Chapter 11
Top-Down Causation Without Levels

Jan Voosholz

Abstract The paper addresses a question concerning George Ellis's theory of top-down causation by considering a challenge to the "level-picture of nature" which he employs as a foundational element in his theory. According to the level-picture, nature is ordered by distinct and finitely many levels, each characterised by its own types of entities, relations, laws and principles of behavior, and causal relations to their respective neighbouring top- and bottom-level. The branching hierarchy that results from this picture enables Ellis to build his model of modular hierarchical structure for complexity, his account of same-level, bottom-up and top-down causation, of emergence, equivalence-classes and multiple realisability. The three main arguments for the level-picture in Ellis's works are reconstructed and shown to face serious problems. Finally, the paper presents a possible solution to this challenge by introducing a reformulation of certain fundamental points of Ellis's theory that does without the level-picture of nature. This allows us to preserve all of his central claims about the model of complexity, the three types of causation, emergence, equivalence-classes and multiple realisability. Any problems pertaining to the level-picture can be remedied in the context of Ellis's theory of top-down causation.

11.1 Introduction

This paper sets out, firstly, to present a challenge to George Ellis's theory of top-down causation by questioning what will be called the "level-picture of nature", and, secondly, to argue that Ellis's theory can answer this challenge if it adapts certain of its ontological premises. The level-picture consists in the claim that nature is ordered by distinct and finitely many levels, each with its own types of entities, relations, laws and principles of behavior, and causal relations to its

J. Voosholz (✉)
Institute for Philosophy, University of Bonn, Bonn, Germany
e-mail: voosholz@uni-bonn.de

J. Voosholz, M. Gabriel (eds.), *Top-Down Causation and Emergence*,
Synthese Library 439, https://doi.org/10.1007/978-3-030-71899-2_11

respective neighbouring top- and bottom-levels. The branching hierarchy that results from this picture enables Ellis to build his model of modular hierarchical structure for complexity, his account of same-level, bottom-up and top-down causation, of emergence, equivalence-classes and multiple realisability. The level-picture is, on the face of it, a cornerstone of Ellis's account.[1] Yet this paper argues for three conclusions:

(C1) Ellis's theory of causation hinges—in its current form—on eight essential assumptions concerning ontologically robust levels of nature. Section 11.2. of this paper assembles these eight crucial points via a reconstruction of key passages of Ellis's work. This serves to substantiate the notion of the level-picture.

(C2) Ellis's arguments in favour of the level-picture stemming from the debates on reductionism and emergence are not fully convincing—at least in the context of the current set-up of his theory. Upon closer inspection, various counterarguments reveal the level-picture to be questionable. Section 11.3. presents this dialectic, yet argues that a central insight can nevertheless be preserved: Ellis successfully establishes that there are constitutional hierarchies for physical entities and that there is a gradient of complexity of systems and structures in the universe.

(C3) Therefore, the apparent failure of the level-picture of nature does not pose a fatal problem for Ellis. My argument for this claim consists in presenting eight alternative foundational hypotheses, which can stand in for the eight assumptions of the level-picture. I further argue that these hypotheses can be reconciled with the critical challenge outlined in the preceding section. Moreover, they can support the key elements of Ellis's theory of causation, such as his modular hierarchical structure-model of complexity, the three types of top-down, bottom-up and same-system causation, equivalence-classes and multiple realisability.

The paper will also stress that this result ought not to be seen as a setback by proponents of Ellis's theory: One line of criticism against his theory and other accounts of top-down causation, strong emergence and genuine complexity tries to exploit the gap in the argument stemming from the level-picture in order to suggest, explicitly or implicitly, that such accounts of causation are wrong. These kinds of criticism could be avoided if it were possible to establish either that a weaker epistemic notion of levels in the universe would suffice for Ellis's position or even that his position remains tenable without any notion of distinct levels at all, instead simply using a gradient of complexity and constitutional hierarchies. I will

---

[1] This definition applies to many different theories of causation; it is neither limited to Ellis nor meant to imply that his proposal originated the level-picture. The latter can be traced back at least to Anderson (1972), who introduced this kind of formulation into the debates about emergence and causation. Thus defined, the level-picture is a broad concept with a long history in thinking about nature. Yet in every different theory of emergence and causation its defining features change considerably. Any critique has to focus on the particular analysis at issue. George Ellis's theory is a perfect case in point, not only due to its substantial weight in the debate on causality in the sciences and philosophy, but on account of its clarity regarding the role of levels. For other literature on the level-picture in the aforementioned debates see (Kim 1992, 1999, Emmeche et al. 1997, 2000, Hulswit 2005, Paolini Paoletti and Orilia 2017).

suggest that the latter is indeed the case. To argue for a complete theory of top-down causation and genuine complexity without distinct levels is a task yet to be achieved.

11.2 The Level-Picture of Nature

Before we look at the arguments for and against, we first have to establish what use of "levels of nature"—the level-picture—Ellis makes in the scope of his theory (step 1), how his account of complexity and, in consequence, his different types of causation, rest on modelling physical systems as or within modular hierarchical structures by presupposing the level-picture (step 2), and finally, how his critics see the underlying assumptions about levels of nature at the foundation of his theory (step 3).

11.2.1 Ellis's Theory of Causation, Complexity and Emergence

George Ellis's account of causation, put forward in his book *How Can Physics Underlie the Mind?* (Ellis 2016) and elsewhere (Ellis 2013, 2012, 2009, 2008, 2006; Ellis and Gabriel 2021; Ellis and Drossel 2019, 2018; Ellis et al. 2012, 2008), is generating increasing interest in discussions of causation in philosophy and the natural sciences. Ellis's influential proposal employs the following definition of cause and effect:

> Causes are separated from effects by searching for correlations between phenomena such that manipulation of one (the cause) can be shown, in a specific context, to reliably result in specific changes in the other (the effect) at a later time. (Ellis 2016, p. 8)

According to the definition, a cause is a factor in the net of dependencies in the world, while an effect is a dependent knot, for if the cause is altered its effects differ too. Ellis's definition closely resembles Jim Woodward's so-called *manipulability theory* of causation (Woodward 2003, 2007, 2021). The basic idea is similar to other counterfactual theories of causation (Paul 2009, p. 166): If the cause were to differ or not obtain at all, the effect would also differ or not obtain at all. Usual problems with counterfactual theories resulting from over-determination, pre-emption and vicious circularity are amended by Woodward through tying causality to the idea that causes are manipulable in principle. He defines *total causation* as follows:

> X is a total cause of Y if and only if under an intervention that changes the value of X (with no other intervention occurring) there is an associated change in the value of Y. (Woodward 2007, p. 73)

Woodward offers corresponding definitions for (type-level) direct causes and contributing causes (Woodward 2003, p. 59). Contrary to other notions of intervention (like Price and Menzies 1993), Woodward tries hard to steer clear of human agency and subjectivity in his explication of an intervention, as the intervention by an act of free will is seen as just any other natural instance of an intervention from one system on another (Woodward 2003, pp. 94–112).

The similarity between Woodward and Ellis is uncanny. The counterfactual, interventionist, non-subjectivist, ontological definitions of cause and effect by Ellis and Woodward are supposed to hold in every case, so it seems to suggest a *monistic* theory of causation, meaning the concept of causality picks out one single type of relation between entities.[2] Yet, according to Ellis, even if we can define causes and effects in Woodward's and his way, we find different kinds or types of causation at work in nature. Ellis works with three main distinctions:

1. Causation on the same level in nature (for instance, two comets that mutually attract each other in their gravitational fields when they closely pass by one another);
2. bottom-up causation across two neighbouring levels of the natural hierarchy, where the cause is at the lower level and the effect on the higher level (for instance, cell necrosis due to a high lead concentration causes the disfunctioning of an organ);
3. top-down causation across two neighbouring levels of the natural hierarchy, where the cause is at the higher level and the effect on the lower level (for instance, the running program in a digital computer determins the actions of the gates and memory registers).

We will revisit these three types of causation later in more detail. First, it is important to note that Ellis uses all three types of causation to present a theory of the emergent complexity of entities on higher levels from lower-level parts or modules. According to Ellis, genuine complexity cannot

> emerge in a bottom-up way alone. [. . .] Like bottom-up causation, top-down causation should be seen as an inter-level phenomenon between neighbouring levels in the hierarchy. Just as bottom-up causation does not (clearly) imply the existence of a clearly identifiable bottom level, top-down causation does not necessarily imply existence of a clearly identifiable topmost level. (Ellis et al. 2012, p. 2)

If we were to assume that Ellis's top-down causation is a real phenomenon, which we can find in many different domains of nature, then it exists independently of our scientific practices and of whether or not we register it in our theories. Since Ellis is quite clear that this is indeed the case, he holds a realistic and ontologically robust view of top-down causation. The question of how this is possible in light of his manipulability theory-style definition of cause and effect, is not addressed in Ellis's works. Similar positions have been defended by Woodward (2008) and Siriwardena (2019) against criticism by Strevens (2007) and Price (2017) respectively.

The following analysis does not attempt to pursue this discussion, but aims to show that Ellis's three basic types of causation are ontologically distinct in so far as

---

[2] There is an extensive ongoing discussion in the philosophy of causation on the issue of pluralism and monism, see (Longworth 2006; Williamson 2006; Cartwright 2006, 2007; Psillos 2010; de Vreese 2010; Strevens 2013).

they exhibit categorically different features. We will see that Ellis is able to make this distinction due to the placement of the causes and effects on different levels of nature in the cases of bottom-up and top-down causation. Should we be inclined to follow Ellis's assessment that top-down causation is a real and ontologically robust phenomenon, we must accept that bottom-up and same-level causation are too. Furthermore, we would then need to concede that there are—in an ontologically robust sense, independently of our registration of them in our theories—different levels in nature.

Ellis's pluralist picture of causation is incomplete until we take his thoughts on the nature of emergence, genuine complexity and the specifics of interlevel causation into account: Ellis compares cases of *weak emergence*, like gas in a container, in which the complexity is no more than the sum of the effective parts, with instances of *genuine complexity*, for which *strong emergence* and *top-down causation* are necessary: Examples include the interconnected neurons in a brain, which form structures inexplicable by bottom-up causation alone.[3]

Aside from same-level causation, we also find interlevel causation, naturally subdivided in two types: Firstly, we have bottom-up causation, the cases where lower level parts add up to form, support or effect higher level entities. The parts form or constitute wholes in a hierarchical fashion, and so "each lower level underlies what happens at each higher level in terms of structure and causation" (Ellis 2016, p. 88). Secondly, we find top-down causation, where higher level systems and objects set constraints for lower level causation by providing a context in the form of specific structure.[4] We can now see how, according to the theory, the three basic types of causation rest entirely on what I will call the MHS-model of causation and complexity ("MHS" for "modular hierarchical structure"). This in turn utilises the picture of levels of nature, which play the key roles of constituting parallel and intersecting hierarchies and so delivering the structure of nature. Both will be the topics of discussion in the next subsection.

---

[3] "Genuine complexity can only emerge from interlevel causation (both bottom-up and top-down) in modular hierarchical structures. [. . .] The structure is emergent from lower level entities, but is much more than the parts. It is the patterns of structuring that count. This is a higher level property of the system: its description requires variables that relate to more than just the properties of the components. [. . .] In addition to the properties of the units themselves, it is the set of relations between units [. . .] that is crucial to building up complexity. These aspects cannot be reduced to lower level variables. [. . .] Higher level structural patterns channel causation at lower levels in the system, breaking symmetry and so constraining what happens at those levels" (Ellis 2016, pp. 85–87).

[4] Ellis can support this type of causation with any example where the whole determines movements of and changes to its parts, like the changes in the cardiovascular system due to heavy body movements of the whole organism. This seems to suggest that for Ellis a top-cause must always be deterministic for the bottom-effect. On the question of how top-down causation and the question of indeterminism fit together, see (Beebee 2014).

11.2.2 The MHS-Model and the Levels of Nature

Regarding emergence, Ellis states that the

> basis of complexity is *modular hierarchical structures*, leading to emergent levels of structure and function based on lower level networks. Each of these aspects ('modular', 'hierarchical' and 'structure') is crucial in the emergence of complexity out of interactions between simpler units. (Ellis 2012, p. 126)

This means any given complex system *S* can be differentiated into modules based on their function in *S*. Take the examples from the last subsection: If *S* is an organ, say a human heart, its parts (modules, subsystems) can be distinguished by their respective function (heart valves, chambers, etc.) and these parts are in turn comprised of cells.

This perspective suggests that there is a hierarchy of structures. *S* is supposed to be more complex than any one of its parts (the modules/subsystems), and those in turn should be more complex than their parts (the module's parts/sub-subsystems). If this constitutes a hierarchy, this complexity relation is transitive: *S* is more complex that its sub-subsystems. This means a heart is far more complex than a cell from the papillary muscles. Here, a worry could be raised: Depending on the supposed function of *S* and our description of the subsystems, one could certainly have a model of the heart with less complexity than that of a model of a single cell. This doubt disappears once we realise that Ellis's modular hierarchical structures are located within real world 'target systems', not in our models of them, because the complexity of a (sub)subsystem of *S* always adds to the complexity of *S*: Firstly, due to the transitive nature of the structural hierarchy, if *S* is comprised of two subsystems (S_{U1}, S_{U2}) *S* is minimally as complex as both subsystems combined (due to the principle of weak emergence).

Secondly, according to Ellis, due to the modular structure of these subsystems, we sometimes find the emergence of genuine complexity in *S*, which is not present at the level of the two subsystems, but results only from their structured contextual interaction (principle of strong emergence). This is crucial: Only because of this emergent genuine complexity, can *S* possess properties which are not mere effects of a subsystem but uniquely appear where *S* is concerned as a whole. In this case, the complexity of *S* is higher than the sum of the complexities of all of its parts.

But why must *S* and its subsystems be on ontologically distinct levels in nature? Up until this point, Ellis's picture of nature seems only to indicate a gradient of complexity between subsystems and systems. Ellis associates the hierarchy of complexity with a particular hierarchy of sciences. He does this in a simple table:

> A simplified version of the basic hierarchy of complexity and causality for natural systems (left) and for human beings (right) is given in Table 11.1. [...] This figure gives a simplified representation of this hierarchy of levels of reality (as characterised by corresponding academic subjects) for natural systems (left) and human beings (right). Each lower level underlies what happens at each higher level, in terms of causation. There is no correlation between the left- and the right-hand columns above the level of chemistry, as emergence and causation are quite different in the two cases; but the first four levels are identical (life emerges out of physics!) (Ellis 2012, pp. 126–127).

Table 11.1 The hierarchy of structure and causation

Level	Discipline	Discipline
10	Cosmology	Sociology/economics/politics
9	Astronomy	Psychology
8	Space science	Physiology
7	Geology, earth science	Cell biology
6	Materials science	Biochemistry
5	Physical chemistry	Chemistry
4	Atomic physics	Atomic physics
3	Nuclear physics	Nuclear physics
2	Particle physics	Particle physics
1	Fundamental theory	Fundamental theory

From Ellis (2016, p. 6)

Ellis also has a deeper view on the question regarding the point at which genuine emergence occurs:

> The way they [the modules; J.V.] interact with each other at a specific level can be characterised by an interaction network showing which modules interact with which other modules through various possible interaction modes (inter-module forces, or matter, energy and information exchange); this is the structure of the system. Emergence of a higher level system from the lower level modules takes place when reliable higher level behaviour arises out of the lower level actions taking place in the context of this structure, with lower level units grouped together to form higher level modules that one can identify as meaningful entities that persist over time and have identifiable laws of behaviour. (Ellis 2012, pp. 126–127)

The table and the descriptions combined are extremely suggestive: According to Ellis, the main target systems studied by the respective disciplines all seem to be on one distinct level of complexity in the modular hierarchical structure of nature.[5]

But is the whole of our solar system (as one level studied by solar system sience) more complex than the whole of the planet Earth (conceived as one system on its own level)? A critic could repeat this point time and again. Ellis answer to the

---

[5] Take solar system science: It studies our solar system. Modular subsystems on the same level of complexity and which causally interact with one another include the planet Earth, the moon Callisto, the dwarf planet Ceres, the Mars trojans, Halley's Comet and many others. The causal relations form a net with an internal hierarchy (Ceres is part of the asteroid belt, moons orbit their planets, trojans are tied to their planets or moons and so on). Yet even at this one distinct level of nature, there are modular subsystems with greater complexity than others: Jupiter seems to be more complex than Halley's Comet if we take Jupiter not only to have more functionally distinct parts (an outer and inner atmosphere with sub-parts like cloud layers, possibly a rocky core, etc. (Guillot et al. 2004)) than Halley's Comet but also, due to the interaction of those parts, more genuine complexity, such as weather events (take the Great Red Spot). According to Ellis's view, interaction between these two modules, Jupiter and Halley's Comet, in the whole of the solar system still has to be conceived as same-level causation, even if there can be a difference of complexity between modules.

question is: The solar system is more complex because it is a *natural systems with subsystems,* where the complexity of a part can never exceed the complexity of the whole it is part of, due to the relations I have outlined before: The complexity of a system is the sum of the weakly emergent complexity of its parts plus the strongly emergent, genuine complexity resulting from their interactions and relations with one another.[6] This is the definition of all complexity according to the MHS-model.

Yet, why should there be *distinct* levels in this ontologically robust way? Why not speak of a gradient of complexity (or rather multiple gradients of complexity) in any direction where natural systems form wholes and parts without clear borders?

> Each of the different levels of the hierarchy function according to laws of behaviour appropriate to that level, and are describable only in terms of language suited to that level [. . .]. Higher level entities, such as plans and intentions, have causal power in their own right, which partially determine what happens at lower levels in the hierarchy (billions of atoms and molecules move in accord with our intentions when we raise our arm). Here, we characterise a level as 'higher' when it can be shown to influence another level ('a lower' level) by setting a context that guides the way the lower level actions take place. (Ellis 2012, p. 127)

This does invite another reason to think of these levels as *not distinct* or *not real*: If an event, process, state, system, etc. on the higher level sets the context, for instance by setting certain boundary conditions (the orbit of Earth around the sun sets certain boundary conditions for Earth's climate), it is easy to see gradual differences between individual systems and subsystems. Yet it leaves completely open how one might justify a hard level distinction. Ellis is very clear on this point:

> Each higher level physical element, created by structured combinations of lower level elements, has different properties from the underlying lower levels—the entities at each level show behaviours characteristic of that level. [. . .]
>
>> **Essential Differences Between Levels**. Hierarchical structures have different kinds of order, phenomenological behaviour, and descriptive languages that characterise each level of the hierarchy.
>
> It is sometimes queried whether these levels actually exist "out there", or are rather impositions of the mind. My position is that different kinds of causation do indeed exist at the different levels as characterised here, and the mind recognises these distinctions which actually exist. They are not just inventions of the mind. Atoms are different from molecules, whether characterised as such by a mind or not. (Ellis 2016, p. 89)
>
>> **Existence**. The different levels are all real, each existing with causal powers in its own right, because [. . .] they each have determinable effects on the levels above and below them. No level is more real than any other. (Ellis 2016, p. 90)

---

[6] The option to define a natural system by way of analysing whether its complexity actually exceeds and not only equals the complexity of the sum of its real functional parts (which would supposedly be the case for any mereological sum of the complexity of arbitrary parts, say the planet Mars and the exoplanet TrES-3b regarded as one part) is not further pursued in Ellis's work as far as I know. If this were possible it would offer a way to search for cases of strong emergence and top-down causation (understood as feedback loops).

Table 11.2 The hierarchy of structure and causation with example target systems

	Discipline	Target systems	Discipline	Target systems
Level	Inanimate matter		Living matter	
10	Cosmology	Universe	Sociology/economics/politics	Global society
9	Astronomy	Galaxies, star clusters	Psychology	Minds
8	Space science	Solar systems/interstellar space	Physiology	Human/animal/plant bodies
7	Geology, earth science	Planets	Cell biology	Cells/cell organelles
6	Materials science	Metals/semiconductors/ceramics/polymers	Biochemistry	Biomolecules
5	Physical chemistry	Molecules	Chemistry	Molecules
4	Atomic physics	Atoms	Atomic physics	Atoms
3	Nuclear physics	Nuclei/protons/neutrons	Nuclear physics	Nuclei/protons/neutrons
2	Particle physics	Quarks/leptons/bosons/quantum fields	Particle physics	Quarks/leptons/bosons/quantum fields
1	Fundamental theory	Strings/etc.	Fundamental theory	Strings/etc.

From Ellis (2016, p. 6), target systems added by J.V.

The apparent reason to assume levels exist in an ontologically robust, real way, independently of their registration, is that the three different forms of *causation* hinge on the existence of different levels, and we can make empirical discoveries about causation once we accept the Ellis-Woodward definition. The exact arguments for this will be analysed in Sect. 11.3 below.

Yet before we come to that, I want to bring out another consequence resulting from the long quote from Ellis above: The scientific disciplines characterising the 10 levels in Table 11.1 are of course practices of human animals on planet Earth, with their respective histories and so on. As all levels exist independently of any registration of them, they exist independently of the respective discipline(s) that researches them. Ellis characterisation seems to be merely instrumental: They are in reality constituted by the entities with causal effects on other entities (on the same or neighbouring levels). In other words, the target systems of the respective sciences are naturally organised in distinct levels. Table 11.2 gives important example target systems for every level of nature.

On Table 11.2, we can now see that Ellis employs the level-picture of nature to categorise ("real") systems as belonging to certain levels depending on their complexity. He then utilises this division to introduce his concepts of emergence and same-level, top-down, and bottom-up causation. Before we turn to the arguments for and against the level-picture, it is important to explicate the eight main points that underlie the levels picture on this reading of Ellis's theory. It will later become apparent that each of these eight points can be appropriately reformulated, even if we abandon the level-picture of nature.

11.2.3 The Eight Points on Levels

According to the interpretation given above, Ellis holds the following view of levels, which he exploits within his theory of top-down causation and emergence:

L1 Levels are structures in nature and real in the sense of existing independently of our (or any) knowledge of them.

L2 Any fully or partly concrete (meaning: existing in spacetime) object or system is situated on a certain level in nature and behaves according to a set of principles and laws effective on this level, which can possibly be discovered by science.

L3 Levels are ordered by the degree of complexity of the objects and systems that are found on them. It is unclear whether there is either a top or a bottom level (or which ones they would be).

L4 To describe different levels (and the associated objects, systems, structures, laws and principles), we need to deploy different scientific languages and different variables suited to any one level. This results in the disunity of sciences, where for every level we have one (or more) science studying it.

L5 For any system S on level n we can identify parts of it either on level n or n^{-1}. For any object O, located on level n, considered as a whole, there are proper parts of it either on level n or n^{-1}.

L6 Levels are distinct entities with clear boundaries, which are marked by emergent objects and systems with corresponding properties, structures, principles, laws, and variables, which can neither be found at, nor explained by, nor fully accounted for by the next lower or higher level.

L7 For any level n (apart from a theoretically existing bottom level) there exists exactly one lower level n^{-1}, which constitutes the complete set of basic parts for all systems and objects on n. In this sense the lowest physical levels are the basis for everything else in nature.

L8 For any level n (apart from a theoretically existing top-level) there exists one higher level n^{+1}, which might be branched. This means it is possible that different higher levels of complexity are constituted by the same lower level. This is the case for levels 4 and 5, where 4 – the level described by atomic physics – forms the basis for level 5, on which the branch of life departs from the branch of inanimate nature, as represented in the sciences of chemistry and physical chemistry.

In the next section, I will critically question L6 and so raise a doubt regarding L1 by first presenting Ellis's arguments for the level-picture and show that they do not fully support L1 and L6. Following this step, instead of arguing against Ellis's account of top-down causation, I will propose a reformulation of L1–L8 in order to show that his theory can be stated in terms that do not impose a level-ontology onto nature. In particular, the idea of modular hierarchical structures, which can be parallel to or intersect one another, will be preserved. Furthermore, his idea of a plurality of causal modes transecting systems of different complexity and the borders of animate and inanimate parts of nature, minded and non-minded reality or the realm of the concrete and the abstract will be left intact by my criticism of the level-picture of nature.

11.3 Arguments for and against the Level-Picture

We can now reconstruct the arguments speaking for and against the level-picture as outlined above (3.1). Though they remain partly implicit, it is possible to distinguish three arguments in Ellis. In a further step (3.2), we consider the possibility of giving up the level-picture. Fortunately, there are good reasons to hope for a superior replacement for the level-picture in antireductionist theories of causation in general and Ellis's in particular. So even if a critic does manage successfully to argue against the level-picture, she does not thereby refute Ellis.

11.3.1 The Arguments for the Level-Picture

We will review Ellis's three arguments for the level-picture in turn. The first argument rests on Ellis's account of genuine complexity and laws. We will see how this argument has the disadvantage of being unable to justify the level-picture as such, because it leaves us stranded half way. The second argument hinges on the apparent disunity of science. We can characterise the problem here as a categorical one: If this argument were successful, it would only prove an epistemic notion of levels derivative from contingent languages and scales. But this limited conclusion may also be attacked by the critic in turn, because the argument might beg the question. The third argument is the strongest: It takes the different local ontologies for different contexts in nature and establishes the existence of whole-part-relationships between entities from different local ontologies. Two consecutive counterarguments by a potential critic of Ellis's third argument point in the direction of an alternative to the level-picture.

11.3.1.1 The Argument from Genuine Complexity, Principles and Laws

Ellis's first argument tries to individuate distinct levels in nature in light of the facts that, firstly, laws of nature (or—to put it more cautiously—reliable principles of behavior) affect only certain entities at a certain scale (of complexity, mass, length, etc.) and, secondly, there are intrinsically higher level variables of an abstract nature and determined by a higher-level logic, which are cases of genuine complexity above the physical laws of nature. Ellis expresses this in the following way:

> Each of the different levels of the hierarchy function according to laws of behaviour appropriate to that level [. . .]. (Ellis 2012, p. 127)
>
> Each higher level, created by structured combinations of lower level elements, has different properties from the underlying lower levels—the entities at each level show behaviours characteristic of that level. (Ellis 2006, p. 80)
>
> Specific higher level variables characterise the macroscopic state of the system at a specific level, and occur in effective laws of behaviour at that level. (Ellis 2016, p. 12)
>
> Higher level variables [e.g. for properties of systems, J.V.] may be emergent from the lower level variables [. . .]. However, there are some kinds of higher level variables that are not emergent: they are intrinsically higher level variable [which in this paper are called *strongly emergent*; J.V.]. (Ellis 2016, p. 105)
>
> Though they are realised in various lower level physical substrates, they are determined by higher level logic, and so are intrinsically of higher level nature:
>
> 1. **Algorithms**. Examples are quicksort or the Google search algorithm [. . .].
> 2. **Codified Laws of Physics**. Our mental representation of physical interactions, such as Newton’s equations or Maxwell’s equations, the foundation of mechanical and electrical engineering, respectively.
> 3. **Social Agreements**. Examples are the rules of football, [. . .] the constitution of an organisation, or exchange rates for money.
> 4. **Conceptual Plans**. Examples are the plans for a building, a town, an aircraft, or for a musical concert, a company, or a physics experiment.
>   [. . .] **Intrinsically Higher Level Variables**. These are not physical variables, and there is no way to obtain them by any kind of coarse-graining process. Rather they are of mental or abstract nature. However, they are certainly causally effective. (Ellis 2016, p. 108)

This argument is twofold, but both lines of reasoning are intertwined: Firstly, there are principles and laws which only hold for a specific context. Examples include the gas laws, which are applicable only to gases, the Friedmann equations, which apply only to expanding space, generalised Lotka–Volterra equations, which apply only to species in competition and trophic relationships in a closed ecosystem. Take the level of statistical mechanics (level 4, inanimate branch in Table 11.2), with target systems such as confined gases in experimental chambers or mesoscopic natural systems of gases: only these systems, gases, can be modelled with the variables behaving according to the gas laws on that particular level of physics. The same holds for the Friedmann equations, which are valid only for the extending space of the universe on the level of cosmology (level 10, inanimate branch in Table 11.2), and for Lotka–Volterra equations, which only describe the behavior of populations of species in direct competition or trophic relationships in one ecosystem over time on the level of ecology (possibly level 10, animate branch in Table 11.2).

In all these cases, the subsystems of the three example systems (single molecules in the gas, stars or galaxies within the universe, individual animals or plants) do not show similar behavior. They do not even possess the properties which are important for law-like behavior (a single molecule has no temperature or pressure, etc.). Therefore, the variables of the equations cannot be applied. So, the first step of the argument clearly shows that we have to cross a distinct border of levels in nature if we go from subsystem to system in these (and other) cases. Ellis concludes that we here have instances of hierarchical complexity and (what we have called in this paper weak or strong) emergence: New laws, principles and properties arise in the order of nature and, depending on our chosen scale, we have to apply different variables which stand in different relationships with one another.

Now we can analyse the next step of the argument: as Ellis's examples of algorithms, social agreements, etc. indicate, the higher levels of natural organisation somehow allow for cases of genuine complexity in which intrinsically higher-level variables appear: Fully abstract variables which behave according to rules of logic or discourse. Social agreements are reached in historic processes of discourse—rational argumentation, group negotiation, exercise of power, etc.—, algorithms are executed by a local conventional logic (binary code in a digital computer, a formula in a formal language on paper). Neither can be accounted for by physical properties or entities alone, even though both can be realised in physical systems (an actual computer or an actual group of people and so on).

Both steps of the argument are intertwined because the first step introduces constitutional hierarchies with distinct levels via level-specific properties and laws of behavior on which the second step is dependent, while the second step places our codified laws of nature and their variables as abstract entities in nature, to allow the first step to take place. With Ellis, we can resist any criticism of this as begging the question: This does not involve a vicious circle, because only our discoveries of and epistemic access to level-specific properties and laws of behavior are dependent on pre-existing discourse and logic (which can be justified and explained with other arguments), and not the entities with said properties or their behavior according to laws or principles.

Yet both parts of the argument, if we grant the premises, suffice only to establish that we have particular constitutional hierarchies with weak and strong emergence as well as cases of genuine complexity, and that these constitutional hierarchies have different levels in the sense that (weakly) emergent variables or (strongly emergent) intrinsically higher-level variables for higher-level properties can be found at the level of the complete system, while these properties are absent in all subsystems. It does not prove that these particular constitutional hierarchies can all be integrated into the general branching hierarchy Ellis shows us in Table 11.1 or 11.2, where all systems fall into one of ten levels. Certainly, it does not convince the critic that the outcome of such an integration would be the discovery of real levels, existing in the ontologically robust sense defined by L1/L6 in the last section. This conclusion needs another step, which would take us from constitutional hierarchies with levels of emerging properties and abstract variables to L1 and L6. This first argument has the disadvantage of being unable to justify the level-picture as such, because it leaves us stranded half way.

11.3.1.2 The Argument from the Disunity of Scientific Language

When Ellis writes about levels, appropriate laws and genuine complexity, he almost always speaks about the different languages of different sciences. Could this help to bring us the rest of the way to the level-picture? Ellis's second argument seems to pick up the thread right where the first argument left off:

> Each higher level physical element, created by structured combinations of lower level elements, has different properties from the underlying lower levels [. . .]. Each level is described in terms of concepts relevant to that level of structure (particle physics deals with quarks and gluons, chemistry with atoms and molecules, and so on), so a different descriptive language applies at each level. Different levels of the hierarchy function according to laws of behavior appropriate to that level, and are describable only in terms of language suited to that level. One cannot even describe higher levels in terms of lower level languages because a different phenomenological description of causation is at work at the higher levels, which may be described in terms of different causal entities. Ideas applicable to lower level causation do not by themselves succeed in explaining the higher level behaviors, for the concepts employed are simply not appropriate to the higher level kinds of causation [. . .].
>
> **Essential Differences Between Levels**. Hierarchical structures have different kinds of order, phenomenological behavior, and descriptive languages that characterise each level of the hierarchy.
>
> [. . .] Thus the hierarchy on the life science side is in term of function and causation rather than the scale of physical entities. The hierarchy is determined by finding out what entities—physical or otherwise—exert constraints or set conditions so as to channel interactions between elements which have their own laws of interaction at their own level, for any environment acts in this way on any system it contains. Together with a careful analysis of what more complex elements emerge from simpler ones, this defines which is a higher level and which a lower level in the hierarchy. (Ellis 2016, p. 89)

Here, Ellis takes the constitutional hierarchies with different levels and the intrinsically higher-level variables as granted and adds an observation regarding the apparent disunity of science. We have different sciences exploring different contexts in nature: Sometimes they look at a specific length or energy scale like radioactive atoms, sometimes they focus on a functional context like different ecosystems, sometimes they concentrate on a type of entity like cells. In order to understand these contexts of nature, the scientists develop their own languages suited to their fields. Apart from science-specific notions ("proton; niche; vacuole") and kinds of description, they often introduce scales with their own units and variables with an associated formal notation. Using these, it is often possible to identify the relevant entities of the context, their properties, relations, laws of behavior, and to explain phenomena and patterns, make predictions, design experiments, control the context and so on.

Ellis notes that it is not possible to take the scientific language (with its concepts, scales and variables) of one context and use it to describe another. One cannot use nuclear physics to explore ecosystems or ecology to investigate cells. He associates this contention with the conclusion of his first argument for constitutional hierarchies with different levels and intrinsically higher-level variables, equating these levels with the contexts of nature researched by different special sciences.

This is the crucial point of this second argument, and it also explains why Ellis uses scientific disciplines to designate ontologically robust levels of nature in his Table 11.1: If the countless constitutional hierarchies of different systems could all be integrated into a single hierarchy formed by (a limited number of) different context of nature, which in turn would comply with the structure of present-day scientific disciplines and their languages, scales and variables (Table 11.1), then the level-picture would be justified.

Whilst Ellis is right to acknowledge the disunity of the sciences, the critic may well regard his further step of associating the many constitutional hierarchies of natural systems with the context under investigation by the various sciences as unmotivated and problematic. Take a human body: A physician can study the organisation of the organs and the role the liver plays within it, while a cell biologist can concentrate on liver cells and a biochemist can look into certain sets of molecules in these cells. All three scientists use different languages, and they investigate different levels of the same constitutional hierarchy. Yet this does not prove that all constitutional hierarchies are integrable into one global hierarchy. It shows that the contexts investigated by these scientists overlap insofar as they study three systems standing in a whole-part-sub-part relationship. The overlap is reflected by the language of the disciplines: The concept of "hepatic lobiles" is used in physiology and cell biology alike, the concept of "DNA" in cell biology and biochemistry. There is no hard disunity here, no distinct borders, only a particular constitutional hierarchy.

Thus, the one further conclusion we can draw here is that the present-day sciences have different languages, scales and variables suited to contexts of phenomena and that these contexts often overlap because the systems under investigation stand in whole-part relationships. It is neither clear that the contexts of phenomena are necessarily all "naturally distinct" and have sharp borders, nor—even if they were—that these contexts form a single branching hierarchy of distinct levels, which could integrate every constitutional hierarchy of natural systems. Just because we might call the context of the phenomena investigated by ecology the "level of ecology", it does not follow that it is equivalent to a distinct level within all constitutional hierarchies of natural systems.

The only possibility with which the critic leaves us is to show that the merely epistemic and operational level of ecology is equivalent to a distinct context in nature by proving that all the systems in context (all ecosystems) share one similarity in their respective constitutional hierarchies: Having the *same general levels of nature* as a single meta-structure for all respective *levels in the constitutional hierarchies* to appear on. But the critic will remark that this would beg the question: We would need to assume general levels in nature to make the argument work in the first place. So even if one could link every science to a distinct context in nature, the argument would still be circular.

11.3.1.3 The Local Ontologies Argument

If we concede to the critic that the first two arguments are not able to prove the level-picture, we can still try to reconstruct an argument that avoids the main problems: We cannot start with differences in scientific languages, and the constitutional hierarchies and their (weakly or strongly emergent) levels do not cut any ice by themselves. But what if we concentrate on the different "local" ontologies of the numerous contexts of nature directly? This is the path taken by the third argument. Albeit more implicit than the first two, it is contained in another explicit argument. Let us therefore consider Ellis's antireductionist argument, which provides another reason to think of levels as distinct and real:

> The core of the reductionist view is that everything can be explained by such bottom-up mechanisms based on the laws of physics, with no remainder. [. . .] The phenomenon of emergent order is when higher levels display new properties not evident at the lower levels. *More is different*, as famously stated by Anderson. Emergence of complexity takes place where quite different laws of behavior hold at the higher levels than at the lower levels. These properties are characterised by named higher level variables, and it is the symbolic naming of these variables that enables us to contemplate their nature. (Ellis 2016, pp. 99–100)

Ellis continues this line of thought by considering the example of the structure of a forest (how the trees, rocks, plants, rivers, etc. are grouped in the landscape) as a formal cause (by setting constraints to effective causal agents such as wind or fire) of the type of top-down causation: The structure as a top-down cause is not operative without an effective causal agent, yet it has causal power over lower level entities (such as the movement of animal bodies). The problem for the reductionist is now that on her picture, it seems impossible even for the lower level entities (rocks, rivers, fire etc.) to do any work in explaining what happens, as they are made up of lower level elements and so do not possess real causal power or effectiveness. They do not exist, strictly speaking.

> In fact nothing does, except vibrations of superstrings, if they exist, which may or may not be the case. Lower level causality vanishes into unknown and untestable regions.
>
> The only sensible way to handle this is to take an interlevel view, i.e., forget the bottommost level and assign real causal power to the lower level with respect to its immediate upper level, and to do this for every pair of levels:
>
> **Interlevel Causation**. For every pair of levels (N, N+1), the lower level "does the work", but the higher level is able to influence what work is to be done by setting constraints on the lower level operations.
>
> This is the basis for regarding every level as real: each is able to do real work. If we don't take this view, then genes and neurons are not able to do real work, as they are not the lowest level [. . .]. (Ellis 2016, p. 112)

The argument is directed at strict reductionism. In Ellis's view, its conclusion establishes the reality of distinct levels. How? By explicitly delivering a reason for thinking that entities on all levels have real causal power and/or effectiveness and implicitly introducing the premise that *every entity depends upon its appearing on one distinct level in nature*. If that premise were true, it would follow that the distinct levels were necessary conditions for any entity to "do work", to cause at all.

The reason for the premise—every entity depends on its appearing on one distinct level in nature—can be explicated by following the thrust of the overall argument. I will call it the local ontologies argument: The ontologies (and not only the languages) of sociology, cell biology, astronomy and particle physics are radically different (they name different entities which are connected by different laws, principles and relations), yet their respective target systems all do "real work", i.e., possess causal power and/or can be causally effective. If the ontologies of these sciences devoted to different respective contexts—let us call them local ontologies—are necessary for their entities to stand in (causal) relations to one another and to exhibit law- and principle-guided behavior at all, then they have to constitute a distinct level. As the antecedent is the case—the entities sometimes stand in (causal) relations and exhibit such behavior—distinct levels are real.

The critic will of course object that Ellis again fails fully to establish what he purports to, even if we grant all the premises. If we do grant them, the main conclusion is that all the different target systems are real (see Table 11.2) and possess causal power and/or effectiveness. This does represent a partial success; it proves that the different local ontologies of different disciplines are equally correct. What is missing is an argument establishing that if there are (correct) local ontologies with sets of (real) entities, then they necessarily form distinct levels, one level per local ontology. One can think of two arguments against this assumption. The first can be posed as a question: If nature really falls into a limited number of distinct levels, what would follow? The second counterargument states that even if it were possible to operate with a multitude of levels, these levels would not be distinct in the way Ellis seems to believe.

Counter 1: Infinite Hierarchies

Let us start with a premise Ellis and many others share. For many entities (say, a particular cell) we can find a constitutional hierarchy: It is composed of organelles, which are in turn composed of bio-molecules, etc, all while the cell itself might be a part of an organ, which is part of an animal etc. The levels of this hierarchy (of cell A), running both upwards and downwards, can be very different from the levels of another hierarchy (of cell B, which is a Acidithiobacillia): This other hierarchy has different organelles and is not part of a larger organ or life-form etc. The critic could make a similar argument for two molecules, one a N_2H^+-molecule in a gas cloud in the star-forming region ρ Oph A (Liseau et al. 2015, p. 2), the other a molecular iron complex in the steel of a sword. The hierarchies of cell A and B (and the two molecules) seem to involve very different levels.

Someone defending Ellis could counter that this is only partially true. Cells A and B might have different upwards hierarchies but the same downwards hierarchy, and it is this fact (the identical downward hierarchy) that identifies two entities as belonging on the same level in nature. This solution seems plausible, since the structurally different cell organelles[7] are on one level with the whole cells (see

---

[7] In cell A we find a cell nucleus and in cell B we do not.

Table 11.2), yet all the sub-subsystems (bio-molecules) and sub-sub-subsystems (atoms) are on the same lower levels. But why should cells and cell organelles be entities on one level, yet bio-molecules on another level? Since some organelles could be just macro-biomolecule, or (if we exclude this by defining a cell organelle as having at least two molecules as parts) the part-whole relationship between cells and cell organelles holds just as well as that between organelle and bio-molecules. In any case, it is unclear why the distinct level border between cell–organelle–bio-molecule should be drawn so as to split the first part-whole relation or the second and not the other. Defending Ellis, we could introduce an extra level (and so a second border) between cell and bio-molecule, namely the newly minted level of organelles. But we have to repeat this manoeuvre for every example where we find part-whole relationships that seem to occur on the same level. Soon, we would have discovered a multitude of distinct levels.

The critic might find different constitutional hierarchies for every "natural kind" or type of target system, maybe in fact for every physical entity. This would mean that there are (almost) infinitely many hierarchies, each possessing many more levels in turn. This would not only make the level-picture difficult to manage within the scope of Ellis's theory; it would also make it impossible to support the missing premise in the local ontologies argument: Even for a small local ontology with a limited set of entities, one has to deal with a vast number of levels. If the local ontology is that of standard-model particle physics, meaning there would be a reason to be concerned only with the downward hierarchies, it might be possible to accommodate all levels. Yet for cell biology or physical chemistry, as the examples above show, it is unfeasible. Certainly, there is no one-to-one correspondence between level and local ontology, yet this is exactly what would be needed to argue for the missing premise.

Counter 2: Against Distinctness

From the critic's perspective, another argument is more damning for the level-picture: Let us, for the sake of the local ontologies argument, assume that a multitude of levels would be manageable—say one hundred levels. Let us further stipulate that we could find a hundred suitable local ontologies, which could be grasped by various disunified sciences. Why should these 100 levels be distinct, with clear cut boundaries, such that for every abstract or physical entity with causal power/effectiveness (electrons, the structure of a forest, the Google search algorithm, Acidithiobacillia cells, a N_2H^+-molecule in a gas cloud, the plan to build an aeroplane, to cite just some examples) we would have a corresponding level number from 1 to 100?

From the reconstructed argument above, the only answer seems to be that the entities at each level always fall within one (and only one!) level-specific local ontology (which Ellis in fact—for simplicity's sake, as he writes—associates with twenty-first century scientific disciplines, while the 100-natural-level-hierarchy we are contemplating is associated with the fictitious scientific disciplines of an ideal future). Yet some entities seem to appear simultaneously in different ontologies: Electrons show up in explanations in material science, particle physics, and physiology (in explanations of neural activity for instance); N_2H^+-molecules are referenced by astronomy and chemistry alike; the Google search algorithm appears in sociology and computer science—and all within Ellis's simplified model. It will certainly become a lot more intractable if we try to differentiate between 100 levels.

One solution could be that the *Google search algorithm* designates a different entity in sociology and computer science (so would *electron*, etc.). This is implausible, however, because even if different sciences with different local ontologies investigate the different properties of different target systems, not only does the algorithm preserve key properties (being an abstract entity, giving a probability distribution for links/web-pages, etc.), but its properties in one local ontology are only fully understandable against the background and within the context of the other local ontologies: The development of the algorithm (and any changes that were made) makes sense for a computer scientist only with its social function in mind, while its success as a tool in global society can only be explained by reference to its mathematical qualities which cause its relative reliability. Even if this problem could be resolved, the solution of confining entities to one level would still fail: Ellis's theory insists that the systems of one level make a reappearance on the next higher level as possible parts.

Might this last point hint at a resolution of the puzzle: systems of level 4 reappear on all higher levels as subsystems and sub-subsystems? Unfortunately not. The Google search algorithm is a proper object of study for sociology and computer science alike, as is fear for physiology, psychology and political science, or water in chemistry, geology and solar system science. All of these appear as proper entities in different ontologies, i.e. target systems in different sciences, not only as parts or subsystems. It is also not true that a system appears only in the neighbourhood of a few levels: the electron and water are telling examples. If we apply the 100-natural-level-hierarchy to these examples (with a hundred levels and sciences) it results in even more problematic cases.

Both counterarguments clearly show that the local-hierarchies-argument does not establish the existence of distinct levels of nature. This failure does not harm the main conclusion of Ellis's argument against reductionism, because this does not require the full-fledged level-picture. All of Ellis's partially implicit arguments for the level-picture seem to fail; but in all three cases, many of the desired conclusions for which Ellis mainly argues follow anyhow: The argument for constitutional hierarchy and genuine complexity, the argument for the disunity of sciences and different phenomenal contexts, and the argument against strict reductionism all succeed—without proving the level-picture.

The failure of the local ontologies argument provides a good reason to think of Ellis's ten distinct levels as a hypothesis for his theory. It is neither proven empirically in cosmology that all systems of the universe fall into distinct levels, nor do science studies of any kind suggest that all scientific ontologies need to operate with a finite number of distinct levels to account for their respective entities, nor is it a result of fundamental theory. In my view, the level-picture seems to be a reductionist premise which was spared by Ellis's critique and utilised for the threefold distinction of same-level, bottom-up and top-down causation and to grant every emergent natural entity a context in which it can "really" exist with causal powers/effectiveness. If that were true, there could be a possibility for Ellis to drop the level-picture but keep his theory intact nevertheless.

11.3.2 Why We Should Give Up the Level-Picture

From our discussion of these arguments, it follows that Ellis's theory of causation and emergence does not have good reasons to assume the level-picture. Yet the same arguments suggest that it is not so much a cornerstone of the theory as it is a contingent addendum to the theory. Therefore, we can apply Occam's razor and see if we might dispense with this hypothesis. As will be shown below in section 4. in more detail, we can find alternative, weaker assumptions which can take over the role of the level-picture. The first argument is that it would thus simply be good scientific practice to jettison the level-picture in accordance with the principle of scientific parsimony.

Another important reason to abandon the level-picture derives from our considerations of Ellis's antireductionist argument above. It is not only a bone of contention for reductionists but a possible weakness to be exploited by them in arguments against Ellis. The level-picture works in favour of reductionism. It states that the basic ontological structure of nature is a single branching tree with fundamental physics at its root, only two branches (the animate and the inanimate), and a finite, neat number of distinct steps in complexity from the most basic components to the most elaborate systems, such as the universe, global society or the human mind. And moreover, it invites the idea of total reduction, because so many systems in one distinct level are just weakly emergent from subsystems on the next lower level (Luu and Meißner 2021; Castellani 2002).[8]

---

[8] Ellis's argument stands in a long line of antireductionist arguments that follow Anderson's classic criticism of the *constructionist hypothesis*, that is "the ability to start from those [fundamental] laws and reconstruct the universe", because "at each new level of complexity entirely new properties appear". (Anderson 1972, p. 393) Many philosophers and scientists with the same agenda follow his lead (see Meyr 1988; Nagel 1998; Castellani 2002; Ellis 2006; Chalmers 2006; Corning 2012). Thomas Nagel's argument against reductionism, for example, introduces levels in nature as joints in nature (Nagel 1998, pp. 7–12): Our explanations and theories always focus on a specific level, depending on which scale we are looking at. According to Nagel, not all natural kinds and laws can be defined in this reductionist way, functional properties in biology being a case in point. One consequence of his argument is that levels in nature are real because the natural kinds of the higher level of biology cannot even in principle be defined by the fundamental level of physics. So we have at least two distinct levels: The fundamental level of physics without functional properties and the higher level of biology with them. Only an ontological reductionist could claim the absence of distinct levels of nature. Nagel does not find a joint in nature which separates the level of larger entities with functional properties and the level of smaller entities without them. (Though even this

The third argument for abandoning the level-picture is that multiple parallel and intersecting modular hierarchies of structure and a gradient of complexity are more compatible with the local ontologies of disunified and specialised sciences with overlapping contexts of interest. The transdisciplinary and intersecting structure of different scientific disciplines and sub-disciplines is captured more accurately and easily without the level-picture. A rigorous proof of this would have to venture deep into both the philosophical meta-theory of current science and empirical studies of research programmes, which are tasks for another time. Such a proof would likely constitute a further, positive argument for favouring a picture of nature without levels.

11.4 Top-Down Causation Without Levels

In this section, our question is what becomes of Ellis's theory of (top-down) causation and emergence if we would decide to abandon the level-picture of nature. It will be answered in two steps: Step one revisits the eight fundamental points, as reconstructed from Ellis's theory in Sect. 11.2.3 above. The aim is to reformulate all eight in light of the criticism of the level-picture, and subsequently prove their value as a foundation for the core concepts of Ellis's theory. Accordingly, the second step is to sketch this possibility.

11.4.1 The Eight Points Revisited

To reformulate all eight points so that they function as a foundation for Ellis's theory without using the level-picture of nature, we must keep in mind three things: (1) We cannot eliminate the notion of levels completely, because the MHS-model requires the "H", the numerous hierarchies, the whole-part relationships of natural system. Without hierarchies, there will be no bottom-up or top-down causation, which is perhaps the key feature of Ellis's account. (2) At the same time, we have to be careful not to reintroduce the level-picture by a sleight-of-hand or by accident. We thus have to, firstly, ensure levels are not construed as necessarily distinct, and secondly, accommodate multiple (possibly infinitely many) hierarchies and therefore levels, and thirdly, secure any other notion which might play the role of levels, by making sure it is not necessarily distinct or only conceived in a finite number either. (3) It would be best if we would not introduce new concepts into Ellis's theory but employ the conceptual means already at our and his disposal, i.e. concepts used in the reconstruction of his position and arguments.

---

is debatable; for a definition of every physical property as dispositional and second-order functional property see Yates 2012.)

I will first present the reformulated eight points, then revisit L1–L8 and compare them with their new counterparts. Here are the alternative ("A" instead of "L") eight points:

A1 Levels are layers in the constitutional hierarchy of fully or partly concrete (meaning existing in spacetime) objects or systems and real in the sense of existing independently of our (or any) knowledge of them. There is an infinite number of parallel and intersecting hierarchies for all such entities and therefore infinitely many levels (from now on designated as CH-levels, for "constitutional hierarchy").

A2 Any fully or partly concrete entity can appear in a context with its own local ontology and behaves according to a set of principles and laws effective in that context, which can potentially be discovered by science.

A3 CH-Levels are ordered by the part-whole relationships of the hierarchy to which they belong. It is unclear if there is either a top or a bottom CH-level (or also which ones that would be) in any hierarchy. All fully or partly concrete entities in and across contexts are ordered by their degree of complexity, their properties and their relations with one another, which are all subjects of the local ontology of a context or of an intersectional ontology of an intersection of many contexts.

A4 To describe different local or intersectional ontologies (and the associated entities they contain—the types of objects, systems, structures, properties, relations, laws and principles) we need to deploy different scientific languages, scales and variables suited to any one local or intersectional ontology. This results in a disunity of sciences where for every local or intersectional ontology we have one (or more) science studying it. It is important to note that, depending on the local ontology, scales and variables can be structures in SV-levels (for "scale and variable", they do not have to be structured in this way), which are epistemic and should not be confused with CH-levels.

A5 For any system S in context n we can identify parts of it either in context n or an intersection i of different contexts, constituting the next lower level in the constitutional hierarchy of S. For any concrete object O, located in context n, considered as a whole, there are proper parts of it either in context n or context-intersection i.

A6 CH-Levels are marked by emergent objects and systems with corresponding properties, structures, principles, laws, and assigned variables which can neither be found at, nor explained by, nor fully accounted for by the next lower or higher CH-levels.

A7 For any CH-level l (apart from a theoretically existing bottom-level) in any hierarchy H_s of a physical system S there exists a finite number of neighbouring lower CH-levels $l^{-1a}, l^{-1b}, \ldots$, which constitute the complete set of basic parts for all (sub)systems and objects on l. In this sense, the lowest physical CH-levels in any constitutional hierarchy of concrete objects or systems are the basis for everything else in nature.

A8 It is possible that many different lower CH-level states $S_{l^1}^1, S_{l^1}^2, \ldots$ of system S constitute S with exactly the same properties in its context n_s.

It is important to note that A1–A8 are formulated solely with conceptual resources from Ellis (2016) or my reconstruction of his position and arguments, so this step does not introduce any problematic notions replacing the level-picture. In all eight reformulated points, the reference to the level-picture is replaced, and some are slightly changed:

A1 preserves the necessary constitutional hierarchies with levels allowing same-level, top-down and bottom-up causation and Ellis's realist notion of them. The notion of CH-level is less demanding than that of the level-picture and can be proven by Ellis's first and third arguments. By contrast, A2 makes no reference to CH-levels but instead uses the notions of context and local ontology to replace the level-picture. This reflects the counterarguments. Otherwise, as with the relation of A1 to L1, A2 is very similar to L2 in all other respects. While L3 presented the order of levels, A3 presents both the order of CH-levels and the order of entities in different contexts. The word "ordered" regarding entities in A3 does not mean ordered into distinct levels, quite the contrary: There are overlaps between different contexts, meaning there is no total distinctness. While some local ontologies might be described as being structured into levels, other contexts can be organised in ways which are radically different.

A4 states the reason for the disunity of science (as well as different languages and variables) in the same manner as L4 but more convincingly: The mystery of why one level can have different sciences associated with it disappears; every divergent epistemic grasp of a different context merits a new discipline. Even closer to L5 is A5: it substitutes context for level in the sense of A2. Regarding A5, it is noteworthy what would follow if we were to postulate an absolute bottom level in the hierarchy of an object O, say a level of vibrating one-dimensional strings: Any such string would not be a system or a concrete object composed of elements. It would undoubtedly be some kind of object (because it would possess properties like being one-dimensional), and it would be located in spacetime. Yet A5 does not deny this: It denies the string is a system (because it lacks the internal structural requirements) and identifies the only proper part of the string as identical to the whole string, located in the given context.

While for the first five points, the L's match the A's relatively closely, this is not true for L6 and A6: One of the key features of the level-picture was that levels are distinct entities with clear boundaries, and this is just not so in all cases of CH-levels. The means of demarcation, on the other hand, has changed only ever so slightly: CH-levels can have more than one higher or one lower neighbouring CH-level. This last point is also the main difference between L7 and A7, which preserves Ellis's hypothesis that physics is fundamental for every concrete object in nature.[9] L8 and A8 are now completely different: Instead of following L8 and claiming a branching hierarchy of animate and inanimate nature with clear distinctions, based on the underlying idea that two higher levels can rest on one lower level, A8 is more modest, simply stating for CH-levels the key feature of equivalence-classes

---

[9] It would of course be possible to consider A7*, which would not make this claim, because it neither follows from the first sentence of A7 nor does the first sentence require the second as a premise. But with the second sentence, A7 is closer to L7.

and multiple-realisability with regard to the context and local ontology of the entity thus constituted. This can function as a good starting point for showing that the key features of Ellis's theory of causation can be defended without the level-picture, instead using A1–A8. This cannot be carried out in any detail here, but should suffice to assure us that the level-picture is at least not a prerequisite for such a theory *per se.*

11.4.2 Ellis's Account Reformulated

After revisiting both versions of the eight points regarding levels, CH-levels and contexts, we can see that Ellis's theory of causation, emergence and complexity can be restated without the level-picture. We can base it on an alternative set of key concepts and see that they will continue to work. To state the theory of Ellis (2016) in these new terms without any reliance on the level-picture is of course an ongoing task.

Definition of Cause and Effect The Woodward-Ellis definition of cause and effect is not harmed by the changes: If anything, the notion of *context* is better defined, and those of variables and intervention/manipulation are as unproblematic as in the level-picture.

Three Basic Types of Causation All basic types of causation can be introduced in the alternative picture. Same-level causation occurs on the same CH-level of any constitutional hierarchy of a given entity in a single context. It could also be called same-system causation between any two entities which are not connected by a mutual downward hierarchy. Let us take an example: Jupiter and Halley's Comet do not share subsystems, so they have separated downward hierarchies. It can remain open whether or not these two entities are causally connected by being modules in the same system on one CH-level (which Jupiter and Halley's Comet are as parts of our solar system). More often than not, there will be a gradient of complexity between two entities (as is the case for Jupiter and Halley's Comet): Yet we should not therefore be misled into thinking of this as a case of top-down causation. It is not a difference in complexity between two entities that constitutes their causal relationship as top-down or bottom-up; rather, it is the constitutive role one entity plays for the other entity as one of the modules in its structural hierarchy or their connection in a context or intersection. Top-down causation is accordingly the name for causal relation from a higher to a lower CH-level in a hierarchy, of the context or intersection on an entity, or both, or from a wider to a smaller context. Finally, bottom-up causation is the reverse of the previous case: from a lower to a higher CH-level, from an entity on a context or intersection or both, or from a smaller to a wider context. This also illuminates interesting subdivisions in bottom-up and top-down causation which could be connected to the five different types of top-down causation.

Genuine Complexity and the MHS-Model All cases of genuine complexity can be preserved and modelled in the alternative picture: Genuine complexity is an effect within the hierarchies of different systems and accords with the MHS-model. The MHS-model itself can be derived from the alternative points: It is *modular* due to the whole-part relationships and the fact that context can appear within contexts. The *hierarchies* are the constitutional hierarchies of the systems, and the structure is a result of how the different hierarchies are parallel to or intersect with one another within a single or many intersection contexts. All the further points (basis of physics, disunity of sciences, different laws, variables, scales, etc.) can all be taken more or less directly from the alternative picture.

Equivalence-Classes and Multiple Realisability These key features of Ellis's theory were not presented thus far. Suffice it to say that they can be formulated on the alternative foundations:

> **Multiple Representation**. In general, many lower [CH-]level states correspond to a single higher [CH-]level state [. . .], because a higher [CH-]level description H_1 is arrived at by ignoring the micro-differences between many lower [CH-]level states L_i, and throwing away a vast amount of lower [CH-]level information (coarse-graining). (Ellis 2016, p. 102)

Examples include "numerous microstates of particle positions and velocities correspond[ing] to a single macrostate of nitrogen gas with a pressure of one bar and a temperature of 20 K in a volume of 1 L." (Ellis 2016, p. 103) The idea is to prove multiple realisability, meaning that a higher-level state or property is realised by different heterogeneous states or properties at a lower level (Green and Batterman 2021), via equivalence-classes of physical systems: Almost infinitely many microstates correspond to a single macro state. In the alternative picture it becomes apparent that these are both hypotheses about CH-levels in a single context.

We could continue this list or take this paper as a reason to give a partial reformulation of Ellis's account on a larger scale. For now, the criticism of the level-picture and the evident functionality of the alternative must suffice.

11.5 Conclusion: New Realist Ontologies and Top-Down Causation

We can now see that how the conclusions presented in the introduction are justified:

(C1) Ellis's theory of causation hinges on eight essential assumptions concerning ontologically robust levels of nature. All eight were reconstructed and their function within his framework explored. Ellis has reasons to employ the level-picture, but:

(C2) Ellis's own arguments in favour of the level-picture fail, at least in the context of Ellis's theory. We contemplated various counterarguments and revealed the level-picture to be untenable. In doing so, we saw that this might not be all too damaging for Ellis because:

(C3) The failure of the level-picture of nature is not a fatal problem for Ellis's theory of (top-down) causation, complexity and emergence. To show this, we presented eight alternative foundational hypotheses that dispense with the level-picture and went on to sketch how one might use them to arrive at all of the key features of Ellis's account.

Finally, I want to make a suggestion, which cannot be substantiated here: It seems likely that Ellis's theory of causation can be merged with a pluralist ontology or metaphysics provided the latter (1) were a thorough realism (which includes causal powers of abstract, mental or social objects) and (2) were to accommodate—at least in my alternative interpretation—local hierarchies which cannot be integrated into a global hierarchy or an overarching context. Both seem to be key features of Markus Gabriel's fields of sense ontology (Gabriel 2015). If Ellis insists on a monist philosophy of nature, which allows the possibility of a global hierarchy and an overarching context—the level-picture—he might be at odds with the general framework of Gabriel and more likely to take up a position like that elaborated by Quentin Meillassoux (Meillassoux 2008).

Appendix

I would like to thank Markus Gabriel and Otávio Bueno for the encouragement and opportunity to write this paper and helpful remarks, Alex Englander and Max Kistler for their comments on an earlier version of this paper, and Noemi Stelzig for all her support during the process of writing.

References

Anderson, P. W. (1972). More is different: Broken symmetries and the nature of the hierarchical structure of science. *Science, 177*, 393–396.

Beebee, H. (2014). Radical Indeterminism and top-down Causation. *Res Philosophica, 91*, 537–545.

Cartwright, N. (2006). Where is the theory in our "Theories" of causality? *Journal of Philosophy, 103*(2), 56–66.

Cartwright, N. (2007). *Hunting causes and using them. Approaches in philosophy and economics.* Cambridge: Cambridge University Press.

Castellani, E. (2002). Reductionism, emergence, and effective field theories. *Studies in History and Philosophy of Modern Physics, 33*, 251–267.

Chalmers, D. (2006). Strong and weak emergence. In P. Clayton & P. Davies (Eds.), *The Re-emergence of emergence. The emergentist hypothesis from science to religion* (pp. 244–256). Oxford: Oxford University Press.

Corning, P. (2012). The Re-emergence of emergence, and the Causal role of synergy in emergent evolution. *Synthese, 185*(2), 295–317.

de Vreese, L. (2010). Disentangling causal pluralism. In R. Vanderbeeken & B. D'Hooghe (Eds.), *Worldviews, science and Us: Studies of analytical metaphysics: A selection of topics from a methodological perspective* (pp. 207–223). Singapur: World Scientific Publication.

Ellis, G. F. R. (2006). On the nature of emergent reality. In P. Clayton & P. Davies (Eds.), *The Re-emergence of emergence. The emergentist hypothesis from science to religion* (pp. 79–107). Oxford: Oxford University Press.

Ellis, G. F. R. (2008). On the nature of causation in complex systems. *Transactions of the Royal Society of South Africa, 63*, 69–84.

Ellis, G. F. R. (2009). Top-down causation and the human brain. In N. Murphy, G. F. R. Ellis, & T. O'Connor (Eds.), *Downward causation and the neurobiology of free will* (pp. 63–81). Berlin: Springer.

Ellis, G. F. R. (2012). Top-down causation and emergence: Some comments on mechanisms. *Interface Focus, 2*, 126–140.

Ellis, G. F. R. (2013). The arrow of time and the nature of spacetime. *Studies in History and Philosophy of Science Part B: Modern Physics, 44*, 242–262.

Ellis, G. F. R. (2016). *How can physics underlie the mind? Top-down causation in the human context.* Berlin: Springer.

Ellis, G. F. R., Auletta, G., & Jaeger, L. (2008). Top-down causation by information control: from a philosophical problem to a scientific research program. *Interface, 5*(27), 1159–1172.

Ellis, G. F. R., Noble, D., & O'Connor, T. (2012). Top-down causation: An integrating theme within and across the sciences? *Interface Focus, 2*, 1–3.

Ellis, G. F. R., & Drossel, B. (2018). Contextual wavefunction collapse: An integrated theory of quantum measurement. *New Journal of Physics, 20*, 113025.

Ellis, G. F. R., & Drossel, B. (2019). How downwards causation occurs in digital computers. *Foundations of Physics, 49*(11), 1253–1277.

Ellis, G. F. R., & Gabriel, M. (2021). Physical, logical, and mental top-down effects. In M. Gabriel & J. Voosholz (Eds.), *Top-down causation and emergence* (vol. 439, pp. 3–38). Cham: Springer.

Emmeche, C., Køppe, S., & Stjernfelt, F. (1997). Explaining emergence—Towards an ontology of levels. *Journal for General Philosophy of Science, 28*, 83–119.

Emmeche, C., Køppe, S., & Stjernfelt, F. (2000). Levels, emergence, and three versions of downward causation. In P. Bøgh Andersen, C. Emmeche, N. Ole Finnemann, & P. Voetmann Christiansen (Eds.), *Downward causation. minds, bodies and matter* (pp. 13–34). Århus: Aarhus University Press.

Gabriel, M. (2015). *Fields of sense. A new realist ontology.* Edinburgh: Edinburgh University Press.

Green, S., & Batterman, R. W. (2021). Causal slack and Top.Down causation. universality and functional equivalence in physics and biology. In M. Gabriel & J. Voosholz (Eds.), *Top-down causation and emergence* (vol. 439, pp. 39–63). Cham: Springer.

Guillot, T., Stevenson, D. J., Hubbard, W. B., & Saumon, D. (2004). The interior of jupiter. In F. Bagenal, T. E. Dowling, & W. B. McKinnon (Eds.) *Jupiter: The planet, satellites and magnetosphere* (pp. 35–57). Cambridge: Cambridge University Press.

Hulswit, M. (2005). How causal is downward causation? *Zeitschrift für allgemeine Wissenschaftstheorie, 36*(2), 261–287.

Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research, 52*, 1–26.

Kim, J. (1999). Making sense of emergence. *Philosophical Studies, 95*, 3–36.

Liseau, R., Larsson, B., Lunttila, T., Olberg, M., Rydbeck, G., Bergman, P., et al. (2015). Gas and dust in the star-forming region ρ Oph A. The dust opacity exponent β and the gas-to-dust mass ratio *g2d*. *A&A, 578*, A313.

Longworth, F. (2006). Causation, pluralism and responsibility. *Philosophica, 77*, 45–68.

Luu, T., Meißner, U.-G. (2021). On the topic of emergence from an effective field theory perspective. In M. Gabriel & J. Voosholz (Eds.), *Top-down causation and emergence* (vol. 439, pp. 101–114). Cham: Springer.

Meillassoux, Q. (2008). *After finitude. An essay on the necessity of contingency.* New York: Continuum.

Meyr, E. (1988). The limits of reductionism. *Nature, 331*, 475.

Nagel, T. (1998). Reductionism and antireductionism. In G. R. Bock, J. A. Goode (Eds.), *The limits of reductionism* (pp. 3–14). Chichster: Wiley.

Paolini Paoletti, M., & Orilia, F. (2017). Downward causation: An opinionated introduction. In M. Paolini Paoletti & F. Orilia (Eds.), *Philosophical and scientific perspectives on downward causation* (pp. 1–21). London, New York, USA: Routledge.

Paul, L. A. (2009). Counterfactual theories. In H. Beebee, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook on causation* (pp. 159–184). Oxford: Oxford University Press.

Price, H. (2017). Causation, intervention and agency—Woodward on Menzies and Price. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a difference* (pp. 73–98). Oxford: Oxford University Press.

Price, H., & Menzies, P. (1993). Causation as a secondary quality. *British Journal of Philosophy of Science, 44*, 187–203.

Psillos, S. (2010). Causal pluralism. In R. Vanderbeeken & B. D'Hooghe (Eds.), *Worldviews, science and Us: Studies of analytical metaphysics: A selection of topics from a methodological perspective* (pp. 131–151). Singapur: World Scientific Publication.

Siriwardena, S. (2019). Old problems for the agency theory of causal discourse. *Erkenntnis, 84*(4), 939–951.

Strevens, M. (2007). Essay review of Woodward, "Making Things Happen". *Philosophy and Phenomenological Research, 74*, 233–249.

Strevens, M. (2013). Causality reunified. *Erkenntnis, 78*, 229–320.

Williamson, J. (2006). Causal pluralism versus epistemic causality. *Philosophica, 77*, 69–96.

Woodward, J. (2003). *Making things happen. A theory of causal explanation.* Oxford: Oxford University Press.

Woodward, J. (2007). Causation with a human face. In H. Price & R. Corry (Eds.), *Causation and the constitution of reality* (pp. 66–105). Oxford: Oxford University Press.

Woodward, J. (2008). Response to Strevens. *Philosophy and Phenomenological Research, 77*, 193–212.

Woodward, J. (2021). Downward causation defended. In M. Gabriel & J. Voosholz (Eds.), s *Top-down causation and emergence* (vol. 439, pp. 217–251). Cham: Springer.

Yates, D. (2012). The essence of dispositional essentialism. *Philosophy and Phenomenological Research, 87*, 93–128.