

Christian List

Levels: descriptive, explanatory, and ontological

Article (Accepted version) (Refereed)

Original citation:

List, Christian (2018) *Levels: descriptive, explanatory, and ontological*. Noûs. ISSN 0029-4624

DOI: [10.1111/nous.12241](https://doi.org/10.1111/nous.12241)

© [10.1111/nous.12241](https://doi.org/10.1111/nous.12241)

This version available at: <http://eprints.lse.ac.uk/87591/>

Available in LSE Research Online: April 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Levels: Descriptive, Explanatory, and Ontological

Christian List*

London School of Economics

March 2016 / December 2017

To appear in *Noûs*

Abstract

Scientists and philosophers frequently speak about levels of description, levels of explanation, and ontological levels. In this paper, I propose a unified framework for modelling levels. I give a general definition of a *system of levels* and show that it can accommodate descriptive, explanatory, and ontological notions of levels. I further illustrate the usefulness of this framework by applying it to some salient philosophical questions: (1) Is there a linear hierarchy of levels, with a fundamental level at the bottom? And what does the answer to this question imply for physicalism, the thesis that everything supervenes on the physical? (2) Are there emergent properties? (3) Are higher-level descriptions reducible to lower-level ones? (4) Can the relationship between normative and non-normative domains be viewed as one involving levels? Although I use the terminology of “levels”, the proposed framework can also represent “scales”, “domains”, or “subject matters”, where these are not linearly but only partially ordered by relations of supervenience or inclusion.

*This paper articulates a framework that has been implicit—to varying degrees—in some of my previous works, but has never been spelt out in full. Relevant works include List and Menzies (2009, 2010), List and Pettit (2011), List and Spiekermann (2013), List (2014), List and Pivato (2015a,b), and Dietrich and List (2016, Section 8). I wish to record my intellectual debt to all of my co-authors of these works: the late Peter Menzies, Philip Pettit, Kai Spiekermann, Marcus Pivato, and Franz Dietrich. I am also grateful to Campbell Brown, Jeremy Butterfield, Johannes Himmelreich, Marcus Pivato, Wlodek Rabinowicz, and an anonymous referee for detailed written comments, and to Richard Bradley, Elizabeth Coppock, Meir Hemmo, George Musser, John Norton, Mike Otsuka, Bryan Roberts, Orly Shenker, and Laura Valentini for helpful discussions or feedback. Finally, I thank conference and seminar audiences at the University of Munich, Lund University, the University of Cambridge, and the LSE. My work was supported by a Leverhulme Major Research Fellowship.

1 Introduction

Scientists as well as philosophers frequently use notions such as *levels of description*, *levels of explanation*, and *ontological levels*.¹ Although it is widely held that everything in the world is ultimately the product of fundamental physical processes, it is also widely recognized that, for many purposes, the right level of description or explanation is not the fundamental physical one, but a “higher” level, which abstracts away from microphysical details. Chemistry, biology, geology, and meteorology would all get bogged down with an informational or computational overload if they tried to explain the phenomena in their domains by modelling the behaviour of every elementary particle, instead of focusing on “higher-level” properties and regularities. For instance, it would be hopeless to try to understand an organism or an ecosystem at the level of the billions of elementary particles of which it is composed, rather than at the macroscopic level of its biological functioning.

Similarly, cognitive scientists tend to assume that human psychology is better understood at the level of the mind (the cognitive-psychological level) than at the level of the brain (the neuro-physiological level),² just as it is easier to understand a word processor at the software level than at the hardware level, where gazillions of electrons flow through microchips. And for many social-scientific purposes, the right level of description is not the “micro”-level of individuals, but a social level, involving “macro”-variables. Economists and political scientists would have a hard time modelling the economy or political systems if they tried to describe the behaviour of every individual market participant or every single citizen.³

Given the ubiquity of higher-level descriptions in science, some philosophers ask whether the world itself might be “layered” or “stratified into levels”, where different levels are organized hierarchically, perhaps with a fundamental level at the bottom.⁴ According to a levelled ontology, the levels in question are not just *levels of description* or *explanation*, but *levels of reality* or *ontological levels*. When we employ different descriptive or explanatory levels in science, on this picture, these correspond to different ontological levels: they are “epistemic markers” of something “ontic”. Historically, a levelled ontology was central to the British Emergentism of the 19th and early 20th centuries.⁵

¹E.g., Oppenheim and Putnam (1958), Fodor (1974), Owens (1989), Beckermann, Flohr, and Kim (1992), Dupré (1993), Kim (1998, 2002), Schaffer (2003), Floridi (2008), and Ellis, Noble, and O’Connor (2012).

²For a classic discussion, see Putnam (1967). See also Bechtel (1994) and Davies (2000).

³This is despite the popularity of methodological individualism, the view that social phenomena should be explained at the level of individuals. On the level of explanation in the social sciences, see, e.g., Kincaid (1986), Sawyer (2002, 2003), List and Pettit (2011), List and Spiekermann (2013), and Haslanger (2016).

⁴For discussions of this “layered” or “stratified” view of the world, see Kim (2002) and Schaffer (2003).

⁵For a review, see, e.g., O’Connor and Wong (2015).

How should we think about levels? Are notions such as *levels of description*, *levels of explanation*, or *ontological levels* mere metaphors, or can we make literal sense of them? Jaegwon Kim, for instance, suggests that “talk of levels may turn out to be only a figure of speech, a harmless but suggestive metaphor”, but he challenges us to “begin [our analysis] by taking the levels talk and its attendant metaphysics seriously” and to “try to make sense of it—or as much sense as we can”.⁶ The aim of this paper is to respond to this challenge, by proposing a unified framework for modelling levels, whether interpreted epistemically or ontically.

I introduce an abstract definition of a *system of levels* and look at a number of examples of systems of levels that fit this definition. Some of them capture descriptive or explanatory notions of levels, others ontological ones. The framework can also make sense of the idea that a level of description may be a marker of an ontological level. My examples build on recent discussions of levels in the literature; the general definition that subsumes them is new and is inspired by category theory.⁷ Crucially, the framework can accommodate the possibility (also raised by Jaegwon Kim) that levels do not form a linear hierarchy, but are only partially ordered by the “higher than” relation. Although I use the conventional terminology of “levels”, the framework can equally be interpreted as capturing the relationship between different “scales”, “domains”, or “subject matters”.⁸

I will illustrate the usefulness of the framework by bringing it to bear on some salient philosophical questions. I will consider questions about fundamentality, emergence, reduction, and normativity. On fundamentality: is there a fundamental level on which all other levels depend?⁹ And what does the answer to this question imply for physicalism, the thesis that everything supervenes on the physical? On emergence: are there emergent higher-level properties that are not explainable in lower-level terms? On reduction: are higher-level descriptions reducible to lower-level ones? On normativity: can we represent the relationship between normative and non-normative domains as one involving levels?

My aim is not to offer comprehensive discussions of all these questions. It would be impossible to do so in a single paper. My aim is rather to show how the proposed framework allows us to frame the relevant debates in helpful ways.

⁶See Kim (2002, p. 3).

⁷The closest precursors to the present work are Butterfield (2011), List (2014), and List and Pivato (2015a,b). Himmelreich (2015, Appendix B) also explicates the idea of levels, building on List and Pivato (2015a). Category theory goes back to Eilenberg and MacLane (1945). For a philosophical survey, see Marquis (2015). Halvorson (2016) gives an account of scientific theories as categories. Category theory has also been suggested as a framework for thinking about levels of description in cognitive and brain science (Gómez Ramirez 2014), though in a different way from the one proposed here.

⁸These terms can be found, e.g., in Wilson (2010), Kim (2002), and Lewis (1988), respectively.

⁹This is the question asked by Schaffer (2003). See also Block (2003).

2 A System of Levels: An Abstract Definition

I begin by giving an abstract definition of a system of levels. In the next section, I discuss some more concrete instances of this definition. In some cases, levels have an epistemic or explanatory interpretation, in others an ontological one.

A *system of levels* is a pair $\langle \mathcal{L}, \mathcal{S} \rangle$, defined as follows:

- \mathcal{L} is a class of objects called *levels* (which will be given more structure later), and
- \mathcal{S} is a class of mappings between levels, called *supervenience mappings*, where each such mapping σ has a *source level* L and a *target level* L' and is denoted $\sigma : L \rightarrow L'$,

such that the following conditions hold:

- (S1) \mathcal{S} is closed under *composition* of mappings, i.e., if \mathcal{S} contains $\sigma : L \rightarrow L'$ and $\sigma' : L' \rightarrow L''$, then it also contains the composite mapping $\sigma \bullet \sigma' : L \rightarrow L''$ defined by first applying σ and then applying σ' (where composition is *associative*¹⁰);
- (S2) for each level L , there is an *identity mapping* $\mathbf{1}_L : L \rightarrow L$ in \mathcal{S} , such that, for every mapping $\sigma : L \rightarrow L'$, we have $\mathbf{1}_L \bullet \sigma = \sigma = \sigma \bullet \mathbf{1}_{L'}$;
- (S3) for any pair of levels L and L' , there is at most one mapping from L to L' in \mathcal{S} .

Interpretationally, when the mapping $\sigma : L \rightarrow L'$ is contained in \mathcal{S} , this means that level L' *supervenies* (or *depends*) on level L . We then call L' the *supervenient* (or *higher*) level and L the *subvenient* (or *lower*) level, according to σ . Note that the lower level is on the left-hand side of the arrow in the expression “ $\sigma : L \rightarrow L'$ ”, and the higher level is on the right-hand side.

Supervenience is usually understood as a relation of determination or necessitation: one set of facts (e.g., the facts at level L') is said to “supervene” on a second set (e.g., the facts at level L) if the second set of facts determines or necessitates the first, i.e., a change in the first set of facts is impossible without any change in the second.¹¹ So, we might also call σ a *determination or necessitation mapping*. There can be different notions of supervenience, corresponding to different modes of determination or necessitation,

¹⁰Formally, $\sigma \bullet (\sigma' \bullet \sigma'') = (\sigma \bullet \sigma') \bullet \sigma''$ whenever $\sigma : L \rightarrow L'$, $\sigma' : L' \rightarrow L''$, and $\sigma'' : L'' \rightarrow L'''$.

¹¹On the concept of supervenience, see, e.g., Kim (1984, 1987) and Horgan (1993). The concept can be traced back to G. E. Moore’s and R. M. Hare’s discussions of the relationship between moral properties and non-moral properties (though Moore did not use the term) and to Donald Davidson’s discussion of the relationship between mental and physical properties. The British Emergentists used the term as well, but not in the modern technical sense. See also McLaughlin and Bennett (2014).

for example metaphysical or nomological ones.¹² The formal framework is compatible with different such interpretations. Moreover, while I here adopt an interpretation in terms of “supervenience”, the framework is also compatible, in principle, with other interpretations of the mappings in \mathcal{S} under which conditions (S1) to (S3) are satisfied.¹³

The present way of formalizing the relationship between levels is in line with Kim’s suggestion to define the “higher than” relation in terms of supervenience: for any two levels L and L' , “ L' is *higher than* L iff L' supervenes on L but L does not supervene on L' ”.¹⁴ As Kim also notes, it may be that levels “are not always comparable to each other in terms of higher and lower”.¹⁵ I return to this point later.

The three conditions on a system of levels I have introduced capture some familiar properties of the notion of supervenience. Condition (S1) entails that supervenience is *transitive*: if L'' supervenes on L' , and L' supervenes on L , then L'' also supervenes on L . Condition (S2) entails that every level supervenes on itself; trivially, supervenience is *reflexive* (though nothing of substance hangs on this). Condition (S3) entails that, whenever L' supervenes on L , the relation in which L and L' stand is *unique*; this is in keeping with the idea of supervenience as a relation of determination or necessitation.¹⁶ The three conditions jointly entail a fourth condition:

(S4) if \mathcal{S} contains a mapping $\sigma : L \rightarrow L'$ and a mapping $\sigma' : L' \rightarrow L$, then $\sigma \bullet \sigma' = \mathbf{1}_L$.¹⁷

Informally, if two levels supervene on one another (which might perhaps never happen if these levels are distinct), then the composite of the relations in which they stand must be the identity relation.

¹²One set of facts supervenes metaphysically on a second if a change in the first set of facts is metaphysically impossible without a change in the second. One set of facts supervenes nomologically on a second if this is nomologically impossible, i.e., impossible relative to the appropriate laws of nature.

¹³For instance, a class of “levels” or “domains”, together with a class of “derivability”, “scrutability”, or “reducibility” relations, might also form a system $\langle \mathcal{L}, \mathcal{S} \rangle$ with properties (S1) to (S3). Obviously, the resulting applications of the framework would have to be interpreted somewhat differently than the ones considered here. I thank Bryan Roberts and an anonymous referee for prompting me to clarify this point. The term “scrutability” comes from Chalmers (2012).

¹⁴The quoted definition comes from Kim (2002, p. 9), with the labels “ L' ” and “ L ” substituted for “M” (as in “mental”) and “N” (as in “neural”), respectively.

¹⁵See Kim (2002, p. 10).

¹⁶By contrast, grounding relations—recently much discussed—have slightly different formal properties. While grounding is arguably transitive, it is irreflexive and asymmetric. See, e.g., Schaffer (2009).

¹⁷To see this, suppose that \mathcal{S} contains a mapping $\sigma : L \rightarrow L'$ and a mapping $\sigma' : L' \rightarrow L$. By (S1), \mathcal{S} will then also contain the composite mapping $\sigma \bullet \sigma' : L \rightarrow L$. By (S2), \mathcal{S} contains an identity mapping $\mathbf{1}_L : L \rightarrow L$ for level L . By (S3), \mathcal{S} contains at most one mapping from L to L . Since both $\sigma \bullet \sigma'$ and $\mathbf{1}_L$ are mappings from L to L , they must coincide; otherwise \mathcal{S} would contain more than one mapping from L to L . This establishes (S4).

In algebraic terms, the pair $\langle \mathcal{L}, \mathcal{S} \rangle$, subject to conditions (S1) and (S2), is a structure called a “category”. Generally, a *category* is a pair consisting of a class of objects and a class of mappings between objects, often called “arrows” or “morphisms”, where conditions (S1) and (S2) hold. In the present context, the “objects” are levels, and the “arrows” or “morphisms” are supervenience mappings.¹⁸ Categories that also satisfy condition (S3), as in the case of $\langle \mathcal{L}, \mathcal{S} \rangle$, are called “posetal categories”. The category-theoretic way of representing systems of levels allows us to express structural relationships between different such systems.

For example, one system of levels, $\langle \mathcal{L}, \mathcal{S} \rangle$, is a *subsystem* of another, $\langle \mathcal{L}', \mathcal{S}' \rangle$, if

- $\mathcal{L} \subseteq \mathcal{L}'$ and $\mathcal{S} \subseteq \mathcal{S}'$, and
- composition and identity in $\langle \mathcal{L}, \mathcal{S} \rangle$ are defined as in $\langle \mathcal{L}', \mathcal{S}' \rangle$.¹⁹

To illustrate, one might imagine a disagreement among scientists or philosophers on how many levels of reality there are. Some think there are more levels than recognized by others. Perhaps some postulate the existence of an additional fundamental level below the commonly recognized microphysical level, while others deny it. Then we have a situation in which some people think that the correct system of levels is a subsystem of the one assumed by others: the two systems of levels are exactly alike, except that one contains more levels than the other.

More generally, there can be structure-preserving mappings between different systems of levels. These are called *functors*. A *functor*, F , from one system of levels, $\langle \mathcal{L}, \mathcal{S} \rangle$, to another, $\langle \mathcal{L}', \mathcal{S}' \rangle$, is a mapping which

- assigns to each level L in \mathcal{L} a corresponding level $L' = F(L)$ in \mathcal{L}' , and
- assigns to each supervenience mapping $\sigma : L \rightarrow L'$ in \mathcal{S} a corresponding supervenience mapping $\sigma' = F(\sigma)$ in \mathcal{S}' , where $\sigma' : F(L) \rightarrow F(L')$,

such that F preserves composition and identity.²⁰ The existence of a functor from one system of levels to another means that we can map the first system into the second in a way that preserves supervenience relationships. In the case in which one system

¹⁸This differs from Gómez Ramirez’s (2014) proposal, where each level—as opposed to a *system* of levels—is represented by a category (e.g., the category of neurons for the neuronal level, with synaptic paths playing the role of morphisms) and there are no supervenience mappings as morphisms.

¹⁹Note that $\langle \mathcal{L}, \mathcal{S} \rangle$ and $\langle \mathcal{L}', \mathcal{S}' \rangle$, qua systems of levels, must each satisfy (S1) to (S3).

²⁰Formally, for any two supervenience mappings σ and σ' in \mathcal{S} , where the target level of σ coincides with the source level of σ' , we have $F(\sigma \bullet \sigma') = F(\sigma) \bullet F(\sigma')$; and for any identity mapping $\mathbf{1}_L$ in \mathcal{S} , we have $F(\mathbf{1}_L) = \mathbf{1}_{F(L)}$, where $\mathbf{1}_{F(L)}$ is the identity mapping in \mathcal{S}' for level $F(L)$.

is a subsystem of another, there trivially exists a functor from the first system to the second, but there can also be functors in less trivial cases. For instance, I later consider the possibility that there may be a functor from a system of *levels of description* to a system of *ontological levels*. If two systems of levels, $\langle \mathcal{L}, \mathcal{S} \rangle$ and $\langle \mathcal{L}', \mathcal{S}' \rangle$, admit functors in both directions (i.e., from $\langle \mathcal{L}, \mathcal{S} \rangle$ to $\langle \mathcal{L}', \mathcal{S}' \rangle$ and vice versa) which are inverses of each other, then this indicates that the two systems are structurally equivalent, at least as far as supervenience relations are concerned. The attraction of the present formalism is its generality, as I will now illustrate.

3 Four Instances of Systems of Levels

I will consider four instances of the general definition just given. Some of them are best interpreted as capturing levels of description or explanation; others are best interpreted as capturing ontological levels; and in some cases they admit both interpretations.

3.1 Levels of grain

I begin with a very simple example of a system of levels, which is generated by different ways of partitioning an underlying non-empty set Ω of possible worlds.²¹ (Alternatively, Ω could be a set of other items, such as the options a decision maker might be faced with.) Let \sim denote an equivalence relation on Ω (a reflexive, symmetrical, and transitive relation). Any such relation \sim partitions Ω into some non-empty, pairwise disjoint, and jointly exhaustive equivalence classes, each of which consists of worlds (or items) that are equivalent with respect to \sim . Let Ω_{\sim} be the resulting set of equivalence classes. We call Ω_{\sim} a *partition* of Ω . We obtain the finest partition if \sim is the identity relation; here, each element of Ω forms a singleton equivalence class by itself. We obtain the coarsest partition if \sim is the total relation, under which all elements of Ω fall into the same equivalence class. Non-trivial partitions lie in between these two extremes.

For any two partitions Ω_{\sim} and Ω_{\approx} , we say that Ω_{\sim} is *at least as fine-grained* as Ω_{\approx} if each equivalence class in Ω_{\approx} is a union of equivalence classes in Ω_{\sim} . The relation “at least as fine-grained as” partially orders partitions. Whenever Ω_{\sim} is at least as fine-grained as Ω_{\approx} , we define a function $\sigma : \Omega_{\sim} \rightarrow \Omega_{\approx}$ that assigns to each equivalence class in Ω_{\sim} the equivalence class in Ω_{\approx} in which it is included.

It is easy to see that we get a system of levels if we define the pair $\langle \mathcal{L}, \mathcal{S} \rangle$ as follows:

²¹On the idea of identifying levels with partitions, see Himmelreich (2015, Appendix B). He develops a version of this idea adapting the framework of worlds-as-histories from List and Pivato (2015a).

- \mathcal{L} is some non-empty set of partitions of Ω , perhaps the set of all logically possible partitions;
- \mathcal{S} consists of every function $\sigma : \Omega_{\sim} \rightarrow \Omega_{\approx}$ under the definition just given, where Ω_{\sim} and Ω_{\approx} are elements of \mathcal{L} such that Ω_{\sim} is at least as fine-grained as Ω_{\approx} .

A *level*, here, is a particular way of partitioning the underlying space of possibilities (worlds or items) into equivalence classes, such that we do not distinguish between members of the same equivalence class.

The most natural interpretation of such *levels of grain* is an epistemic one: different levels correspond to different ways of perceiving or representing the world. In decision theory, for example, an agent’s *level of awareness* is often modelled in this way.²² Someone’s *awareness* is defined in terms of the distinctions he or she is able to draw. The agent is *aware* of some feature of the world (or a feature of some item) if and only if he or she is able to distinguish worlds (or items) with that feature from ones without it. The more features an agent is aware of, the more distinctions between worlds (or items) he or she is able to draw. Greater awareness thus corresponds to the adoption of a more fine-grained partition of the space of possibilities, lesser awareness to the adoption of a more coarse-grained partition.

Another interpretation of a system of levels of grain can be found in David Lewis’s notion of a “subject matter”.²³ Informally, a *subject matter* is—or at least picks out—a part of the world: the part that has to do with that subject matter. As an illustration, Lewis says: “the 17th Century is a subject matter, and also a part of this world... [T]wo worlds are alike with respect to the 17th Century iff their 17th Centuries are exact intrinsic duplicates”.²⁴ More generally, two worlds are alike with respect to a particular subject matter if and only if their relevant parts coincide. Demography, biology, and climate are all subject matters in this sense. Formally, Lewis defines a *subject matter* as an equivalence relation on the set of possible worlds, i.e., as a partition of Ω as discussed here. So, the subject matter “demography” partitions the set Ω into equivalence classes of worlds that are alike with respect to demography. Similarly, the subject matters “biology” and “climate” partition Ω , respectively, into equivalence classes of biologically indistinguishable worlds and into equivalence classes with the same climate. Lewis also defines the notion of “inclusion of subject matters”: one subject matter *includes* another

²²See, e.g., Modica and Rustichini (1999). In a recent working paper, Dietrich (2017) has proposed a model of decision-making under uncertainty in which an agent’s subjective conceptualization of outcomes and states takes the form of appropriate partitions of some underlying space of possibilities.

²³See, e.g., Lewis (1988).

²⁴He adds “or neither one has a 17th Century”. See Lewis (1988, p. 161).

if the former—understood as an equivalence relation—is at least as fine-grained as the latter. It should be evident that any set of Lewisian subject matters, together with the associated inclusion relations, forms a system of levels of grain as I have defined it.

3.2 Ontological levels

As already noted, it is sometimes suggested that the world itself is stratified into levels.²⁵ I will here explicate such a levelled ontology in terms of the idea that there is not just a single set of possible worlds (“possible worlds *simpliciter*”), but different such sets, which encode facts at different levels. Specifically, on this picture, each level is associated with its own set of all possible *level-specific* worlds.

To explain this, let me begin by revisiting the standard notion of a possible world. As usually defined, a *possible world* is a full specification of the way the world might be. It encodes all the facts that obtain at that world. To capture the idea that facts may belong to particular levels—physical, chemical, biological, etc.—we must amend this definition, so as to make the level-specificity of facts explicit. We can define a *possible world at a particular level* as a full specification of the way the world might be *at that level*. Worlds at the physical level thus encode the totality of physical facts; worlds at the chemical or biological levels encode the totality of chemical or biological facts; worlds at the psychological or social levels encode the totality of psychological or social facts; and so on. Then each level is associated with the set of all possible worlds at that level.

To the extent that the same higher-level facts are compatible with different lower-level facts (e.g., the biological or chemical facts leave some physical facts open), we can think of higher-level worlds as “thinner”—less richly specified—than lower-level worlds, and we can think of lower-level worlds as “thicker”. From a lower-level perspective, a higher-level world thus looks like a partial or incomplete specification of the world, which leaves certain facts (namely lower-level ones) indeterminate. But it is important to stress that, *at the higher-level itself*, a higher-level world encodes the totality of facts *at that level*.

The relationship between lower-level and higher-level worlds is one of *supervenience with multiple realizability*.²⁶ Supervenience means that each lower-level world determines

²⁵Recall the discussions in Kim (2002) and Schaffer (2003).

²⁶This way of thinking about levels is consistent with Block’s (2003). He suggests that different sets of properties characterize different levels (p. 142): the psychological properties characterize the psychological level; the neurophysiological properties characterize the neuroscientific level; and so on. “These distinct branches of science and their associated families of properties plausibly form a supervenience hierarchy”; this yields a “notion of level ... keyed to branches of science” (ibid.). I will later show that we need not assume a linear hierarchy. For an influential discussion of supervenience and multiple realizability (especially in the relationship between the physical and mental levels), see Yablo (1992).

a corresponding higher-level world: the lower-level facts, say the physical ones, determine the higher-level facts, say the chemical ones. By fixing all physical properties, we necessarily fix all chemical properties, in this example. Multiple realizability means that different lower-level worlds can correspond to the same higher-level world: different configurations of lower-level facts can give rise to or necessitate the same higher-level facts. For instance, many different states of the individual water molecules in a flask can instantiate the same aggregate state of the water. Similarly, different configurations of physical properties can instantiate the same chemical or biological properties.

To be clear: I am not defending any particular account of what makes some fact physical, chemical, or biological, and so on. Rather, I am making a structural point: whatever criteria we employ for demarcating the facts at a particular level, a world at that level is simply a total specification of the level-specific facts. Since the totality of facts at a higher level, say the psychological, is typically thinner than the totality of facts at a lower level, say the physical, we can coherently distinguish higher-level worlds from the underlying lower-level worlds (despite the dependency of the former on the latter).

For any given level, the set of all possible worlds at that level has the formal properties that we would expect such a set to have: its members are mutually exclusive, and they jointly exhaust all the possible ways the world could be at that level. Crucially, while in a “flat” ontology there is a single set Ω of all possible worlds, interpretable as the set of possible worlds *simpliciter*, the levelled picture suggests that there are several such sets, Ω , Ω' , Ω'' , and so on: one for each level.²⁷

We can formalize this ontological picture as a system of levels $\langle \mathcal{L}, \mathcal{S} \rangle$, where:

- \mathcal{L} is some non-empty class of sets of “level-specific” worlds (with each set of level-specific worlds non-empty);
- \mathcal{S} is some class of surjective (“onto”) functions of the form $\sigma : \Omega \rightarrow \Omega'$,²⁸ where Ω and Ω' are elements of \mathcal{L} , such that \mathcal{S} satisfies (S1), (S2), and (S3).

Each element Ω of \mathcal{L} can be interpreted as an *ontological level*: it is the set of possible worlds at that level. For example, \mathcal{L} might contain a set Ω corresponding to the physical level, a set Ω' corresponding to the chemical level, a set Ω'' corresponding to the biological level, and so on. A physical-level world—an element of Ω —settles all physical facts; a chemical-level world—an element of Ω' —settles all chemical facts; and so on.

²⁷Crucially, higher-level worlds need not be *defined* as equivalence classes of lower-level worlds; it only matters that they are each *realizable* by an equivalence class of lower-level worlds. See Section 3.3.

²⁸A function $\sigma : \Omega \rightarrow \Omega'$ is *surjective* if, for every $\omega' \in \Omega'$, there exists some $\omega \in \Omega$ such that $\sigma(\omega) = \omega'$.

To say that the chemical level supervenes on the physical, or that the biological supervenes on the chemical, is to say that the class \mathcal{S} of supervenience mappings contains one such mapping, $\sigma : \Omega \rightarrow \Omega'$, from the relevant lower level to the higher one, where σ maps each lower-level world $\omega \in \Omega$ to the higher-level world $\omega' \in \Omega'$ to which it corresponds. We then call ω a *lower-level realizer* of ω' .²⁹ The surjectivity of σ means that there are no possible worlds at the higher level that lack a possible lower-level realizer. For example, for a world to be chemically possible—i.e., contained in Ω' —it must have a possible physical realizer: there must be some $\omega \in \Omega$ such that $\sigma(\omega) = \omega'$. An instance of *multiple realizability* occurs when the function $\sigma : \Omega \rightarrow \Omega'$ is “many-to-one”: several distinct elements of Ω can realize the same element of Ω' .³⁰

We can use the present framework to express not only the idea that higher-level *worlds* in Ω' supervene on lower-level *worlds* in Ω , but also the idea that specific higher-level *facts* supervene on specific lower-level *facts*. Let $E' \subseteq \Omega'$ represent some higher-level fact, namely the fact that the higher-level world falls inside the set E' . We write $\sigma^{-1}(E')$ for the *inverse image* of E' under the supervenience mapping σ , defined as the set of all lower-level worlds that are mapped by σ to some element of E' . Formally, $\sigma^{-1}(E') = \{\omega \in \Omega : \sigma(\omega) \in E'\}$. We can then interpret $E = \sigma^{-1}(E')$ as the (weakest) *supervenience base* of E' . It consists of all the possible lower-level realizers of E' . Whether the higher-level fact E' obtains (i.e., the higher-level world falls inside E') depends on whether the lower-level fact E obtains (i.e., the lower-level world falls inside E). For instance, whether someone is in pain (a psychological fact) supervenes on whether this person’s brain is in a pain-generating state, such as “C-fibres firing” (a neuro-physiological fact).

While I have adopted a *world-based* understanding of ontological levels, it is worth contrasting this with a more traditional *entity-based* understanding. This defines ontological levels as *levels of entities*, where lower-level entities are the building blocks of higher-level entities.³¹ On that approach, “higher than” is defined not in *modal* terms, using the notion of supervenience, but in *mereological* terms, using notions such as composition and part-whole relationships.³² The microphysical level, for example, is the level

²⁹On my understanding of the term, we can speak not only of *lower-level properties realizing higher-level properties* (e.g., when certain molecular micro-properties realize a particular macro-property of a liquid), but also of *lower-level worlds realizing higher-level worlds* (namely when a lower-level world ω determines, via σ , a higher-level world ω'). Since level-specific worlds can be viewed as full specifications of level-specific properties, these two usages are entirely compatible with each other.

³⁰“Many-to-one” is the negation of injectivity. A function $\sigma : \Omega \rightarrow \Omega'$ is *injective* (“one-to-one”) if, for any $\omega, \omega' \in \Omega$, $\sigma(\omega) = \sigma(\omega')$ implies $\omega = \omega'$.

³¹For discussion, see, e.g., Schaffer (2003) and Kim (1993, 2002).

³²Similarly, Block (2003, pp. 141/142) distinguishes between “a notion of level keyed to objects” and one “keyed to relations among properties” and defends the latter. Himmelreich (2015, Appendix B) also

of elementary particles, while the macrophysical level is the level of larger aggregates. In their classic article on the unity of science, Paul Oppenheim and Hilary Putnam describe this idea as follows: “Any thing of any level except the lowest must possess a decomposition into things belonging to the next lower level.”³³ A related idea can also be found in the 19th and early 20th century British Emergentism. For instance, C. Lloyd Morgan wrote: “Each higher entity in the ascending series is an emergent ‘complex’ of many entities of lower grades, within which a new kind of relatedness gives integral unity.”³⁴

As Kim has observed, the entity-based understanding of levels faces some difficulties. Part-whole relationships do not neatly match up with “lower than”/“higher than” relationships as they are plausibly understood: “not every ‘complex’ of ‘lower-grade’ entities will be a higher entity; there is no useful sense in which a slab of marble is a higher entity than the smaller marble parts that make it up”.³⁵ Furthermore, we may not be able to assign every entity to a unique level. For instance, one and the same entity might be said to have both lower-level and higher-level properties; think of a computer, which might be said to have both physical and computational properties. Finally, two distinct entities need not always be comparable in “lower than”/“higher than” terms, though such non-comparability is not restricted to the entity-based understanding of levels, but applies to the world-based understanding too; see Section 4.1.

Although I prefer the world-based understanding to the entity-based one, it is worth noting that the world-based understanding does not preclude characterizing level-specific worlds in a way that refers to level-specific entities. One might define a level-specific world as a specification of all relevant facts about the level-specific entities, provided—and this is non-trivial—one can come up with a criterion for assigning entities to levels. A world at the fundamental physical level, for example, might be defined as a full specification of the states of all elementary particles across space and time. However, a level-specific world would then be not merely a collection of level-specific entities, but a specification of all facts about them. In line with the world-based understanding, we would still be able to associate each level with the set of possible worlds at that level and explicate the relationship between levels in terms of supervenience. The mereological “part-whole” distinction must not be confused with the “subvenient-supervenient” distinction.³⁶

distinguishes between a mereological understanding of levels and a world/state-based understanding and argues for the latter. Norton (2014) distinguishes between two ways of delineating lower and higher levels in physics. On one criterion, the contrast is that between molecular micro-states and thermodynamic macro-states; on another, it is that between a system with few components and one with many.

³³See Oppenheim and Putnam (1958, p. 9).

³⁴This quote comes from Kim (2002, p. 10).

³⁵See Kim (2002, p. 11).

³⁶I thank an anonymous referee for prompting me to clarify this point.

The picture I have defended, in which worlds and levels are defined primarily in terms of facts rather than entities, is consistent with Wittgenstein’s famous dictum:

“The world is everything that is the case. The world is the totality of facts, not of things.”³⁷

Similarly, we may say: *a level-specific world is everything that is the case at that level; it is the totality of level-specific facts, not of level-specific things.*

3.3 An interlude: how ontological levels relate to levels of grain

From a purely formal perspective, a system of levels of grain as defined in Section 3.1 can be viewed as a special case of a system of ontological levels as defined in Section 3.2, although the primary interpretation of levels of grain was not ontological. If, in the earlier definition, each coarsened partition of the underlying set Ω is re-interpreted as a set of higher-level worlds, then the given *levels of grain* will fit the formal definition of *ontological levels*. Indeed, if levels of grain are understood as Lewisian subject matters, this re-interpretation is not too far-fetched. Thus, for each system of levels of grain, we can define a structurally equivalent system of ontological levels. Recall that structural equivalence means that there are functors, in both directions, between the two systems, where these functors are inverses of each other.

Importantly, however, the definition of ontological levels is formally more general than the definition of levels of grain. Higher-level worlds, as defined in Section 3.2, need not be *identified* with equivalence classes of lower-level worlds, as defined in Section 3.1. They merely *pick out* such equivalence classes.³⁸ For example, the definition of ontological levels permits the inclusion in \mathcal{L} of two distinct levels (or “domains”) Ω' and Ω'' that each supervene on some lower level Ω and pick out the same equivalence classes of worlds in Ω . In fact, Ω' and Ω'' could supervene on each other and be viewed as distinct but isomorphic.³⁹ By contrast, if levels are simply partitions of some underlying set of worlds, no two distinct levels could ever supervene on one another.

Although a system of ontological levels is formally more general than a system of levels of grain, there exists a functor from any system of ontological levels to some system

³⁷See Wittgenstein (1922, 1 and 1.1).

³⁸For any supervenience mapping $\sigma : \Omega \rightarrow \Omega'$ in \mathcal{S} , each world $\omega' \in \Omega'$ picks out the equivalence class of those worlds $\omega \in \Omega$ such that $\sigma(\omega) = \omega'$.

³⁹This means that \mathcal{S} contains a supervenience mapping $\sigma : \Omega \rightarrow \Omega'$ and also a supervenience mapping $\sigma' : \Omega' \rightarrow \Omega$. It follows from our definitions that each of these mappings must then be *bijective* (i.e., injective and surjective). Surjectivity follows from the definition of \mathcal{S} . If injectivity were violated, we would not have $\sigma \bullet \sigma' = \mathbf{1}_L$, thereby contradicting condition (S4) in Section 2.

of levels of grain. That system will then mirror some (though not necessarily all) of the structure of the given system of ontological levels. We arrive at this functor in two steps.

- First, identify a lowest level: a level on which all levels supervene. If the system $\langle \mathcal{L}, \mathcal{S} \rangle$ has a lowest level, then this step is straightforward. But if it has no lowest level—an issue to which we return in Section 4.1—we must construct a hypothetical level on which all levels supervene, a so-called “inverse limit”. For a posetal category such as $\langle \mathcal{L}, \mathcal{S} \rangle$, an inverse limit can be constructed, though we need not interpret it as anything more than a mathematical construct. Formally, this involves extending $\langle \mathcal{L}, \mathcal{S} \rangle$ to a larger category which contains this inverse limit.
- In the second step, associate each level in \mathcal{L} with the partition of lowest-level worlds that the given ontological level picks out. In this way, we can map the given system of ontological levels to some system of levels of grain.⁴⁰

3.4 Levels of description

Regardless of whether we consider a levelled ontology independently plausible, it is undeniable that we use different levels of description to think and speak about the world. In fundamental physics, we describe the world in different terms than in the special sciences, such as chemistry, biology, psychology, or the social sciences. And within each of these sciences, there are debates about which level of description is appropriate for the phenomena of interest: the level of molecular descriptions versus that of aggregate descriptions in physics and chemistry; the level of biochemical and cellular descriptions versus that of broader systemic ones in biology, focusing on organisms or ecosystems; the level of the brain versus that of the mind in psychology; and the level of individual behaviour versus that of broader social patterns in the social sciences. The notion of a level of explanation is closely related to that of a level of description. An explanation at a particular level—say, a macroeconomic explanation—is an explanation that uses descriptions at that level. As Kim notes,

“[t]alk of levels of organization, descriptions or languages, of analysis, of explanation, and the like is encountered everywhere—it has thoroughly permeated primary scientific literature in many fields, in particular, various areas of psychology and cognitive science, systems theory, and computer science—as well as philosophical writings about science.”⁴¹

⁴⁰I am indebted to Marcus Pivato for suggesting the inverse-limit construction.

⁴¹See Kim (1998, p. 16, emphasis removed).

To define a system of levels of description formally, I begin by introducing the notion of a *language* that we may use to talk about the world.⁴² I define a *language*, \mathbf{L} , as a set of formal expressions—called *sentences*—which is endowed with two things:

- a *negation operator*, denoted \neg , such that, for each sentence $\phi \in \mathbf{L}$, there exists a corresponding negated sentence, $\neg\phi \in \mathbf{L}$;
- a notion of *consistency*, which deems some sets of sentences consistent and the remaining sets of sentences inconsistent.⁴³

An example of a language is standard propositional logic. Here \mathbf{L} is the set of all well-formed sentences that can be constructed out of some atomic sentences and the standard logical connectives (“and”, “or”, “not”, “if-then”, etc.), and we call a set of sentences *consistent* if all its members can be simultaneously true. Other examples of languages are more expressive logics, such as predicate, modal, conditional, and deontic logics.⁴⁴

Crucially, any language \mathbf{L} , as I have defined it, induces a corresponding “ontology”, understood as a minimally rich set of worlds $\Omega_{\mathbf{L}}$ such that each world in $\Omega_{\mathbf{L}}$ “settles” everything that can be expressed in \mathbf{L} . To *settle* a sentence is to assign a determinate truth-value to it: either “true” or “false”. If one takes the sentences in \mathbf{L} to have truth-conditions, then one is, in effect, committed to positing such an ontology. The set $\Omega_{\mathbf{L}}$ can be interpreted as the set of all possible ways the world could be such that

- (i) everything that is expressible in \mathbf{L} is settled, and
- (ii) nothing else is settled that is not entailed by what is expressible in \mathbf{L} .

One cannot take the language \mathbf{L} at face value (be a realist about the things expressible in it) without assuming that there is a fact of the matter as to which world in $\Omega_{\mathbf{L}}$ is the actual one. In that sense, the language \mathbf{L} is a “marker” of the associated ontology $\Omega_{\mathbf{L}}$.

For modelling purposes, the easiest way to define the set $\Omega_{\mathbf{L}}$ is to take it to be the set of all maximal consistent subsets of \mathbf{L} . A *maximal consistent* subset of \mathbf{L} is a consistent set of sentences to which no further sentences can be added without undermining

⁴²I borrow the present abstract definition of a language from Dietrich (2007), who introduced it in a different context, namely that of judgment-aggregation theory.

⁴³The notion of consistency must satisfy three minimal conditions (Dietrich 2007): (i) any sentence-negation pair $\{\phi, \neg\phi\}$ is inconsistent; (ii) any superset of any inconsistent set is inconsistent; (iii) the empty set is consistent and every consistent set has a consistent superset containing a member of each sentence-negation pair in \mathbf{L} .

⁴⁴Even a Boolean algebra may formally qualify as a language in the present sense, where this algebra is a set \mathcal{A} of subsets (perhaps *all* subsets) of some non-empty set Ω of possible worlds, such that \mathcal{A} is closed under intersection, union, and complementation. Here the role of “sentences” is played by elements of \mathcal{A} , and any set of such elements is *consistent* if their intersection (a subset of Ω) is non-empty.

consistency.⁴⁵ Alternatively, if taking worlds to *be* maximal consistent subsets of \mathbf{L} is too artificial, we only need to assume that the worlds in $\Omega_{\mathbf{L}}$ *correspond* to the maximal consistent subsets of \mathbf{L} .⁴⁶ We say that a sentence $\phi \in \mathbf{L}$ is *true* at a world $\omega \in \Omega_{\mathbf{L}}$ if the maximal consistent subset of \mathbf{L} to which ω corresponds contains ϕ ; the sentence is *false* otherwise. For each sentence $\phi \in \mathbf{L}$, we write $[\phi]$ to denote the set of worlds in $\Omega_{\mathbf{L}}$ at which ϕ is true; we call this the *propositional content* of ϕ .

We call any pair consisting of a language \mathbf{L} and the induced set of worlds $\Omega_{\mathbf{L}}$ a *level of description*. We can now define a system of levels of description $\langle \mathcal{L}, \mathcal{S} \rangle$ as follows:

- \mathcal{L} is some non-empty class of levels of description, each of which is a pair $\langle \mathbf{L}, \Omega_{\mathbf{L}} \rangle$;
- \mathcal{S} is some class of surjective functions of the form $\sigma : \Omega_{\mathbf{L}} \rightarrow \Omega_{\mathbf{L}'}$, where $\langle \mathbf{L}, \Omega_{\mathbf{L}} \rangle$ and $\langle \mathbf{L}', \Omega_{\mathbf{L}'} \rangle$ are levels of description in \mathcal{L} , such that \mathcal{S} satisfies (S1), (S2), and (S3).

For example, \mathcal{L} may contain levels corresponding to fundamental physics, chemistry, biology, psychology, and the social sciences. Each such level is a pair of an appropriate level-specific language and the induced set of level-specific worlds. The supervenience mappings capture the idea that chemical-level worlds supervene on physical-level worlds, biological-level worlds supervene on chemical-level worlds, and so on. In this way, a system of levels of description can capture the different levels corresponding to the different special sciences; the supervenience mappings between them capture the relationships between levels. I return to those relationships in Section 4.3, where I discuss whether supervenience entails reducibility, in a sense to be made precise.⁴⁷

Note that by focusing just on the sets of worlds induced by each level-specific language, we can map a system of levels of description to a corresponding system of ontological levels, as defined earlier. Technically, there is a functor from any system of levels of description to the induced system of ontological levels. This captures the sense in

⁴⁵The present construction, in which maximal consistent sets of sentences play the role of “worlds”, can also be found in the “Linguistic Ersatzism” discussed by Lewis (1983, section 3.2). Maximal consistent sets prominently go back to the work of Leon Henkin (1949).

⁴⁶Formally, there exists a bijection from $\Omega_{\mathbf{L}}$ to the set of maximal consistent subsets of \mathbf{L} .

⁴⁷It is worth mentioning one special case of a system of levels of description. Here, each pair $\langle \mathbf{L}, \Omega_{\mathbf{L}} \rangle$ in \mathcal{L} is of the following form: $\Omega_{\mathbf{L}}$ is some partition Ω_{\sim} of an underlying non-empty set Ω of possible worlds as in Section 3.1 (where \sim is the equivalence relation generating that partition), and \mathbf{L} is some canonical algebra $\mathcal{A}_{\Omega_{\sim}}$ over Ω_{\sim} (in the simplest case, the set of all subsets of Ω_{\sim}). We can then define the supervenience mappings in \mathcal{S} as in Section 3.1, i.e., two levels $\langle \mathcal{A}_{\Omega_{\sim}}, \Omega_{\sim} \rangle$ and $\langle \mathcal{A}_{\Omega_{\approx}}, \Omega_{\approx} \rangle$ are related by a mapping $\sigma : \Omega_{\sim} \rightarrow \Omega_{\approx}$ if and only if Ω_{\sim} is at least as fine-grained as Ω_{\approx} . Under this definition, higher-level algebras, such as $\mathcal{A}_{\Omega_{\approx}}$, may be interpreted as subalgebras of lower-level algebras, such as $\mathcal{A}_{\Omega_{\sim}}$. In particular, each element of $\mathcal{A}_{\Omega_{\approx}}$ may have an inverse image (with respect to σ) in $\mathcal{A}_{\Omega_{\sim}}$. As will become clear in Section 4.3, this feature is not shared by levels of description in general.

which our levels of description pick out corresponding ontological levels. A system of levels of description, however, encodes more information than the induced system of ontological levels: different systems of levels of description, involving different level-specific languages, could induce structurally equivalent systems of ontological levels. It should be no surprise that different languages can in principle be used to describe the same sets of level-specific worlds, while describing them differently.

3.5 Levels of dynamics

Ever since the development of statistical mechanics, there has been considerable interest in the dynamics of physical and other systems at different levels (or “scales”): microscopic and macroscopic. A coin-tossing system can be studied at a microscopic level, where the focus is on the precise details of the coin’s trajectory as it is being tossed. Alternatively, the system can be studied at a macroscopic level, where the coin is viewed as a simple Bernoulli-distributed stochastic process with only two possible outcomes, “heads” or “tails”. Similarly, the weather, the climate, or the economy can each be studied at a micro-level, where the focus is on detailed processes at a fine-grained resolution, or at a macro-level, where the system is specified in terms of certain aggregate variables. To mark this contrast, we often speak of a system’s “higher-level dynamics” and its “lower-level dynamics”. Practically any interesting dynamic system can be studied at multiple levels, and as we will see later, the dynamic properties of such a system—for instance, whether it is deterministic or not—may depend on the level in question.⁴⁸ I will now briefly explain how such “levels of dynamics” fit into the present framework. Consistently with the formalism to be presented, we can interpret levels of dynamics either epistemically (as corresponding to different ways of describing a system) or ontically (as corresponding to different ontological levels).

I begin with a simple definition of a *temporally evolving system*.⁴⁹ Let T be the set of all points of *time*, which is linearly ordered by the “before” relation. Let X be the set of all possible *states* in which the system could be at any time; we call X the system’s *state space*. A *history* of the system is a trajectory of the system through its state space, formally a function from time into the space space, $h : T \rightarrow X$, which assigns to each time $t \in T$ the system’s state at that time, denoted $h(t)$. We write Ω to denote the set of all histories that are permitted by the laws of the system. We can think of these histories

⁴⁸See, e.g., Butterfield (2011), List and Pivato (2015a), and Werndl (2009). On higher-level probabilistic dynamics, see also Ismael (2009), Glynn (2010), Strevens (2011).

⁴⁹The definition of a temporally evolving system and its analysis at different levels are based on List and Pivato (2015a,b). For a more basic analysis without a specification of probabilities, see List (2014).

as the *nomologically possible* ones. The set Ω plays the role of the set of possible worlds. Collections of histories are called *events*.

To capture the fact that the system may be stochastic, we further require the notion of a *conditional probability function*. This is a function Pr that assigns to each pair of events $E, F \subseteq \Omega$ a real number $Pr(F|E)$ between 0 and 1, interpreted as the conditional probability of F , given E , with standard properties. For instance, to determine the probability of the event F in history h at time t , we need to conditionalize on the event E that we have reached time t in history h .⁵⁰ We call the pair $\langle \Omega, Pr \rangle$ a *temporally evolving system*.

To see how we can model temporally evolving systems at multiple levels, let such a system be given, and interpret its state space X as a set of lower-level states, for example microphysical states of some coin-tossing system. Now assume that each state in X determines (realizes, instantiates) some higher-level state in some other set X' , a *higher-level state space*. In the coin-tossing example, X' could be a set of aggregate states such as “heads” and “tails”. For each higher-level state (an element of X'), there is an equivalence class of lower-level states (a subset of X) that may realize that higher-level state. For instance, different micro-states of the coin can each correspond to the same macro-state, such as landing “heads”. Similarly, each micro-state of the billions of water molecules in a flask determines a corresponding macro-state of the water.

Let $\sigma : X \rightarrow X'$ be the function that assigns to each lower-level state the resulting higher-level state. We call σ a *supervenience mapping* from the given lower-level system $\langle \Omega, Pr \rangle$ to a resulting higher-level system $\langle \Omega', Pr' \rangle$ if it has the following properties:

- (i) every higher-level state in X' has at least one possible lower-level realizing state in X according to σ , where X and X' are the two systems' state spaces; formally, the function $\sigma : X \rightarrow X'$ is surjective;
- (ii) the set Ω determines the set Ω' via σ ; formally, σ induces a surjective mapping from Ω to Ω' : for each history $h \in \Omega$, $\sigma(h) = h'$, where, for each $t \in T$, $h'(t) = \sigma(h(t))$;
- (iii) the conditional probability function Pr (for lower-level events) determines the conditional probability function Pr' (for higher-level events) via σ .⁵¹

⁵⁰So, the probability is $Pr(F|E)$, where E is the set of all histories $h' \in \Omega$ that coincide with h up to time t .

⁵¹Formally, Pr' assigns to each pair of higher-level events $E', F' \subseteq \Omega'$ the conditional probability $Pr'(F'|E') = Pr(F|E)$, where E and F are the inverse images of E' and F' under σ . As before, the *inverse image* of any $E' \subseteq \Omega'$ under σ is $\{h \in \Omega : \sigma(h) \in E'\}$. The assumption that higher-level probabilities are determined by lower-level probabilities is a simplifying assumption. One could also

We are now in a position to define a system of levels, $\langle \mathcal{L}, \mathcal{S} \rangle$:

- \mathcal{L} is some non-empty class of “level-specific” temporally evolving systems, where, for each set Ω , \mathcal{L} contains at most one system $\langle \Omega, Pr \rangle$ whose set of histories is Ω (i.e., each Ω is endowed with a unique conditional probability function);
- \mathcal{S} is some class of supervenience mappings with the properties just defined, such that \mathcal{S} satisfies (S1), (S2), and (S3).

So, when two level-specific temporally evolving systems in \mathcal{L} are related via a supervenience mapping σ in \mathcal{S} , the dynamics of the higher-level system is determined by the dynamics of the lower-level system: higher-level states (in X') supervene on lower-level states (in X), and higher-level histories as well as higher-level probabilities are determined by lower-level ones.⁵² We are then able to consider how the dynamics at different levels relate to one another; I return to this issue in Section 4.2.

Again, the present example of a system of levels can be related to one of our earlier examples. By focusing just on the sets of level-specific histories, where each history is interpreted as a possible world, we can map any system of level-specific dynamic systems to a corresponding system of ontological levels, as defined in Section 3.2. We have thereby constructed a functor from a system of levels of the present kind to a system of the earlier kind. Of course, level-specific probabilistic information is lost under this functor. A system of level-specific dynamic systems encodes more information than a system of ontological levels under the earlier definition.

4 Illustrative Applications

As we have seen, the present framework can capture a variety of instances of systems of levels, where levels may be interpreted either as levels of description or as ontological levels, and in some cases in both ways. I will now illustrate how the framework can be brought to bear on some salient philosophical questions. My illustrative questions concern the topics of fundamentality, emergence, reduction, and normativity. As already noted, my aim is not to offer comprehensive discussions of these questions, but merely to show how the relevant debates can be framed in helpful ways.

take the view that, while each level is endowed with a level-specific conditional probability function Pr , there need not be a determination relationship between the probability functions at different levels. This would amount to dropping condition (iii). On probabilities at different levels, see, e.g., Albert (2000) and List and Pivato (2015).

⁵²As noted in List and Pivato (2015b), the higher-level system is a *factor system* of the lower-level one.

4.1 Is there a linear hierarchy of levels, with a fundamental level at the bottom?

A positive answer to this question is widely assumed, but seldom carefully defended, as Jonathan Schaffer notes. In particular, “[t]alk about ‘the fundamental level of reality’ pervades contemporary metaphysics.”⁵³ Similarly, Kim comments:

“The Cartesian model of a *bifurcated* world has been replaced by that of a *layered* world, a hierarchically stratified structure of ‘levels’ or ‘orders’ of entities and their characteristic properties. It is generally thought that there is a bottom level, one consisting of whatever microphysics is going to tell us are the most basic physical particles out of which all matter is composed (electrons, neutrons, quarks, or whatever).”⁵⁴

A locus classicus of this assumption is Oppenheim and Putnam’s article on the unity of science, in which they write:

“There must be a unique lowest level (i.e., a unique ‘beginner’ under the relation ‘potential micro-reducer’).”⁵⁵

They then describe a linear hierarchy in which the level of “elementary particles” is at the bottom, followed by the levels of “atoms”, “molecules”, “cells”, “(multicellular) living things”, and “social groups”. Note the entity-based understanding of levels adopted by Oppenheim and Putnam.

David Owens, too, acknowledges a hierarchical structure of science with a bottom level, though he refers to levels of *explanation*:

“[T]he levels metaphor naturally suggests itself as a way of visualizing the structure of science. According to this picture, there is a hierarchy made up of different levels of explanation. Physics is at the base of this hierarchy and the rest of the structure depends upon it.”⁵⁶

Are we justified in assuming that there is a linearly ordered hierarchy of this kind, with a fundamental level at the bottom? Let us begin by defining the relevant properties:

- A system of levels $\langle \mathcal{L}, \mathcal{S} \rangle$ is *linear* if the levels in \mathcal{L} are totally ordered by the supervenience mappings, i.e., supervenience is (i) transitive, as already defined;

⁵³See Schaffer (2003, p. 498).

⁵⁴See Kim (1993, p. 337), quoted in Schaffer (2003).

⁵⁵See Oppenheim and Putnam (1958, p. 9); the original quote ends with a semicolon.

⁵⁶See Owens (1989, p. 59).

(ii) antisymmetric, which means that if L supervenes on L' and L' also supervenes on L , then $L = L'$; and (iii) complete, which means that, for any two levels L and L' , either L supervenes on L' , or L' supervenes on L (or, in special cases, both).

- The system $\langle \mathcal{L}, \mathcal{S} \rangle$ has a *fundamental level* if there is some level in \mathcal{L} on which every level supervenes.

The first thing to note is that linearity is independent of the existence of a fundamental level. A system of levels could have either of these properties without the other. Moreover, a system of levels might lack both properties. Linearity and the existence of a fundamental level obtain only in special cases, as I will now explain.

As already noted, a system of levels is, in general, only partially ordered.⁵⁷ A totally ordered system, as envisaged by Oppenheim and Putnam, is a special case. While the relation of supervenience is always transitive (and reflexive), it need not be complete or antisymmetric. As far as completeness is concerned, it is not generally true that, for any two levels, one of them supervenes on the other. Even if we define levels in the simplest way, as levels of grain, the resulting system of levels is only partially ordered by the relation “at least as fine-grained as”. Although the terminology of “levels” is conventional, the terms “scales”, “domains”, or “subject matters” might better capture the lack of a complete ordering: there can be many different scales or domains, which need not always be comparable in “higher-lower” terms. To illustrate, consider biology and geology. While both biological and geological facts presumably supervene on physical facts, biological facts do not plausibly supervene on geological facts or vice versa. Similarly, economic facts do not supervene on geological facts, or vice versa.

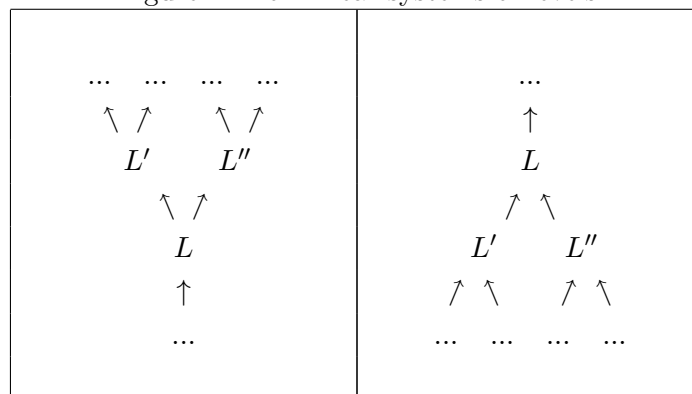
Even antisymmetry need not be satisfied by supervenience relations. While supervenience is antisymmetric for levels of grain (no two distinct levels of grain can mutually supervene on each other), it need not be antisymmetric for ontological levels or levels of description. Two distinct such levels could in principle supervene on one another, consistently with the definitions I have given. An example would be a case in which two distinct descriptive languages coincide in their ontological commitments, so that we have two distinct, but mutually supervenient levels (or domains) of description.

A partially ordered system of levels, in turn, could include more than one linearly ordered chain, where different such chains meet at most in some places. For example, the system of levels could look like a tree, or even an upside-down tree, where some levels supervene on—or subvene—others that do not themselves stand in any supervenience relation relative to each other. Kim makes a similar remark: “If a comprehensive levels

⁵⁷See Kim (2002). See also List and Pivato (2015a, fn. 41).

ontology is wanted, a tree-like structure is what we should look for; it seems to me that there is no way to build a linear system like ... Oppenheim-Putnam’s that will work.”⁵⁸ One scenario is a tree with a fundamental physical level at the bottom and the various special-science levels branching out in sometimes incomparable ways, as in the example of biology and geology. Another scenario, equally coherent though perhaps more hypothetical, is one in which, for each level L , there are two distinct levels L' and L'' such that L supervenes on each of them, while L' and L'' stand in no supervenience relation relative to each other.⁵⁹ Figure 1 illustrates both scenarios.

Figure 1: Non-linear systems of levels



Next consider the question of a fundamental level. From a formal perspective, there need not exist a level on which every level supervenes. If there were a fundamental level, say with a set Ω of “bottom-level” worlds, then any world in Ω would determine not only all facts at the fundamental level, but also all higher-level facts, via the appropriate supervenience mappings. We might then interpret an element of Ω as a “world *simpliciter*”. However, a partially ordered system of levels could have multiple distinct lowest levels: one for each linearly ordered chain. Each descending chain of levels would have its own bottom level, but there would be no overarching fundamental level. Furthermore, a system of levels, whether totally ordered or not, could have infinitely descending chains, in which each level supervenes on an even lower level. Technically, a system of levels need not be *well-founded*. This is the scenario of a “metaphysic of infinite descent”, as

⁵⁸See Kim (2002, pp. 17-18).

⁵⁹Formally, if levels are given by level-specific sets of worlds, then $L = \Omega$ might supervene on each of $L' = \Omega \times \Phi$ and $L'' = \Omega \times \Psi$, where Φ and Ψ are disjoint. Here, each world at level L' is an ordered pair consisting of an element of Ω and an element of Φ , and each world at level L'' is an ordered pair consisting of an element of Ω and an element of Ψ . The supervenience mappings from L' to L and from L'' to L then map each L' -level world or each L'' -level world to their first component.

described by Schaffer.⁶⁰

In science, the jury is arguably still out on whether there is a fundamental level. Throughout the history of science, human representations of the world have become ever more fine-grained, and this trend does not seem to have come to a halt yet. Each time people thought they had discovered a bottom level—often understood in mereological terms, say in terms of atoms or subatomic particles—new discoveries eventually led to even more fine-grained representations. To be sure, it would be too speculative to run a “pessimistic meta-induction argument” against the existence of a fundamental level on the basis of the history of science.⁶¹ But, as Ned Block remarks:

“The hypothesis that there is no bottom level ... appears to be an open question, not a mere philosopher’s possibility like the possibility that the world was created 5 seconds ago complete with the evidence of an ancient provenance.”⁶²

An interesting mathematical observation is the following. As noted in Section 3.3, for any system of ontological levels, one can theoretically construct an “inverse limit” in the category-theoretic sense: a hypothetical lowest level, on which every level supervenes. So, any system of ontological levels without a fundamental level can be mathematically viewed as a subsystem of an augmented system in which there is a fundamental level. In a formal sense, we can thus embed an ontological picture in which there is an infinite descent within an enriched ontological picture in which there is a fundamental level. Of course, the latter need not be anything more than a mathematical construct.

What does all this imply for physicalism, the thesis that everything supervenes on the physical? The exact meaning of “physicalism” depends, among other things, on how we define “the physical”.⁶³ We could either define “the physical” in terms of our *current* best account of what the fundamental physical facts are. Or, alternatively, we could define “the physical” in terms of the best future account of those facts, whatever it turns out to be. Regardless of the definition, however, the thesis that everything supervenes on the physical can be true only if there is *some* level on which all other levels supervene, i.e., a fundamental level. If there is no such level, physicalism seems in trouble. At best, physicalism could be true as a claim about a certain subclass of levels—those that do in fact supervene on a designated physical level. Perhaps the meteorological and chemical levels supervene on some underlying physical level, for example. But as an all-encompassing thesis (asserting the supervenience of *all* levels on some fundamental

⁶⁰See Schaffer (2003, p. 499). For discussion, see also List and Pivato (2015a).

⁶¹For a critical discussion, see Callender’s (2001) response to Schaffer.

⁶²See Block (2003, p. 138).

⁶³On the notion of physicalism, see, e.g., Stoljar (2010).

physical level), physicalism would be structurally false if a particular kind of bottomless levelled ontology were vindicated. Alternatively, we could accept a weaker notion of physicalism, according to which some subclass of levels, say $\mathcal{L}_{\text{phys}} \subseteq \mathcal{L}$, counts as “physical”, where every level in \mathcal{L} supervenes on some level in $\mathcal{L}_{\text{phys}}$, but where $\mathcal{L}_{\text{phys}}$ need not contain a lowest level. This, however, is a watered-down notion of physicalism, even if some physicalists might still be happy to embrace it.

Interestingly, this weakened notion of physicalism faces a special challenge, which does not affect the notions of physicalism that are committed to a fundamental level. As argued by Ned Block, Kim’s “causal exclusion argument” has a particularly radical implication for a bottomless ontology.⁶⁴ Let me first sketch the exclusion argument in a standard form. Suppose one accepts a “lower-level causal closure principle” which says that any putative higher-level cause for some effect must be underwritten by some subvenient lower-level cause, so that we cannot posit, for instance, social or biological causes that are not underwritten by subvenient physical causes. Suppose, further, one accepts an “exclusion principle” which prohibits the simultaneous ascription of *distinct* lower-level and higher-level causes for the same event. Then one can argue that all causal powers must be located at the lowest level; lower-level causes exclude higher-level ones. Now this raises a challenge for the “bottomless” version of physicalism, as Block notes:

“If there is no bottom level, and if every (putatively) causally efficacious property is supervenient on a lower ‘level’ property (Call it: ‘endless subvenience’), then (arguably) Kim’s Causal Exclusion Argument would show ... that any claim to causal efficacy of properties is undermined by a claim of a lower level, and thus that there is no causation.”⁶⁵

Block’s own response, to which I am sympathetic, is to reject the exclusion principle. Another response would be to reject some version of “endless supervenience”.⁶⁶ I need not commit myself to any particular response to this problem here, other than to note that the question of whether or not there is a fundamental level has significant repercussions.

⁶⁴See, e.g., Block (2003) and, on the exclusion argument, Kim (1998) and Bennett (2003).

⁶⁵See again Block (2003, p. 138).

⁶⁶The response that Block attributes to Kim is to argue that “the causal efficacy of the mental [level] (outside of consciousness) does indeed drain down to the physiological, but the physiological doesn’t drain down any further—at least not in any sense that makes the physiological inefficacious—because of an identity between the physiological and the next ‘level’ down, say the biochemical. And the biochemical doesn’t drain down any further because of an identity between the biochemical and say the atomic-physical. The mental is causally unreal, but the physiological is causally real and the draining stops there” (Block 2003, p. 145). For further detail, see Kim (1998, esp. pp. 81 ff.). For a critical discussion of the exclusion principle, see List and Menzies (2009).

In Section 4.4, we will further see that, even if there exists a fundamental physical level on which all the special-science levels supervene (such as the chemical, biological, and geological levels etc.), it is far from clear whether this level also provides a supervenience base for facts beyond the scientific realm, such as normative facts. The present framework can thus be used to analyze the limits of a physicalist worldview.

4.2 Are there emergent properties?

In a levelled ontology, properties are instantiated, not at worlds *simpliciter*, but at worlds *at a particular level*. For example, fundamental physical properties belong to the fundamental physical level, while chemical, biological, or psychological properties belong to the relevant higher levels. Paul Humphreys describes such a picture in terms of the assumption that “[t]here is a hierarchy of levels of properties $L_0, L_1, \dots, L_n, \dots$ of which at least one distinct level is associated with the subject matter of each special science” (for present purposes, we may replace the word “hierarchy” with “system”).⁶⁷ Similarly, Block takes different kinds of properties to characterize different levels: “The family of mental properties can be used to characterize the level of psychology, the family of neurological properties can be used to characterize the level of neuroscience, the family of elementary particle properties can be used to characterize the level of elementary particle physics.”⁶⁸ In the present framework, one might formalize this by taking each level-specific world ω to correspond to a list of property values of the form $\langle p_1, p_2, p_3, \dots \rangle$ (or more generally, a family $\langle p_i \rangle_{i \in I}$, with i ranging over an appropriate index set I), where p_i is the value of the i^{th} property among the level-specific properties.⁶⁹ In the case of binary properties, p_i would specify whether the i^{th} property is satisfied or not satisfied at world ω . In the case of non-binary properties, p_i could take more than two values.

Now if properties are tied to levels, then there is a trivial sense in which higher-level properties “emerge”: they are instantiated only at the relevant higher levels, not at lower ones, simply by virtue of their level-specificity. This, however, is not yet an interesting sense of “emergence”. An *emergent* property, in a more interesting sense, is one that is not only *not instantiated* at lower levels, but also *not explainable* in terms of—or, alternatively, *not reducible* to—lower-level properties. Of course, this is only a

⁶⁷See Humphreys (1997, p. 5), italics added. He continues: “ L_j cannot be reduced to L_i for any $i < j$ ”. On reducibility, see Section 4.3.

⁶⁸See Block (2003, p. 142), also quoted in footnote 26.

⁶⁹For example, the physical-level worlds in Ω might correspond to families of the values of physical properties, while the higher-level worlds in Ω' , Ω'' , and so on, might correspond to families of the values of the appropriate higher-level properties. The relevant supervenience mappings would map each family of lower-level property values to the resulting families of higher-level property values.

definition scheme. We obtain more precise definitions of emergence, of varying strengths, depending on how we fill in the “non-explainability” or “non-reducibility” requirement.⁷⁰

On a very weak definition, a higher-level property would count as “emergent” simply if it is not accompanied by some type-equivalent lower-level property. On a more demanding definition, a property would count as “emergent” if it is “*not predictable* even from the most complete and exhaustive knowledge of [its] emergence base”, as Kim puts it.⁷¹ “Not predictable” might, in turn, be spelt out further by adding qualifications such as “practically speaking”, “for any finite knower”, or “for even an ideal knower”.⁷² In short, an emergent property is one whose value p'_j cannot be predicted, in the relevant sense, from a complete list of lower-level property values $\langle p_1, p_2, p_3, \dots \rangle$. Alternatively, a property might count as “emergent” if it is a “systemic feature ... governed by true, law-like generalizations within a special science that is irreducible to fundamental physical theory for conceptual reasons” (quoting Timothy O’Connor and Hong Yu Wong).⁷³

As is widely recognized, on all but the most demanding definitions, emergence is consistent with supervenience. Even among the British Emergentists, who famously defended strong notions of emergence, one can arguably identify some views according to which emergent properties can be supervenient on lower-level properties, despite being irreducible to them in some sense. For example, Samuel Alexander—unlike J. S. Mill and C. D. Broad—may be interpreted as holding a view close to contemporary non-reductive physicalism: the view that higher-level properties supervene on lower-level properties while being non-identical to them and having their own causal powers.⁷⁴

Since there is a large literature on emergence, I will here review only one example of how an emergent property can be represented in the present framework. For this example, I adopt a weak, but intuitive, notion of emergence, according to which a higher-level property is “emergent” if, despite supervening on lower-level properties, it is neither accompanied by a matching lower-level property, nor readily predictable from a lower-level perspective. I discuss a more demanding notion of irreducibility in the next section.

The example is that of emergent indeterminism: a dynamic system can display indeterminism at a higher level, in the presence of lower-level determinism.⁷⁵ Consider a temporally evolving system (as defined in Section 3.5) whose histories across five time

⁷⁰I thank an anonymous referee for helping me to clarify this definition. On emergence, see the papers in Bedau and Humphreys (2008) as well as Kim (1999) and O’Connor and Wong (2015).

⁷¹See Kim (1999, p. 6).

⁷²The quoted qualifications come from O’Connor and Wong (2015).


⁷³See O’Connor and Wong (2015).

⁷⁴This interpretation is defended by O’Connor and Wong (2015).

⁷⁵For earlier analyses of this phenomenon, see, e.g., Butterfield (2011), List (2014), List and Pivato (2015a), and Werndl (2009). Relatedly, see Ismael (2009), Glynn (2010), and Strevens (2011).

periods are as shown in Figure 2.⁷⁶ Each dot represents one state in the state space X , and each path through the figure from bottom ($t = 1$) to top ($t = 5$) represents one possible history. The set of all such paths is Ω . Clearly, all histories in Ω are *deterministic*, in the sense that any initial segment of any history admits only a single possible continuation in Ω .⁷⁷

Figure 2: Lower-level histories



(reproduced from List 2014)

Now consider the temporally evolving system at a higher, more macroscopic level. We interpret the original states in X and the histories in Ω as lower-level histories, and introduce a higher-level state space X' that results from X via a many-to-one supervenience mapping. Specifically, suppose that any lower-level states that lie in the same rectangular cell in Figure 2 realize the same higher-level state. So, the supervenience mapping treats all lower-level states within the same cell in the grid as belonging to the same equivalence class. Figure 3 shows the image of Figure 2 under this supervenience mapping. The thick dots represent higher-level states, and the possible paths through the figure from bottom to top represent possible histories, this time at the higher level. Note that the set Ω' of higher-level histories is the image of Ω under the given supervenience mapping. Clearly, higher-level histories are *indeterministic* here: the initial segment of any higher-level history admits more than one possible continuation in Ω' .⁷⁸


The example shows that a temporally evolving system may display indeterminism at a higher level, consistently with determinism at a lower level. One may not even be able to ask meaningfully whether a system is deterministic or indeterministic *simpliciter*.

⁷⁶The example comes from List (2014).

⁷⁷Formally, for any history $h : T \rightarrow X$, the *initial segment* of h up to time t , denoted h_t , is the restriction of the function h to all points in time up to t . History h' is a *continuation* of an initial segment h_t if $h'_t = h_t$. History h is *deterministic* if every initial segment of h has only one continuation in Ω , namely h itself. History h is *indeterministic* if some initial segment of h has more than one continuation in Ω .

⁷⁸Formally, determinism and indeterminism are defined as they are at the lower level, except that the quantification is now over higher-level histories (in Ω') rather than lower-level histories (in Ω).

Figure 3: Higher-level histories



(reproduced from List 2014)

The answer to this question depends entirely on the level at which we are considering the system.⁷⁹ We must define determinism and indeterminism as level-specific properties. So, to avoid ambiguity, it is perhaps best to speak of “ Ω -determinism” and “ Ω' -indeterminism”. Moreover, prior to adopting a higher-level perspective, we might not expect to find the kind of higher-level indeterminism we have encountered. Its discovery involves an element of surprise, in line with the idea that emergent properties are “not readily predictable” from a lower-level perspective.

Arguably, there are many real-world instances of emergent properties in this not-too-demanding sense. For example, social or economic properties such as *unemployment* and *being a monetary exchange* are neither accompanied by matching lower-level properties, nor—it seems—readily predictable in lower-level terms. As Jerry Fodor notes,

“an immortal econophysicist might ... find a predicate in physics that was, in brute fact, co-extensive with ‘is a monetary exchange’. If physics is general—if the ontological biases of reductivism are true—then there must be such a predicate. But (a) to paraphrase a remark Donald Davidson made in a slightly different context, nothing but brute enumeration could convince us of this brute co-extensivity, and (b) there would seem to be no chance at all that the physical predicate employed in stating the co-extensivity is a natural kind term, and (c) there is still less chance that the co-extension would be lawful (i.e., that it would hold not only for the nomologically possible world that turned out to be real, but for any nomologically possible world at all).”⁸⁰

In light of (a), (b), and (c), social and economic properties may well qualify as candidates for emergent properties, while being supervenient on lower-level properties. Similar

⁷⁹For earlier arguments for this claim, see List and Pivato (2015a) and relatedly Werndl (2009).

⁸⁰See Fodor (1974, p. 104).

remarks apply to certain psychological properties, such as *being in a certain belief-and-desire state*, which—it is often argued—cannot be readily inferred from the underlying extremely complex neural brain state, even though it supervenes on it. Interestingly, Fodor concedes that if we consider a higher-level predicate that picks out a higher-level property of interest, then—if physicalism is true—“there must be [a co-extensive physical-level] predicate”. He hastens to add, however, that the existence of such a lower-level predicate is of little explanatory use. Nonetheless, we might wonder whether supervenience does entail some form of *in-principle* explanatory reducibility. I will now turn to this question.

4.3 Are higher-level descriptions reducible to lower-level ones?

It is sometimes assumed that because higher-level phenomena *supervene* on lower-level ones, we should—at least in principle—be able to *explain* higher-level phenomena in terms of lower-level ones. The idea, in short, is that supervenience implies explanatory reducibility. If this were right, it would rule out any strong form of emergence, where higher-level properties supervene on lower-level properties but are not explainable in terms of them at all. Because chemical phenomena supervene on physical ones, for example, we should ultimately be able to explain chemical phenomena in physical terms. Similarly, because social phenomena supervene on the interactions of a large number of individuals, we should ultimately be able to explain social phenomena in individual-level terms. If those claims were true, then higher-level descriptions would be dispensable, at least in principle. They would be shorthand notations for certain things that we could equally express at a lower level.

However, as several philosophers have argued in relation to domains ranging from the philosophy of mind to the philosophy of social science, supervenience does not imply explanatory reducibility.⁸¹ Fodor’s quote above illustrates this view. If we were to dispense with higher-level descriptions, we would fail to see some higher-level patterns in the world.⁸² Systematic regularities are not confined to the physical level or some other lower level; rather, they occur at higher levels too. For example, the causes that make a counterfactual difference to some effect may sometimes be higher-level properties, rather than their lower-level realizers.⁸³ The difference-making cause of a decrease in

⁸¹See, e.g., Putnam (1967), Fodor (1974), Owens (1989), Sawyer (2002, 2003), List and Pettit (2011), List and Spiekermann (2013), and Dietrich and List (2016, Section 8).

⁸²See, e.g., Dennett (1991).

⁸³See, e.g., List and Menzies (2009, 2010) as well as Yablo’s defence of mental causation (1992). On higher-level causation, see also Glynn (2013).

inflation may be the increase in the interest rate by the central bank—a higher-level property—rather than its individual-level or physical-level realizers. While inflation may systematically co-vary with the interest rate, it need not equally systematically co-vary with any particular lower-level realizing properties on which the interest rate supervenes.

I will now show that the present framework lends further support to the claim that higher-level descriptions may be irreducible to lower-level ones, even in the presence of supervenience. To discuss this claim, consider a system of levels of description, as introduced in Section 3.4, and assume, for the sake of argument, that there is a fundamental level at the bottom, on which all other levels supervene.⁸⁴ Let this fundamental level be given by the pair $\langle \mathbf{L}, \Omega_{\mathbf{L}} \rangle$, where \mathbf{L} is the fundamental-level language—say a physical one—and $\Omega_{\mathbf{L}}$ is the fundamental-level set of worlds. I will make two assumptions. First, $\Omega_{\mathbf{L}}$ is an infinite set. This is a reasonable assumption; plausibly, infinitely many worlds are compatible with the fundamental laws of nature, provided we allow variations in initial conditions. Second, the language \mathbf{L} is countable; i.e., it admits infinitely many expressions but no more than there are natural numbers. This is also a reasonable assumption; all familiar languages are countable, from propositional logic to English.

Now let us turn to some higher level of description, given by the pair $\langle \mathbf{L}', \Omega_{\mathbf{L}'} \rangle$, where \mathbf{L}' is some higher-level language and $\Omega_{\mathbf{L}'}$ is the corresponding set of higher-level worlds. Let σ be the supervenience mapping from the fundamental level to this higher level; formally, σ is a surjective function from $\Omega_{\mathbf{L}}$ to $\Omega_{\mathbf{L}'}$. We want to know whether all higher-level descriptions are “reducible” to corresponding lower-level descriptions, so that we would then be able to translate higher-level explanations into lower-level ones.

Call a higher-level sentence $\phi' \in \mathbf{L}'$ *reducible* to a lower-level sentence from \mathbf{L} if it is *equivalent in content* to some sentence in \mathbf{L} (modulo supervenience): there is some sentence $\phi \in \mathbf{L}$ whose propositional content (in $\Omega_{\mathbf{L}}$) is the inverse image, under σ , of the propositional content (in $\Omega_{\mathbf{L}'}$) of ϕ' ; formally, $[\phi] = \sigma^{-1}([\phi'])$. It is important to note that, if the language \mathbf{L} is closed under conjunction and disjunction, as in the case of standard propositional logic, then reducibility to a single sentence is equivalent to reducibility to any finite logical combination of sentences (because any such logical combination can be expressed as a single composite sentence). If every higher-level sentence were reducible to a corresponding lower-level sentence (or a finite combination of lower-level sentences), then we might indeed consider higher-level descriptions dispensable, at least in principle. Higher-level explanations would be translatable into lower-level ones.

However, there is an important combinatorial reason why reducibility is the exception

⁸⁴If there is no fundamental level, then *all* levels are higher levels relative to some other levels, and so the attempt to “reduce away” higher-level descriptions cannot get off the ground.

rather than the rule, as I will now explain.⁸⁵ The supervenience of the higher level on the lower level entails that:

- (i) the inverse image, under σ , of the propositional content of any higher-level sentence is some subset of $\Omega_{\mathbf{L}}$.

However, it does not follow that:

- (ii) this subset can be described by some sentence (or by a finite combination of sentences) from the lower-level language \mathbf{L} .

For reducibility, both (i) and (ii) are needed; (i) alone is not enough. Why might (ii) fail? The set $\Omega_{\mathbf{L}}$, being infinite, has uncountably many subsets, of which \mathbf{L} allows us to describe only countably many. Since \mathbf{L} is countable, the set of subsets of $\Omega_{\mathbf{L}}$ that can be finitely described in \mathbf{L} is also countable. Therefore, almost all subsets of $\Omega_{\mathbf{L}}$ —namely all except a countable number—do *not* admit a finite description in \mathbf{L} . For almost every subset of $\Omega_{\mathbf{L}}$, then, there exists no sentence in \mathbf{L} (or even a finite combination of sentences) whose propositional content is that subset. Of course, it is logically possible that two levels of description are so well aligned that all higher-level sentences can be reduced to equivalent lower-level sentences, in that the inverse images of their propositional contents are the contents of some lower-level sentences. But, from a combinatorial perspective, this is a highly special case. It would be surprising—a “cosmic coincidence”—if the levels corresponding to different special sciences turned out to be so well aligned.

The bottom line is that, despite the existence of a supervenience relation between a lower and a higher level, higher-level descriptions are not generally reducible to lower-level descriptions, even in principle.⁸⁶ And in those special cases in which a higher-level sentence *can* be reduced to some lower-level sentence (or a finite combination of sentences), this translation may be so cumbersome and uneconomical as to be of little practical use. For instance, the lower-level sentences might be unmanageably long; and all of Fodor’s concerns expressed in the quote above would still apply. So, higher-level descriptions may be irreducible in practice, whether or not they are irreducible in principle.

4.4 Can we represent the is-ought relationship in a levelled framework?

Inspired in part by G. E. Moore’s and R. M. Hare’s influential works on the is-ought gap, there has been much discussion about the relationship between normative and

⁸⁵The present argument draws on List and Pivato (2015a, Section 7).

⁸⁶As an anonymous referee has pointed out, the idea that there can be supervenience without reducibility is somewhat similar to a point made by concept dualists, namely that certain higher-level concepts are distinct from, and not reducible to, lower-level concepts. See, e.g., Papineau (2002, ch. 2).

non-normative domains of discourse, especially on whether “the normative” supervenes on “the descriptive” or not.⁸⁷ In this final section, I want to show that the present framework can help to clarify what is at stake in this debate.

For my discussion, I will take the normative domain to be represented by language involving “obligation” and “permission” operators (“ought” and “may”), and I will take the non-normative domain to be represented by language that is free from such operators. For the sake of argument, I will further assume that the sentences we express in normative language can be true or false. That is: there are truth-conditions for the normative domain, just as there are truth-conditions for the descriptive domain. Some people, notably non-cognitivists, reject this assumption, and even among those who accept it, there is little agreement on what the truth-conditions for the normative domain are, or on how they relate to the truth-conditions for the descriptive domain. Naturalists, for instance, think that normative truths supervene on non-normative ones, while non-naturalists disagree.⁸⁸ My aim is to show that we can helpfully think about the relationship between the normative and non-normative domains by representing them as two different levels of description, with their associated ontologies.

Let \mathbf{L} denote some descriptive, non-normative language, and let $\Omega_{\mathbf{L}}$ be the associated set of worlds. We can think of \mathbf{L} as our non-normative “base language”, and we can think of the worlds in $\Omega_{\mathbf{L}}$ as fully specifying all descriptive facts. We now augment the language by introducing normative terms. For present purposes, these will be the deontic operators “it is obligatory that” and “it is permissible that”. I assume that these operators have the standard properties assumed in deontic logic. (My arguments could equally be developed if the normative language were specified differently.) Let \mathbf{L}^+ be the “normatively augmented” language. While the original language \mathbf{L} can express only descriptive discourse, the augmented language \mathbf{L}^+ can express *both* descriptive *and* normative discourse. I suggest that we can think of \mathbf{L} and \mathbf{L}^+ as giving rise to two different levels of description.

Our central question is this: what set $\Omega_{\mathbf{L}^+}$ of worlds is associated with \mathbf{L}^+ , and how does it relate to $\Omega_{\mathbf{L}}$, the set of worlds associated with \mathbf{L} ? In other words, how is the normatively augmented level, $\langle \mathbf{L}^+, \Omega_{\mathbf{L}^+} \rangle$, related to the non-normative one, $\langle \mathbf{L}, \Omega_{\mathbf{L}} \rangle$?

⁸⁷For a recent discussion of the is-ought gap, see Brown (2014).

⁸⁸A weaker form of non-naturalism would accept that the normative truths supervene on the non-normative ones, but assert that they are not reducible to them, in the sense of Section 4.3. We can define *supervenience naturalism* as the view that the normative truths supervene on the non-normative ones. Consistently with supervenience naturalism, we can then distinguish between the view that the normative truths are reducible to non-normative truths (*reductive naturalism*) and the view that they are not (*non-reductive naturalism*). I am grateful to Wlodek Rabinowicz for pressing me on this point.

In particular, are there supervenience mappings connecting the two levels? Does one of these levels supervene on the other? I will argue that, while there trivially exists a supervenience mapping from $\Omega_{\mathbf{L}^+}$ to $\Omega_{\mathbf{L}}$, the converse (which is of greater interest to naturalists) holds only in special cases.

The first thing to note is that since \mathbf{L}^+ is an augmented version of \mathbf{L} , every sentence from \mathbf{L} is also contained in \mathbf{L}^+ . And so, since any world in $\Omega_{\mathbf{L}^+}$ settles all sentences from \mathbf{L}^+ , it must also settle all sentences from \mathbf{L} . This shows that, to each world ω^+ in $\Omega_{\mathbf{L}^+}$, there corresponds a world ω in $\Omega_{\mathbf{L}}$ such that ω^+ and ω assign the same truth-values to all sentences from \mathbf{L} . Let σ^+ be the function that maps each $\omega^+ \in \Omega_{\mathbf{L}^+}$ to the corresponding $\omega \in \Omega_{\mathbf{L}}$. If we assume that any consistent subset of \mathbf{L} remains consistent when viewed as a subset of \mathbf{L}^+ , it follows further that, for every world $\omega \in \Omega_{\mathbf{L}}$, there exists at least one world $\omega^+ \in \Omega_{\mathbf{L}^+}$ with $\sigma^+(\omega^+) = \omega$. Therefore σ^+ is surjective. This establishes (unsurprisingly) the existence of a supervenience mapping $\sigma^+ : \Omega_{\mathbf{L}^+} \rightarrow \Omega_{\mathbf{L}}$.

What about the converse? Does there exist a supervenience mapping $\sigma : \Omega_{\mathbf{L}} \rightarrow \Omega_{\mathbf{L}^+}$, as naturalism would require? To answer this question, let me begin by observing that the sentences in \mathbf{L}^+ , unlike those in \mathbf{L} , may involve the operators “it is obligatory that” and “it is permissible that”, abbreviated “O” and “P”. In line with standard deontic logic, we assume that $O\phi$ is true at some world if and only if ϕ is true at *all* worlds that are permissible relative to that world. The permissible worlds, in turn, are specified by a *selection function*, f , which assigns to each world a set of *permissible worlds* relative to the given world. Similarly, $P\phi$ is true at some world if and only if ϕ is true at *some* world that is permissible relative to it. But over which worlds should we quantify in this definition? Should we quantify over the worlds in $\Omega_{\mathbf{L}^+}$ or over the worlds in $\Omega_{\mathbf{L}}$?

Clearly, if we somehow treat the selection function f as given, we do not need to quantify over worlds from $\Omega_{\mathbf{L}^+}$. Given f , all the information needed to define the truth-values of $O\phi$ and $P\phi$ is already encoded in the “non-normative” worlds from $\Omega_{\mathbf{L}}$. So, any world $\omega \in \Omega_{\mathbf{L}}$, together with the selection function f , settles all the sentences from \mathbf{L}^+ , including those that go beyond \mathbf{L} . In other words, if the selection function f is held fixed in the background, we can identify $\Omega_{\mathbf{L}^+}$ with $\Omega_{\mathbf{L}}$, and, by implication, there is a supervenience mapping $\sigma : \Omega_{\mathbf{L}} \rightarrow \Omega_{\mathbf{L}^+}$. We can then think of \mathbf{L}^+ and \mathbf{L} as two different languages whose associated ontologies are essentially the same, despite the fact that \mathbf{L}^+ is expressively richer than \mathbf{L} . On this picture, normative and non-normative forms of discourse refer to the same world, just described differently. This, I think, is the picture that proponents of naturalism about normativity have in mind.⁸⁹

⁸⁹Note that this picture will be shared by all proponents of supervenience naturalism, whether or not they hold that the “the normative” is reducible to the “the non-normative”. The question of reducibility

By contrast, if we do not take the selection function f as given, the supervenience of the normatively augmented level on the non-normative one breaks down. Without the function f , the worlds in $\Omega_{\mathbf{L}}$ are insufficient to settle everything that can be expressed in \mathbf{L}^+ . The truth-values of sentences such as $O\phi$ and $P\phi$ are then left open. So, worlds in $\Omega_{\mathbf{L}}$ encode strictly less information than worlds in $\Omega_{\mathbf{L}^+}$. Note that, for something to qualify as a world in $\Omega_{\mathbf{L}^+}$, it must settle everything that can be expressed in \mathbf{L}^+ .

Let F be the set of all possible selection functions: each element of F , in effect, encodes a particular criterion of permissibility. In order to settle everything that can be expressed in \mathbf{L}^+ , we require not only a world $\omega \in \Omega_{\mathbf{L}}$, but also a selection function $f \in F$. We may thus define $\Omega_{\mathbf{L}^+}$ as the set of all possible pairs of the form $\langle \omega, f \rangle$, formally $\Omega_{\mathbf{L}^+} = \Omega_{\mathbf{L}} \times F$ (or a suitable subset thereof). Any world in $\Omega_{\mathbf{L}^+}$, formally a pair $\langle \omega, f \rangle$, will then suffice to settle all sentences in \mathbf{L}^+ . Clearly, there will continue to be a supervenience mapping from $\Omega_{\mathbf{L}^+}$ to $\Omega_{\mathbf{L}}$ (it simply maps each pair $\langle \omega, f \rangle$ to ω alone), but the converse is no longer true, assuming there can be more than one possible selection function. The relationship between $\Omega_{\mathbf{L}}$ and $\Omega_{\mathbf{L}^+}$ will be one-to-many: several distinct normatively augmented worlds will share the same non-normative part.

However, there is one limiting case in which there exists a supervenience mapping $\sigma : \Omega_{\mathbf{L}} \rightarrow \Omega_{\mathbf{L}^+}$. This is the case in which the set F is assumed to be singleton, i.e., in which there is only a single possible selection function. In this case, $\Omega_{\mathbf{L}}$ and $\Omega_{\mathbf{L}^+}$ are isomorphic, so that they can essentially be identified with one another. This is again the picture that naturalists have in mind.

Recall that there are different notions of supervenience. Earlier, when I was not talking about normative matters, I gave the examples of metaphysical and nomological supervenience. We can now observe the following: *relative to a given selection function*, the normatively augmented level does indeed supervene on the non-normative one. This is a form of *relativized supervenience*: supervenience relative to some further parameter, here the selection function f . Without such a relativization, the normatively augmented level does not supervene on the non-normative one. So, the hidden assumption underlying naturalism seems to be a particular choice of selection function: the normative supervenes on the descriptive relative to this hidden parameter, but not without it.

5 Concluding Remarks

I have given a general definition of a system of levels, inspired by ideas from category theory, and I have discussed several instances of this definition, which show that the pro-

can in principle be analyzed as in Section 4.3.

posed framework can capture descriptive, explanatory, and ontological notions of levels. I have further suggested that the framework can be helpfully brought to bear on some familiar philosophical questions, such as questions about fundamentality, emergence, reduction, and normativity. Another possible application, which I have omitted due to space constraints, concerns the relationship between first-personal and third-personal phenomena in the study of subjectivity and consciousness. Arguably, one may helpfully introduce the notion of a first-personal level, as distinct from the third-personal level, so that first-personal descriptions as well as first-personal properties can be associated with a different level than third-personal ones. In sum, since references to “levels” are ubiquitous in science and philosophy, I hope that, by going beyond metaphorical uses of this idea and offering a precise and unified formalization, the present framework will clarify the conceptual terrain and lead to further useful applications.

6 References

- Albert, David. 2000. *Time and Chance*. Cambridge, MA: Harvard University Press.
- Bechtel, William. 1994. Levels of Description and Explanation in Cognitive Science. *Minds and Machines* 4(1): 1–25.
- Beckermann, Ansgar, Hans Flohr, and Jaegwon Kim, eds. 1992. *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. Berlin: de Gruyter.
- Bedau, Mark A., and Paul Humphreys, eds. 2008. *Emergence: Contemporary Readings in Philosophy and Science*. Cambridge, MA: MIT Press.
- Bennett, Karen. 2003. Why the Exclusion Problem Seems Intractable, and How, Just Maybe, to Tract It. *Noûs* 37(3): 471–497.
- Block, Ned. 2003. Do Causal Powers Drain Away? *Philosophy and Phenomenological Research* 67(1): 133–150.
- Brown, Campbell. 2014. Minding the Is-Ought Gap. *Journal of Philosophical Logic* 43(1): 53–69.
- Butterfield, Jeremy. 2011. Laws, Causation and Dynamics at Different Levels. *Interface Focus* 2(1): 101–114.
- Callender, Craig. 2001. Why be a Fundamentalist: Reply to Schaffer. PhilSci-Archive <<http://philsci-archive.pitt.edu/215/>>.

- Chalmers, David. 2012. *Constructing the World*. Oxford: Oxford University Press.
- Davies, Martin. 2000. Interaction without reduction: The relationship between personal and sub-personal levels of description. *Mind and Society* 1(2): 87–105.
- Dennett, Daniel. 1991. Real Patterns. *Journal of Philosophy* 88(1): 27–51.
- Dietrich, Franz. 2007. A generalised model of judgment aggregation. *Social Choice and Welfare* 28(4): 529–565.
- Dietrich, Franz. 2017. Savage’s Theorem under changing awareness. Working paper, Paris School of Economics.
- Dietrich, Franz, and Christian List. 2016. Mentalism versus behaviourism in economics: A philosophy-of-science perspective. *Economics and Philosophy* 32(2): 249–281.
- Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Eilenberg, Samuel, and Saunders MacLane. 1945. General Theory of Natural Equivalences. *Transactions of the American Mathematical Society* 58(2): 231–294.
- Ellis, George F. R., Denis Noble, and Timothy O’Connor, eds. 2012. Top-down causation. Theme issue of *Interface Focus* 2(1): 1–140.
- Floridi, Luciano. 2008. The Method of Levels of Abstraction. *Minds and Machines* 18(3): 303–329.
- Fodor, Jerry. 1974. Special Sciences (or: The Disunity of Science as a Working Hypothesis). *Synthese* 28(2): 97–115.
- Glynn, Luke. 2010. Deterministic Chance. *British Journal for the Philosophy of Science* 61(1): 51–80.
- Glynn, Luke. 2013. Is There High-Level Causation? PhilSci-Archive <<http://philsci-archive.pitt.edu/9823/>>.
- Gómez Ramirez, Jaime. 2014. *A New Foundation for Representation in Cognitive and Brain Science: Category Theory and the Hippocampus*. Dordrecht/Heidelberg: Springer.
- Halvorson, Hans. 2016. Scientific Theories. In Paul Humphreys, ed., *The Oxford Handbook of Philosophy of Science*, 585–608. Oxford: Oxford University Press.

- Haslanger, Sally. 2016. What is a (social) structural explanation? *Philosophical Studies* 173(1): 113–130.
- Henkin, Leon. 1949. The Completeness of the First-Order Functional Calculus. *Journal of Symbolic Logic* 14(3): 159–166.
- Himmelreich, Johannes. 2015. *Agency as difference-making: Causal Foundations of Moral Responsibility*. PhD thesis, LSE <<http://etheses.lse.ac.uk/3277/>>.
- Horgan, Terence. 1993. From Supervenience to Superdupervenience: Meeting the Demands of a Material World. *Mind* 102(408): 555–586.
- Ismael, Jenann. 2009. Probability in Deterministic Physics. *Journal of Philosophy* 106(2): 89–108.
- Kim, Jaegwon. 1984. Concepts of Supervenience. *Philosophy and Phenomenological Research* 45(2): 153–176.
- Kim, Jaegwon. 1987. “Strong” and “Global” Supervenience Revisited. *Philosophy and Phenomenological Research* 48(2): 315–326.
- Kim, Jaegwon. 1993. The nonreductivist’s troubles with mental causation. In *Supervenience and Mind: Selected Philosophical Essays*, 336–357. Cambridge: Cambridge University Press.
- Kim, Jaegwon. 1998. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: MIT Press.
- Kim, Jaegwon. 1999. Making Sense of Emergence. *Philosophical Studies* 95(1–2): 3–36.
- Kim, Jaegwon. 2002. The Layered Model: Metaphysical Considerations. *Philosophical Explorations* 5(1): 2–20.
- Kincaid, Harold. 1986. Reduction, explanation, and individualism. *Philosophy of Science* 53(4): 492–513.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, David. 1988. Relevant implication. *Theoria* 54(3): 161–174.
- List, Christian. 2014. Free will, determinism, and the possibility of doing otherwise. *Noûs* 48(1): 156–178.

- List, Christian, and Peter Menzies. 2009. Non-reductive physicalism and the limits of the exclusion principle. *Journal of Philosophy* CVI(9): 475–502.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- List, Christian, and Marcus Pivato. 2015a. Emergent Chance. *The Philosophical Review* 124(1): 119–152.
- List, Christian, and Marcus Pivato. 2015b. Dynamic and stochastic systems as a framework for metaphysics and the philosophy of science. PhilSci-Archive <<http://philsci-archive.pitt.edu/12371/>>.
- List, Christian, and Kai Spiekermann. 2013. Methodological Individualism and Holism in Political Science: A Reconciliation. *American Political Science Review* 107(4): 629–643.
- Marquis, Jean-Pierre. 2015. Category Theory. In Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition) <<http://plato.stanford.edu/archives/win2015/entries/category-theory/>>.
- McLaughlin, Brian, and Karen Bennett. 2014. Supervenience. In Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition) <<http://plato.stanford.edu/archives/spr2014/entries/supervenience/>>.
- Menzies, Peter, and Christian List. 2010. The Causal Autonomy of the Special Sciences. In Cynthia McDonald and Graham McDonald, eds., *Emergence in Mind*, 108–128. Oxford: Oxford University Press.
- Modica, Salvatore, and Aldo Rustichini. 1999. Unawareness and Partitional Information Structures. *Games and Economic Behavior* 27(2): 265–298.
- Norton, John. 2014. Infinite Idealizations. In *European Philosophy of Science: Philosophy of Science in Europe and the Viennese Heritage*, Vienna Circle Institute Yearbook, vol. 17, 197–210. Dordrecht/Heidelberg/London/New York: Springer.
- O'Connor, Timothy, and Hong Yu Wong. 2015. Emergent Properties. In Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition) <<http://plato.stanford.edu/archives/sum2015/entries/properties-emergent/>>.
- Oppenheim, Paul, and Hilary Putnam. 1958. Unity of Science as a Working Hypothesis. *Minnesota Studies in the Philosophy of Science* 2: 3–36.

- Owens, David. 1989. Levels of Explanation. *Mind* 98(389): 59–79.
- Papineau, David. 2002. *Thinking about Consciousness*. Oxford: Oxford University Press.
- Putnam, Hilary. 1967. Psychological predicates. In W. H. Capitan and D. D. Merrill, eds., *Art, Mind, and Religion*, 37–48. Pittsburgh: University of Pittsburgh Press.
- Sawyer, R. Keith. 2002. Nonreductive Individualism. Part I—Supervenience and Wild Disjunction. *Philosophy of the Social Sciences* 32(4): 537–559.
- Sawyer, R. Keith. 2003. Nonreductive Individualism. Part II—Social Causation. *Philosophy of the Social Sciences* 33(2): 203–224.
- Schaffer, Jonathan. 2003. Is there a fundamental level? *Noûs* 37(3): 498–517.
- Schaffer, Jonathan. 2009. On What Grounds What. In David Chalmers, David Manley, and Ryan Wasserman, eds., *Metametaphysics: New Essays on the Foundations of Ontology*, 347–383. Oxford: Oxford University Press.
- Stoljar, Daniel. 2010. *Physicalism*. London: Routledge.
- Strevens, Michael. 2011. Probability out of Determinism. In Claus Beisbart and Stephan Hartmann, eds., *Probabilities in Physics*, 339–364. Oxford: Oxford University Press.
- Werndl, Charlotte. 2009. Are deterministic descriptions and indeterministic descriptions observationally equivalent? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 40(3): 232–242.
- Wilson, Mark. 2010. Mixed-Level Explanation. *Philosophy of Science* 77(5): 933–946.
- Wittgenstein, Ludwig. 1922. *Tractatus Logico-Philosophicus*. London: Kegan Paul.
- Yablo, Stephen. 1992. Mental Causation. *The Philosophical Review* 101(2): 245–280.