

This is a Pre-Print Draft – Please Do Not Quote This Version

The Ontology of Organisms: Mechanistic Modules or Patterned Processes?

Christopher J. Austin

Abstract

Though the realm of biology has long been under the philosophical rule of the mechanistic magisterium, recent years have seen a surprisingly steady rise in the usurping prowess of process ontology. According to its proponents, theoretical advances in the contemporary science of evo-devo have afforded that ontology a particularly powerful claim to the throne: in that increasingly empirically confirmed discipline, emergently autonomous, higher-order entities are the reigning explanantia. If we are to accept the election of evo-devo as our best conceptualisation of the biological realm with metaphysical rigour, must we depose our mechanistic ontology for failing to properly “carve at the joints” of organisms? In this paper, I challenge the legitimacy of that claim: not only can the theoretical benefits offered by a process ontology be had without it, they cannot be sufficiently grounded without the metaphysical underpinning of the very mechanisms which processes purport to replace. The biological realm, I argue, remains one best understood as under the governance of mechanistic principles.

There’s no doubt that one of the most trending topics in the philosophy of science is the so-called ‘new mechanism’ movement. In the philosophy of biology in particular, the movement is truly a metaphysics *en vogue*: it represents a conceptual schema which appears to more than adequately capture the framework within which a wide variety of empirical data is commonly interpreted, and within which the experimental practitioners of that field carry out their work. According to the new mechanists, the biological realm is a *mechanical* realm, and its denizens – organisms – are machines *par excellence*. And although it’s undeniably the case that biology is a science of mechanisms in one sense or another, the pertinent question at hand is whether and to what extent the particulars of this now popular ontology properly carve at the same joints that our best contemporary biological models do.

If we’re going to answer that question, a plausible place to do so is within the conceptual remit of evolutionary developmental biology (evo-devo), a research programme whose fruit has been the reliable delineation of the various ontogenically and evolutionarily salient modular sub-systems which compose the meta-systems we recognise as organisms. What we want to know then is whether the ontology of evo-devo is an ontology of mechanisms – that is, whether our best model of the composition of organisms is one capable of being constructed mechanistically. One would think that, given the past successes of a broadly mechanistic characterisation of the biological realm, this is a question that is likely to cause very few any pause – but there has rather recently arisen a dissenting voice claiming that the ontology of organisms evo-devo presents us with is not – and indeed, cannot be – one of “entities and activities”, but is rather one consisting of activities *alone*.

According to ‘process ontology’, the familiar entities which our scientific theories describe and quantify over are nothing more than particularly stable processes whose temporal persistence (and our subsequent ability to track them) consists in their repeated patterned-based behaviour. Of course, no one – not even those who favour an ontology of mechanisms – will deny the importance of characterising certain classes of biological phenomena in terms of processes, but it is important to note that adopting a

process ontology amounts to a much more robust (and radical) claim about the nature of organisms: on this ontology, organisms *just are* a set of interrelated, ever-fluctuating processes, and the parts and pieces of the mechanisms we commonly individuate within them are *nothing more* than the interlocking instantiations of particular patterns of activity. For its proponents, this revisionary ontology is no metaphysical extravagance – its application to our conceptual understanding of the nature of organisms is in fact demanded by the framework and findings of evo-devo: a mechanistic ontology simply cannot adequately model the characteristic features of an ontology centred on the systems we now understand to primarily construct organisms.

Choosing an ontology that best captures the concept of organism that contemporary evo-devo research presents us with is a timely and necessary task for the philosophy of biology. In this paper, I examine the intricacies of the conflict between a mechanistic and a processual ontology in the context of that decision. To do so, I first explicate the current ontological landscape of evo-devo, focusing on the principal players – developmental modules composed of highly integrated genetic regulatory networks, responsible for the generative specificity of phenotypic development *via* homologue determination. I draw special attention to a unique feature of these modular systems – namely, their *generative robustness* – and discuss how the concept it appears to entail – namely, *higher-order multiple realisability* – may not only sit uneasily within a mechanistic schema, but lend itself to a proper categorisation within a process ontology. By specifically detailing the causal-*cum*-explanatory structure of these central, “organism-building” genetic regulatory networks with respect to ontogenesis, I argue that even if one is willing to pay the revisionist price of building an organism from the materials afforded by a processual ontology, it is nonetheless an unnecessary indulgence.

Why Might Biology Require a Process Ontology?

The main motivation, as far as I am able to discern, for thinking that contemporary biology requires a process ontology is borne from taking seriously certain lessons from evolutionary developmental biology (evo-devo). But what sorts of lessons are these, and what consequences have they for our ontology? As John Dupré, an outspoken advocate of a biological process ontology, and Eric Baptiste rather simply put it: “...as [evo-devo] has emphasised, an organism is a developmental process...[and] as evolution is uncontroversially a process, an evolutionary ontology will quite naturally be processual...[an] evolutionary ontology of the living world *should distinguish the real evolutionary players, the units with causal powers resulting from or contributing to evolutionary processes*”.¹ I take it the claim is that if we properly understand the nature of these centrally important “units”, we will understand the utility, and perhaps the necessity, of conceptualising them according to a process ontology. What then are these “real players”, according to evo-devo?

Plausibly, the units which occupy the main-stage of contemporary evo-devo ontology are so-called *developmental modules*, discrete sub-systems which are responsible for the specified development of a particular morphological structure in a developing organism. These modular sub-systems might be considered central (in the aforementioned sense) to an evo-devo ontology on account of their being ontogenically explanatory with respect to the most contemporarily important “units” of evolution – *homologues*, discrete morphological features whose broad phenotypic similarity among their various instances is underwritten by shared structural-*cum*-causal developmental mechanisms which exhibit a traceable phylogenetic lineage.² For if the evolutionary process can be conceived as the successive propagation and progression of homologue variation and canalisation, then these highly integrated genetic regulatory networks (**GRNs**) embedded in the “bottleneck” of ontogenesis (Galis & Metz 2001; Kalinka *et*

¹ Baptiste & Dupré (2013: 380-381); my emphasis.

² In order to retain sufficient generality, this definition encapsulates both the *phylogenetic* and *developmental* concepts of ‘homologue’.

al. 2010), in virtue of their exerting downstream spatial and temporal regulatory control *via* the production of transcription factors whose patterns of expression causally specify the particularised developmental pathways of those morphological structures, are surely prime candidates for being the “real players” in an adequate evo-devo ontology (Brandon 1999; Brigandt 2007; McCune & Schimenti 2012; Wagner 2014).

Importantly, it has become increasingly common to treat these homologue-specifying developmental modules as “higher-order” entities, defined functionally *via* their morphological end-states (that is, the homologues whose development they causally control). For the explanatorily fruitful research programme of uncovering the specificities of the regulatory architecture underlying various particular homologues which has illuminated our understanding of molecular-based phylogeny (Carroll, Grenier, and Weatherbee 2001; Wilkins 2002) has at the same time presented us with a view of homologues which are importantly *dissociated*, or *autonomous* with respect to their underlying networks (Müller 2003; Wagner 2014). This is due to the phenomena of *robustness*, now widely acknowledged to be ubiquitous in the biological realm (Greenspan 2001; Mason 2010). Because these modules’ production of their associated morphological structures represents a highly *robust* process, one underwritten by their regulatory architecture consisting of both redundant network motifs and degenerately functional elements, they are able to maintain their generative competence with respect to that structure in the face of intra-species mutational variations of their component elements and epigenetic variations on their regulatory structure (Davidson 2001; Carroll 2008; Wagner & Lynch 2010).³

The ability to maintain their generative integrity throughout the developmental process in the face of genetic and regulatory variation has resulted in the *canalisation* of these modules in such a way that, over time, and in successive generations, they have gained a kind of independence from their (original) underlying genetic constituents and their accompanying regulatory structure (Müller & Newman 1999; Müller 2003)⁴. The now classic paragon of this is Owen’s (1848) own – that of the *tetrapod limb* which, although performing a variety of distinct functions in its multiple instances throughout evolutionary history, is readily identifiable in innumerable species despite its specific generative competency being developmentally anchored in a wide variety of diverse underlying GRNs throughout those instances (Zuniga 2015).⁵ Due to the co-variational disconnect between the mechanistic composition and the morphological product of these modules, it has been increasingly conceptually advantageous to *functionally* individuate them with respect to their generatively specific end-states – that is, according to their generative capacities in establishing structurally specific morphospaces, the varied permutations of which represent species-specific instances of their respective homologues (Rieppel 2005; Brigandt 2007; Love 2009; Wagner 2014). Individuating the modules which are causally responsible for homologues *functionally* allows them to feature in higher-order explanations of developmental phenomena, ones which operate at a higher “causality horizon” (Salazar-Ciudad & Jernvall 2013), or at explanatory levels “above” the workings

³ It may even be the case, as has recently been argued, that the generative robustness inherent in developmental systems is in fact a *necessary* requirement for their ability to evolve. See Edelman & Gally (2001), and Whitacre & Bender (2010).

⁴ As Rosenberg (2001) argues, if the process of natural selection is in a certain sense “blind” to *structure*, operating instead primarily on *function*, it’s not at all surprising that one and the same functional homologue might, over time, and in successive generations, be underpinned by distinct sets of GRNs (that is, given mutational robustness, and the accumulation of cryptic variation, etc.)

⁵ Recently, there’s been a surge of compelling evidence that homologues in a wide range of taxa might be principally underwritten by shared, “core” GRN elements – see Davidson & Erwin’s (2006) ‘kernel’ concept, and Wagner’s (2014) concept of ‘character identity networks’. It’s important to note that even if a particular homologue does share some central GRN elements in all of its instances, the generative capacity of each instance to produce its specific variation on that homologous structure must also be grounded in that GRN being highly integrated with both upstream and downstream regulatory cassettes, and these elements will be substantially variable across these instances.

of their molecular constituents – indeed, their ability to function at this level is the *sine quo non* of the now prominent explanatory project of *dynamical systems theory* (which I discuss in more detail later).

So, at first blush, the “real players” of an ontology which has any hope of being empirically adequate with respect to evo-devo are, as it turns out, of rather an odd breed. Due to the phenomenon of robustness, and thus their existential persistence throughout various alterations in their constitutive regulatory structures, these modules are commonly understood as *emergent* – that is, as novel, “higher-order” entities whose unique features ensure that they are incapable of being either ontologically or explanatorily reducible to any particular set of constitutive, “lower-level” entities (Wimsatt 2000; Callebaut *et al.* 2007; Mitchell 2012; Walsh 2013; Brigandt 2015). Typically, emergent entities are understood as resisting these forms of reduction due to their possession of novel (often downwardly directed) causal-*cum*-explanatory power which is neither had by, nor combinatorially attainable by the linear composition of their lower-level constituents (Andersen *et al.* 2000; Ellis *et al.* 2012).

And because the persistence conditions and particular explanatory prowess of these modules are unfettered, as it were, from any particular molecular mooring, they are also understood as being *multiply realisable*: capable of being “realised” by a number of distinct underlying structures, but incapable of being strictly *identified* with any particular such structure. Given that the set of distinct regulatory architectures which are capable of performing the higher-order, functional role (of producing a particular developmental end-state) which metaphysically individuates a module is a heterogeneous collection (over developmental, and eventually evolutionary time-scales), although each member of that set may *realise* the *same* module, the latter cannot be ontologically identified with the former. In this way, these modules are seemingly best understood – ontologically – as being “...supported by [various] biological components, but not *composed* of them” (Cahoone 2013: 141; my emphasis), where ‘composition’ is understood as constitutive of metaphysical *identity*.⁶

These strikingly peculiar features of developmental modules are, as I understand it, the ones which the process theorist takes to be uncapturable by, and incompatible with an ontology of mechanisms. The short of the argument is: if *these* are the sorts of entities which are going to occupy the centre stage of our ontology, we must abandon *mechanisms* and embrace *processes*. Let us see why.

Why Mechanisms Won’t Do

Evo-devo is routinely heralded as a properly ‘mechanistic science’, and for an important reason: while population genetics has focused primarily on mapping the dynamics of allele frequencies over evolutionary time-scales, the research programme of evo-devo consists in the attempt to lay bare the local, organism-level developmental systems which causally underwrite that topology (Wagner *et al.* 2000; Hall 2003; Canestro *et al.* 2007; Laubichler 2010). That being said, whether the scientific community regards the latter discipline as properly *mechanistic* is of course irrelevant to the question of whether its requisite ontology consists, as a matter of fact, of mechanisms *proper*. In order to answer that second question, we require a firm grip on the concept of ‘mechanism’, and so we ought to look to how that concept is understood in the contemporary philosophical literature.

There are quite a few fairly nuanced definitions of ‘mechanism’ in the ever-growing literature on the topic, but there is a discernible common core. On a basic level, a mechanism can be defined as “a set of entities and activities organised such that they exhibit the phenomenon to be explained” (Craver 2007:

⁶ The description just offered will no doubt strike a fair few as a fundamentally unfair reading of the ontological picture painted by the data of evo-devo – and, to my mind, it certainly is. However, in taking a naïve view of that picture, the aim of this part of the paper is to understand the *prima facie*, conceptual motivations one might have for finding the process theory attractive (or the mechanistic picture unattractive). The critique of this naïve view, in defence of the mechanistic ontology, is the focus of the rest of the paper.

5), where ‘exhibit’ is understood as those entities and activities being “productive of regular changes from start or set-up conditions to finish or termination conditions” (Machamer, Darden, and Craver 2000: 3). A more sophisticated rendering may define a mechanism as “a structure performing a function in virtue of its component parts, component operations, and their organisation....responsible for one or more phenomenon” (Bechtel & Abrahamsen 2005: 423), where ‘organisation’ is understood as the component parts being “appropriately located, structured, and oriented, and the activities in which they engage [having] a temporal order, rate, and duration” (Machamer, Darden, and Craver 2000: 3). This is enough to give us the fundamental picture: a mechanistic ontology consists, simply enough, of *entities* and *activities* – and a particular ‘mechanism’ consists of a specific set of spatio-temporally arranged entities (read: parts) which are connected by a specific set of activities (read: causal operations). Importantly then, mechanisms are ontologically *defined/individuated* by their unique four-fold structure – that is, their (1) type and number of entities and (2) their spatial organisation, and their (3) type and number of connective activities and (4) their spatial and temporal organisation. Thus, *what it is to be* a particular mechanism is to be an instance of a specific four-fold structure: distinct permutations of the values of (1) – (4) constitute distinct mechanisms, and so in any particular case the alteration of any of those values amounts to the effective dissolution of that mechanism.

Describing the GRNs which underwrite developmental modules *via* a mechanistic ontology no doubt seems rather natural. After all, we know quite a lot about the types of entities which compose regulatory networks and their typical organisational structures, as well as their typical causal connectives (motifs) and the “logic” of those connectives, etc., and rather complex schematic mappings of these entities and their activities, representing the end product of a great deal of intricate research, are ten a penny in the relevant scientific literature.⁷ Why then ought we think that our mechanistic ontology might not be up to the task of adequately modelling such developmental systems? One *prima facie*, quite general worry that’s been raised is whether, given the overwhelming integrative complexity of the GRNs underwriting the relevant developmental systems, we can properly individuate the *parts* of, and subsequently demarcate the *boundaries* of those networks. Even when aided by the conceptually powerful ‘manipulation theory’ (Woodward 2002; Craver 2007), detailed precisifications of the concept of ‘part-hood’ and ‘modularity’ have arguably been fairly unsuccessful in this context: there are *tant* causal knots here, not easily untied by counterfactual discrimination (McManus 2012).

Important as these foundational issues are, I’m going to set them aside for now – let’s assume they can be adequately answered.⁸ Even still, there is a perhaps more specifically important issue at hand with respect to the focus of this paper, already hinted at in the discussion of the previous section, and at the heart of contemporary objections to a mechanistic ontology: the ontological consequences of *emergence* and *multiple realizability*.⁹ Recall that these phenomena arise from the *robustness* of developmental systems – that is, their ability to maintain their generative integrity with respect to their associated end-state amidst mutational and regulatory perturbations. Typically, the regulatory networks which compose these systems exhibit this homeostatic behaviour either in virtue of their possessing a number of ‘redundant’ elements

⁷ See Alon (2006) for an excellent overview of the “regulatory logic” used to construct models of these networks.

⁸ As far as I’m concerned, these difficulties are being adequately addressed outside of a strictly conceptual, philosophical account, within empirical studies on ‘modularity’ in developmental systems. There, a sub-system is considered ‘modular’ just in case it satisfies certain criteria on regulatory connectivity. This type of modularity is no doubt central to the evolutionary process (Callebaut & Rasskin-Gutman 2005), and recent evidence with respect to the processes involved in shaping the *developmental hourglass* lends credence to this conception (Raff 1996, Galis & Metz 2001, Kalinka *et al.* 2010).

⁹ For specific critiques of a mechanistic ontology which stem the homeostatic properties of developmental systems, see Woese (2004), McManus (2012), and Dupré (2013). It’s interesting to note that even Woodward (2013), a card-carrying mechanist, thinks that the phenomenon of robustness (in an important class of cases) lies outside the explanatory remit of a mechanism ontology.

(gained perhaps through various duplication events) which can take-up the slack of missing/mutated/disabled elements or, more intriguingly perhaps, in virtue of their ability to “re-wire” their regulatory architecture such that its non-isomorphic elements are able to become isofunctional, causally mirroring the required role within the perturbed network - this is the phenomenon of *degeneracy*.

The quandary for an ontology of mechanism should be fairly plain: due to the generative robustness of the developmental systems in question, and their subsequent ability to existentially *persist* through substantive alterations in their constitutive elements, their metaphysical individuation *via* the entities and activities which comprise their aforementioned four-fold structure looks at best problematic, and at worst impossible; note that this isn’t merely an epistemological worry – in order for it to be the case that mechanistic structures truly “carve at the joints” of the world, they must do so *at* that four-fold structure. In developmental systems whose robustness is derived from some measure of redundancy, we are confronted with systems whose parts (read: coding genes, proteins, etc.) are potentially transient, able to be removed or functionally disabled, with novel instances of functional connections arising from their replacements (Zhenglong *et al.* 2003; MacNeil & Walhout 2011). In other words, these are systems whose existential persistence isn’t perturbed by their *loss of entities* – a fact which clearly does not sit well with the individuation requirements central to a mechanistic ontology wherein a change in the four-fold structure of any particular mechanisms is tantamount to the dissolution of that mechanism. In the more extreme case, where developmental systems’ robustness is derived from their degenerative capacities, we are confronted with systems whose parts *and* causal connectives are transient, where one part capable of performing multiple functional roles can causally correlate downstream activity towards an end-state in a novel fashion (Edelman & Gally 2001; Mason 2010). In other words, these are systems whose existential persistence isn’t perturbed by their *loss of entities* and their *acquisition of activities* – another, perhaps more worrying instance where the ontological individuation requirements of a mechanistic ontology seem ill-fitted to capture the nature of these systems.

It’s easy to see why, if we want to be able to capture the fact that developmental systems undergo, and persist through, the homeostatic activity that typifies the phenomenon of robustness (as indeed we *must*, if we have ontological aims), characterising these systems as *emergent*, and *multiply realisable* is both natural and appealing: they can be individuated by a definable higher-order activity – namely, the specified production of their associated morphological structures – which remains invariant (over developmental, and eventually evolutionary time-scales) across a large set of substitutionary permutations in a series of particularised instances of a four-fold mechanistic structure. But while this characterisation is certainly a boon for theoretical simplicity, the facts upon which it is grounded are at the same time the bane of a mechanistic ontology: the robustness, and thus multiple realisability of developmental modules makes any hope of metaphysically individuating, and therefore defining the persistence conditions of these systems with respect to a single and stable four-fold structure appear fundamentally misplaced.¹⁰

In this way, the dynamic, homeostatic capacities of developmental systems – and their accompanying ontological consequences - lie at the heart of contemporary critiques of the application of a mechanistic ontology in the biological realm. Thus the eminent Carl Woese urged:

Let's stop looking at the organism purely as a molecular machine. The machine metaphor certainly provides insights, but these come at the price of overlooking much of what biology is...Machines are stable and accurate because they are designed and built to be so. The stability of an organism lies in resilience, the homeostatic capacity to reestablish itself. While a machine is a mere collection of parts, some sort of ‘sense of the whole’ inheres in the

¹⁰ As Dupré (2013: 27) puts it, the worry is that, for many cases “...there is no unique and definitive sequence of molecular events” by which these mechanisms might be individuated.

organism, a quality that becomes particularly apparent in phenomena such as regeneration in amphibians and...in the homeorhesis exhibited by developing embryos (2004: 176)

The question is, however: if we are convinced that we cannot look at the developmental systems which compose organisms as machines – how else ought we to view them? One sharply contrasting perspective which has been increasingly favoured in the context of these of questions is afforded by *process ontology*, wherein *entities* are eschewed in favour of *activities*: according to this ontology, “[w]hat are stable and robust in biology are not things, but processes” (Dupré 2013: 30). For the process theorist, the denizens of the biological world are fundamentally a set of interrelated, ever-fluctuating though meta-stabilised activities, and the entity-like parts and pieces of the developmental systems which compose them are themselves *nothing more* than the interlocking instantiations of further stabilised patterns of activity. The *stability* and *discreteness* which we have historically associated with the mereological constituents of modular developmental systems is, on this view, little more than an abstraction, affording only an incomplete view of “...a particular time slice from [the] developmental process” (Baptiste & Dupré (2013: 380-381) - to think otherwise is to commit what Whitehead (1925), the figurehead of modern process ontologies, colourfully called “the fallacy of misplaced concreteness”. In other words, the ‘entities’ which populate the mechanistic ontology – be they parts of a mechanism, or whole mechanisms – are nothing more than particularly stable processes whose temporal persistence (and our subsequent ability to track them) consists in their repeated, patterned-based dynamical behaviour.

For the process theorist then, what ontologically *persists* (and thus what truly *exists*) in the biological realm are higher-order *patterns of activity*, and the underlying flux of seemingly *static entities* which variously instantiate these patterns are only conceptual artefacts which, while perhaps heuristically powerful, are unable to properly carve that realm at its joints. Accordingly, on this perspective, to capture the *dynamical flow* of a system is to capture its *essence*, defined by the “...dynamical interactions among its constituents, not the constituents *per se*” (Jaeger & Monk 2015). The process theorist will, for instance, grant ontological primacy to developmental *pathways*, rather than to any specific configurations of entities which compose those pathways in any particular case, or at any particular time (Gilbert & Bolker 2001): on their view, it is the dynamic features *of the pathway* – the spatio-temporally ordered, continuously connected stages of its causal structure¹¹ – which accurately capture the ontological facets of biological objects, rather than any static set of entities in a particular four-fold configuration which might enter into, or satisfy that structure.¹²

Even from this rather simple picture of *what it is to be* a ‘process’, it’s not difficult to see why one might find this ontology to be an attractive framework for conceptualising the biological realm – but here I want to focus on two reasons which seem to be the most compelling, and which have subsequently received the most attention in the literature. The first is relatively straightforward, given the previous discussion: a process ontology can easily accommodate the *emergent* nature of developmental modules. For note that processes, being higher-order, *multiply realisable* entities, are neither defined by nor existentially dependent upon the stability of any particular “lower-level” structures – indeed, the fact that those structures are in constant flux is a simple consequence of the ontology. On that ontology, *stability* is located at a higher-order: the parts and pieces below that level may come and go as they please, so long as they collectively instantiate a particular higher-order dynamic ‘pattern of activity’. In this way, it’s clear that

¹¹ These ‘pathway features’ are the comparative basis of so-called *developmental homologies*; see Gilbert & Bolker (2001), Rosa & Etxeberria (2011), and Nathan and Borghini (2013) for recent discussions of the conceptual distinctions between phylogenetic and developmental homology.

¹² One might also consider the higher-order, patterned-based nature of Goodwin’s (Webster & Goodwin 1996) ‘morphogenetic fields’ as another instructive example. See Levin (2012) for some interesting recent work inspired by this approach.

one and the same process is capable of existentially persisting throughout (lower-level) structural variation. Thus, if developmental modules are conceived as processes, the phenomenon of generative robustness is no longer a threat to their persistence: as their ontological continuance will consist solely in their producing (or their capability to produce) a particular morphological feature, variation in precisely *how* that role is performed – that is, variation in which specific four-fold configuration satisfies that dynamic structure – over developmental (and eventually, evolutionary) time-scales is of little existential consequence.

The second reason is more complicated, and perhaps more compelling. Plausibly, if the biological world *is* correctly carved-up into processes, we should expect these higher-order entities to possess an accompanying higher-order explanatory prowess, being causally capable in a fashion that is independent and irrespective of their lower-level constituents. In other words, in line with the widely adopted ‘Eleatic Principle’, if we’ve no evidence of emergent, higher-order causal competence, we’ve no evidence of a process ontology.¹³ As it happens, such competence is precisely what we find in the models of cutting-edge evo-devo research, where the explanatory burden with respect to the development of particular morphological traits rests upon a high “causality horizon”, operating above the particularities of any specific lower-level network-regulatory configurations. In particular, this is what we find in *dynamic systems theory* (DST), a now prominent research project.¹⁴

Rather than attempting to model the morphological development of a particular module *via* the conceptual decomposition of its GRN into a precisely detailed explication of its every elemental constituent and their variously interconnected causal intricacies, the guiding methodology of DST is to take a more *holistic* approach. Instead of conceptualising a module’s development as its step-wise causal progression through a series of discrete state-changes in a collection of isolatable elements, DST views the developmental process as a continuous series of temporally successive transitions between *entire system states*. On this perspective, one can model the development of a module as the tracing of a temporal pathway through an abstract *state space* – a multi-dimensional map whose coordinate points represent possible whole system states, or the complete ‘genetic expression profile’ of the entire imaginal disc.

Importantly, because the *transitions* between states within that space are governed by the regulatory architecture of that module in virtue of it determining how the expression of particular sets of the module’s constitutive genetic elements at one time affect other sets at later times (*via* enhancing or repressing their transcription rates, for instance), we can assign to each state a relative *stability* value – e.g. a value representing the likelihood the system will shift to another state within that space (Kim & Wang 2007; Bhattacharya *et al.* 2011). If we then represent the stability of each state of the system within that space along another, additional dimension by assigning it an *elevation* value (where a *higher* elevation represents a *higher* level of regulatory instability), our abstract state space becomes a structured topology, complete with peaked hills and low-lying valleys. Utilising this abstract three-dimensional representation of the collection of possible system-wide states of a module, DST models the morphological development of these systems as a kind of kinetically-constrained temporal traversal across the topological curvatures of this epigenetic landscape, from the high peaks of the system-wide instability of its initial conditions “downward” toward a state of regulatory stability (Wang *et al.* 2011; Huang 2012; Davila-Velderrain *et al.* 2015).¹⁵ More picturesquely, on DST the state of a developmental system at any time is conceptualised as a frictionless orb, and the temporal succession of various distinct states of that system throughout its morphological development is modelled as the dynamic trajectory of that orb through a pathway

¹³ A principle derived from Plato’s *Sophist*, but more recently articulated by Armstrong (1997)

¹⁴ There are now a number of specialist journals which focus on holistic treatments of developmental phenomena – see, for instance, *Molecular Systems Biology* and *BMC Systems Biology*.

¹⁵ Contemporary forms of DST represent the theoretical union of Waddington’s (1957) geometrical-dynamical models and Kauffman’s (1969) Boolean network configurations.

geometrically constrained by the topological ridges and valleys of the system's Boolean regulatory configuration.

Accordingly, in this framework a developmental system's morphological end-state is conceptualised as a particular collection of system-wide states of regulatory stability, and thus its area within state space functions as an *attractor* whose wide, low-lying *basin of attraction* dynamically constrains the system's various possible causal trajectories to follow a pathway toward those states within the sloping walls of its surrounding topology. Thus it is the *higher-order*, formal properties of a developmental system – that is, the *dynamic flow* of the system's topology, represented by the vectorisation of its state space and defined by its regulatory architecture of Boolean network connectives – which are explanatory with respect to *why* that system reaches a particular end-state: the dynamics of attractor states (and their metric-bending basins) possess the relevant causal-*cum*-explanatory power.¹⁶ In contrast, the system's *non-formal* properties – that is, the other characteristics of whatever it is to which those Boolean values belong (read: the underlying “entities and activities”) which *secure* the stability measures of this topology – are explanatorily irrelevant, and seemingly causally impotent: the flow of causation in the process of development is directed solely by the contours of their collective phase space, its path effectively oblivious to any lower-order, non-geometric features.

This is of course a gross simplification of the rather complex formalism of DST, but it adequately illustrates the relevant point – namely, that developmental systems exhibit emergent, higher-order causal competence, precisely as predicted by (and required for) a process ontology. Furthermore, understanding the nature of that competence according to DST lays bare, in a more formal fashion, the *multiple realisability* afforded by a process ontology: one and the same process, defined by a “pattern of activity”, or a directedness toward a particular end-state *via* a specifically contoured dynamic landscape, can be *realised* by any number of distinct sets of underlying constituents which instantiate that topology (Gilbert & Bolker 2001; Dupré 2013; Jaeger & Monk 2015). All of this suggests that in order for our ontology to correctly and fully capture the developmental capacities, and perhaps even the existential persistence of these modules, it requires the inclusion of multiple realisable, uniquely causally potent processes.¹⁷ The resources of a mechanistic ontology, wherein a system's causal activity is constituted merely by the ordered orchestrations of miniscule manipulations in an unchanging set of inflexible clockwork cogs, just don't look up to the task.

Mechanisms, Emergence, and Explanation

It's clear that although a process ontology undoubtedly represents a picture of the world which is scientifically unorthodox, it isn't empirically unmotivated – it may be theoretically costly, in that its framework demands that our fundamental categories be substantially restructured, but perhaps, if we wish to truly understand the empirical data of contemporary biological science, that's simply the cost of admission. Be that as it may, it's wise to exercise caution before investing in such revisionary stock and in what follows, I make a case for that caution. On my view, adopting a process ontology amounts to buying in to an inflated market – the theoretical benefits, I argue, have been overvalued.

Importantly, I don't want to argue that the benefits gained from higher-order causal modelling of the sort provided by DST aren't worth paying for: as far as I'm concerned, these models afford us genuine and wide-reaching explanatory power with respect to everything from sub-organismal cell-fate

¹⁶ Note that this is true even in the case of *non-autonomous* systems (Corson & Siggia 2012) where that topology changes over developmental time.

¹⁷ That process ontology and DST are natural bedfellows has been pointed out before. Waddington (1969) himself professed to being deeply influenced by Whitehead, the founder of modern process ontology, as Gilbert & Bolker (2001) note. More recently, Hall (2013) has characterised the contemporary, mathematical models of DST as having a natural home within a Whiteheadian ontology.

(Bhattacharya *et al.* 2011; Verd *et al.* 2014) to the evolvability of organism populations (Streider 1998; Jaeger & Monk 2014).¹⁸ What I do want to argue however is that these are benefits that our current mechanistic ontology already affords us, and at no extra charge. In the literature, one often encounters just the opposite view functioning as an implicit premise in an argument for process ontology: if we want to account for the higher-order nature of developmental systems (and we certainly do, and must), then we have no choice but to pay the price in adopting an ontology of processes.¹⁹ However, my argument is that this implicit transcendental claim – namely, that a process ontology functions as a *necessary condition for the possibility* of the existence of higher-order phenomena – is simply false. That claim is false for a simple, but powerful reason: the *true* transcendental claim, so far as we know or have any reliable evidence of, concerns an ontology of mechanisms. For although the claim that a mechanistic ontology functions as the *logically* necessary condition for that possibility is certainly too strong, that ontology is, at the very least, *one* such condition – and importantly, it is one whose transcendental role can be, and has been laid bare by empirically testable, predictive models.

In order to see this, consider first the case of the higher-order explanatory power of developmental systems as captured by the models of DST. My claim is that the causal competence exhibited in those models is *mechanistically grounded*, and thus is *mechanistically explicable*. Importantly, this is *not* the claim that the explanatory prowess of those models is illusory – no, those models *are* genuinely explanatory. Nor is it the claim that the causal structure inherent in these models is *mechanistic* – no, the activity captured by those models operates *without* “entities and activities”. But while the higher-order causal models of the sort offered by DST provide *genuine* explanations of the phenotypic products of developmental modules *without* the use of mechanisms, their ability to do so is yet underwritten by the *mechanistic constitution* of those modules. For though the explanatory weight of those models rests *solely* upon the dynamic pathway of the orb of development flowing down the curvatures in a landscape of various carved-out trajectories, the character of that landscape – the geometrical mapping which defines its topology – is shaped by the mechanistic architecture of an underlying GRN.

Note that, in DST modelling, a system’s dynamic landscape from which we derive the relevant higher-order explanatory power is constructed from two elements – an exhaustive collection of possible system states which defines its state space, and a surjective assignment of elevation values on that space. Capturing this landscape then first requires the definition of a specific state space, each unique point of which represents a distinct system-wide *expression profile* – that is, a quantifiable measure on the concentration and activity of a particular proteome (Wang *et al.* 2011; Huang 2012). Of course, the specification of this profile cannot be had without a corresponding specification of a particular set of causally responsible coding regions embedded in a particular genotype which are *expressed*. In this way, the collection of states in the state space which defines a particular developmental module *just is* a collection of all the possible configuration-outputs of the expression of a specific set of components in its constitutive GRN. Note then that this entails that our ability to have even a simple, one-dimensional plane upon which to subsequently place the orb of development in our dynamic model *depends* upon our correctly mapping the causal capabilities of a certain set of *mechanistic* resources – for surely the production of the proteome *via* the transcription and translation of coding regions is mechanistically explicable if anything is (Darden & Craver 2002).

¹⁸ There is a lively and interesting debate concerning whether higher-order, dynamic models can be *genuinely* explanatory – see Brigandt (2015) and Kaplan (2015) for opposing views. Even if one isn’t convinced that they *are* so, the point can be granted for the sake of argument.

¹⁹ See for instance Dupré’s (2013) argument that a process ontology is required for the phenomenon of robustness, Jaeger & Monk’s (2015) claim that a process ontology is the required to utilise the explanatory power of DST with respect to regulatory networks, or Gilbert & Bolker’s (2001) argument that a process ontology is required to understand the embryonic organising prowess of homologous structures.

Perhaps the process theorist will not find this point convincing, replying that we ought to focus on what it is that's *doing the higher-order causal work* – namely, the topological features of a dynamic quasi-potential energy landscape: it is here that the aforementioned underlying mechanisms are fundamentally ignored. This brings us to the aforementioned second requirement in the construction of our landscape – the assignment of an elevation value to each state within that landscape – as these values collectively determine the “flow” of the system which constitute its unique causal structure. As we have seen, the dynamics which ground the higher-order explanatory power of DST are represented in a mapping of measures of *stability* over a state space, and these measures are inversely reflected in the elevation values of each state within that space – the *higher* the elevation, the *less* stable the state.²⁰ Once we have a complete elevation assignment, we have a structured topology and a set of trajectories throughout that landscape which converge on a particular state: we have the system dynamics, and thus, its higher-order causal structure.

However, the important question for our purposes is: what *determines* the elevation value-mapping of such a system? For while no one will deny that all we need in order to give a higher-order causal explanation of a system's development is this topologically defining connective structure, there yet remains the further important question as to *what kind* of structure is required, and *what determines* that structure. An immediate answer is that this mapping represents a functional assignment of system *stability*. True enough – but what *determines* the stability of each system state? Again, an immediate answer is that the Boolean value of each state is determined by a function which specifies how a set of inputs (read: other state's values) are connected to a set of outputs (read: other state's values) – thus the *stability* of a state is a reflection of the causally connective structure of its system. Again, true enough – but what determines *that* structure?

In the case of biological systems, that structure is constructed from *regulatory* relations specified by the architecture of GRNs: it is this structure which *shapes* and *constrains* the dynamic evolution of system state values during development. In other words, the elevation values assigned to each state within a system are a reflection of its underlying network logic: *stability* is a function of the system's expression profile and a particular set of regulatory relations governing state-transitions. Importantly then, in any particular developmental system, these relations are determined by, and are therefore *ontologically dependent upon*, the character of its constitutive GRN which encodes a specific set of available protein-protein and gene-protein interactions, their directionality (constituting upstream and downstream expression control), and their interaction modalities (activation, repression, etc.) This is evidenced by the fact that mutational changes in both the protein-coding and *cis*-regulatory regions of a system's genome are capable of substantially modifying its dynamical landscape: these alterations effect both the *topology* (the number, kind, and relative placement of attractors and their associated basins) and the *geometry* (the position, shape and/or size of attractors and their associated basins) of that landscape (Kim & Wang 2007; Jaeger *et al.* 2012; Verd *et al.* 2014). Thus, the regulatory structure of a particular genome which determines the specific elevation mapping, and thus the higher-order dynamics of any particular developmental system is, in the end, mechanistically explicable: the regulatory logic governing the stability of its states is one whose content is cashed-out by the familiar and well-studied causal role of transcription factors at *cis*-regulatory regions – a role which has been quite naturally explicated *via* a mechanistic ontology.

This strongly suggests that the argument from ‘novel explanatory power’ to ‘novel entities’ doesn't work: a process ontology doesn't look to be a necessary condition for the possibility of the existence of higher-order explanatory power. The proponent of a mechanistic ontology then is not required to regard the higher-order explanatory prowess of the sort offered by DST modelling as illusory

²⁰ Elevation levels can also be conceptualised as (inversely) representing probability measures – see Huang (2009), and Zhou *et al.* (2012).

if she is unwilling to countenance the existence of higher-order entities to which that power belongs: taking the former seriously does not commit one to treating the latter likewise. Indeed, as we have seen, so far as we have any evidence of, assigning the proper theoretical weight to the former is a job which requires the conceptual resources afforded by a mechanistic ontology. For while the *holistic, dynamical explanations* of the causal activity of developmental systems offered by DST take place at a higher-level, abstracted away from any particular mechanistic underpinning, the specification of those dynamics *ontologically depends upon* the particular workings of a set of underlying mechanisms, and their role in transcendentally grounding the possibility of such higher-order explanatory power is one we now have a fairly firm empirical grip upon (Davila-Velderrain *et al.* 2015).²¹

But what of the other higher-order phenomenon associated with developmental modules – namely, their *multiple realisability*: if we are to account for this, must we adopt a process ontology? Recall that what we want to capture is the higher-order *stability* of those modules, one which reigns over and above the underlying flux of its part-like constituents. Due to their generative robustness, we require an ontology which can account for the *persistence* of those modules throughout substantial variation in their constitutive collection of entities and activities. While this phenomenon poses no problems for a process ontology, the same cannot be said for an ontology of mechanisms – for, as we have seen, mechanisms are ontologically individuated with respect to those collections, and precisely so: the aforementioned four-fold structure which defines any particular mechanism consists of a *specific* set of entities and activities.²²

How is the defender of a mechanistic ontology to respond? A popular reply to this worry is to insist that it simply misses the mark in virtue of the fact that the *static* conception of ‘mechanism’ it trades on is outmoded and should be abandoned: entities and activities are no doubt central to our contemporary concept of ‘mechanism’, but so is a particular *dynamism* among those elements, and the *correct* concept must reflect that (Bechtel & Abrahamsen 2010; Brigandt 2015; Kaplan 2015). Here the defender of a mechanistic ontology is essentially calling for a reform of her concept of ‘mechanism’. This isn’t a move without precedent – our ‘mechanism’ concept has been reformed in important ways before, after all. We no longer consider spatial linearity as constitutive of that concept, for instance – the entities and activities which constitute a mechanism need not follow a uni-directional connective causal pathway in producing their effect, but are acknowledged to bring about them about *via* a complex, branching web of interactions. The discovery of the importance of auto-regulatory elements and feedback loops has informed our ‘mechanism’ concept too: strict temporal linearity is no longer a requirement on a system qualifying as properly mechanistic. Perhaps it’s open then to the defender of a mechanistic ontology to further reform her concept of ‘mechanism’ to account for a certain fluidity in their four-fold structures in such a way that the persistence of a particular mechanism need not depend so rigidly upon its being constituted by a static set of entities performing a static set of activities.

²¹ One could hold a hybrid view according to which both mechanisms and processes are understood as ontologically robust, and in which no process is without some mechanistic configuration which operates as its metaphysical ground by specifying its higher-order dynamic profile. The argument just offered was not meant to take this sort of view off the table, but it may very well serve to undercut its motivation: if the conceptual resources of a mechanistic ontology are capable of performing this grounding role, and higher-order explanatory power need not be attributed to any higher-order entities, what do we stand to gain by positing the existence of (ontologically robust) processes? Thanks to an anonymous reviewer for raising this point.

²² It’s worth noting that it’s conceptually open to the mechanist to simply deny that any such stability exists, understanding it instead as a heuristic abstraction from the underlying metaphysical state of novel mechanisms popping in and out of existence whenever reconfiguration events take place. I don’t think this is an attractive view, for a variety of reasons, but notice that adopting it doesn’t require any novel ontology: even if *which* mechanisms underlie a particular phenomenon becomes a matter of relative indiscernibility, this doesn’t commit the ontology to countenancing *anything other than* mechanisms as the performers of that causal work.

Notably, Bechtel and Abrahamsen (2010: 323) have recently offered this sort of amendment: “a mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism, *manifested in patterns of change over time in properties of its parts and operations*, is responsible for one or more phenomena”. On this type of definition, a mechanism is not only individuated with respect to its entities and activities, but also with respect to its “orchestrated functioning”. The specification of this ‘orchestration’ will undoubtedly contain a certain set of entities and activities, though importantly accompanied by the ways in which those components might differ both compositionally and functionally over a certain time-scale.

However, even if this redefinition solves some of the mechanist’s conceptual difficulties, it’s easy to see why the defender of a process ontology might remain rather unimpressed: this conceptual reformation may appear to amount to little more than an *ad hoc* reshuffling of the goal-posts. Thus, in discussing the individuation of mechanisms in this fashion, Dupré (2013: 28) laments that “...it may well be possible to shoehorn descriptions of biological systems into talk of mechanisms if one is sufficiently determined...[but] there is a serious danger of vacuity...in which it seems that mechanisms are just whatever explains whatever happens”. Even if one doesn’t find the sort of definition given above conceptually empty, the point is nonetheless worth heeding: if we are to allow the mechanistic ontology such a central conceptual redefinition, it had better be on account of something more than its defenders’ desire to discharge a particularly difficult objection to it.

Plausibly though, in this case, there are independent reasons for permitting such a move. Note firstly that adequately modelling the development of organisms is a job which is going to require capturing the coming about/production of novel entities and activities – this is a simple consequence of the nature of ontogenesis. And because whatever it is that’s ultimately causally responsible for the production of an organism’s morphology must *itself* be formed and shaped *within* the process of development, any ontology attempting to capture that process must necessarily be flexible with respect to its “stock of entities” (and so, its “stock of activities”). With that in mind, a refusal to allow the mechanistic ontology a redefinition of the aforementioned sort looks akin to disallowing *ab initio* even the *possibility* of that ontology correctly modelling that process: this strikes me as being at the very least rather uncharitable. But perhaps a more compelling reason for permitting the mechanist’s proposed redefinition is that it accords with the ‘mechanism’ concept as it is utilised in contemporary biology; this is, after all, a stated goal of that ontology’s proponents (Darden 2007; Bechtel & Abrahamsen 2010; McManus 2012).

When one examines the literature, it’s a trivially easy task to find a large swathe of widely accepted and well understood *mechanistic* models of developmental processes which feature the sort of “flexibility” the redefinition calls for. There are mechanisms which feature transient entities and activities, such as those responsible for the development of distal digits which involve the *removal* of entities and their associated activities (Tickle 2003), as well as those which feature non-standard, specialised entities and activities, such as those responsible for cell-fate determination *via* morphogen gradients, where a certain *concentration* of entities – fleeting though its individual members at any time may be – are causally relevant (Tabata 2001). This being the case, it strikes me as rather difficult to accuse the defenders of a mechanistic ontology of an *ad hoc* revision of their central concept - this adjustment might plausibly be understood instead as a simple reflection of the tried and true methodology of contemporary scientific classification.

Still, I grant that some will view the argument for the dynamical redefinition of the concept of ‘mechanism’ from the practice of contemporary science as misguided, and that’s fair enough - we can’t always reliably infer ontology from methodology. There is however another, and I think more decisive reason that a more dynamic definition of ‘mechanism’ ought to be allowed and accepted, and it is this: the variational flexibility which typifies *robust* developmental systems is a phenomenon for which we already possess a plethora of richly explanatory *mechanistic* models. The phenomenon of developmental robustness is, after all, no conceptual outlier in contemporary biology – it is widely acknowledged within the discipline of evo-devo to play a central role in the evolutionary process (Edelman & Gally 2001; Kitano

2004; Whitacre 2010). It's no surprise then that the phenomenon has been extensively studied and comprehensively modelled *within* the prevailing mechanistic paradigm – a project which has borne considerable explanatory fruit, both at an abstracted level of generality and in a wide variety of specific cases.

We now know that the robustness of developmental modules can be understood in an abstracted sense as a function of certain thematic connective motifs in the “logic” of their constitutive GRNs – by means of their regulatory architecture featuring, for instance, ‘integral feedback control’, wherein particular downstream elements have feedback loops which send the time integral of system “errors” back into the upstream input elements (Yi *et al.* 2000; Alon 2007). That robustness can also be understood in a more specific sense as a function of certain complex causal roles of its constitutive network elements. One prominent mechanism which is understood to undergird the robustness of GRNs is the causal “cross-talk” established by pleiotropic transcription factors which functions as a redundancy “safety net” when system perturbations occur (Frankel *et al.* 2010; Whitacre & Bender 2010); rather tellingly, studies have shown that a system’s robustness is significantly compromised when the coding regions for these highly pleiotropic transcription factors are rendered non-functional (Wagner 2005). Another instance of this is exhibited by the multi-faceted control features of *cis*-regulatory elements where, for instance, two structurally non-isomorphic proteins might naturally perform a collective regulatory role for which neither is necessary, and either is sufficient (Edelman & Gally 2001; MacNeil & Walhout 2011). Without getting too far into any particular details, the important point is this: all of these various well-understood ways by which the robustness of developmental systems is secured are underwritten by the operational capacities of *mechanisms*, and are thus *mechanistically explicable*; we have a well-defined set of participating entities and their governing activities responsible for protein binding, *cis*-regulatory control, acyclical enhancer/repressor relations, etc.

The fact that the robustness of developmental systems is explicable within a mechanistic ontology is not trivial, remembering that, in the current context, to account for *robustness* is to account for *multiple realisability*. This is because, as we have seen, the stability of a system throughout alterations in its four-fold mechanistic structure which underlies the requirement that it be individuated at a “higher-level” is itself a reflection of that system’s generative robustness.²³ In this way then, the phenomenon of multiple realisability ought to pose no threat to an ontology of mechanisms, it being a phenomenon for which we have multiple mechanistic explanations – that is, explanations which posit and appeal to discrete collections of separable entities, their intrinsic causal capacities, and the interplay among them. That being the case, why *shouldn't* we permit the aforementioned dynamic redefinition of ‘mechanism’, given that the reason such a move has been deemed necessary is to account for a phenomenon which is *itself* mechanistically explicable?

Of course, it may be the case that the entities and activities responsible for a mechanism’s robustness are not included in any particular *model* of that mechanism, as those elements typically don’t directly contribute to the mechanism’s explanatory prowess with respect to its relevant explanandum (read: the production of a particular morphological feature), but that isn’t a reflection on its ontological make-up: there is not – nor should we expect there to be – a bijective mapping of ‘definitional elements’ from metaphysic to model.²⁴ Nevertheless, widening the ‘mechanism’ concept to reflect the dynamism entailed by the phenomenon of generative robustness inherent in developmental systems seems both well-motivated and necessary: any sufficiently precise mechanistic definition meant to capture the underlying *ontology* of those systems must include the “entities and activities” causally responsible for that

²³ Even on evolutionary time-scales, the higher-order stability of developmental systems *qua* the canalization of phenotypic traits is a long-term reflection of their mechanistically explicable network robustness. See Flatt (2005).

²⁴ For a more full discussion of the relationship between ‘ontology’ and ‘models’, see Bechtel & Abrahamsen (2005)

phenomenon. And in doing so, as we have seen, that definition will at the same time capture the ontological (and ultimately, mechanistic) basis for the higher-order persistence, and thus the multiple realisability of those systems.

There is a final plea to be made to those who remain yet unconvinced that the defenders of a mechanistic ontology ought to be allowed their redefinition on these grounds, and it is based on the following theoretical consideration: if whichever ontology we choose must be capable of capturing a particular phenomenon, and we have two competing ontologies which are able to do so, we ought to prefer the one which does not merely *weather* that phenomenon, but rather *forecasts* it.²⁵ For note that it is one thing for the phenomenon of multiple realisability to simply “fall out”, as a natural consequence, of an ontology, and quite another for an ontology to possess the conceptual resources to *explain* and *predict* the occurrence, empirical limitations, and general structural features of that phenomenon.²⁶ As we’ve already seen, one can easily account for the phenomenon of multiple realisability from within a process ontology, as its fundamental entities are *causally dynamic patterns* whose individuation is ontologically untethered from any patterned particulars.²⁷ In this respect, having that phenomenon functioning at the “ground floor” of that ontology is certainly a virtue: on a process ontology, where that phenomenon is taken as primitive, it is simply the “of the nature” of processes to be multiply realisable.

But in another respect, this primitiveness is a vice: for, in virtue of being ontologically *fundamental*, the character of that phenomenon is incapable of being given a satisfyingly robust *analysis*, or explication *by* that ontology. This is in sharp contrast to an ontology of mechanisms where, as we have just seen, there exist a great variety of models which explain *for which reasons* and *in which respects* developmental systems are multiply realisable, and subsequently predict *under which conditions* and *to which degrees* those systems are so. A mechanistic ontology then is capable not only of weathering that phenomenon, but also forecasting it: utilising its models allows one to get a firm explanatory grip on *how* and *why* developmental systems are multiply realisable.²⁸ *Prima facie*, for the aforementioned reasons, the same cannot be said for an ontology of processes.²⁹ Importantly however, even if it *can*, the onus is decidedly on the defender of that ontology to show that it has the conceptual resources to do so: accounting, in a suitably rich sense, for the multiple realisability of robust developmental systems is a theoretical burden on process ontology, not an ontology of mechanisms. In light of the fact that a mechanistic ontology possesses the conceptual resources required in order to offer a *properly explanatory* account of the system dynamics which constitute the phenomenon of generative robustness, and hence the phenomenon of multiple realisability, allowing an

²⁵ Of course, this preference is only relevant if it’s the case that the phenomenon in question *is* able to be given some further metaphysical analysis, *and* if only one of our competitors *is* capable of doing so. Below, I argue that both of these conditions are met.

²⁶ Cf. Mitchell’s (2012) excellent discussion of accounting for the phenomenon of ‘emergence’ in the physical sciences.

²⁷ This is true at least in principle. In practice, the only examples I know of that *explicitly* discuss the identity conditions of particular processes focus on comparing ‘homologous processes’ (in Gilbert & Bolker 2001, and Rosa & Etzeberria 2011) – but their homology seems to implicitly rely on a shared mechanistic underpinning. In Gilbert & Bolker (*ibid.*: 3-5) for instance, the ‘Wnt pathway’ is conceived as a “modular unit of process” which has many distinct instances, but each of these instances are acknowledged to feature homologous sets of genes, proteins, and a certain set of interactions between them: this looks like individuation *via* mechanistic constitution, as already discussed.

²⁸ For comparison, consider the analogous point which Dispositionists raise against their Humean opponents in the debate on the ‘laws of nature’: while regularity-based accounts get the *right laws*, they can’t *explain how or why* the *these* are the ‘laws’. See Mumford (2004) for an excellent in-depth discussion.

²⁹ It’s worth noting that there are DST models of robustness as well – see, for instance, Kitano (2004) and Huneman (2010). Though as I have already argued, even if they possess explanatory power in this respect, it is ultimately (ontologically) derived from the features of underlying operative mechanisms. For related discussions, see Kaplan & Craver (2011), Brigandt (2015), and Kaplan (2015).

appropriately dynamic definition of ‘mechanism’ ought to be understood as not only acceptable, but only natural.

A Few Final Thoughts

Throughout this paper I’ve argued that, contrary to claims of the process ontologists, a mechanistic ontology has the conceptual resources to not only accommodate, but importantly also to explicate the “higher-order” nature of the developmental modules which form the centre of the contemporary ontology of evo-devo. This is because, in short, we have conceptually plausible, empirically testable models according to which the phenomena of *higher-order explanatory prowess* and *multiple realisability* are grounded in the causal workings of sets of entities and their activities. Thus the implicit transcendental claim at the heart of the motivation for accepting a process ontology is false, and the cost of its revisionary metaphysics gratuitous: the theoretical benefits we gain from its purchase can be had by our current metaphysic, and at no extra cost.

As a final thought, one question worth asking is: why *has* there recently been a rise in the promoters of process ontology within the philosophy of biology? I suspect that a central reason has its origins in a muddling of the various senses of ‘mechanism’ floating around in the literature (Allen 2005; Nicholson 2012). Indeed, the contemporary dialectic against organisms *qua* ontologically composed of *mechanisms* appears to often be focused on objections to organisms *qua machines* (Woese 2004; Dupré & O’Malley 2007; Dupré 2013; Jaeger & Monk 2015). ‘Machines’, being exhaustively dissectible into sets of entities whose activities can be studied in isolation, are strongly associated with the philosophical project of *reductionism*, and the rejection of *holistic phenomena*: after all, their 17th century origin as the centrepiece of the scientific revolution directly pitted them against the irreducible and immaterial forces of *vitalism*. No one ought to dispute the claim that *machines* of this sort are unable to ground the holistic, emergent phenomena which we now know characterises much of the biological realm: they are ontologically outmoded in, and their associated reductionist programme outmatched by our contemporary science, itself actively engaged in studying and classifying these phenomena (Callebaut *et al.* 2007; Mason 2010).

But of course, no one does dispute this, least of all those who endorse an ontology of mechanisms – a point on which they have been quite clear (Darden 2007; Craver 2007; Bechtel & Abrahamsen 2008). As we have seen, unlike *machines*, the activity of contemporary *mechanisms* does not consist wholly in step-wise successions through linear series of clockwork-like connectives among their parts, but instead regularly feature spatially and temporally complex causal motifs such as self-replication, self-regulation, etc. And unlike *machines*, contemporary *mechanisms*, as we have seen, allow for holistic phenomena arising from the collective activity of parts which imbue entire systems with higher-order predictive, explanatory power.

A contemporary *mechanistic* ontology is *not* an ontology of *machines*, and the conflation of the two seems to function as an implicit cause of much of the process theorists’ ire. Indeed, many of their arguments can be read as various attempts to show that such a conflation constitutes a justified conceptual collapse. Understood in this way, the aim of this paper has been to illustrate the illegitimacy of this move: the higher-order, emergent features of developmental systems which are inexplicable within the machine ontology of old are the very ones which our contemporary mechanistic ontology is able to both accommodate and comprehensively account for.

Acknowledgements

I am grateful for the generous support of the Analysis Trust.

References

- Allen, G. (2005). Mechanism, Vitalism and Organicism in Late Nineteenth and Twentieth-Century Biology: The Importance of Historical Context. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 36(2), 261-283.
- Alon, U. (2006). *An Introduction to Systems Biology: Design Principles of Biological Circuits*. Taylor & Francis.
- Alon, U. (2007). Network Motifs: Theory and Experimental Approaches. *Nature Reviews Genetics*, 8(6), 450-461.
- Andersen, P., Emmeche, C., Finnemann, N., & Christiansen, P. (Eds.). (2000). *Downward Causation: Minds, Bodies and Matter*. Aarhus: Aarhus University Press.
- Armstrong, D. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Bapteste, E., & Dupré J. (2013). Towards a Processual Microbial Ontology. *Biology & Philosophy*, 28(2), 379-404.
- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A Mechanistic Alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 421-441.
- Bechtel, W., & Abrahamsen, A. (2008). From Reduction Back To Higher Levels. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 559-564).
- Bechtel, W., & Abrahamsen, A. (2010). Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science. *Studies in History and Philosophy of Science*, 41, 321-333.
- Bhattacharya, S., Zhang, Q., & Andersen, M. (2011). A Deterministic Map of Waddington's Epigenetic Landscape for Cell Fate Specification. *BMC Systems Biology*, 5(85), 1-11.
- Brakefield, P. (2011). Evo-Devo and Accounting for Darwin's Endless Forms. *Philosophical Transactions of the Royal Society*, 366(1574), 2069-2075.
- Brandon, R. (1999). The Units of Selection Revisited: The Modules of Selection. *Biology & Philosophy*, 14(2), 167-180.
- Brigandt, I. (2007). Typology Now: Homology and Developmental Constraints Explain Evolvability. *Biology & Philosophy*, 22(5), 709-725.
- Brigandt, I. (2015). Evolutionary Developmental Biology and the Limits of Philosophical Accounts of Mechanistic Explanation. In P. A. Braillard, & C. Malaterre (Eds.), *Explanation in Biology: An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences* (pp. 135-173). Springer.
- Cahoone, L. (2013). Ordinal Pluralism as Metaphysics for Biology. In B. Henning, & A. Scarfe (Eds.), *Beyond Mechanism: Putting Life Back Into Biology* (pp. 133-146). Lexington Books/Rowan & Littlefield.
- Callebaut, W., & Rasskin-Gutman, D. (Eds.). (2005). *Modularity: Understanding the Development and Evolution of Natural Complex Systems*. Cambridge: MIT Press.
- Callebaut, W., Müller, G., & Newman, S. (2007). The Organismic Systems Approach: Evo-Devo and the Streamlining of the Naturalistic Agenda. In R. Sansom, & R. Brandon (Eds.), *Integrating Evolution and Development: From Theory to Practice* (pp. 25-92). Cambridge: MIT Press.
- Canestro, C., Yokoi, H., & Postlethwait, J. (Nature Reviews Genetics). Evolutionary Developmental Biology and Genomics. 2007, 8(12), 932-942.
- Carroll, S. (2008). Evo-Devo and an Expanding Evolutionary Synthesis: A Genetic Theory of Morphological Evolution. *Cell*, 134(1), 25-36.
- Carroll, S., Grenier, J., & Weatherbee, S. (2001). *From DNA to Diversity: Molecular Genetics and the Evolution of Animal Design*. Oxford: Blackwell Science.
- Corson, F., & Siggia, E. D. (2012). Geometry, Epistasis, and Developmental Patterning. *Proceedings of the National Academy of the Sciences of the United States of America*, 109(15), 5568-5575.
- Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press.
- Darden, L. (2007). Mechanisms and Models. In D. Hull, & M. Ruse (Eds.), *The Cambridge Companion to The Philosophy of Biology* (pp. 139-159). Cambridge: Cambridge University Press.
- Darden, L., & Craver, C. (2002). Strategies in the Interfield Discovery of the Mechanism of Protein Synthesis. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 33, 1-28.
- Davidson, E. (2001). *Genomic Regulatory Systems: In Development and Evolution*. London: Academic Press.
- Davidson, E., & Erwin, D. (2006). Gene Regulatory Networks and the Evolution of Animal Body Plans. *Science*, 311(5762), 796-800.
- Davila-Velderrain, J., Martinez-Garcia, J. C., & Alvarez-Buyila, E. R. (2015). Modeling the Epigenetic Attractors Landscape: Toward a Post-Genomic Mechanistic Understanding of Development. *Frontiers in Genetics*, 6(160), doi: 10.3389/fgene.2015.00160.
- Dupré, J. (2013). Living Causes. *Proceedings of the Aristotelian Society Supplementary Volume*, 87(1), 19-38.
- Dupré, J., & O'Malley, M. (2007). Metagenomics and Biological Ontology. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 38(4), 834-846.
- Edelman, G., & Gally, J. (2001). Degeneracy and Complexity in Biological Systems. *Proceedings of the National Academy of the Sciences*, 98(24), 13763-13768.
- Ellis, G. (2012). Top-Down Causation and Emergence: Some Comments on Mechanisms. *Interface Focus*, 2(1), 126-140.
- Flatt, T. (2005). The Evolutionary Genetics of Canalization. *The Quarterly Review of Biology*, 80(3), 287-316.

- Frankel, N., Davis, G., Vargas, D., Wang, S., Payre, F., & Stern, D. (2010). Phenotypic Robustness Conferred by Apparently Redundant Transcriptional Enhancers. *Nature*, *466*, 490-493.
- Galis, F., & Metz, J. (2001). Testing the Vulnerability of the Phylotypic Stage: On Modularity and Evolutionary Conservation. *Journal of Experimental Zoology*, *291*(2), 195-204.
- Gilbert, S., & Bolker, J. (2001). Homologies of Process and Modular Elements of Embryonic Construction. *Journal of Experimental Zoology*, *291*(1), 1-12.
- Greenspan, R. (2001). The Flexible Genome. *Nature*, *2*(5), 383-387.
- Hall, B. (2003). Evo-Devo: Evolutionary Developmental Mechanisms. *International Journal of Developmental Biology*, *47*(7-8), 491-495.
- Hall, B. (2013). Epigenesis, Epigenetics, and the Epigenotype: Toward An Inclusive Concept of Development and Evolution. In B. Henning, & A. Scarfe (Eds.), *Beyond Mechanism: Putting Life Back Into Biology* (pp. 345-368). Lexington Books/Rowan & Littlefield.
- Huang, S. (2009). Reprogramming Cell Fates: Reconciling Rarity with Robustness. *Bioessays*, *31*(5), 546-560.
- Huang, S. (2012). The Molecular and Mathematical Basis of Waddington's Epigenetic Landscape: A Framework for Post-Darwinian Biology? *Bioessays*, *34*(2), 149-157.
- Huneman, P. (2010). Topological Explanations and Robustness in Biological Sciences. *Synthese*, *177*(2), 213-245.
- Jaeger, J., & Monk, N. (2014). Bioattractors: Dynamical Systems Theory and the Evolution of Regulatory Processes. *Journal of Physiology*, *592*(11), 2267-2281.
- Jaeger, J., & Monk, N. (2015). Everything Flows: A Process Perspective On Life. *EMBO Reports*, *16*(9), 1064-1067.
- Jaeger, J., Irons, D., & Monk, N. (2012). The Inheritance of Process: A Dynamical Systems Approach. *Journal of Experimental Zoology*, *318*(8), 591-612.
- Kalinka, A., Varga, K., Gerrard, D., Preibisch, S., Corcoran, D., Jarrells, J., et al. (2010). Gene Expression Divergence Recapitulates the Developmental Hourglass Model. *Nature*, *468*, 811-814.
- Kaplan, D. (2015). Moving Parts: The Natural Alliance Between Dynamical and Mechanistic Modeling Approaches. *Biology & Philosophy*, *30*(6), 757-786.
- Kaplan, D., & Craver, C. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science*, *78*(4), 601-627.
- Kauffman, S. A. (1969). Metabolic Stability and Epigenesis in Randomly Constructed Nets. *Journal of Theoretical Biology*, *22*(3), 437-467.
- Kim, K., & Wang, J. (2007). Potential Energy Landscape and Robustness of a Gene Regulatory Network: Toggle Switch. *PLoS Computational Biology*, *3*(3), 0565-0577.
- Kitano, H. (2004). Biological Robustness. *Nature Reviews Genetics*, *5*(11), 826-837.
- Laubichler, M. D. (2010). Evolutionary Developmental Biology Offers a Significant Challenge to the Neo-Darwinian Paradigm. In F. Ayala, & R. Arp (Eds.), *Contemporary Debates in the Philosophy of Biology* (pp. 199-212). Malden: Wiley-Blackwell.
- Levin, M. (2012). Morphogenetic Fields in Emryogenesis, Regeneration, and Cancer: Non-Local Control of Complex Patterning. *Biosystems*, *109*(3), 243-261.
- Love, A. (2009). Typology Reconfigured: From the Metaphysics of Essentialism to the Epistemology of Representation. *Acta Biotheoretica*, *57*(1-2), 51-75.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking About Mechanisms. *Philosophy of Science*, *67*(1), 1-25.
- MacNeil, L., & Walhout, A. (2011). Gene Regulatory Networks and the Role of Robustness and Stochasticity in the Control of Gene Expression. *Genome Research*, *21*(5), 645-657.
- Mason, P. (2010). Degeneracy at Multiple Levels of Complexity. *Biological Theory*, *5*(3), 277-288.
- McCune, A., & Schimenti, J. (2012). Using Genomic Networks and Homology to Understand the Evolution of Phenotypic Traits. *Current Genomics*, *13*(1), 74-84.
- McManus, F. (2012). Development and Mechanistic Explanation. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *43*, 532-541.
- Mitchell, S. (2012). Emergence: Logical, Functional and Dynamical. *Synthese*, *185*(2), 171-186.
- Müller, G. (2003). Homology: The Evolution of Morphological Organization. In G. B. Müller, & S. A. Newman (Eds.), *Origination of Organismal Form: Beyond the Gene in Developmental and Evolutionary Biology* (pp. 51-69). Cambridge: MIT Press.
- Müller, G., & Newman, S. A. (1999). Generation, Integration, Autonomy: Three Steps in the Evolution of Homology. *Novartis Foundation Symposia*, *222*, 65-73.
- Mumford, S. (2004). *Laws in Nature*. London: Routledge.
- Nathan, M., & Borghini, A. (2014). Development and Natural Kinds: Some Lessons from Biology. *Synthese*, *191*(3), 539-556.
- Nicholson, D. (2012). The Concept of Mechanism in Biology. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *43*(1), 152-163.
- Owen, R. (1848). *On the Archetype and Homologies of the Vertebrate Skeleton*. London: John van Voorst.

- Raff, R. A. (1996). *The Shape of Life: Genes, Development, and the Evolution of Animal Form*. Chicago: University of Chicago Press.
- Rieppel, O. (2005). Modules, Kinds, and Homology. *Journal of Experimental Zoology*, 304B(1), 18-27.
- Rosa, L., & Etcheberria, A. (2011). Pattern and Process in Evo-Devo: Descriptions and Explanations. In H. de Regt, S. Hartmann, & S. Okasha (Eds.), *EPSA Philosophy of Science: Amsterdam 2009* (pp. 263-274). Springer.
- Rosenberg, A. (2001). On Multiple Realization and the Special Sciences. *The Journal of Philosophy*, 98(7), 365-373.
- Salazar-Ciudad, I., & Jernvall, J. (2013). The Causality Horizon and the Developmental Bases of Morphological Evolution. *Biological Theory*, 8(3), 286-292.
- Striedter, G. (1998). Stepping Into the Same River Twice: Homologues as Recurring Attractors in Epigenetic Landscapes. *Brain, Behavior and Evolution*, 52(4-5), 218-231.
- Tabata, T. (2001). Genetics of Morphogen Gradients. *Nature*, 2(8), 620-630.
- Tickle, C. (2003). Patterning Systems: From One End of the Limb to the Other. *Developmental Cell*, 4(4), 449-458.
- Verd, B., Crombach, A., & Jaeger, J. (2014). Classification of Transient Behaviours in a Time-Dependent Toggle Switch Model. *BMC Systems Biology*, 8(43), 1-19.
- Waddington, C. H. (1957). *The Strategy of the Genes*. London: George Allen & Unwin.
- Waddington, C. H. (1969). The Practical Consequences of Metaphysical Beliefs on a Biologist's Work: An Autobiographical Note. In C. H. Waddington (Ed.), *Towards a Theoretical Biology 2: Sketches* (pp. 72-81). Edinburgh: Edinburgh University Press.
- Wagner, A. (2005). Distributed Robustness Versus Redundancy as Causes of Mutational Robustness. *BioEssays*, 27(2), 176-188.
- Wagner, G. (2014). *Homology, Genes, and Evolutionary Innovation*. Princeton: Princeton University Press.
- Wagner, G., & Lynch, V. (2010). Evolutionary Novelty. *Current Biology*, 20(2), R48-R52.
- Wagner, G., Chiu, C., & Laubichler, M. D. (2000). Developmental Evolution as a Mechanistic Science: The Inference from Developmental Mechanisms to Evolutionary Processes. *American Zoologist*, 40(5), 819-831.
- Walsh, D. (2013). Mechanism, Emergence, and Miscibility: The Autonomy of Evo-Devo. In P. Huneman (Ed.), *Functions: Selection and Mechanisms* (pp. 43-65). Springer.
- Wang, J., Zhang, K., Xu, L., & Wang, E. (2011). Quantifying the Waddington Landscape and Biological Paths for Development and Differentiation. *Proceedings of the National Academy of Sciences of the United States of America*, 108(20), 8257-8262.
- Webster, G., & Goodwin, B. (1996). *Form and Transformation: Generative and Relational Principles in Biology*. Cambridge: Cambridge University Press.
- Whitacre, J. (2010). Degeneracy: A Link between Evolvability, Robustness and Complexity in Biological Systems. *Theoretical Biology and Medical Modelling*, 7(6), 7:6.
- Whitacre, J., & Bender, A. (2010). Networked Buffering: A Basic Mechanism for Distributed Robustness in Complex Adaptive Systems. *Theoretical Biology and Medical Modelling*, 7(20), 1-20.
- Whitehead, A. N. (1925). *Science and the Modern World*. Cambridge: Cambridge University Press.
- Wilkins, A. (2002). *The Evolution of Developmental Pathways*. Sunderland, MA.: Sinauer Associates, Inc.
- Wimsatt, W. (2000). Emergence as Non-Aggregativity and the Biases of Reductionisms. *Foundations of Science*, 5(3), 269-297.
- Woese, C. (2004). A New Biology for a New Century. *Microbiology and Molecular Biology Reviews*, 68(2), 173-186.
- Woodward, J. (2002). What is a Mechanism? A Counterfactual Account. *Philosophy of Science*(3), S366-S377.
- Woodward, J. (2003). Mechanistic Explanation: Its Scope and Limits. *Proceedings of the Aristotelian Society Supplementary Volume*, 87(1), 39-65.
- Yi, T., Huang, Y., Simon, M., & Doyle, J. (2000). Robust Perfect Adaptation in Bacterial Chemotaxis through Integral Feedback Control. *Proceedings of the National Academy of Sciences of the United States of America*, 97(9), 4649-4653.
- Zhenglong, G., Steinmetz, L., Gu, X., Scharfe, C., Davis, R., & Li, W.-H. (2003). Role of Duplicate Genes in Genetic Robustness Against Null Mutations. *Nature*, 421(6918), 63-66.
- Zhou, J. X., Aliyu, M. D., & Huang, S. (2012). Quasi-Potential Landscape in Complex Multi-Stable Systems. *Journal of the Royal Society Interface*, 9(77), 3539-3553.
- Zuniga, A. (2015). Next Generation Limb Development and Evolution: Old Questions, New Perspectives. *Development*, 142(22), 3810-20.