Levels of Description and Levels of Reality: A General Framework

Christian List
Munich Center for Mathematical Philosophy, LMU Munich

Working Paper, March 2022 (with minor revisions in August 2022)

Abstract: This expository paper presents a general framework for representing levels and inter-level relations. The framework is intended to capture both epistemic and ontological notions of levels and to clarify the sense in which levels of explanation might or might not be related to a levelled ontology. The framework also allows us to study and compare different kinds of inter-level relations, especially supervenience and reduction but also grounding and mereological constitution. This, in turn, enables us to explore questions such as whether supervenience implies explanatory reducibility and whether there can be irreducible higher-level explanations or even "emergent" higher-level properties.

1. Introduction

An important feature of science is its organization into different domains of enquiry. In different such domains, we focus on different phenomena and use different concepts and categories to describe and explain those phenomena. Some areas of science focus on larger-scale phenomena – think of astronomy, ecology, or macroeconomics, for instance – while others focus on smaller-scale phenomena, such as particle physics, molecular biology, or microeconomics. We then say that these areas of science operate at different "levels of description" or different "levels of explanation". Some operate at what we call a "micro-level", while others operate at a "macro-level".

But what are "levels"? Although talk of "levels", such as "levels of description", "levels of explanation", or even "levels of reality", is very common in both science and philosophy,[1] and there are many debates on what the right level of explanation is for certain phenomena, such as for social, psychological, or biological ones, this talk of levels is often criticized for being too metaphorical and imprecise. As Jaegwon Kim writes, "talk of levels may turn out to be only a figure of speech, a harmless but suggestive metaphor".[2]

We may have an intuitive grasp of what it means to say that macroeconomics operates at a higher level than microeconomics, or that systems biology operates at a higher level than cell biology, but there is no consensus among scientists and philosophers on how to make those

---

[1] See, among others, Oppenheim and Putnam (1958), Fodor (1974), Owens (1989), Beckermann, Flohr, and Kim (1992), Dupré (1993), Bechtel (1994), Kim (1998, 2002), Schaffer (2003), Floridi (2008), Ellis, Noble, and O'Connor (2012), and Knox (2016).

[2] See Kim (2002, p. 3).

claims precise. Further, there is no consensus on how higher-level phenomena or explanations are related to lower-level ones, and whether the former are somehow "reducible" to the latter, at least in principle. Finally, there is no consensus on whether "levels" should be understood only in epistemic terms, as a feature of how we think about the world, or also in ontic terms, as a feature of reality itself.

The aim of this paper is to present a general framework for representing levels and inter-level relations. The framework is intended to capture both epistemic and ontological notions of levels and to clarify the sense in which levels of explanation might or might not be related to a levelled ontology. Moreover, the framework is intended to allow us to study and compare different kinds of inter-level relations, especially supervenience and reduction but also grounding and mereological constitution. This, in turn, will enable us to explore questions such as whether supervenience implies explanatory reducibility and whether there can be irreducible higher-level explanations or even "emergent" higher-level properties.

I will first review several salient uses of the idea of levels, beginning with levels in the epistemic sense (Section 2), followed by levels in the ontic sense (Section 3). I will then show how to accommodate these different notions in a unified framework (Section 4). Next, I will use the framework to address some key questions about the relationship between epistemic and ontic notions of levels (Section 5). And finally, I will briefly mention some other theoretical payoffs and applications (Section 6).[3]

2. Levels in the epistemic sense

I will begin with an account of levels in the epistemic sense, i.e., levels of description or levels of explanation, which seems to be the least controversial sense, and only subsequently turn to levels in the ontic sense, i.e., levels of reality, which seems to be more controversial.

As already noted, we use different concepts and categories when we describe and explain the phenomena in different domains. For example, fundamental physics speaks of particles, fields, and forces; biology speaks of cells, organisms, and ecosystems; psychology speaks of mental states, intentionality, and cognition; and the social sciences speak of institutions, norms, and conventions.

We have very good explanatory reasons for following this differentiated explanatory practice. Different *explananda* – different phenomena to be explained – require different explanatory concepts and categories, which enable us to recognize different patterns and regularities in the world. It should be evident, for instance, that explaining the movement of the planets in physics

---

[3] This paper builds on some material in, and is a sequel to, List (2019a), where the proposed unified framework was first introduced. The present exposition is, however, new and updated in some respects. Among other things, I explicitly consider a greater variety of inter-level relations than I did in that earlier work. I am grateful to Jan Borner, Neil Dewar, and the participants of a seminar at the Munich Center for Mathematical Philosophy in February 2022 for helpful comments and discussion.

or photosynthesis in biology requires very different conceptual resources than explaining inflation in economics or voting behaviour in politics.

"Operating at a different level of description or different level of explanation" simply means describing and explaining the world through the lens of a different system of concepts and categories. A *level of description* or a *level of explanation* can thus be informally defined as a particular system of concepts and categories through which one might describe and explain the world. For instance, the fundamental physical level is defined by the concepts and categories of fundamental physics, such as particles, fields, and forces, and the psychological level is defined by a system of psychological concepts and categories, such as beliefs, desires, and other mental states and processes.

Now, there are at least two ways in which such an epistemic understanding of levels can be made more precise.

- *The coarse-graining understanding*: This is based on the idea that each level of description corresponds to a particular way of partitioning some underlying set of possibilities into equivalence classes.

- *The linguistic understanding*: This is based on the idea that each level corresponds to a particular level-specific descriptive language.

Let me explain these in turn.

2.1 Levels as equivalence relations

On a "coarse-graining understanding", different levels correspond to different ways of partitioning some underlying set of possibilities – for instance, the set of all possible worlds or the set of all possible states of the world – into equivalence classes. Formally, each level thus corresponds to a particular equivalence relation on that set. An *equivalence relation* on a given set specifies, for any two of its elements, whether they count as equivalent according to the standard encoded by that relation.[4]

Understanding levels as equivalence relations captures the idea that the concepts and categories available at different levels allow us to draw different distinctions in the world and force us to ignore others. Specifically, at each level, the level-specific concepts and categories allow us to distinguish between possibilities that lie in different equivalence classes but not between possibilities that lie within the same equivalence class.

David Lewis already introduced this way of representing levels of description, albeit without using the terminology of "levels". Specifically, he introduced the notion of a "subject matter",

[4] Formally, an equivalence relation is a reflexive, symmetrical, and transitive binary relation on the given set.

which is essentially the same as a level in the present sense.[5] A *subject matter*, for Lewis, picks out a part – or perhaps better: an aspect – of the world, namely the one that has to do with that subject matter. Formally, Lewis takes each subject matter to be representable by an equivalence relation on the set of possible worlds. Physics, biology, and psychology, for instance, are all subject matters under this definition; they each partition possibilities differently, thereby focusing on different distinctions. Two worlds are indistinguishable with respect to physics, or biology, or psychology if and only if they coincide with respect to all physical, all biological, or all psychological properties, respectively.

Lewis also introduces the notion of "inclusion of subject matters".[6] One subject matter is said to *include* another if the equivalence relation representing the former is at least as fine-grained as the equivalence relation representing the latter, i.e., any two possibilities that are distinguished by the latter equivalence relation are also distinguished by the former. So, whenever one subject matter includes another, it is true that any distinction that can be drawn in terms of the latter (the included subject matter) can also be drawn in terms of the former (the including one).

Similarly, in some parts of economics and psychology, an agent's *awareness* is sometimes characterized in terms of the distinctions that this agent is able to draw and formally defined as an equivalence relation on some underlying set of possibilities.[7] The agent is said to be *aware* of some feature of the world if and only if he or she can distinguish worlds with that feature from worlds without it. Greater awareness corresponds to a more fine-grained partition, and lesser awareness to a more coarse-grained one. Awareness growth would involve fine-graining. Levels of awareness can again be related to each other by an inclusion relation, defined as in Lewis's account of subject matters.

Inclusion as defined by Lewis and applicable also to awareness is our first example of an inter-level relation. We can say that one level counts as "higher" than another if the equivalence relation representing the former is strictly more coarse-grained than the equivalence relation representing the latter. The inclusion relation ("at least as fine-grained as") yields a partial ordering over all Lewisian subject matters (or levels as equivalence relations), defined for some underlying set of possible worlds.

At this point, we can already make the first substantive observation: levels in the epistemic sense need not be totally ordered. That is, we shouldn't think of there being a linear hierarchy of levels. Rather, there may be only a partial ordering. Some levels may be comparable in terms of the "higher than" relation, others not. For instance, the levels of biology and geology may each be higher than the level of physics, but neither of them may be higher than the other.

---

[5] See Lewis (1988).
[6] Ibid.
[7] See, e.g., Modica and Rustichini (1999) and Dietrich (2017).

Some people may therefore prefer to speak of "scales", "domains", "conceptual schemes", or indeed Lewisian "subject matters" instead of "levels", but since talk of "levels" is ubiquitous, I propose to retain this terminology, despite the lack of a linear hierarchy.[8]

2.2 Levels as descriptive languages

Let me turn to the second way in which levels in the epistemic sense can be made more precise. Here, different levels correspond to different level-specific languages for describing the world. To provide a simple formalization of this, let me define a *descriptive language*, **L**, as the set of all (indicative) sentences that can be expressed in it (this includes all sentences that express propositional content but excludes, for instance, questions and commands), where this language is endowed with (i) some *logical operations*, at a minimum a negation operator, such that, for each sentence in **L**, its negation is also in **L**, and (ii) a well-behaved *notion of consistency*, which partitions the set of all subsets of **L** into those that are consistent and those that are inconsistent. (The latter, in turn, also allows us to define a notion of *logical entailment*.)[9] The simplest examples of such languages come from standard propositional logic, but we could also consider more expressive languages, which may include not only predicates, but also modal operators (such as "necessarily" and "possibly") and/or non-material conditionals (such as "if X were the case, then Y would be the case").

If different levels correspond to different descriptive languages, we can now also introduce one salient kind of inter-level relation for such levels, namely the reduction relation. One language **L** is *reducible* to another language **L**′ if there exists a *translation function f* from **L** to **L**′ which assigns to each sentence ϕ in **L** an "equivalent" sentence $\phi' = f(\phi)$ in **L**′, where logical properties (such as consistency, inconsistency, and negation) are preserved under translation.

For example, if we had a function that assigns to each sentence expressible in the language of chemistry a content-wise equivalent sentence in the language of physics, then we would have achieved a reduction of chemistry to physics. It is a non-trivial question, however, whether, and under what conditions, such reductions exist, and I will say more about it in Section 5. For the moment, I want to note that even the question of whether chemical descriptions are reducible to physical ones – a familiar example of purported reducibility – is controversial.[10] Again, different levels in the present sense are partially, but not completely, ordered by the given inter-level relation.

[8] These alternative terms appear, for instance, in Wilson (2010), Kim (2002), Davidson (1973), and Lewis (1988).
[9] For this notion of a language, see Dietrich (2007). To count as *well-behaved*, the notion of consistency must satisfy the following conditions: first, any sentence-negation pair is inconsistent; second, any subset of any consistent set is still consistent; third, the empty set is consistent and every consistent set has a consistent superset containing a member of each sentence-negation pair within the language. We can then further say that a set of sentences *logically entails* another sentence of the set together with the negation of the sentence is inconsistent.
[10] See, e.g., Manafu (2015) and Hettema (2012).

We may also ask how the two epistemic notions of levels I have introduced – levels as equivalence relations and levels as descriptive languages – are related to one another, and similarly how their respective inter-levels relations are related. As should become clear, the framework to be presented will offer some formal tools for addressing those questions.

3. Levels in the ontic sense

Let me move on to the ontic understanding of levels. Here the idea is that levels are not merely a feature of our way of thinking about the world and describing it, but a feature of reality itself. According to a levelled ontology, the world is somehow stratified into levels. In line with such a picture, philosophers often invoke notions such as "the fundamental level of reality". And if one speaks of the fundamental level of reality, then presumably it also makes sense to speak of other, higher levels. As Jonathan Schaffer, for instance, observes: "[t]alk about 'the fundamental level of reality' pervades contemporary metaphysics".[11] And Jaegwon Kim writes: "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."[12] As an example, he mentions the "bottom level", "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)."[13]

Again, there are at least two ways in which this can be made more precise:

- *The entity-based understanding*: This is based on the idea that each ontological level corresponds to a particular set of level-specific entities and perhaps their properties.

- *The fact or world-based understanding*: This is based on the idea that each ontological level corresponds to a particular set of level-specific facts and by implication a level-specific way of defining worlds.

I will ultimately endorse only the second of these understandings.[14]

[11] See Schaffer (2003, p. 498).
[12] This passage from Kim (1993, p. 337) is also quoted in Schaffer (2003).
[13] Ibid.
[14] Others have drawn similar distinctions and supported the second understanding. Notably, Block (2003, pp. 141/142) contrasts "a notion of level keyed to objects" and another "keyed to relations among properties" and defends the latter, and Himmelreich (2015, Appendix B) contrasts a mereological understanding of levels and a world/state-based understanding and argues for the second. Relatedly, Norton (2014) distinguishes between different criteria for distinguishing between lower and higher levels in physics. One criterion focuses on the states of a system (distinguishing between micro- and micro-states), while the other focuses on the number of components of a system.

3.1 Levels of entities

The entity-based way of understanding ontological levels is the most conventional one. Its key idea is that, at each level, there are certain level-specific entities, which serve as building blocks of higher-level entities. Recall, for instance, how Jaegwon Kim describes how people conventionally think about the fundamental level: it "consist[s] 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)".[15] On this understanding, higher levels consist of more complex entities, such as molecules in chemistry or cells or organisms in biology.

A version of this understanding of levels can already be found in the writings of some British Emergentists, as for instance in the following quote from C. Llloyd Morgan: "Each higher entity in the ascending series in an emergent 'complex' of many entities of lower grades, within which a new kind of relatedness gives integral unity."[16] The entity-based understanding of levels can also be found in a classic article by Paul Oppenheim and Hilary Putnam, who write: "Any thing of any level except the lowest must possess a decomposition into things belonging to the next lower level."[17]

On an entity-based understanding, inter-level relations are mereological relations, such as *composition* or *parthood* relations. One level is "higher" than another if the entities of the former (higher) level are composites or aggregates of the entities of the latter (lower), or conversely, the entities of the latter (lower) level are the parts or building blocks of the entities of the former (higher). Again, this would yield a partial ordering over levels.

However, as critics such as Jaegwon Kim have pointed out, the entity-based understanding of levels has several shortcomings.[18] First, it is not clear that part-whole relationships always capture lower-versus-higher-level relationships. Only some part-whole relationships seem to do. Plausibly, the elementary particles in physics of which larger entities are composed are associated with a lower level than, say, cells in biology. But it is not plausible, as Kim notes, that "a slab of marble is a higher entity than the smaller marble parts that make it up".[19] Second, it is unclear that every entity can be associated with a unique level. For instance, an organism or a computer might have both physical properties and higher-level ones, such as biological or computational ones. One would then not be able to say which level the organism or computer, *qua* entity, belongs to. Is it low-level, is it high-level, or is it both? Unless we clarify which *properties* of the entity we are interested in, the answer seems unclear.

[15] See Kim (1993, p. 337).
[16] For this quote, see Kim (2002, p. 10).
[17] Oppenheim and Putnam (1958, p. 9).
[18] See, e.g., Kim (2002).
[19] Ibid., p. 11.

3.2 Levels of facts

The shortcomings of the entity-based understanding of ontological levels motivate the alternative, fact- or world-based understanding. On this understanding, it is not entities that are primarily assigned to levels but rather facts (or properties of the world). For example, some facts belong to the fundamental physical level, such as facts about the physical microstate of the universe, while other facts belong to higher levels, such as facts about metabolism in biology, mental states in psychology, or inflation and the exchange rate in economics.

If, for the moment, we run with the idea that different levels can be associated with different level-specific facts, we can see that the notion of "the world" can also be defined in a level-specific way. To introduce this idea, let's begin by recalling the standard notion of a possible world, as we find it, for instance, in Wittgenstein's *Tractatus*: "The world is everything that is the case".[20] On this picture, a *world* is a full specification of all facts that obtain at that world. Moreover, consistently with a fact-based rather than entity-based ontology, Wittgenstein emphasizes that we should think of the world as "the totality of facts, not of things".[21] Now, to incorporate the idea that facts are level-specific, we must amend Wittgenstein's definition. We can define a *possible world at a particular level* as a full specification of the way the world might be at that level. For instance, the world at the microphysical level is the totality of microphysical facts; the world at the biological level is the totality of biological facts; the world at the psychological level is the totality of psychological facts; and so on. Amending Wittgenstein's definition, we can say: "The world at a particular level is everything that is the case at that level."

The world at some higher level, under this definition, will omit certain lower-level facts – for instance, facts about certain microphysical details – that are irrelevant at the higher level. From a lower-level perspective, higher-level worlds may then look like *partial* worlds. However, from a higher-level perspective, this would be the wrong interpretation, since, as far as higher-level facts are concerned, they are complete specifications of those.

Higher-level worlds might also include some other facts which, despite being somehow *determined* by lower-level facts, are not explicitly included in any purely lower-level factual inventory of the world. For instance, if a certain version of non-reductive physicalism is true, psychological-level worlds may include certain mental facts which, despite being supervenient on underlying physical facts, do not themselves qualify as physical. At any rate, at each level, a *possible world at the given level* is a total specification of all level-specific facts.

On the present understanding, we can associate each level with its own level-specific set of possible worlds: the physical level is associated with the set of all possible physical-level worlds; the biological level is associated with the set of all possible biological-level worlds;

[20] See Wittgenstein (1922, §1).
[21] Ibid., §1.1.

and so on. Furthermore, we can think of inter-level relations as supervenience relations between facts or, more globally, between worlds at different levels. Recall that one set of facts (call it the B-facts) *supervenes* on another set of facts (call it the A-facts) if it is impossible for the former (the B-facts) to be any different without the latter (the A-facts) being different too. A standard example is the commonly assumed supervenience of chemical facts on physical facts.

Formally, on this picture, one level counts as "higher" than another if there exists a mapping from the set of worlds associated with the latter (lower) level to the set of worlds associated with the former (higher), where that mapping has the following property:

- *surjectivity*: for each "higher-level" world, there exists at least one "lower-level" world that is mapped to it (a *lower-level realizer* of the higher-level world).

The mapping may also have a second property:

- *many-to-one*: for at least one "higher-level world" (perhaps many), there exists more than one "lower-level" world that is mapped to it (*multiple realizability*).

These are of course standard properties of *supervenience*. Importantly, the "many-to-one" property is optional and should not be built into the definition of supervenience because we can have cases of supervenience mappings that are not many-to-one.

Once more, the present inter-level relation yields a partial ordering over levels. Formally, in accordance with the mathematical notation for functions, we represent a supervenience relation by a function $\sigma : \Omega \rightarrow \Omega'$, where Ω (the domain of σ) is the relevant set of lower-level worlds and Ω' (the co-domain of σ) is the relevant set of higher-level worlds.

Alternatively, levels of facts, or levels of worlds, could be related to each other by grounding relations, such as when we say that the physical facts ground the chemical ones or that the chemical facts ground the biological ones.[22] But for reasons that will become clearer later, I here prefer to focus on supervenience. Importantly, neither supervenience nor grounding, which are suitable inter-level relations on a fact- or world-based understanding, should be confused with the mereological part-whole relations on the entity-based understanding.

4. A unifying framework

So far, I have reviewed four notions of levels, two of an epistemic sort and two of an ontic sort. Since we find each of these notions in some discourse about levels, does this suggest that "levels" talk is inherently diverse and pluralistic, and that there is no hope of unifying or at least reconciling all the different ways of understanding levels? Or can we find something that

---

[22] On grounding, see, e.g., Schaffer (2009) and Rosen (2010).

all these different notions have in common, and/or identify some interesting relationships between them?

What I want to show is that all four ways of thinking about levels and inter-level relations can be subsumed under a single unified framework. This framework further allows us to compare some key aspects of the different notions and to address some additional questions about levels and their relations.

I will proceed by first giving an abstract definition of a *system of levels* and then showing that each understanding of levels defines precisely such a system.[23] We can subsequently compare the systems of levels that are defined by the different understandings.

4.1 A system of levels

A *system of levels* is an ordered pair $<\mathcal{L},\mathcal{M}>$, defined as follows:

- $\mathcal{L}$ is a class of objects called *levels*;
- $\mathcal{M}$ is a class of directed arrows (mappings) between levels in $\mathcal{L}$, called *(inter-level) morphisms*, each of which has a *source level* L and a *target level* L' and is of the form $\mu : L \rightarrow L'$.

A system of levels, I propose, must ideally satisfy three conditions:

(1) *Closure under composition*: If $\mathcal{M}$ contains a mapping from level L to level L' and a mapping from L' to L'', it also contains a composite mapping from L to L''.

(2) *Identity*: For each level L, $\mathcal{M}$ contains an identity mapping from L to itself.

(3) *Uniqueness*: For any pair of levels L, L', $\mathcal{M}$ contains *at most* one mapping from L to L'.

Condition (1) captures the *transitivity* of the inter-level relation encoded by the morphism: whenever L stands in this relation to level L', and level L' stands in this relation to level L'', then level L also stands in the relation to level L''. Condition (2) captures the *reflexivity* of the inter-level relation: each level L stands in the relevant relation to itself (perhaps trivially or vacuously). Condition (3) captures the *uniqueness* of the inter-level relation: whenever two levels L and L' are related by it, then that relation must be unique, though two levels could well be unrelated to each other.

Mathematically speaking, the pair $<\mathcal{L},\mathcal{M}>$ is an algebraic structure called a "category".[24] A *category* is an ordered pair consisting of a class of objects and a class of mappings between objects ("arrows" or "morphisms") satisfying closure under composition (1) and the existence

[23] I first introduced this formalism in List (2019a).
[24] See Marquis (2015).

of an identity map (2). A category whose mappings additionally satisfy the uniqueness property (3) is called a "posetal category". In the present application, the "objects" are levels, and the arrows or mappings are inter-level morphisms.

4.2 The four notions of levels revisited

It should be evident that all four notions of levels give rise to an ordered pair $<\mathcal{L},\mathcal{M}>$. Let's briefly run through them.

First, on the coarse-graining understanding of levels of description,

- the elements of $\mathcal{L}$ are equivalence relations on some underlying set of possibilities, namely one equivalence relation for each level, and
- $\mathcal{M}$ contains precisely one inclusion mapping for every pair of equivalence relations that stand in an inclusion relation to one another.

Second, on the linguistic understanding of levels of description,

- the elements of $\mathcal{L}$ are descriptive languages, namely one language for each level, and
- $\mathcal{M}$ contains precisely one translation function for any pair of such languages that stand in a reducibility relation to one another.

Third, on the entity-based understanding of levels of reality,

- the elements of $\mathcal{L}$ are classes of level-specific entities, namely one class of entities for each level, and
- $\mathcal{M}$ contains precisely one arrow for any pair of such classes where the entities in one of them are parts or building blocks of the entities in the other.

Finally, on the fact- or world-based understanding of levels of reality,

- the elements of $\mathcal{L}$ are sets of level-specific worlds, namely one set of all possible level-specific worlds for each level, and
- $\mathcal{M}$ contains precisely one supervenience mapping for any pair of levels that are related by supervenience.

It should also be evident that, with the possible exception of the entity-based understanding of levels of reality, all of the different understandings define systems of levels satisfying the key category-theoretic conditions of (1) closure under composition, (2) identity, and (3) uniqueness. Specifically, the inclusion relation between equivalence relations, the reducibility relation between languages, and the supervenience relation between facts or worlds are transitive, as required by condition (1). Each of these relations trivially admits identity as a special case, i.e., inclusion, reducibility, and supervenience are each reflexive, as required by condition (2). And each of these inter-level relations, when it exists between two levels, is unique, as required by condition (3). In the case of a parthood or composition relation, the transitivity requirement of condition (1) might still be relatively unproblematic.[25] But it is unclear that parthood or composition are reflexive, so condition (2) may well be violated. Similar remarks would also apply if we were to adopt a fact- or world-based understanding of levels but used grounding instead of supervenience as the inter-level relation. Grounding is not only irreflexive – no fact grounds itself – thereby violating condition (2), but its transitivity is also controversial.[26] Because of its (arguably) neater formal properties, I here prefer to use supervenience rather than grounding as the default inter-level relation for the fact- or world-based understanding of ontological levels. Still, it is worth noting that the present framework, with suitably weakened conditions on a system of levels, could also capture ontological levels that are related via grounding.

4.3 Some broader observations

All four notions of levels I have discussed share the feature that they do not generally give us a linear hierarchy of levels but just a partial ordering. This vindicates a critical remark that Jaegwon Kim made about Oppenheim and Putnam's understanding of levels: "If a comprehensive levels 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."[27] For instance, a system of levels could look like one of the examples in Figure 1.[28]

The category-theoretic perspective confirms that a linearly ordered system of levels, with a fundamental level at the bottom, is just a very special case. Levels could not only be partially rather than totally ordered, but there could also be infinitely descending chains of levels that do not terminate in any bottom level. This, in turn, shows that a "metaphysic of infinite descent", as considered by Jonathan Schaffer in his discussion of whether there is a fundamental level, is coherent, even if a "bottomless ontology" may not ultimately be supported by our best scientific theories of reality.[29] From a historical perspective, however, it is interesting to note that whenever scientists thought that they had hit "rock bottom" and identified the most fundamental building blocks of nature, new discoveries eventually led them to identify even more fine-grained constituents. Think of the move from atoms to electrons, neutrons, and protons, and subsequently to smaller elementary particles, and now to even tinier strings or superstrings of which everything may be composed.

---

[25] For a discussion of arguments for and against the transitivity of parthood, see Varzi (2006).
[26] See Schaffer (2012).
[27] See Kim (2002, pp. 17-18).
[28] This figure is reproduced from List (2019a).
[29] See Schaffer (2003, p. 499). Marcus Pivato and I have also discussed such a scenario in List and Pivato (2015).

Figure 1: Non-linear systems of levels

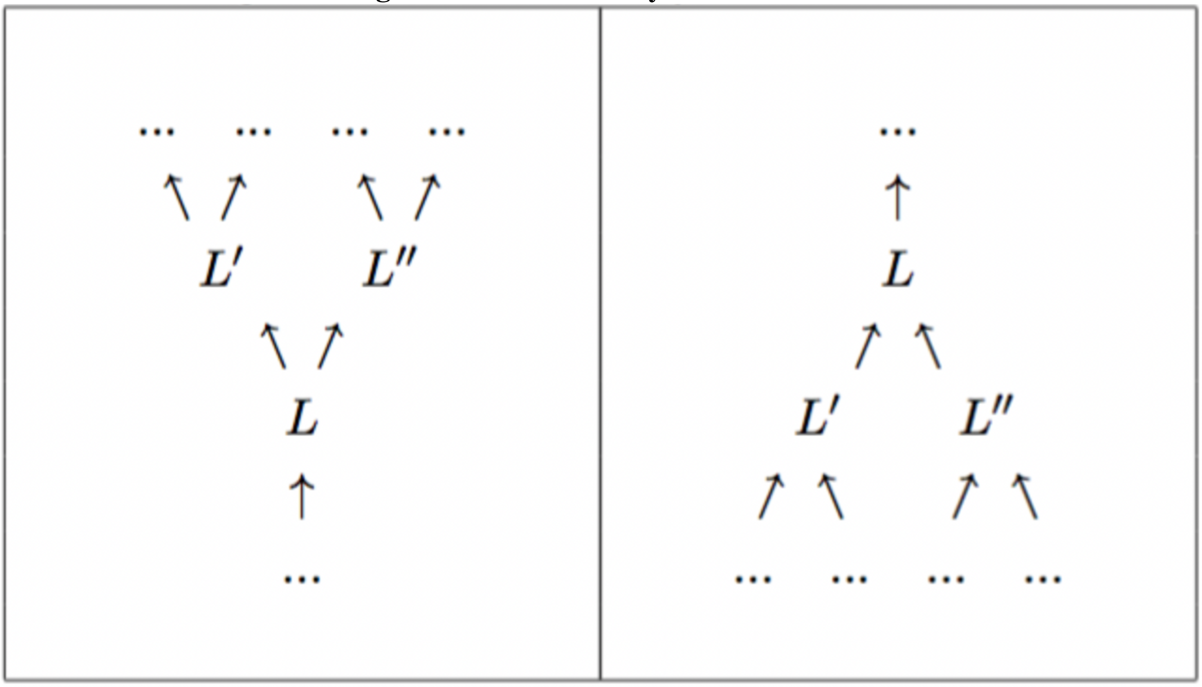

In addition to vindicating the coherence of a non-linear and even bottomless system of levels, the category-theoretic perspective also allows us to study structural relationships between different systems of levels, including relationships between systems of levels of description on the one hand and systems of levels of reality on the other. Technically, it yields a criterion for saying when one system of levels is a *subsystem* of another, and it allows us to identify structure-preserving mappings (so-called *functors*) between different such systems, which can capture structural commonalities between them. We call one system of levels, $<\mathcal{L},\mathcal{M}>$, a *subsystem* of another, $<\mathcal{L}',\mathcal{M}'>$, if $\mathcal{L}$ is a subset of $\mathcal{L}'$ ($\mathcal{L} \subseteq \mathcal{L}'$), $\mathcal{M}$ is a subset of $\mathcal{M}'$ ($\mathcal{M} \subseteq \mathcal{M}'$), and composition and identity in $<\mathcal{L},\mathcal{M}>$ are defined in the same way as in $<\mathcal{L}',\mathcal{M}'>$. If one scientist thinks there are more levels than recognized by another scientist, then the system of levels according to the second scientist is a subsystem of that according to the first. A *functor* is a mapping F from one system of levels, $<\mathcal{L},\mathcal{M}>$, to another such system, $<\mathcal{L}',\mathcal{M}'>$, where F assigns, to each level L in $\mathcal{L}$, a level $L' = F(L)$ in $\mathcal{L}'$, and to each mapping μ in $\mathcal{M}$, a mapping $\mu' = F(\mu)$ in $\mathcal{M}'$ such that composition of mappings and identity are preserved. In the next section, I will give one example of such a structure-preserving mapping, namely between a system of levels of description and a system of ontological levels.

5. The relationship between levels of description and levels of reality

I have raised the question of whether "levels" should be regarded mainly as an epistemic phenomenon, i.e., a feature of how we think about the world, or also as an ontic phenomenon, i.e., a feature of reality itself. And I have asked how levels in the epistemic sense, which we undeniably find in science, relate to levels in the ontic sense.

In response, I will now sketch a technical and a philosophical argument for the thesis that levels of description do indeed correspond to levels of reality, and on the assumption that this thesis is correct, I will consider the relationship between supervenience, which is a key inter-level relation on the ontic side, and reducibility, which is a key inter-level relation on the epistemic side.

5.1 Do levels of description correspond to levels of reality?

I will first sketch a purely formal answer to this question, and I will then suggest a philosophical answer.[30] Formally, I will show that any descriptive language, as defined in Section 2.2, induces a corresponding set of level-specific possible worlds, as defined in Section 3.2. To establish this claim, let **L** be a descriptive language. This allows us to define – at least in formal terms – a corresponding set of worlds, which we may call $\Omega_{\mathbf{L}}$. Specifically, we can identify the elements of $\Omega_{\mathbf{L}}$, the "worlds", with *maximal consistent* subsets of **L**, i.e., sets of sentences from **L** that are consistent but where the addition of *any* further sentence from **L** would introduce an inconsistency. One can think of each element of $\Omega_{\mathbf{L}}$ as a minimally rich world that "settles" everything that can be expressed in **L**. To *settle* a sentence in **L** is to assign a truth-value to it: "true" or "false". A sentence ϕ in **L** is true at some world ω in $\Omega_{\mathbf{L}}$ if ϕ is contained in the maximal consistent subset of **L** representing ω, and ϕ is false if it isn't. Each element of $\Omega_{\mathbf{L}}$ thus picks out a way the world could be (a "possible world") such that everything that can be expressed in **L** is settled and nothing else is settled that isn't entailed by a set of sentences expressible in **L**. Of course, we need not literally think of a maximal consistent subset of **L** as a world, but we can think of it as representing a world. As soon as we are treating the sentences in **L** as having truth-conditions, we are thereby at least implicitly postulating the existence of some world ω in $\Omega_{\mathbf{L}}$ that determines which sentences in **L** are true and which not.

Applying this reasoning to an entire system of levels of description, we can see that each descriptive language **L** in the given system induces a corresponding set of level-specific worlds $\Omega_{\mathbf{L}}$ within a system of ontological levels. Moreover, whenever two descriptive languages **L** and **L′** stand in a reducibility relation, then the worlds in the corresponding sets $\Omega_{\mathbf{L}}$ and $\Omega_{\mathbf{L'}}$ are related by supervenience. To see this, let *f* be the translation function from **L′** (the higher-level language) to **L** (the lower-level language), and consider any lower-level world ω in $\Omega_{\mathbf{L}}$. We need to show that this determines a supervenient higher-level world ω' in $\Omega_{\mathbf{L'}}$. Let ω' be given by the set consisting of every sentence ϕ from the higher-level language **L′** whose lower-level counterpart $f(\phi)$ is true at ω. Since *f* preserves logical properties such as consistency and inconsistency, the set of higher-level sentences thus defined forms a consistent subset of **L′**. To see that it is *maximal* consistent, consider any other sentence ψ from **L′** that is not yet included in it. It follows from the definition of our set that $f(\psi)$ is not true at ω, so its negation $\neg f(\psi)$ is true at ω. Since *f* preserves negation, $f(\neg\psi)$ must also be true at ω, and therefore $\neg\psi$ meets the

[30] My formal answer is based on the analysis in List (2019a), but that paper did not contain an explicit formal argument to the effect that reducibility implies supervenience.

membership criterion for the set of sentences defining ω'. We can then infer that this set together with ψ is inconsistent, and consequently that ω' is indeed maximal consistent. So, the mapping that assigns to each lower-level world ω the higher-level world ω' thus constructed qualifies as a supervenience mapping from $\Omega_{\mathbf{L}}$ to $\Omega_{\mathbf{L}'}$.

In this way, we have arrived at a functor which maps a given system of levels of description, with reducibility as the inter-level relation, to a system of ontological levels, with supervenience as the inter-level relation.

This formal result also suggests a philosophical answer to the question of whether levels of description correspond to levels of reality. It is this: we can take the fact that levels of description are so useful and even indispensable in science as indicative of an underlying levelled ontology of reality. The idea is that if science supports a certain system of levels of description $<\mathcal{L},\mathcal{M}>$ in our best explanations of reality, then this is good evidence – a good indicator – that reality itself contains a corresponding system of ontological levels $<\mathcal{L}',\mathcal{M}'>$. On this picture, each level-specific language $\mathbf{L}$ in $\mathcal{L}$ picks out a corresponding ontological level in $\mathcal{L}'$, given by the set of level-specific worlds $\Omega_{\mathbf{L}}$ derivable from $\mathbf{L}$. Further, $\mathcal{M}'$ consists of supervenience mappings between appropriately related pairs of such sets of level-specific worlds. Minimally, a supervenience relation holds between any two levels for which the corresponding level-specific languages stand in a reducibility relation, but we shall see in the next subsection that reducibility is not necessary for supervenience.

The upshot is that the system of levels of description supported by science might mirror a system of ontological levels "out there in reality".

5.2 Does supervenience imply reducibility?

We have seen that whenever two distinct level-specific languages stand in a reducibility relation, then the facts or worlds at the higher one of the two corresponding ontological levels supervene on the fact or worlds at the lower. But what about the converse? Is it also true that whenever the facts or worlds at some higher level supervene on those at some lower level, then the corresponding higher-level descriptions are reducible to the relevant lower-level ones?[31]

I will now answer this question in the negative: supervenience is not sufficient for reducibility. To establish this, consider two distinct level-specific languages $\mathbf{L}$ and $\mathbf{L}'$, and let $\Omega_{\mathbf{L}}$ and $\Omega_{\mathbf{L}'}$ be the corresponding level-specific sets of possible worlds. Moreover, suppose that there is a supervenience mapping from the lower one of the two levels to the higher, formally a surjective function σ from $\Omega_{\mathbf{L}}$ to $\Omega_{\mathbf{L}'}$. I want to show that, under plausible assumptions, the existence of a translation function for reducing the higher-level language $\mathbf{L}'$ to its lower-level counterpart $\mathbf{L}$ is not guaranteed but rather a very special case. Recall that such a translation function, say f,

[31] The discussion of the possibility of supervenience without reducibility goes back at least to Fodor (1974) and Putnam (1967).

would have to assign to each sentence in the higher-level language $\mathbf{L}'$ an "equivalent" sentence in the lower-level language **L**, where logical properties are preserved under translation. To capture the requirement of "equivalence", in turn, we require that whenever ϕ is a higher-level sentence and $f(\phi)$ is its lower-level counterpart, the set of worlds ω in $\Omega_{\mathbf{L}}$ at which the lower-level sentence $f(\phi)$ is true – call that set $[f(\phi)]$ – is the inverse image, under the supervenience mapping σ, of the set of worlds ω' in $\Omega_{\mathbf{L}'}$ at which the higher-level sentence ϕ is true, denoted $[\phi]$. Formally,

$$[f(\phi)] = \sigma^{-1}([\phi]) = \{\omega \in \Omega_{\mathbf{L}} : \sigma(\omega) \in [\phi]\}.$$

Could there be such a translation function? Suppose that

(1) the set $\Omega_{\mathbf{L}}$ of lower-level worlds is infinite, in line with the assumption that infinitely many distinct initial conditions of the world are at least in principle nomologically possible;
(2) the languages we are considering, including the lower-level language **L**, are countable, in the sense that they permit the expression of as many sentences as there are natural numbers, but no more; this is a feature of practically all familiar formal and natural languages, from standard propositional logic to English.

From assumption (1), it follows that there are uncountably many subsets of $\Omega_{\mathbf{L}}$ (because any infinite set has uncountably many subsets), and from assumption (2), it follows that only countably many of them are describable in the lower-level language **L**, in the sense that there exists some sentence ψ in **L** whose content $[\psi]$ matches the given subset of $\Omega_{\mathbf{L}}$ (because the language admits only countably many sentences). In consequence, *almost all* subsets of $\Omega_{\mathbf{L}}$, i.e., all but a countable number, are *not* describable by a sentence (or equivalently, even by a finite logical combination of sentences) from the lower-level language **L**. This has an immediate implication for our question of whether we can assume the existence of a translation function from the higher-level language $\mathbf{L}'$ to the lower-level language **L**. Take any higher-level sentence ϕ. Given supervenience, it will certainly be true that there exists some set of lower-level worlds that forms the "supervenience base" of the content expressed by ϕ. Formally, the set $\sigma^{-1}([\phi])$ will exist and be a subset of $\Omega_{\mathbf{L}}$. However, since almost all subsets of $\Omega_{\mathbf{L}}$ are not describable by any sentence from **L**, it would be a highly special case if $\sigma^{-1}([\phi])$ were so describable. Therefore, we cannot generally assume that there will exist a sentence ψ in **L** whose content $[\psi]$ is equal to $\sigma^{-1}([\phi])$. And so, the existence of a translation function f from $\mathbf{L}'$ to **L** is the exception rather than the rule, in combinatorial terms. I conclude that supervenience does not imply reducibility.

Of course, one could try to formulate additional conditions under which supervenience does imply reducibility. Notably, Neil Dewar, Samuel Fletcher, and Laurenz Hudetz have proposed two conditions on the two languages **L** and $\mathbf{L}'$ that are jointly sufficient for supervenience to imply reducibility.[32] One condition, called *compatibility*, requires, informally, that if the two languages share some vocabulary, they "agree" with regard to things expressible in the shared vocabulary. The other condition, called *joint characterizability*, requires, in the authors' own informal gloss, that "the union of two levels of description relative to a supervenience map admits of a description itself".[33] Now, compatibility seems to me to be a relatively undemanding condition. Moreover, it does not require the existence of any shared vocabulary between the two languages at all; it only requires that *if* there is some shared vocabulary, its meaning must be matched. Joint characterizability, however, seems much more demanding, as the authors recognize. If we take the example of psychology and fundamental physics, should we really assume that the union of these two levels of description admits a description itself? I take it that there is such a joint description whenever we are able to spell out explicit bridge laws between the levels in question, but often we aren't able to spell out such bridge laws. For this reason, the assumption of joint characterizability seems to me to come close to the assumption that there are explicitly describable bridge laws, in which case it is less of a surprise that this condition is favourable to the existence of a translation function between the two languages. So, I suggest that even though Dewar, Fletcher, and Hudetz have obtained an interesting formal result which may be applicable to some cases of inter-level relations, the cases it covers remain special, and we cannot generally assume that when there is supervenience, there is also reducibility.

As an aside, an analysis similar to the one given in this subsection would also show that if Ω_1 and Ω_2 are distinct ways of coarse-graining some underlying set Ω of possible worlds, representable by distinct equivalence relations on Ω, then the inclusion of the equivalence relation representing Ω_1 within the one representing Ω_2 would not imply the reducibility of a descriptive language we might use to describe Ω_1 to a language we might use to describe Ω_2. This speaks to the question of how the two epistemic inter-level relations mentioned earlier, inclusion of equivalence relations on the one hand and reducibility on the other, relate to one another.

6. Some further payoffs and applications

Arguably, many philosophical problems concern the relationship between phenomena that are intuitively at different levels, and so the present framework offers some resources for thinking about such problems. I will here mention just a few examples.

6.1 The compatibility of free will and determinism

Free-will sceptics often argue that because everything in the world is governed by the fundamental laws of physics, there is no room for free will. Humans might have the illusion that they are able to choose and control their own actions, but in reality everything is

[32] See Dewar, Fletcher, and Hudetz (2019).
[33] Ibid.

determined by underlying physical processes over which we have no control.[34] One way to respond to this kind of free-will scepticism is to note that free will and choice are phenomena at the level of agency rather than at the level of physics. In particular, we can speak about free will and choice only if we use the concepts and categories of psychology and the human sciences. Without those concepts and categories, we would not be able to refer to agents and their actions, let alone ask whether these qualify as free. By contrast, the underlying physical processes, for instance those in the brain and body, are sub-agential phenomena, which belong to the level of physics, biology, or neuroscience. Many of the sceptical arguments fail to recognize the multi-levelled nature of the free-will problem and involve a mixing of levels.

To give just one example, free will plausibly requires the possibility of doing otherwise, i.e., of choosing between alternative actions, and at first sight there seems to be no such possibility if the fundamental laws of physics are deterministic, and determinism has not yet been ruled out by the physical sciences. However, once we carefully distinguish between the level of physics and the level of agency, we can see that each level is endowed with its own modal notions: possibility at the level of agency ("agential possibility") on the one hand, and physical possibility on the other. These are distinct notions, just as chemical possibility, biological possibility, and economic possibility are distinct. This insight, in turn, leaves room for showing that the possibility of doing otherwise at the level of agency can co-exist with determinism at the level of physics. Conditional on the state of the world *as specified at the level of agency*, different courses of action may be open to me and thus *agentially possible* for me, even if there is some sub-agential specification of the state of the world *at the level of microphysics* at which only a single physical trajectory is *physically possible*. There is no contradiction here: at the level of physics, we would not even be able to speak about the choices that I could or could not make; the agential "can" does not belong to the vocabulary of physics. At the level of agency, on the other hand, we would not be able to refer to, or conditionalize on, the detailed physical microstate. So, it would also make little sense to say that "conditional on the physical microstate, it is agentially impossible for me to act otherwise". This claim would mix two different levels of description that do not go together and between which there is arguably no relation of reducibility. Arguments for the incompatibility of free will and determinism, such as van Inwagen's famous consequence argument, tend to draw conclusions about what agents can and cannot do from premises about the constraints that the fundamental laws of physics place on the physical microstate, thereby in effect conflating physical and agential levels.[35]

---

[34] This kind of scepticism is reviewed (with literature references) in List (2019b), which (along with List 2014) is also the source of the response summarized here. Others who have defended free will by arguing, in a variety of ways, that free-will is a higher-level phenomenon rather than a physical-level one include Kenny (1978), Dennett (2003), Siderits (2008), and Carroll (2016). Furthermore, as Koons (2002) has recently pointed out and further elaborated, Wilfrid Sellars, who famously discussed the contrast between what he called the "scientific image" and the "manifest image", held a view on free will that is arguably a precursor to the one sketched here.

[35] See, e.g., van Inwagen (1975) and the response in List (2019c).

6.2 The level-specificity of dynamic properties

As already implicit in my brief discussion of free will, dynamic properties of a system, such as whether the system is deterministic or indeterministic, are best understood as level-relative properties. If we ask whether the world is deterministic or whether there is room for genuine randomness or some other source of indeterminism, the answer can be given only once we are clear about the level at which we are asking those questions. There might well be determinism at one level, say that of microphysics, and indeterminism at another, say that associated with some special science. The contrast between classical and statistical mechanics, where systems are conceptualized as, respectively, deterministic and probabilistic, is a case in point.[36]

Formally, if we think about each possible world as a trajectory the world might take through its state space across time (specifying in which state the world is at each point in time), then *determinism* means that any initial segment of any such trajectory up to any point in time admits only one continuation among the nomologically possible trajectories. *Indeterminism* means that some initial segment of some trajectory up to some point in time admits two or more distinct continuations among the nomologically possible trajectories: there is, at least sometimes, a "fork in the road".

It is easy to see that if macro-level trajectories result from micro-level trajectories via some way of coarse-graining the underlying state space, such as with the help of some equivalence relation on the set of microstates, then the distinction between determinism and indeterminism is level-specific. Low-level trajectories could be deterministic while high-level trajectories could be indeterministic, or it could be the other way round.[37] As Jeremy Butterfield puts it, the micro- and macro-level dynamics of a system need not "mesh".[38] When we move from a lower level of description to a higher one, we might see a kind of "phase transition" from deterministic to indeterministic dynamics or vice versa. Empirical considerations alone would then not allow us to settle the question of whether a particular system is deterministic or not, as Charlotte Werndl has pointed out.[39] The question receives a determinate answer only when we are clear about the level at which we are considering the system. Even a bottomless hierarchy of levels in which there is determinism at even-numbered levels and indeterminism at odd-numbered levels is coherent, albeit a somewhat contrived scenario.[40]

Similarly, one may argue that there can be "emergent" higher-level chance in a system that admits a deterministic lower-level description.[41] A necessary condition for non-trivial objective

[36] For a recent discussion of coarse-graining in the move from classical to statistical mechanics, see Robertson (2020).
[37] For formal versions of this point, see Werndl (2009), Butterfield (2012), Yoshimi (2012), List (2014), and List and Pivato (2015).
[38] See Butterfield (2012).
[39] On the observational indistinguishability of deterministic and indeterministic descriptions, see Werndl (2009).
[40] See List and Pivato (2015).
[41] Ibid. On probability in the context of deterministic physics, see also Ismael (2009).

chance at a given level is merely the presence of the indeterminism at the relevant level, not the presence of indeterminism at some lower level.[42] We can thus see that, while within a given level objective chance is incompatible with determinism, across levels the incompatibility goes away: lower-level determinism is compatible with higher-level objective chance.

6.3 Indexical versus non-indexical and first-personal versus third-personal descriptions

In discussions of indexicality and subjectivity, it is often acknowledged that indexical facts cannot be derived from non-indexical ones and, similarly, that subjective facts cannot be derived from objective ones. David Lewis famously gives the following example:

> "Consider the case of the two gods. They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. Insofar as knowledge is a propositional attitude [with third-personal, non-indexical content], they are omniscient. Still I can imagine them to suffer ignorance: neither one knows which of the two he is."[43]

Each of the two gods has complete third-personal and non-indexical knowledge of the world, and yet lacks knowledge of his own position relative to the world: is he the one on the left or the one on the right, for example?

Similarly, even if we had complete information about the entire trajectory of the physical universe – from the beginning of time *ad infinitum* – we would not be able to infer from this what the present time is, i.e., the location of the "now", or at which spatial coordinates we are positioned, i.e., the location of the "here". In short, the non-indexical facts under-determine the indexical ones.

I suggest that we can think of non-indexical and indexical phenomena as residing on two different levels. Using the present framework, we can identify the non-indexical level with an ordinary set Ω of possible worlds, each of which is a total specification of all non-indexical facts, while we can identify the indexical level with a set of *centred worlds*, a set of ordered pairs consisting of a world ω in Ω and a centre c within that world, which could be a spatio-temporal coordinate or a pointer to a particular individual.[44] Such centred worlds settle indexical as well as non-indexical facts, by including a centre as a kind of location pointer. On this picture, the non-indexical level is the higher, more coarse-grained one, while the indexical level is the lower, more fine-grained one; different centres can be combined with the same total body of non-indexical facts. The non-supervenience of indexical facts on non-indexical ones is an immediate consequence.

---

[42] This point is formally developed in List and Pivato (2015).

[43] See Lewis (1979, p. 520).

[44] On centred worlds, see Quine (1969), Lewis (1979), Liao (2012), and Milano (2018).

Similarly, some philosophers of mind have argued that even if we were to specify the totality of third-personal facts about the world, i.e., those describable by the ordinary sciences, this would leave open the facts about first-personal experience: what it is like for conscious subjects to experience and perceive the world first-personally, or indeed whether there are any first-personal experiences at all.[45] If this is right, then the third-personal facts under-determine the first-personal ones. David Chalmers describes the challenge for a science of consciousness as follows:

> "The task of a science of consciousness ... is to systematically integrate two key classes of data into a scientific framework: *third-person data*, or data about behavior and brain processes, and *first-person data*, or data about subjective experience."[46]

In analogy with my brief discussion of indexicality, I suggest that we can think of first-personal and third-personal facts as residing on two different levels too.[47] We can amend the machinery of centred worlds to capture the idea that the facts of first-person experience hold only at what we may call "first-personally centred worlds", ordered pairs consisting of an ordinary third-personal world ω and a "locus of subjectivity" π, where π encodes a subject's first-person perspective on the world ω. The combination of ω and π will then determine not only all third-personal facts that hold at ω but also all first-personal facts that hold for the relevant subject.

Once more, we have a two-level structure. The first-personally centred level is given by the set of all possible first-personally centred worlds, and the third-personal level is given by the ordinary set Ω of all possible third-personal worlds. Just as, in the case of indexicality, the indexical level is lower (subvenient) and the non-indexical level is higher (supervenient), so the first-personally centred level is lower (subvenient) and the third-personal level is higher (supervenient).

This vindicates the claim, made by Chalmers and others, that the facts about first-personal experience do not supervene on the ordinary physical facts.[48] It further shows that there is a structural parallel between indexicality and subjectivity. Finally, on the present picture, the much-discussed "hard problem of consciousness" is due to the fact that ordinary science only ever delivers third-personal explanations of third-personal phenomena, while the explanation of first-personal experience involves an explanandum that can only be found at a different, more richly specified level, namely the first-personally centred one.[49]

---

[45] Classic discussions of this point include Nagel (1974), Jackson (1982), Levine (1983), and Chalmers (1996).
[46] See Chalmers (2004, p. 1111).
[47] I have discussed this proposal in detail in List (2022).
[48] See, in particular, Chalmers (1996).
[49] For more on this, see List (2022).

6.4 Positive versus normative facts

A final illustrative application of the present framework concerns the relationship between positive and normative facts and the fact-value distinction. Positive facts, sometimes also just called "descriptive facts", are facts such as "H_2O consists of two hydrogen atoms and one oxygen atom", "green plants use light energy to convert water, carbon dioxide, and minerals into oxygen and certain organic compounds", and "increases in the interest rate tend to lead to decreases in inflation, other things being equal". Normative facts – if they exist, as moral realists assume – are facts such as "killing is wrong", "all humans deserve equal moral consideration", and "society ought, or ought not, to be organized in such-and-such a way". Similarly, evaluative facts – again, if they are genuine facts – are facts such "education is good", "freedom is desirable", and "ecosystems are valuable".

Debates about moral naturalism and non-naturalism revolve around the question of how normative or evaluative facts relate to positive or descriptive ones. Do normative or evaluative facts supervene on positive or descriptive ones, or is this not the case? Moreover, if there is supervenience, is there also reducibility, in the sense that normative or evaluative discourse is translatable into positive or descriptive discourse? Or could we have a case of supervenience without reducibility? Normative or evaluative descriptions might be irreducible, even if the facts they express are, or supervene on, natural facts.

While the present framework can obviously not settle these difficult meta-ethical questions, it provides a formalism in which they can be articulated precisely. For a start, we can compare a purely positive and descriptive language with a normative or evaluative language. The latter is, in some ways, richer than the former, insofar as it includes deontic operators such as "ought" and "may" and/or evaluative predicates such "good", "bad", "desirable", and "undesirable", which are absent from the positive and descriptive language. The two languages – call them **L** and **L**′ – clearly define different levels of description in the sense discussed in this paper, and this already allows us to see precisely what it would mean to say that normative or evaluative discourse is reducible to positive or descriptive discourse: there would have to be a translation function from **L**′ to **L** which preserves content and logical properties. Moreover, the two languages, at least when taken at face value, can be thought to induce two corresponding ontological levels: one level would be given by the set of all possible worlds in a positive or descriptive sense, the other by the set of all possible worlds in some normatively or evaluatively augmented sense. A possible world in the latter set explicitly includes – in addition to ordinary positive facts – a specification of all normative or evaluative facts, while a possible world in the former set omits such facts or includes them at most implicitly, in case the hypothesis that they supervene on positive facts is true.

Elsewhere I have suggested that we could model "normatively augmented worlds" as ordered pairs consisting of an ordinary positive or descriptive world ω from some set Ω and a selection function f which assigns to each world ω a set of permissible worlds relative to ω.[50] Any ordered pair of the form $<\omega, f>$ will then be rich enough to settle not only the truth-value of all positive and descriptive sentences but also that of all sentences involving normative operators such as "ought" and "may". For instance, "it is obligatory that p" ("ought p") is true at the normatively augmented world $<\omega, f>$ if and only if p is true at *all* worlds that f deems permissible relative to ω, i.e., which are in the set $f(\omega)$. Similarly, "it is permissible that p" ("may p") is true at $<\omega, f>$ if and only if p is true at *some* worlds in $f(\omega)$.

Under this construction, there exists a many-to-one supervenience mapping from the set of all normatively augmented worlds to the set of positive or descriptive worlds. This mapping, σ, would simply map each ordered pair $<\omega, f>$ to its first component, i.e., $\sigma(<\omega, f>) = \omega$. So, the positive or descriptive level appears to be higher or more coarse-grained, while the normatively augmented level is lower or more fine-grained. This, in turn, would speak against the supervenience thesis entailed by normative naturalism and vindicate the claim that deriving an "ought" from an "is" is indeed a fallacy.[51]

However, if one could somehow show that one and only one selection function f is possible relative to each positive or descriptive world ω, then one might still be able to defend the naturalistic supervenience thesis. In this case, there would be a one-to-one correspondence between the positive or descriptive worlds and the normatively augmented ones. But at least from the perspective of logic, it is hard to see why only one selection function f should be *logically* possible for each ω. This is not the place to discuss these questions in any detail. I simply hope to have shown that the present framework allows us to look at them in a clear and systematic way.

In sum, I have reviewed several salient uses of the idea of levels, in both epistemic and ontic senses, and explained how they can all be accommodated within a unified framework. I have shown that this allows us to shed light on questions such as how levels of description relate to levels of reality and whether supervenience implies reducibility. Finally, I have considered some illustrative applications of this framework, in the hope that they will inspire further applications as well as extensions of the framework itself.

References

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.

Block, Ned. 2003. Do Causal Powers Drain Away? *Philosophy and Phenomenological Research* 67(1): 133–150.

[50] See List (2019a).

[51] For a recent analysis of the is-ought gap, see also Brown (2014).

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.

Carroll, Sean M. (2016). *The Big Picture: On the Origins of Life, Meaning, and the Universe itself*. New York: Dutton.

Chalmers, David. 1996. *The Conscious Mind*. New York: Oxford University Press.

Chalmers, David. 2004. How can we construct a science of consciousness? In Michael S. Gazzaniga, ed., *The cognitive neurosciences III*, 3rd ed., 1111–1120. Cambridge, MA: MIT Press.

Davidson, Donald. 1973. On the Very Idea of a Conceptual Scheme. *Proceedings and Addresses of the American Philosophical Association* 47: 5–20.

Dennett, Daniel. 2003. *Freedom Evolves*. London: Penguin.

Dewar, Neil, Samuel C. Fletcher, and Laurenz Hudetz. 2019. Extending List's Levels. In Marek Kuś and Bartłomiej Skowron, eds., *Category Theory in Physics, Mathematics, and Philosophy*, *CTPMP 2017. Springer Proceedings in Physics*, vol. 235, 63–81. Heidelberg: Springer.

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.

Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

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.

Jackson, Frank. 1982. Epiphenomenal qualia. *Philosophical Quarterly* 32(127): 127–136.

Hettema, Hinne. 2012. *Reducing Chemistry to Physics: Limits, Models, Consequences*. PhD thesis, University of Groningen.

Himmelreich, Johannes. 2015. *Agency as difference-making: Causal Foundations of Moral Responsibility*. PhD thesis, LSE <http://etheses.lse.ac.uk/3277/>.

Ismael, Jenann. 2009. Probability in Deterministic Physics. *Journal of Philosophy* 106(2): 89–108.

Kenny, Anthony. 1978. *Freewill and Responsibility*. London: Routledge.

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. 2002. The Layered Model: Metaphysical Considerations. *Philosophical Explorations* 5(1): 2–20.

Knox, Eleanor. 2016. Abstraction and its Limits: Finding Space For Novel Explanation. *Noûs* 50(1): 41–60.

Koons, Jeremy. 2022. Sellars on compatibilism and the consequence argument. *Philosophical Studies* 179: 2361–2389.

Levine, Joseph. 1983. Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly* 64(4): 354–361.

Lewis, David. 1979. Attitudes de dicto and de se. *The Philosophical Review* 88(4): 513–543.

Lewis, David. 1988. Relevant implication. *Theoria* 54(3): 161–174.

Liao, Shen-yi. 2012. What are centered worlds? *The Philosophical Quarterly* 62(247), 294–316.

List, Christian. 2014. Free will, determinism, and the possibility of doing otherwise. *Noûs* 48(1): 156–178.

List, Christian. 2019a. Levels: Descriptive, explanatory, and ontological. *Noûs* 53(4): 852–883.

List, Christian. 2019b. *Why Free Will is Real*. Cambridge, MA: Harvard University Press.

List, Christian. 2019c. What's Wrong with the Consequence Argument: A Compatibilist Libertarian Response. *Proceedings of the Aristotelian Society* 119(3): 253–274.

List, Christian. 2022. The many-worlds theory of consciousness. *Noûs*, online early.

List, Christian, and Marcus Pivato. 2015. Emergent Chance. *The Philosophical Review* 124(1): 119–152.

Manafu, Alexandru 2015. A Novel Approach to Emergence in Chemistry. In Eric Scerri and Lee McIntyre, eds., *Philosophy of Chemistry: Growth of a New Discipline*, 39–55. Heidelberg: Springer.

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/>.

Milano, Silvia. 2018. *De se beliefs and centred uncertainty*. PhD thesis, London School of Economics <https://doi.org/10.21953/lse.b2wsi6xbghk6>.

Modica, Salvatore, and Aldo Rustichini. 1999. Unawareness and Partitional Information Structures. *Games and Economic Behavior* 27(2): 265–298.

Nagel, Thomas. 1974. What is it like to be a bat? *The Philosophical Review* 83(4): 435–450.

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.

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.

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.

Quine, W.V. 1969. Propositional objects. In *Ontological Relativity and Other Essays*, 139–160. New York: Columbia University Press.

Robertson, Katie. 2020. Asymmetry, Abstraction, and Autonomy: Justifying Coarse-Graining in Statistical Mechanics. *British Journal for the Philosophy of Science* 71(2): 547–579.

Rosen, Gideon. 2010. Metaphysical Dependence: Grounding and Reduction. In Bob Hale and Aviv Hoffmann, eds., *Modality: Metaphysics, Logic, and Epistemology*, 109–135. Oxford: Oxford University Press.

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.

Schaffer, Jonathan. 2012. Grounding, transitivity, and contrastivity. In Fabrice Correia and Benjamin Schnieder, eds., *Metaphysical Grounding: Understanding the Structure of Reality*, 122–138. Cambridge: Cambridge University Press.

Siderits, Mark. 2008. Paleo-Compatibilism and Buddhist Reductionism. *Sophia* 47(1): 29– 42.

Van Inwagen, Peter. 1975. The Incompatibility of Free Will and Determinism. *Philosophical Studies* 27(3): 185–199.

Varzi, Achille C. 2006. A Note on the Transitivity of Parthood. *Applied Ontology* 1(2): 141–146.

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.

Yoshimi, Jeffrey. 2012. Supervenience, Dynamical Systems Theory, and Non-Reductive Physicalism. *British Journal for the Philosophy of Science* 63(2): 373–398.