characterization of top-down experiments (TDE) modifies CGP's condition (CR2*) (2021, p. 8822; see above), by 1) introducing the condition of level-specificity, 2) categorizing each variable as belonging to a specific level, and 3) making the temporal structure of the experiments explicit.

(TDE) In a top-down experiment, 1) a variable I intervenes level-specifically on S 's Ψ_{in} at t_1 , then 2) X_2 's Φ_2 -ing is measured level-specifically at t_2 and 3) S 's Ψ_{out} is measured level-specifically at t_3 .

In the retraction mechanism of *c. elegans*, the manipulation of (i.e., intervention on) Ψ_{in} at t_1 consists in touching the worm's head, X_2 's Φ_2 -ing at t_2 may represent the activity of the ALM sensory neuron for head touch, and S 's Ψ_{out} at t_3 represents the worm's backing up.

The intervention on Ψ_{in} at t_1 , by touching the worm's head, is level-specific for the level of the whole mechanism, simply because the instrument used to touch the worm's head is of the same size as the worm's head. (Adult *c. elegans* worms are around 1 mm long.) The measurement of X_2 's Φ_2 -ing, i.e., of the activity of the ALM sensory neuron at t_2 , is level-specific for

²⁹ On causal specificity, Woodward (2010), Kistler (2017; 2021).

the lower level of neurons: it proceeds by an experimental technique tuned to the level of neurons, e.g., by inserting an electrode into the individual neuron and measuring the neuron's depolarization. Again, the observation of S 's Ψ_{out} at t_3 is level-specific for the level of the whole worm: it is done by a human observer using a magnifying glass or some other tool that is tuned to the observation of events of the size of the whole worm.

A bottom-up experiment (BUE), such as a bottom-up excitatory manipulation, consists in intervening, in a level-specific way, on a variable representing the activity of a part of the organism that is hypothesized to be constitutive of the mechanism. If the target of the BUE is the hypothesis that X_2 's Φ_2 -ing is constitutively relevant for Ψ , it involves the same variables as the second and third variables involved in schema (TDE) of the top-down experiment targeted at that constitution relation. Instead of measuring X_2 's Φ_2 -ing at t_2 , the value of X_2 's Φ_2 is set, in a level-specific way, by the experimenter at t_1 . Then the value of S 's Ψ_{out} is measured a little later, at t_2 .

Given that X_2 is a part of S , if it is found that X_2 's Φ_2 -ing at t_1 causally influences Ψ_{out} at t_2 , that influence is bottom-up.

The following characterizations of excitatory (BUE) and inhibitory (BUI) bottom-up experiments modify CGP's conditions (CRIi) and (CRIe) (2021, p. 8822; see above).

(BUE) In a bottom-up excitatory experiment, 1) a variable I intervenes level-specifically on X_2 's Φ_2 -ing at t_1 , then 2) S 's Ψ_{out} is measured level-specifically at t_2 .

(BUI) In a bottom-up inhibitory experiment, 1) a variable I_1 intervenes level-specifically on S 's Ψ_{in} at t_1 , then 2) a variable I_2 intervenes level-specifically on X_2 's Φ_2 -ing at t_2 and then 3) S 's Ψ_{out} is measured level-specifically at t_3 .

A non-causal relation, such as constitution, cannot be directly detected by a single experiment, but its inclusion in the model can be indirectly justified by experiments that reveal downward and upward causal influences. My suggestion is that such a model is built by integrating the causal relations between variables lying at different levels, which have been established by TDE, BUI and BUE experiments.

Here is a sketch of the reasoning that uses the results of TDE, BUI and BUE experiments to build models that contain constitution relations. Such a model must be built in several steps. In a first series of steps, a partial, purely causal, model is constructed for each hypothetical constituent, at each relevant level: Φ_1 at level n-1, χ_j at level n-2 and so on.

In causation there is always a delay between the time of the cause variable and the time of the effect variable. TDE, BUE, and BUI are all experiments aimed at discovering causal relations. In neuroscience, the delay between the times at which the cause and effect are modified or measured is often very short but it has always a finite non-zero value. The use of specific variables (Spohn 2001a; 2001b; 2006) makes explicit the fact that causal influence takes time, i.e., that the effect variable corresponds to a later time than the cause variable. In a TDE, e.g., X_2 's Φ_2 -ing is observed at t_2 , a few milliseconds after t_1 , at which S 's Ψ_{in} is modified by intervention³⁰.

The following reasoning is then applied separately to each constituent. Let us consider one of them: X_i 's activity of Φ_i -ing. If there are top-down experiments and bottom-up experiments of both types with respect to X_i 's activity of Φ_i -ing³¹, so that (TDE), (BUI), and (BUE) are all satisfied with respect to that activity, it is possible to combine them, by virtue of criterion (MM), which is inspired by Craver's (2007) original criterion of mutual manipulability, to construct a partial model of the causes and effects X_i 's activity of Φ_i -ing.

(MM) X_i 's Φ_i -ing is a constituent activity of S 's mechanism of Ψ -ing, iff there are experiments establishing that X_i 's Φ_i -ing at t_2 is 1) the effect of S 's Ψ_{in} at t_1 (with $t_1 < t_2$) (TDE), 2) the cause of S 's Ψ_{out} at some later time t_3 ($t_2 < t_3$) (BUE), and 2) a causal intermediary between S 's Ψ_{in} at t_1 and S 's Ψ_{out} at t_3 (BUI).

To sum up, the first step of model-building, it results in 1) a list of constituents and 2), for each constituent, a partial model of its causes and effects, which lie in general at different levels than the constituent itself.

³⁰ One of Krickel's conditions on "causation-based constitutive relevance" (2018, p. 64), which corresponds to top-down experiments, spells out the requirement that X is a proper part of S and that the first effect (X 's Φ -ing at t_2) occurs later than the cause (S 's Ψ_{in} at t_1), in terms of temporal and spatial "EIO-parts". Instead of saying that S 's Ψ_{in} at t_1 causes X_2 's Φ_2 -ing at t_2 , Krickel says that "there is a temporal EIO-part of S 's Ψ -ing that is a cause of X 's Φ -ing", given that " X 's Φ -ing is a spatial EIO-part of S 's Ψ -ing" (Krickel, 2018, p. 64). I think that this condition is more awkward but equivalent to my condition in terms of specific variables. Krickel's account is formulated in terms of "actual causation" because "the *relata* are taken to be concrete individuals". This may be a source of confusion. In the context of scientific research on mechanisms, the variables always represent features of types of organisms and their mechanisms although experiments provide knowledge about the values of these variables in particular individuals. This is so whether the research aims at discovering causal relations or constitutive relations.

³¹ This requirement is analogous to CGP's (p. 8822) requirement of "matching". See above.

In a second and final step, the information contained in these purely causal partial models for each constituent is combined into a synthetic model containing both causal and constitution relations. One crucial criterion for constructing such models is spatio-temporal inclusion³². If 1) the location of X_2 lies spatially within the location of S , if 2) in the relevant TDE experiment X_i 's Φ_j -ing at t_2 lies temporally between S 's Ψ_{in} at t_1 and S 's Ψ_{out} at t_3 , if 3), in the relevant BUE experiment, X_i 's Φ_j -ing at t_2 precedes S 's Ψ_{out} at t_3 , and if 4) in the relevant BUI experiment, X_i 's Φ_j -ing at t_2 is causally intermediary between S 's Ψ_{in} at t_1 and S 's Ψ_{out} at t_3 , this justifies the hypothesis that X_i 's Φ_j -ing at t_2 is constitutive of S 's mechanism of Ψ -ing.

CGP are right when they say that “philosophical analysis must keep conceptual, epistemic, and ontological questions about mechanisms and constitutive relevance distinct while recognizing that they are systematically connected to one another” (CGP, p. 35). Both causal and constitution relations are metaphysical posits that are represented in models as relations between variables. The justification of causal relations is more direct, and the causal arrows linking variables can be added in the first step of model construction. Constitution relations can then be added in a second step.

Contrary to CGP's analysis, each variable is explicitly attributed to a specific level in the part-whole hierarchy, with the whole organism at the top level. This makes it possible to interpret the results of various experiments as detecting parts of a multi-level model. True, “each experiment tests a different causal claim” (CGP, p. 26). However, by using variables that explicitly represent the level to which they belong, one can make explicit the fact that top-down and bottom-up experiments provide information that goes beyond the mere fact that some constituent activity lies between the whole mechanism's Ψ_{in} and Ψ_{out} . A top-down experiment doesn't only show that a variable such as X_j 's Φ_j -ing lies “between” S 's Ψ_{in} and S 's Ψ_{out} but also that it lies at a lower level than the mechanism of Ψ -ing itself.

Let me end with a remark on stable constituents. CGP acknowledge as a limitation of their account that “MIE is blind (as was MM) to certain kinds of mechanistic component. [...] [E]very mechanism we know [...]

³² As I said before, what follows is only meant as a raw sketch. One question I cannot tackle in this paper is whether the concept of mechanism that is the object of the present analysis is applicable beyond biological mechanisms, i.e., whether ecosystems, epidemics or climates can be taken to be mechanisms in the sense that their scientific exploration follows the same logic as the exploration of biological mechanisms. If an epidemic is a mechanism, its parts may be spread out far away from each other in space and time. The application of the analysis of the logic of the scientific exploration of such a mechanism would certainly have to be adapted, in particular concerning the relevant spatio-temporal constraints. I thank an anonymous referee for having brought up this issue.

require[s] (relatively) stable structures as standing conditions to work. [...] [T]he mechanism depends on their *not* changing with the causal input.” (CGP, p. 8824; italics in text).

The thesis that mechanisms are just made up of entities and activities, but do not contain stable constituents, is a metaphysical presupposition that CGP adopt without justification, from Machamer, Darden, and Craver (2000). To allow only “activities” but not stable states to be part of mechanistic models corresponds to the traditional (but unjustified) metaphysical prejudice that only changes can be causes, but not standing conditions. It is strange that this posit has not been abandoned when philosophers have begun analyzing mechanisms within the framework of causal models. In order to account for the crucial role that standing conditions play in mechanisms, it suffices to drop the requirement that only activities can be constituents. This makes it possible to see that the experimental demonstration of the constitutive relevance of stable conditions follows the same logic as the detection of activities that are constitutively relevant for a mechanism. The same sort of experiments and the same reasoning serve to establish that standing conditions and activities are constitutively relevant for a mechanism.

Conclusion

Constitution is a non-causal form of dependence. It shouldn't therefore be surprising that it is impossible to directly apply Woodward's (2003) criteria for detecting causal influence to the task of detecting constitutive relevance. We have examined two proposals for finding empirical criteria for constitutive relevance, which are 1) based on experiments that are tailored to the detection of causal influence, but 2) aim at providing necessary and sufficient conditions for constitutive relevance. We found that both provide only a part of the information needed to construct a model of a multi-level mechanism, such as the mechanism of backing up in *c. elegans*.

Baumgartner, Casini and Krickel's (2020) criteria have two limitations. 1) They require the detection of two synchronous effects. This does not correspond to actual experimental practice, which always aims at detecting causal influence between variables that represent states or activities at different times. 2) Their criteria cannot be satisfied in mechanisms that contain activities acting in parallel. Furthermore, 3) in mechanisms that contain causal chains of activities, their criteria can only be met by the very first link in such a chain.

Craver, Glennan and Povich (2021) give up the project of building models that contain representations of non-causal constitution relations.

Given that experimental methods are tailored for the detection of relations of causal influence, they suggest that a model of a mechanism that is empirically justified must be purely causal. However, an account according to which models of mechanisms contain only activities at a hypothetical fundamental level, but no representation of activities at higher levels and, a fortiori, no non-causal constitution relations, is an eliminativist, or at best nominalist, account of mechanisms. Activities at higher levels, and mechanisms themselves, don't exist in reality: They are just names, which we choose to use to group together specific parts of fundamental level causal chains.

I have suggested that CGP's account provides a crucial part of the information that is needed to build a model of a multi-level mechanism. Experiments can be targeted at different levels, using interventions and measurements that are *level-specific*. CGP's analysis of bottom-up and top-down experiments can be modified to represent level-specific interventions at various levels, over and above their own application to a hypothetical fundamental level.

Once chains of causal influence have been identified at different levels, they can be integrated in a separate step into a multi-level model. Constitutive relations can be added to the model on the basis of spatio-temporal constraints. If 1) the location of X_i lies spatially within the location of S , if 2) in the relevant TDE experiment X_i 's Φ_j -ing at t_2 lies temporally between S 's Ψ_{in} at t_1 and S 's Ψ_{out} at t_3 , if 3), in the relevant BUE experiment, X_i 's Φ_j -ing at t_2 precedes S 's Ψ_{out} at t_3 , and if 4) in the relevant BUI experiment, X_i 's Φ_j -ing at t_2 is causally intermediary between S 's Ψ_{in} at t_1 and S 's Ψ_{out} at t_3 , this justifies the hypothesis that X_i 's Φ_j -ing at t_2 is constitutive of S 's mechanism of Ψ -ing³³.

References

- Baumgartner, M. (2018). The inherent empirical underdetermination of mental causation. *The Australasian Journal of Philosophy*, 96, 335-350.
- Baumgartner, M., & Gebharter, A. (2016). Constitutive relevance, mutual manipulability, and fat-handedness. *British Journal for the Philosophy of Science*, 67, 731-756.
- Baumgartner, M., & Casini, L. (2017). An abductive theory of constitution. *Philosophy of Science*, 84, 214-233.
- Baumgartner, M., Casini, L., & Krickel, B. (2020). Horizontal surgicality

³³ I thank my auditors in Buenos Aires and Paris, as well as two anonymous referees for *Análisis Filosófico*, for their helpful critical comments on earlier versions of this paper. I thank Institut universitaire de France (IUF) for funding a partial leave that made the research for this paper possible.

- and mechanistic constitution. *Erkenntnis*, 85, 417-430.
- Bennett, K. (2017). *Making things up*. Oxford University Press.
- Blanchard, T. (2016). Physics and causation. *Philosophy Compass*, 11, 256-266.
- Blanchard, T. (2023). The causal efficacy of composites: A dilemma for interventionism. *Philosophical Studies*, 180, 2685-2706.
- Craver, C. F. (2002). Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science Supplemental*, 69, S83-S97.
- Craver, C. F. (2007). *Explaining the brain*. Oxford University Press.
- Craver, C. F. (2015). Levels. In T. Metzinger & J. M. Windt (Eds.), *Open MIND 8(T)*, MIND Group.
- Craver C. F., & W. Bechtel (2007). Top-down causation without top-down causes. *Biology & Philosophy*, 22, 547-563.
- Craver, C., Glennan, S., & Povich, M. (2021). Constitutive relevance and mutual manipulability revisited. *Synthese*, 199, 8807-8828.
- Eronen, M. (2013). No levels, no problems: Downward causation in neuroscience. *Philosophy of Science*, 80, 1042-1052.
- Eronen, M. (2015). Levels of organization: A deflationary account. *Biology & Philosophy*, 30, 39-58.
- Eronen, M. (2021). The levels problem in psychopathology. *Psychological Medicine*, 51(6), 927-933. <https://doi.org/10.1017/S0033291719002514>
- Field, H. (2003). Causation in a physical world. In M. J. Loux & D. W. Zimmerman (Eds.), *The Oxford Handbook of Metaphysics* (pp. 435-460). Oxford University Press.
- Friend, T. (2019). Can parts cause their wholes? *Synthese*, 196, 5061-5082.
- Goodman, M., & Sengupta, P. (2019). How *Caenorhabditis elegans* senses mechanical stress, temperature, and other physical stimuli. *Genetics*, 212, 25-51.
- Gray, J. M., Hill, J. J., & Bargmann, C. I. (2005). A circuit for navigation in *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences*, 102, 3184-3191.
- Hacking, I. (1983). *Representing and intervening*. Cambridge University Press.
- Harinen, T. (2018). Mutual manipulability and causal inbetweenness. *Synthese*, 195, 35-54.
- Kim, J. (1984). Epiphenomenal and supervenient causation. *Midwest Studies in Philosophy*, 9, 257-270. Reprinted in J. Kim (1993), *Supervenience and mind* (pp. 92-108). Cambridge University Press.
- Kistler, M. (2006). *Causation and laws of nature*. Routledge.
- Kistler, M. (2017). Higher-level, downward and specific causation. In

- M. Paolini Paoletti & F. Orilia (Eds.), *Philosophical and scientific perspectives on downward causation* (pp. 54-75). Routledge.
- Kistler, M. (2021). Models of downward causation. In J. Voosholz & M. Gabriel (Eds.), *Top-down causation and emergence* (Synthese Library, vol. 439, pp. 305-326). Springer.
- Kistler, M. (2022). La constitution en métaphysique des sciences et en métaphysique descriptive. In R. Künstler (Ed.), *Métaphysique et sciences: Nouveaux problèmes* (161-181). Hermann.
- Kiverstein, J., & M. D. Kirchhoff (2021). Diachronic constitution. Preprint. <http://philsci-archive.pitt.edu/19690/>
- Korman, D. Z. (2020). Ordinary objects. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition). <https://plato.stanford.edu/archives/fall2020/entries/ordinary-objects/>
- Krickel, B. (2018). Saving the mutual manipulability account of constitutive relevance. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 54, 10-19.
- Lange, M. (2009). *Laws and lawmakers: Science, metaphysics and the laws of nature*. Oxford University Press.
- Leuridan, B. (2012). Three problems for the mutual manipulability account of constitutive relevance in mechanisms. *British Journal for the Philosophy of Science*, 63, 399-427.
- Leuridan, B., & T. Lodewyckx (2020). Diachronic causal constitutive relations. *Synthese*, 198, 9035-9065.
- Lewis, D. (2004). Causation as influence. In J. Collins, N. Hall & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 75-106). The MIT Press. First publication in *Journal of Philosophy* 97, 182-197 (2000).
- Machamer, P., Darden, L. & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.
- Mellor D. H. (2012). Nature's joints: A realistic defense of natural properties. *Ratio (New series)*, 25, 387-404.
- Merricks, T. (2001). *Objects and persons*. Oxford University Press.
- Norton, J. D. (2003). Causation as folk science. *Philosophers' Imprint*, 3(4). <http://www.philosophersimprint.org/003004/>. Reprinted in H. Price & R. Corry, *Causation and the constitution of reality: Russell's republic revisited* (pp. 11-44). Clarendon Press.
- Prychitko, E. (2021). The causal situationist account of constitutive relevance. *Synthese*, 198, 1829-1843.
- Raatikainen, P. (2010). Causation, exclusion, and the special sciences. *Erkenntnis*, 73, 349-363.
- Romero, F. (2015). Why there isn't inter-level causation in mechanisms. *Synthese*, 192, 3731-3755.

- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton University Press.
- Schaffer, J. (2003). Is there a fundamental level? *Noûs*, 37, 498-517.
- Shapiro, L. (2010). Lessons from causal exclusion. *Philosophy and Phenomenological Research*, 81, 594-604.
- Shapiro, L., & Sober, E. (2007). Epiphenomenalism: The dos and don'ts. In G. Wolter & P. Machamer (Eds.), *Thinking about causes: From Greek philosophy to modern physics* (pp. 235-264). Pittsburgh University Press.
- Spohn, W. (2001a). Bayesian nets are all there is to causal dependence. In M. C. Galavotti et al. (Eds.), *Stochastic causality*, CSLI, and in W. Spohn (2009), *Causation, coherence, and concepts: A collection of essays*. Springer.
- Spohn, W. (2001b). Deterministic causation. In W. Spohn, M. Ledwig & M. Esfeld (Eds.), *Current issues in causation*, Mentis, 21-46.
- Spohn, W. (2006). Causation: An alternative. *British Journal for the Philosophy of Science*, 57, 93-119.
- Suzuki, H., Kerr, R., Bianchi, L., Frokjaer-Jensen, C., Slone, D. et al. (2003). In vivo imaging of *C. elegans* mechanosensory neurons demonstrates a specific role for the MEC-4 channel in the process of gentle touch sensation. *Neuron*, 39, 1005-1017.
- White, J. G., Southgate, E., Thomson, J. N., & Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 314(1165), 1-340.
- Wilson, A. (2018). Metaphysical causation. *Noûs*, 52(4), 723-751.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Woodward, J. (2009). Agency and interventionist theories. In H. Beebe, C. Hitchcock & P. Menzies (Eds.), *Oxford handbook of causation* (pp. 234-262). Oxford University Press.
- Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25, 287-318.
- Woodward, J. (2014). A functional account of causation; or, A defense of the legitimacy of causal thinking by reference to the only standard that matters—usefulness (as opposed to metaphysics or agreement with intuitive judgment). *Philosophy of Science*, 81, 691-713.
- Woodward, J. (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research*, 91, 303-347.

Received 23rd December 2024; accepted 10th February 2025.