*Philosophy of Science* (2025), **00**, 1–18
doi:10.1017/psa.2024.63

ARTICLE

# Closure of Constraints as a Theoretical Model

Campbell Rider

School of Humanities, The University of Sydney, Sydney, Australia
Email: campbell.rider@sydney.edu.au

(Received 21 April 2024; revised 16 October 2024; accepted 01 November 2024)

## Abstract

In this paper I offer a model-theoretic interpretation of Autonomy Theory as defended by Moreno, Mossio, Montévil, and Bich. I address accusations that Autonomy Theory is excessively liberal, such as those made by Garson (2017), arguing that these misunderstand the role of strategic abstractions and generalizations in theory construction. Conceiving of closure of constraints as a model-building effort that emphasizes generality—in the spirit of Levins (1966)—also clarifies its potential for application in empirical contexts.

## 1. Introduction

Numerous organizational theories of life have emerged over the past half-century following fruitful collaborations between scientists and philosophers. Among the most prominent contemporary flagbearers of this tradition are Matteo Mossio, Alvaro Moreno, Maël Montévil, and Leonardo Bich. Their version of the theory, which emphasizes "closure of constraints" in complex living systems, also grounds an account of biological function. Justin Garson's "liberality objection" (2017, 2019) holds that closure of constraints can be applied to many simple, intuitively non-living systems. While I review efforts to resist this objection, I wish to show that assessment of the theory as an analytic definition misunderstands its philosophical basis. Instead, I reconstruct Autonomy Theory in model-theoretic terms, following Levins (1966). A close reading of work produced by Mossio, Moreno, Montévil, and Bich is sufficient to show that closure of constraints can be interpreted according to its capacity to produce models guiding experimental work. I suggest that it is worthwhile to understand the relationship between closure of constraints and theoretical modeling, since this clarifies the role of strategic simplifications in the description of the theory. Since generality, idealization, and abstraction are part and parcel of any modeling endeavor, I argue that evaluation of Autonomy Theory should not depend on how well it explains the family of counterexamples proposed by Garson. Rather, the quality of the theory should be judged by its efficacy in guiding modeling efforts in experimental application.

## 2. Historical overview of organizational accounts of life

There is a rich tradition in theoretical biology that seeks to understand living systems through the concept of self-organization. In the last decade of the eighteenth century, Immanuel Kant and, following him, Friedrich Schelling, developed the earliest systematic accounts of life as self-organizing (Huneman 2006; Cooper 2023; Heuser-Keßler 1992). Claude Bernard's development of the concept of homeostasis in the middle of the nineteenth century also represents an important precursor to this family of theories (Keller 2008, 56). In the twentieth century, Humberto Maturana and Francisco Varela's "autopoiesis" theory of biological organization has been especially influential for shaping this school of thought in contemporary theoretical work (Varela 1979; Maturana and Varela 1980). They introduced the notion of autonomy and closure as necessary for an adequate understanding of biological forms, arguing that the process by which the organism integrates matter into a complex and self-maintaining system demands a holistic understanding of life. "Autopoiesis" or "self-production" denotes the causal and structural form whose realization—especially at the cellular level, but also at larger scales—defines a system as a living and self-enclosed individual (Maturana and Varela 1980).

Ilya Prigogine's work on thermodynamics in the 1960s laid the foundations for the application of self-organization to biochemistry and molecular biology (Keller 2009, 3). These insights were taken up by Stuart Kauffman in the 1990s, whose applications of the idea to the biochemistry of autocatalysis provided a robust conceptual core around which origins of life research could focus its efforts (Kauffman 1993). There are clear resonances between autopoiesis and Kauffman's notion of life as that entity or process which, via the causal self-relation embodied in chemical autocatalysis, can act "on its own behalf" (Kauffman 2000, x). Robert Rosen's work on relational systems in the 1980s and 1990s has also exerted considerable influence on the philosophical tradition, and shares with Varela and Maturana a holistic and formal emphasis on the organizational structure of systems, perhaps at the expense of their specific materiality (Wolfe 2010, 222). Interestingly, Rosen does not cite Varela anywhere in his work—nor does Varela cite Rosen in his later texts (Thompson 2004, 390).

In the 1990s and 2000s philosophers Wayne Christensen and Cliff Hooker (2000, 2001) published a number of articles foregrounding self-regulatory capacities in an effort to describe organisms as self-directed agents, embodying "bio-agency" (Skewes and Hooker 2009). While Hooker and Christensen represent an important intermediary body of work that brought autopoiesis and related concepts of organization and agency to analytic philosophy of biology, Alvaro Moreno, Matteo Mossio, Maël Montévil, and Leonardo Bich are the contemporary standard bearers of this theory in analytic philosophy of biology, and address themselves more directly to its key controversies and disputes. In this respect, they offer the most promising articulation of this view if we are interested in evaluating its relevance for ongoing debates. It should be noted that these theorists avoid the phrase "self-organization," favoring *self-maintenance* and *self-determination*, in order to differentiate their position from work on spontaneous self-assembly in physical systems (Mossio and Bich 2017, 1104). It is their encompassing project of "Autonomy Theory" which produces organizational *accounts* of phenomena such as function and agency.

There is a further, practical reason to highlight the tradition out of which Moreno, Mossio, Montévil, and Bich's thought emerges. Although they now occupy a position in analytic philosophy of biology, they grow out of a long tradition of thought on self-organization which has its roots in German philosophy of nature, and, more recently, the phenomenologically inclined work of Varela and Maturana, who themselves have identified the German phenomenologist Hans Jonas as a precursor (Weber and Varela, 2002). It is important to take note of this fact, as it explains why they employ a more abstract discursive style than the familiar conceptual analysis common to contemporary philosophy of biology. Justin Garson, for instance, mischaracterizes what he calls the "Organizational Theory" in this way: "OT is attractive to a certain kind of philosophical disposition, one that goes in for 'conceptual engineering'" (Garson 2022, 382). I suggest that it is not in fact a view that should be interpreted as concerned with an analytic definition of life or function, proceeding by way of necessary and sufficient conditions. Autonomy Theory must be understood as a fusion of a continental tradition with the methods and techniques of contemporary philosophy of biology and theoretical biology. This means it is entirely appropriate to evaluate it using the tools and methods of analytic philosophy of science, as it remains firmly planted in this tradition as well. This paper therefore attempts to show how, within the context of debates on scientific modeling and representation, it is desirable to understand closure of constraints, and Autonomy Theory generally, as a model-building project.

## 3. Self-determination and function

In Autonomy Theory, "organizational closure" is the causal regime that specifies biological systems, and which distinguishes them from dissipative structures. In the case of a single-celled organism, for example, closure simply refers to the organism's capacity for self-regulation, self-maintenance, and self-repair. It is a system that is able to replenish its own constituents, which undergo constant exchange with the outside world. The important point is that these self-directed processes of replenishment are under the control of the system itself, and form a network of parts that constrain and nourish one another. More formally, closure occurs when the constituents of a system maintain the boundary conditions necessary for its own existence via a series of mutually interacting constraints (Moreno and Mossio 2015, 19–20). Montévil and Mossio (2015) formalize this description as "closure of constraints." The system is organizationally closed but thermodynamically open—the causal regime of closure via self-constraint is the source of its biological autonomy, and also what guarantees that it can consume energy and resources from the outside without the collapse of its constitutive boundary conditions (2015, 23–24). A system which maintains itself through organizational and causal closure does so in virtue of different contributions from a diverse array of subsystems, which each have the global effect of sustaining the system at large and thus maintaining themselves.

On this view, a trait has the function that it does because of its role as a constraint on the boundary conditions of the organizational whole into which it is integrated. Closure "grounds functionality within biological systems: constraints do not exert functions when taken in isolation, but only insofar as they are subject to a closed organisation" (Montévil and Mossio 2015, 186). A trait has its function because, as a constraint, it stands in a causal relation to other constraints such that each one is mutually dependent on the adequate functioning of the others. Contrary to theorists concerned with naturalizing the apparent normativity of functions by locating the appearance of design in a function's selective history, they argue that "functionality is an emergent property of closure" (Moreno and Mossio 2015, 71). The normativity of the system derives from the fact that the parts of an organism must fulfill certain organizational roles to ensure the continuity of the system. The norm of each function is subsumed under the global norm of continuing to exist.

## 4. The role of modeling in Autonomy Theory

In Giere's classic formulation of modeling, there is "no direct relationship between sets of statements and the real world. The relationship is indirect through the intermediary of a theoretical model" (Giere 1990, 82). Models mediate between the complexity of real-world phenomena and the level of abstraction at which all scientific reasoning must operate, making complex empirical phenomena accessible to scientific analysis. This is perhaps the sense of Montévil and Mossio's (2015, 182) statement that "closure is at the core of the very constitution of biological phenomena as scientific objects." I suggest that the Autonomy Theory described by Moreno, Mossio, Montévil, and Bich can be positioned within the contemporary literature on scientific representation. While their own work fails to provide an explicit statement of the status of the theory as a kind of modeling, I propose an interpretation or rehabilitation of the theory along these lines.

First, it is important to address one paper in which Mossio and Bich distinguish their own "conceptual approaches" from what they call "model-based approaches." They write that modeling approaches:

> characterize self-determination by developing a formal system (the model) [ . . . ] with the objective of establishing a correspondence between the derived formal structure and the causal structure of the natural system. Within the model, it is then possible to infer the behavior of the system from certain initial conditions: these inferences constitute predictions about the behavior of the natural system we wish to study. Whereas conceptual approaches attempt to clarify "what self-determination is," model-based approaches attempt to explain how a system achieves self-determination by explicitly describing the causal regime at work. (Mossio and Bich 2014, 149, translation my own)

The terminology used here is somewhat counterintuitive. In my view the above passage is not sufficient evidence to show that Mossio and Bich are in fact opposed to modeling: instead, they oppose a specific kind of modeling which attempts an explicit and direct characterization of a living system. What they call a modeling approach, in other words, appears to imply what we might prefer to call an *application* of a model, such as in the simulation of a specific living system. What they call a "conceptual" approach, I suggest, corresponds with *highly general* theoretical models, which do not seek to directly represent or make predictions about a particular system.

In contrast to more direct forms of representation, they see their "conceptual" attitude as one step removed from concrete real-world systems, thus circumventing the explanatory overreach they regard as problematic in certain formal models. But they are just as keen to avoid the pitfalls of excessive abstraction and formalization, which they see as the principal limitation of Robert Rosen's work. Rosen's framework "remains too abstract, and therefore hardly applicable as a guiding principle for biological theorising, modelling and experimentation" (Montévil and Mossio 2015, 180). This statement certainly implies that they regard the possibility of application in modeling and experimentation as a key advantage of their own theory.

Montévil and Mossio (2015) refer to the notion of "tendency to closure" as an "operational tool," which is represented graphically in a "toy example" (189). They admit that their description of closure is "very general and schematic, and unable to capture the complexity of its actual realisations by biological systems" (Montévil and Mossio 2015, 186). This seems to concede that their view only *indirectly* represents living systems, and exists at a level of generality that requires it to be operationalized in the context of experiment or simulation. There is an ambivalence in their description, however, since they also appear to *contrast* their project with modeling:

> It is important to underline that our purpose is by no means to provide a model of closure which would adequately capture the complexity of real biological systems. Rather, we conceive this paper as a contribution to characterise in precise terms some of the general features of closure, which might subsequently be used to develop models of biological organisation. (Montévil and Mossio 2015, 180)

Montévil and Mossio regard their theory as deliberately idealized, and at a conceptual distance from the specificity of "real biological systems." But do they understand the theory as a *model*? In his later work on mathematical models, and in his collaboration with Ana Soto (discussed below), Maël Montévil appears to demonstrate his understanding of Autonomy Theory as a modeling effort (Montévil 2018; Montévil and Soto 2024). But the other philosophers associated with Autonomy Theory have not made an explicit or unified statement as to the model-theoretic status or potential of the framework, and, as we have seen above, use somewhat idiosyncratic language on those occasions where they have acknowledged the relationship between theory and modeling. Autonomy Theory in the abstract may not itself be called a model, but I suggest the encompassing project should still be interpreted as an effort to generate principles that contribute to the description of more concrete systems via a modeling relation. I therefore propose to offer a *reinterpretation* of Moreno, Mossio, Montévil, and Bich according to the model-theoretic conception of science. I will argue that even if they do not, in their published writings, have an explicit stance on whether Autonomy Theory is a modeling effort, they *should* see themselves as model-builders.

It will first be necessary to understand the status of modeling in contemporary philosophy of biology. The literature on scientific modeling offers a sizable typology of models (Frigg and Hartmann 2020). My aim below is to articulate some of the key attributes of the theory employed by Moreno, Mossio, Montévil, and Bich and assess how its key characteristics position it relative to the various theoretical alternatives available in the literature on modeling. In particular, I aim to show that the model of closure of constraints can be profitably understood within the framework of the "semantic" conception of theory structure.

The semantic conception stands opposed to the "syntactic" conception. As Thompson (1989) explains, in the syntactic conception a theory is an interpreted formal system comprising a "deductively-related set of statements" (32). Within a scientific theory, this network of statements defines a correspondence between theoretical terms and an observational vocabulary, such that any experimental procedure can be redescribed as a rigorous logical operation. Thompson summarizes this as an "axiomatic-deductive structure," which can be contrasted with the semantic conception's "model-theoretic structure" (1989, 32). The semantic view emphasizes that scientific interpretation of the world is indirect, and mediated by *models*. The key explanatory role of a given theory is therefore to define the class or family of models most apt for simulating and describing a target system. Empirical work then explores the effectiveness of a model or cluster of models in their explanatory role. Multiple models may be deployed to explain diverse aspects of a target system at different scales, and differing explanatory emphases can be regarded as complementary efforts to define the family of models suitable for a plurality of interpretive tasks. Indeed, a key point of the semantic view is that it captures the diversity of interpretive and hermeneutical conceptions which may comprise a scientific theory. Much contemporary work on modeling, whether or not it explicitly aligns itself with the semantic conception, has been influenced by this emphasis on the role of models. Hartmann (1996), for example, has described the use of simulations as a form of modeling which enables theorists to pre-select hypotheses for experiment and to analyze alternative experimental setups. Frigg (2010) and Godfrey-Smith (2006a, 2009) endorse the view that there is an aesthetic dimension to the modeling relation, insofar as models possess a "fictional" quality that enables imaginative contemplation conducive to theory-construction.

In Autonomy Theory we are clearly dealing with a highly *mediated* interpretation of living phenomena. Closure of constraints, on my view, illuminates important features of living systems via a modeling relation. If Moreno, Mossio, Montévil, and Bich are to be interpreted as contributing to a model-theoretic understanding of biology, then we must ask more concretely what *kind* of modeling this is. One avenue toward interpreting their remarks is to simply situate the principles of Autonomy Theory and closure of constraints within the influential framework elaborated by Levins (1966). Levins defines three dimensions of model construction: generality, realism, and precision. Modeling involves trade-offs between these factors: "The multiplicity of models is imposed by the contradictory demands of a complex, heterogeneous nature and a mind that can only cope with few variables at a time; by the contradictory desiderata of generality, realism, and precision [...]" (Levins 1966, 431).

Levins claims that a model can only maximize any two of these variables at the expense of the third: for example, a theory can be general (covering more cases) and realistic (accurately capturing the structure of the modeled entity), but it cannot be precise at the same time (able to produce predictions). Understood in these terms, the status of modeling in the work of Moreno, Mossio, Montévil, and Bich becomes clearer: they perhaps regard certain modeling attempts as failing to follow through on the trade-offs necessary to attain a model with the right sort of abstract generality. Rosen's work may be guilty of this: although his work is highly general, he also regards it as exceedingly realistic in its articulation of the structure of living organization. For our purposes, Levins' framework usefully defines the term "generality" as a characteristic that can be contrasted with realism and precision. I suggest that this is the sense of generality invoked by Autonomy Theory when, for example, constraints closure is described as "very general and [ . . . ] too schematic to capture the complexity of its actual realisations in biological systems" (Moreno and Mossio, 2015, 21).

Two questions may be asked at this stage. First, as I have described closure of constraints, its model-theoretic status is related to its generality and abstractness. But just about any form of scientific representation is in some way partial or incomplete—so what makes this modeling effort different from, for example, abstract direct representation? Secondly, it may be asked whether there is in fact anything distinct about the model described by Autonomy Theory when compared with other minimal models for the origins of life. We have already seen how Montévil and Mossio distinguish themselves from Rosen, but what about Ganti's "chemoton" or the "hypercycle?"

According to Weisberg (2007), models construct an idealized situation that deliberately misrepresents key features of the target system, whereas abstract direct representation refers itself explicitly to the target system and provides an abstract description of its relationships and processes. Weisberg claims that Volterra's celebrated model of the predator–prey relation exemplifies the former, while Mendeleev's derivation of the periodic table is an example of the latter (Weisberg 2007, 215). What differs is whether the theorist begins from a real-world phenomenon or an abstract model. Whereas in abstract direct representation, the theorist is in a sense "working backwards" from the real-world phenomena they intend to describe and explain, a modeler begins with an abstract and highly idealized description from which they undertake work of analysis and manipulation. There is therefore a final step that requires actual coordination of the model with real-world phenomena, in order to determine whether and to what extent the model fits. In abstract direct representation, this final step is not included. Weisberg states that if "the theorist is analyzing a representation that is directly related to a real phenomenon, anything she discovers in her analysis of the representation is a discovery about the phenomenon itself, assuming that it was represented properly. There is no extra stage where the theorist must coordinate the model to a real phenomenon" (Weisberg 2007, 227). Autonomy Theory exists at a degree of a generality that certainly requires this additional work of coordination.

Peter Godfrey-Smith (2006a) adopts Weisberg's distinction between abstract direct representation and modeling in his comparison of the use of models in two key texts on evolutionary transitions in individuality. In contrast to Maynard Smith and Szathmáry's *The Major Transitions in Evolution*, where "the currency of theoretical argument at each stage is the model," Godfrey-Smith observes that in *The Evolution of Individuality* Leo Buss does not employ models of any kind: his "entire argument is based on the causal roles and consequences of actual cellular machineries, actual environmental circumstances, and actual developmental sequences" (2006a, 731–32). Buss's work resembles abstract direct representation, and accordingly there is "no significant role for deliberate consideration of fictional, idealized, or merely schematic organisms, and the distinction between cautious exposition and deliberate fiction is a crucial one here" (2006a, 731).

Certainly, Autonomy Theory comes closer to Maynard Smith and Szathmáry's use of hypercycle and chemoton models in this respect. But Autonomy Theory has more lofty ambitions. It not only aspires to describe a class of chemical or physical loops that may have led to the emergence of life, but to theorize a category of organizational structure instantiated by every living thing, and which explains other attributes of living systems: autonomy, agency, and functionality. For example, the question of how RNA might have replicated itself in the absence of other cellular machinery is a concrete puzzle that requires a solution detailing a possible physical or chemical process. But if that resulting process is regarded as living, or as being instantiated in all living forms, this will only be because it is a special case of a more general organizational structure—that is, closure of constraints.

Letelier et al. (2011) provide a useful framework for distinguishing Autonomy Theory from other models used in origins of life studies. In their view, one tradition focuses on "design" (chemical makeup) as opposed to "principles" (metabolic dynamics). The latter strategy is associated with the classic example of the Carnot engine, since "the emphasis is not the production (even on paper) of an actual metabolic network, but a search for general principles" (Letelier et al. 2011, 104). The Carnot engine is an early example of what we might now call a "minimal" model. Importantly, it was not a schematic for an engine Carnot intended to construct: instead, it served as a theoretical tool enabling an understanding of the dynamics of work, energy expenditure, and efficiency in real engines and other work cycles. Since the hypercycle and chemoton models typically aim to account for the specific chemical properties of a hypothetical minimal life form, Letelier et al. see them as distinct from the modeling tradition exemplified by the Carnot engine—that is, the tradition to which Autonomy Theory belongs.

Godfrey-Smith (2006b) argues that there is a close analogy between model-based science and the strategies of abstraction and idealization used by analytic philosophers in their metaphysical system-building. If we think of a model as "an imagined or hypothetical structure that we describe and investigate in the hope of using it to understand some more complex, real-world 'target' system," he argues, then much work in metaphysics might be fruitfully understood as an effort of this kind (2006b, 7). For instance, "Humean supervenience and Armstrong's anti-Humeanism are both toy models of the universe" which offer "a particular kind of problem-solving and systematizing power" (2006b, 16–17). Godfrey-Smith observes that these philosophers are likely unaware of the extent to which their theories can be conceptualized in model-theoretic terms, since "a person can be engaged in modeling while having a different self-understanding" (2006b, 9). Indeed, someone with more robustly metaphysical ambitions may resist such an analogy, and even suspect that a model-theoretic interpretation threatens an excessively instrumentalist reading of their position. But if we wish to understand what is attractive or persuasive about certain metaphysical frameworks, Godfrey-Smith claims, it is useful to acknowledge the role played by strategic abstractions and generalizations.

It is on this final point that Godfrey-Smith's remarks might be brought to bear on our preceding discussion of Moreno, Mossio, Montévil, and Bich. Certainly, the theory of organizational closure attempts an "ontic" rather than merely "heuristic" characterization of the causal structure of living systems (Desmond and Huneman 2020). But to make the complexity of living organization tractable, they have deployed techniques of generalization, idealization, and abstraction, and to the extent that this model-theoretic grounding recontextualizes the abstract generality of the theory as an advantage rather than a weakness, it is surely in the interests of these theorists that the role of these techniques be explicitly acknowledged.

## 5. Autonomy Theory in practice

Since the model of organizational closure exists at such a high level of generality, Montévil and Mossio (2015) explain that any application of the model will need to specify clear parameters, particularly concerning scale. For example, the model requires a distinction between "processes" and "constraints." While a red blood cell may be considered a stable "constraint" if one is studying intracellular processes, from the perspective of the circulatory system at large it may be assimilated to the "processual" activity of blood flow, constrained by arteries and capillaries. Since by the logic of Autonomy Theory all system components contribute to the living process and are mutually constraining, to even describe a subsystem such as the vascular system we must specify what, for the purposes of the model, will be considered the "processes" and what will be considered the "constraints." Since this is true at every scale and for every functional system, "a complete characterisation of the whole set of mutually dependent constraints is usually not available, and constitutes a sort of 'theoretical horizon' of biological explanation" (Montévil and Mossio 2015, 187). Any analysis under this theory must make provisional assumptions. As such, "in practice, any actual description of closure in biological systems is a partial one" (Montévil and Mossio 2015, 187). We must therefore idealize any given model in order to posit an underlying stability against which the process or mechanism of interest can be assessed. For instance, in an organizational model of the circulatory system, the heart and blood vessels must be treated as stable entities if what we are interested in is the generation of blood cells, oxygenation, and so on. This is typical of any modeling relation: it begins at a level of abstraction that distorts or omits features that will necessarily constrain any real instantiation of the represented system, so that in application the model proceeds by "de-idealization," that is to say, the "successive removal of the distorting assumptions" (Frigg and Hartmann 2020).

The theory elaborated by Moreno, Mossio, Montévil, and Bich has many of the hallmarks of an effort at scientific modeling. Their closure model is well suited to stimulating experimental design, with the ultimate purpose of guiding empirical practice, and, as we shall see, also enabling limitations of certain experimental setups to become clear. Let us consult three examples.

### *5.1. Modeling morphogenesis in mammary tissues*

In collaboration with Mossio, Montévil, and Bich, cancer researcher Ana Soto has endorsed Autonomy Theory as a theoretical resource for the development of experimental models. Soto regards Autonomy Theory as critically important to the project of constituting biology as an independent domain of inquiry, with its own principles and norms of experimentation. She takes seriously the view that "biological objects are agents capable of creating their own norms," and argues that conceiving of them in this way "opens up the possibility of anchoring mathematical modelling on properly biological principles" (Soto et al. 2016, 81). To demonstrate the applicability of these organicist principles, Bich, Mossio, and Soto (2020) have produced an organizational model of glycemia regulation, emphasizing its advantages over feedback loop models, as well as a model based on Soto and Sonnenschein's (2011) "tissue organisation field theory" of carcinogenesis.

This latter model, constructed in collaboration with Maël Montévil, examines tumor formation in mammary tissue (Montévil and Soto 2024; Montévil et al. 2016). Applied to mammary cells suspended in a three-dimensional fibrillar culture, Montévil and Soto's model recognizes five interacting layers of constraints, including cellular, nutritive, and inhibitory constraints, in addition to mechanical forces (Montévil et al. 2016, 64). The model also presumes four possibilities for cellular behavior: exertion of force on neighboring cells; generation and division; random movement; and death. Montévil and Soto explain that organizational closure comes into play at two different levels, first via the mutual interactions that occur directly between cells, and second through the cells' alteration of the physical matrix which organizes them in space, which indirectly impacts other cells. Here is how they describe the advantage to the experimenter of modeling tissue morphogenesis as the complex result of interacting levels of constraint closure:

> Biological meaning is construed by applying similar constraints to those which operate in vivo and which seem to play a role in the determination of the phenomenon. In this way, we can "reduce" the number of constraints to those necessary to answer our specific question [. . .] In vitro 3D models allow researchers to manipulate constraints beyond the range operating in vivo. That is, constraints are determined by the organism and its parts, while in the in vitro model the researcher also plays a direct role in modifying these constraints and parameterizing them. (Montévil et al. 2016, 67)

This example is instructive because it highlights the possibility for experimenters to undertake processes of "de-idealization." By defining the constraints, they can be successively added or subtracted in order to isolate their distinctive causal contributions. In Montévil and Soto's paper, multiple interacting models are at play: a theoretical model (Autonomy Theory in the abstract), an experimental model (the in vitro three-dimensional cell culture), in addition to a mathematical description of the theoretical model, for the purposes of programming yet another model, via computer. Furthermore, these overlapping models draw from differing domains of biological knowledge but interact collaboratively to guide how experimenters proceed. This layering and self-reinforcement of models is examined under the heading of "robustness analysis" in the contemporary philosophical literature, following Levins' (1966, 431) suggestion that "a satisfactory theory is usually a cluster of models" which generate "robust theorems" (see Orzack and Sober 1993; Weisberg 2006; Schupbach 2018). Here, a model is said to be robust when it tracks a phenomenon across different iterations which vary in their underlying assumptions, idealizations, or simplifications. In this respect, the above example demonstrates yet another typical modeling procedure afforded by the generality of Autonomy Theory.

### 5.2. Computational models of metabolism

Using the example of *computational* models of metabolism, Moreno and Ruiz-Mirazo (1999, 48) show how thinking in terms of organizational closure can guide actual modeling and experiment. First of all, developing formal models of a minimal metabolic organization enables the dynamics of a metabolic system to be abstracted away from specific chemical pathways, reactions, and enzymatic action. This can give rise to fruitful model systems, and Moreno and Ruiz-Mirazo offer the examples of the computer simulations produced by Varela and McMullin and separately by Stuart Kauffman in the 1990s. For Varela and McMullin, the central aim of the simulation is to demonstrate that when certain conditions are met, a spatial boundary separating system from environment emerges. For Kauffman, computer simulations provide some of the principal evidence for autocatalytic sets as a model for the chemistry of early life: in a simulation, the parameters of large chemical networks can be manipulated so as to demonstrate what variables affect whether a system will reach a threshold where it can no longer grow, and under what circumstances it pursues open-ended (autocatalytic) growth. In both cases, there are two levels of abstraction: first, the theoretical model (autopoiesis and autocatalysis), and second, the computer simulation which realizes the model. The former prescribes the structure of the latter. But as Moreno and Ruiz-Mirazo point out, an excessively abstract theoretical model will carry over some of its limitations to the more “concrete” simulation: “the conception of metabolism as a purely relational-constructive system disregards the material causal interrelations in it, which are associated to processes that take place in intrinsic times and with intrinsic energies. So the processes that computational metabolisms reproduce occur in times and with energies which are arbitrary” (Moreno and Ruiz-Mirazo 1999, 59).

What the computer simulations lack is an awareness that the formal structure (metabolic closure) of the system is enmeshed in a cycle of thermodynamic openness and energy transfer that cannot be excluded from the formal account. This is not to say that the model is inadequate if it does not model all of the chemical reactions actually taking place in a system. But if the model does not recognize that different energetic processes have different reaction times and that the metabolic process involves synchronization of these reaction times, or that the organism as a whole must be able to independently store energy and access it later (Moreno and Ruiz-Mirazo 1999, 51), then the model has abstracted too far, to the point of also generalizing away a key formal element: the autonomous control and direction of energetic flow. This example demonstrates once again that while Autonomy Theory emphasizes generality, this does not prevent it from illuminating cases where too much realism has been “traded-off.” This case is illustrative as it shows that the generality of the model does not imply an empty formalism: concrete thermodynamic considerations continue to play a role in the model. At the same time, this is a clear-cut example of how modeling relates to experimental practice, including the development of simulations. The model guides, informs, and provides a basis for the critique of a given experimental setup.

### 5.3. Modeling ecosystem function

Systems of organizationally closed individuals may interact symbiotically or parasitically in such a way that closure may be ascribed to the encompassing system itself. Montévil and Mossio (2015) have suggested a measure of "tendency to closure," where closure is not conceived as a binary state which a system may or may not possess, but rather a characteristic that can come in degrees and vary across levels of a system. El-Hani et al. (2024) and Nunes-Neto et al. (2014) have used this notion to extend the organizational account of function to ecosystems:

> [I]t seems justifiable to ascribe ecological functions to the organisms constituting the system, even if the system does not show fully-fledged functional integration or constraints closure [ . . . ] We think we can extend [the notion of tendency to closure], however, to conceive of subsystems or modules, generally speaking, which show a relatively large degree of internal cohesion but yet depend on other modules in a given network. (El-Hani et al. 2024, 302)

Different modules or components in an ecological network can relate to each other as dependent or enabling, exerting influence according to the structure of self-maintenance under mutual constraint. As such they are "functionally coupled" and can be described as "elements within a set or network of modules" (El-Hani et al. 2024, 302). The notion of tendency to closure permits a "somewhat more relaxed notion of internal cohesion that makes it clear how the scope of the organizational theory of ecological functions is substantially broader than just a limited number of cases showing fully-fledged closure" (El-Hani et al. 2024, 302). They offer the example of a local ecosystem that can be found in species of the bromeliad *Quesnelia arvensis*, explicitly describing it in model-theoretic terms: "We cannot account here for all the details and complexity found in real bromeliad ecosystems. Our strategy will be, rather, to build a simple model of a possible (idealized) ecosystem [ . . . ] The most important feature of this model system is that it allows us to think clearer about function in ecological terms" (Nunes-Neto et al. 2014, 133). By tracking the flow of energy through the system, the model helps to determine the precise enabling/dependency relationships between producers and consumers, autotrophs and heterotrophs. Described in this way, the health of the system can then be assessed according to the criteria of functionality given by the organizational account.

I have offered some examples of the formal and empirical applications of Autonomy Theory. I have emphasized that this tradition leans toward the abstract and conceptual (highly "general" in Levins' framework), but without disregarding empirical applicability. The examples given above—studies of mammary morphogenesis, computer simulations, and ecosystem models—illustrate how a model exhibiting a high degree of generality may be progressively "de-idealized" for application to a variety of cases. Moreno, Mossio, Montévil, and Bich do not intend to represent, however abstractly, a specific living system. They aspire to describe an ideal class of systems to which the concrete structure, processes, and organizational regime embodied by real living systems may be referred. The above case studies suggest an emerging interest in operationalizing this theory to generate models for application in more specific, empirical contexts. To this end, an explicit model-theoretic interpretation of Autonomy Theory seems a worthwhile preliminary to this sort of work.

## 6. Garson's criticisms of organizational functions

In this section, I outline Garson's criticisms of organizational accounts and survey possible responses on behalf of Moreno, Mossio, Montévil, and Bich. My intention is not to robustly defend the Autonomy Theory against the charges of liberality, but to show that Garson misrepresents the depth and complexity of the theory. As such, I only offer a brief discussion. In the next section, I argue that the strength of Autonomy Theory should be measured by its contribution to theoretical modeling of biological phenomena, and that what Garson calls liberality—in our terms, generality—is an unremarkable by-product of any effort to construct a model of this kind.

Garson (2019, 1148) argues that an effective account of biological function ought to do at least two things. Firstly, it should be able to distinguish function from dysfunction. Secondly, it must be able to distinguish functions from accidents. Garson (2017, 526) also includes the requirement that a theory of function should make sense of the idea that entities with functions exist *because* they have those functions. With these conditions in mind, Garson argues that organizational theories succumb to "liberality" objections of the sort advanced by Boorse (1976). Central to Garson's argument is a simplification of the organizational account of function to the following definition: "The basic organizational theory holds, roughly, that trait T has function F either because T's doing F contributes to the intragenerational persistence of T or because T's doing F contributes to the intergenerational multiplication of T" (Garson 2017, 1094).

On this definition, a trait has its function in virtue of its contribution to the persistence of the system to which it belongs. A trait's function derives from its role in the self-maintenance of the organism. Garson argues that these accounts are too liberal because the proposed definitions include cases which our intuitions tell us are not examples of biological function. Garson suggests that we consider a gas leak in a laboratory, which ensures its own persistence by knocking out any scientist who wishes to seal it up (Garson 2017, 1095). It clearly does not, for this reason, have the "function" of self-preservation. Another process that would appear to have functional status on Garson's reading of the organizational account is a panic disorder: panic disorders entrench physiological and psychological conditions that contribute to the recurring likelihood of panic attacks. Each panic attack leads to heightened nervousness, adrenal changes, and so on, which increase the chances of subsequent panic attacks, entrenching the disorder (Garson 2017, 1097). Garson's point is that even though panic disorders are "prima facie dysfunctional states," they can be aptly described as functional traits according to organizational theories insofar as they instantiate a bootstrap-like process of self-reproduction (1097). Mossio and Moreno (2010, 277), Moreno and Mossio (2015), and Mossio and Bich (2017) include conditions that are overlooked by Garson. However, in the following I intend to only provide a brief assessment of these potential responses. As I will argue, additional conditions can only give rise to further counterexamples and marginal cases—the strength of Autonomy Theory should not depend on its capacity to answer such objections.

Garson's assessment of dysfunctional states can be answered by the account of dysfunction offered by Moreno and Mossio in *Biological Autonomy* (2015), which distinguishes first-order "constitutive" processes from second-order "regulatory" processes (Moreno and Mossio 2015, 83–84). A mechanism which is functional when considered at the first-order level may be interpreted as dysfunctional when considered relative to second-order regulatory norms imposed on it by signaling and feedback channels. According to this framework, what we call "panic disorder" is a state in which the adrenal, circulatory, and nervous systems interact in such a way as to lead to heightened respiration, anxiety, and fear. As Moreno and Mossio would remind us, the constituent parts of the panic disorder are not an isolated system. If the sense organs in conjunction with the nervous system are detecting danger when there is none, then this system—which stands in a higher-level regulatory relation to the adrenal and circulatory subsystems it coordinates—is in a state of dysfunction. Montévil and Mossio write that "no subsystem of collectively dependent constraints that can be shown to belong to an encompassing closed system can be said to realise closure" (2015, 188). Since closure is the basis of functionality, a given subsystem is functional only insofar as it contributes to the closure of the higher-level system to which it belongs.

These considerations motivate a distinction between mere cycles and genuine cases of closure. As Mossio and Bich (2017, 1106) explain, there is a difference between "a causal regime in which some effects happen to circularly contribute to generate their own (material) causes, and a causal regime which itself plays a role in determining the conditions under which the effects contribute to generating their own causes." Mossio and Moreno (2010, 277) clarify this point with their example of a frictionless billiards table, where billiard balls bounce off each other in a self-sustaining loop. On their view, this is just a case of positive feedback and self-reinforcing energy transfer, and therefore merely a cycle rather than a case of closure: the loop does not in fact produce, replicate, or maintain the actual components of the system. According to the principles of Autonomy Theory, the panic disorder is an example of a cycle that does not independently realize closure, and to which functions cannot be ascribed.

It may be asked why, if Moreno, Mossio, Montévil, and Bich have model-theoretic ambitions, they nonetheless include conditions in their description of closure that appear to shore it up against accusations of liberality. It should not be surprising that even a theory that wants to avoid the game of definitions and counterexamples would wish to avoid being *excessively* liberal. More importantly, however, the above qualifications in the definition of closure are not included primarily to exclude counterintuitive cases; rather, they appear as natural consequences of core theoretical commitments: the view that living organization is autonomous, or that function is contextualized by the norm of continued viability at the level of the whole organism.

## 7. The liberality objection

Above, I have identified some avenues for Autonomy Theory to respond to Garson's counterexamples. Rather than treating these as counterarguments in their own right, my aim has been to show that objections of this kind tend to rely on a somewhat oversimplified account of Autonomy Theory. I wish to argue, following the discussion of modeling in preceding sections, that the theory can be better defended by recasting itself as a model-building effort that is able to formalize general principles and guide empirical research.

A somewhat analogous discussion can be found in the penultimate chapter of *The Intentional Stance*, where Daniel Dennett addresses similar "liberality" objections leveled at the computational theory of mind, such as Ned Block's famous variation on the Chinese room. Block asks whether consciousness will be instantiated at the scale of a large nation (such as China) if all its citizens signal each other by two-way radio, simulating the action of neurons:

> Like any effort at scientific modeling, AI modeling has been attempted in a spirit of opportunistic oversimplification. Things that are horribly complicated may be usefully and revealingly approximated by partitionings, averagings, idealizations, and other deliberate oversimplifications [ . . . ] it is by no means obvious that any united combination of the sorts of simplified models and subsystems developed so far in AI can approximate the perspicuous behavior of a normal human brain [ . . . ] but that still does not impeach the research methodology of AI, any more than their incapacity to predict real-world weather accurately impeaches all meteorological oversimplifications as scientific models. (Dennett 1998, 330–31)

We can just as easily imagine a scenario in which members of a nation—still communicating via two-way radio—are responsible for constructing vital infrastructure projects that facilitate the lifestyle of other members. The distributed communication system means that all members are reciprocally dependent on directing each other to their respective tasks. What happens if a breakdown in communication occurs, disrupting construction projects and leading to widespread disarray? In the style of Block's classic example, we might say that in this scenario faulty buildings and communications infrastructure are not obviously "tumorous," "oncogenetic," or capable of metastasis. But this does not impugn the theoretical model of closure, such as the Soto–Montévil model, which defines these terms according to failed coordination of mutually constraining components in a tissue field.

As Dennett explains, the original China-brain thought experiment parodies the simplifications made by researchers who strategically disregard the biophysical materiality of the mind in order to develop digital or computational models of cognition. It aims to show that it is difficult to imagine such a system instantiating consciousness; as such, computational models which operate via the same principle should not be taken seriously. Dennett's response to the above sort of objection just as adequately addresses criticisms of Autonomy Theory. When interpreted as an analytic definition, the organizational account of functions might be open to similar objections. Understood model-theoretically, however, it becomes clear that arguments of this sort misleadingly target the very characteristics that make models effective: generality, abstraction, and idealization.

## 8. Conclusion

Generality, abstraction, and idealization are the signature procedures of any attempt to construct a theoretical model. These are certainly features of Autonomy Theory, regardless of the degree to which Moreno, Mossio, Montévil, and Bich conceive of it model-theoretically. Following Godfrey-Smith (2006a), I have suggested that it is useful to explicitly state the advantages afforded to a theorist who is engaged in model-building. Critics of Autonomy Theory may claim that it fundamentally relies on a picture of recursive self-propagation which is not restrictive enough to distinguish trivial non-living cases—mere cycles and loops—from genuine organizational closure. I have argued that this fact does not impair the capacity of the theory to enhance our understanding of experimental setups and simulations. While a defense may be mounted against accusations of liberality, Autonomy Theory does not need to justify itself in these terms when it is understood as model-building project. Where generality is a given, what must be evaluated is the effectiveness and applicability of the model itself.

**Acknowledgements.** I am grateful to Paul Griffiths and the Theory and Method in Biosciences group at the University of Sydney for their helpful comments. I would also like to thank Matteo Mossio for his feedback on an earlier version of this paper.

## References

Bich, Leonardo, Matteo Mossio, and Ana M. Soto. 2020. “Glycemia regulation: From feedback loops to organizational closure.” *Frontiers in Physiology* 11. https://www.frontiersin.org/articles/10.3389/fphys.2020.00069.

Boorse, Christopher. 1976. “Wright on functions.” *The Philosophical Review* 85 (1):70–86. https://doi.org/10.2307/2184255.

Christensen, Wayne. D., and Cliff. A. Hooker. 2000. “An interactivist-constructivist approach to intelligence: Self-directed anticipative learning.” *Philosophical Psychology* 13 (1):5–45. https://doi.org/10.1080/09515080050002717.

Christensen, Wayne. D., and Cliff. A. Hooker. 2001. “Self-directed agents.” *Canadian Journal of Philosophy Supplementary Volume* 27:18–52. https://doi.org/10.1080/00455091.2001.10715995.

Cooper, Andrew. 2023. “Kant and biological theory.” In *Life, Organisms, and Human Nature: New Perspectives on Classical German Philosophy*, edited by Luca Corti and Johannes-Georg Schülein, 21–38. New York: Springer. https://doi.org/10.1007/978-3-031-41558-6_2.

Dennett, Daniel C. 1998. *The Intentional Stance*. Cambridge, MA: MIT Press.

Desmond, Hugh, and Philippe Huneman. 2020. “The ontology of organismic agency: A Kantian approach.” In *Natural Born Monads*, edited by Andrea Altobrando and Pierfrancesco Biasetti, 33–64. Berlin: De Gruyter. https://doi.org/10.1515/9783110604665-003.

El-Hani, Charbel N., Felipe Rebelo Gomes De Lima, and Nei De Freitas Nunes-Neto. 2024. “From the organizational theory of ecological functions to a new notion of sustainability.” In *Organization in Biology*, edited by Matteo Mossio, 285–328. New York: Springer. https://doi.org/10.1007/978-3-031-38968-9_13.

Frigg, Roman. 2010. “Models and fiction.” *Synthese* 172 (2):251–68. https://doi.org/10.1007/s11229-009-9505-0.

Frigg, Roman, and Stephan Hartmann. 2020. “Models in science.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford, CA: Stanford University Press. https://plato.stanford.edu/archives/spr2020/entries/models-science/.

Garson, Justin. 2017. “Against organizational functions.” *Philosophy of Science* 84 (5):1093–103. https://doi.org/10.1086/694009.

Garson, Justin. 2019. “There are no ahistorical theories of function.” *Philosophy of Science* 86 (5):1146–56. https://doi.org/10.1086/705472.

Garson, Justin. 2022. “What are functions good for?” *Australasian Philosophical Review* 6 (4):374–85. https://doi.org/10.1080/24740500.2024.2370624.

Giere, Ronald N. 1990. *Explaining Science: A Cognitive Approach*. Chicago, IL: University of Chicago Press.

Godfrey-Smith, Peter. 2006a. “Theories and models in metaphysics” *The Harvard Review of Philosophy* 14 (1):4–19. https://doi.org/10.5840/harvardreview20061411.

Godfrey-Smith, Peter. 2006b. “The strategy of model-based science.” *Biology & Philosophy* 21 (5):725–40. https://doi.org/10.1007/s10539-006-9054-6.

Godfrey-Smith, Peter. 2009. "Models and fictions in science." *Philosophical Studies* 143 (1):101–16. https://doi.org/10.1007/s11098-008-9313-2.

Hartmann, Stephan. 1996. "The world as a process." In *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, edited by Rainer Hegselmann, Ulrich Mueller, and Klaus G. Troitzsch, 77–100. Dordrecht: Springer. https://doi.org/10.1007/978-94-015-8686-3_5.

Heuser-Keßler, Marie-Luise. 1992. "Schelling's concept of self-organization." In *Evolution of Dynamical Structures in Complex Systems*, edited by Rudolf Friedrich and Arne Wunderlin, 395–415. Berlin: Springer. https://doi.org/10.1007/978-3-642-84781-3_21.

Huneman, Philippe. 2006. "From the critique of judgment to the hermeneutics of Nature: Sketching the fate of philosophy of Nature after Kant." *Continental Philosophy Review* 39 (1):1–34. https://doi.org/10.1007/s11007-006-9017-2.

Kauffman, Stuart A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Kauffman, Stuart A. 2000. *Investigations*. Oxford: Oxford University Press.

Keller, Evelyn Fox. 2008. "Organisms, machines, and thunderstorms: A history of self-organization, part one." *Historical Studies in the Natural Sciences* 38 (1):45–75. https://doi.org/10.1525/hsns.2008.38.1.45.

Keller, Evelyn Fox. 2009. "Organisms, machines, and thunderstorms: A history of self-organization, part two. Complexity, emergence, and stable attractors." *Historical Studies in the Natural Sciences* 39 (1):1–31. https://doi.org/10.1525/hsns.2009.39.1.1.

Letelier, Juan-Carlos, María Luz Cárdenas, and Athel Cornish-Bowden. 2011. "From *L'Homme Machine* to metabolic closure: Steps towards understanding life." *Journal of Theoretical Biology* 286:100–113. https://doi.org/10.1016/j.jtbi.2011.06.033.

Levins, Richard. 1966. "The strategy of model building in population biology." *American Scientist* 54 (4):421–31. https://www.jstor.org/stable/27836590.

Maturana, Humberto R., and Francisco J. Varela. 1980. *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.

Montévil, Maël. 2018. "A primer on mathematical modeling in the study of organisms and their parts." In *Systems Biology*, edited by Mariano Bizzarri, 41–55. New York: Springer. https://doi.org/10.1007/978-1-4939-7456-6_4.

Montévil, Maël, and Matteo Mossio. 2015. "Biological organisation as closure of constraints." *Journal of Theoretical Biology* 372:179–91. https://doi.org/10.1016/j.jtbi.2015.02.029.

Montévil, Maël, and Ana M. Soto. 2024. "Modeling organogenesis from biological first principles." In *Organization in Biology*, edited by Matteo Mossio, 263–83. New York: Springer. https://doi.org/10.1007/978-3-031-38968-9_12.

Montévil, Maël, Lucia Speroni, Carlos Sonnenschein, and Ana M. Soto. 2016. "Modeling mammary organogenesis from biological first principles: Cells and their physical constraints." *Progress in Biophysics and Molecular Biology*, 122 (1):58–69. https://doi.org/10.1016/j.pbiomolbio.2016.08.004.

Moreno, Alvaro, and Matteo Mossio. 2015. *Biological Autonomy*. Dordrecht: Springer. https://doi.org/10.1007/978-94-017-9837-2.

Moreno, Alvaro, and Kepa Ruiz-Mirazo. 1999. "Metabolism and the problem of its universalization." *Biosystems* 49 (1):45–61. https://doi.org/10.1016/S0303-2647(98)00034-3.

Mossio, Matteo, and Leonardo Bich. 2014. "La circularité biologique: Concepts et modèles." In *Modéliser et Simuler : Epistémologies et Pratiques de La Modélisation et de La Simulation*, Vol. 2, edited by Franck Varenne, Marc Silberstein, and Sébastien Dutreuil, 165–200. Cork: Éditions Matériologiques. https://doi.org/10.3917/edmat.varen.2014.01.0137.

Mossio, Matteo, and Leonardo Bich. 2017. "What makes biological organisation teleological?" *Synthese* 194 (4):1089–114. https://doi.org/10.1007/s11229-014-0594-z.

Mossio, Matteo, and Alvaro Moreno. 2010. "Organisational closure in biological organisms." *History and Philosophy of the Life Sciences* 32 (2/3):269–88.

Nunes-Neto, Nei, Alvaro Moreno, and Charbel N. El-Hani. 2014. "Function in ecology: An organizational approach." *Biology & Philosophy* 29 (1):123–41. https://doi.org/10.1007/s10539-013-9398-7.

Orzack, Steven Hecht, and Elliott Sober. 1993. "A critical assessment of Levins's *The Strategy of Model Building in Population Biology* (1966)." *The Quarterly Review of Biology* 68 (4):533–46. https://doi.org/10.1086/418301.

Schupbach, Jonah N. 2018. “Robustness analysis as explanatory reasoning.” *The British Journal for the Philosophy of Science* 69 (1):275–300. https://doi.org/10.1093/bjps/axw008.

Skewes, Joshua. C., and Cliff. A. Hooker. 2009. “Bio-agency and the problem of action.” *Biology & Philosophy* 24 (3):283–300. https://doi.org/10.1007/s10539-008-9135-9.

Soto, Ana M., Giuseppe Longo, Paul-Antoine Miquel, Maël Montevil, Matteo Mossio, Nicole Perret, Arnaud Pocheville, and Carlos Sonnenschein. 2016. “Toward a theory of organisms: Three founding principles in search of a useful integration.” *Progress in Biophysics and Molecular Biology*, 122 (1):77–82. https://doi.org/10.1016/j.pbiomolbio.2016.07.006.

Soto, Ana M., and Carlos Sonnenschein. 2011. “The tissue organization field theory of cancer: A testable replacement for the somatic mutation theory.” *BioEssays* 33 (5):332–40. https://doi.org/10.1002/bies.201100025.

Thompson, Evan. 2004. “Life and mind: From autopoiesis to neurophenomenology. A tribute to Francisco Varela.” *Phenomenology and the Cognitive Sciences* 3 (4):381–98. https://doi.org/10.1023/B:PHEN.0000048936.73339.dd.

Thompson, Paul. 1989. *The Structure of Biological Theories*. New York: State University of New York Press.

Varela, Francisco J. 1979. *Principles of Biological Autonomy*. New York: North Holland.

Weber, Andreas, and Francisco J. Varela. 2002. “Life after Kant: Natural purposes and the autopoietic foundations of biological individuality.” *Phenomenology and the Cognitive Sciences* 1 (2):97–125. https://doi.org/10.1023/A:1020368120174.

Weisberg, Michael. 2006. “Robustness analysis.” *Philosophy of Science* 73 (5):730–42. https://doi.org/10.1086/518628.

Weisberg, Michael. 2007. “Who is a modeler?” *The British Journal for the Philosophy of Science* 58 (2):207–33. https://doi.org/10.1093/bjps/axm011.

Wolfe, Charles T. 2010. “Do organisms have an ontological status?” *History and Philosophy of the Life Sciences* 32 (2/3):195–231.

---

**Cite this article:** Rider, Campbell. 2025. “Closure of Constraints as a Theoretical Model.” *Philosophy of Science*. https://doi.org/10.1017/psa.2024.63