Extending List's Levels

Neil Dewar, Samuel C. Fletcher, and Laurenz Hudetz

Abstract Christian List [24] has recently proposed a category-theoretic model of a system of levels, applying it to various pertinent metaphysical questions. We modify and extend this framework to correct some minor defects and better adapt it to application in philosophy of science. This includes a richer use of category theoretic ideas and some illustrations using social choice theory.

1 List's descriptive, explanatory, and ontological levels

In general, a *system of levels* for List [24] is a preordered class: that is, a class $\mathscr{L}$ equipped with a reflexive and transitive binary relation $\leq$. The elements of $\mathscr{L}$ are to be interpreted as *levels*, and the binary relation $\leq$ as the relation of *supervenience*: $L \leq L'$ means L' supervenes on L. Thus, requiring that $\leq$ be a preorder amounts to assuming that every level supervenes upon itself, and if the level L_1 supervenes on the level L_2, and L_2 on the level L_3, then L_1 supervenes on L_3. In requiring only these characteristics, List is deliberately opting for a fairly weak conception of supervenience. For instance, it is possible to have two distinct levels within the system, neither of which supervene upon the other, or to have two distinct levels both of which supervene upon the other.

Neil Dewar
LMU Munich, Geschwister-Scholl-Platz 1, 80539 Munich, Germany
e-mail: neil.dewar@lrz.uni-muenchen.de

Samuel C. Fletcher
University of Minnesota, Twin Cities, 271 19th Ave S, Minneapolis, MN 55455, USA
e-mail: scfletch@umn.edu

Laurenz Hudetz
London School of Economics, Houghton Street, WC2A 2AE, London, UK
e-mail: l.hudetz@lse.ac.uk

This framework gives us the resources to consider certain relationships between systems of levels. Given two systems of levels $\mathscr{L}$ and $\mathscr{L}'$, a function $f : \mathscr{L} \to \mathscr{L}'$ is a *monotonic* map if it preserves the order relation: i.e., for any levels $L_1, L_2 \in \mathscr{L}$, if $L_1 \leq L_2$ then $f(L_1) \leq f(L_2)$. If there are monotonic maps $f : \mathscr{L} \to \mathscr{L}'$ and $f' : \mathscr{L}' \to \mathscr{L}$ that are mutually inverse to one another, then $\mathscr{L}$ and $\mathscr{L}'$ are said to be *isomorphic*, or structurally equivalent. If $\mathscr{L} \subseteq \mathscr{L}'$ and there is some map $f : \mathscr{L} \to \mathscr{L}'$ such that $L_1 \leq L_2$ if and only if $f(L_1) \leq f(L_2)$, then $\mathscr{L}$ is a *subsystem* of $\mathscr{L}'$.

Such is List's general framework. There is one significant difference between our presentation of this framework so far and his: we have not yet mentioned *categories* at all. Although we will start to use category-theoretic language below, we have avoided it so far to make clear that the abstract framework in his [24] can be understood without category theory.

Let us now turn to more specific kinds of systems of levels. First, consider the case of a system of *ontological* levels. In such a system, each level is associated (or identified) with a set of *possible worlds* for that level. If one level supervenes upon another, then there is a (unique) *supervenience map* from the subvenient level to the supervenient level; this map is required to be surjective.[1] We also impose the following two requirements:

1. for any level L, the identity map is the supervenience map from L to itself; and
2. if σ is the supervenience map from L_1 to L_2, and σ' is the supervenience map from L_2 to L_3, then $\sigma' \circ \sigma$ is the supervenience map from L_1 to L_3.

More compactly stated, a system of ontological levels forms a *concrete posetal category* in which every function is surjective. To say that it is a *concrete category* means that it is a class of sets, equipped with functions between those sets that are closed under composition and include all identity functions. To say that it is a *posetal* category means that between any two sets, there is at most one function. Although this compact description uses category-theoretic apparatus, it only does so in the form of appeal to concrete categories (i.e., categories of sets and functions). This means that the category-theoretic language provides a convenient way to express our requirements, rather than being an indispensable tool; if desired, we could do everything purely in set-theoretic terms, as our initial description of a system of levels in terms of a preordered class evinces.[2]

Two specific kinds of systems of ontological levels are worth mentioning. The first is what List calls a *system of levels of grain*. For this, we take as given a set Ω of "possible worlds". Each level in the system is given by some partition of Ω: one such partition supervenes upon another just in case every cell of the former is a union of cells of the latter, with the supervenience map taking each "fine-grained" cell to the "coarse-grained" cell of which it is a part. Since any such map is guaranteed to

[1] In section 3 we will consider whether this requirement is justified.

[2] Of course, there is a sense in which this is true for any application of category-theoretic apparatus, at least insofar as one can represent a category as a set-theoretic structure. But this trivial sense is not the one we have in mind here.

be surjective, it follows that a system of levels of grain is an ontological system of levels.

List [24, p. 10] claims that even though "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 of levels of grain." In case the system of ontological levels has a lowest level—a level on which all others supervene—the desired functor maps that level bijectively onto a class representing the lowest-level possible worlds, and all other levels are mapped to partitions of these worlds induced by the supervenience maps of the system of ontological levels. In case there is no lowest level, List proposes to construct one formally using an inverse limit of the ontological levels. "For a posetal category," he writes, "an inverse limit can always be constructed, though we need not interpret it as anything more than a mathematical construct."

This last statement must be qualified, however. Inverse limits are defined only for posetal categories each pair of whose elements has a greatest lower bound [1, p. 194],[3] so an inverse limit can be constructed in a system of ontological levels when any two levels have a common level on which they both supervene that itself supervenes on all such common subvenient levels. A simple example of a system of levels in which this does not occur consists of two levels, L_1 and L_2, that only supervene on themselves. No inverse limit can be constructed for this system, and it does not mirror any system of levels of grain.

Perhaps there is some other construction that allows one to exhibit a functor from a large class of—if not *any*—system of ontological levels to some system of levels of grain. In the previous simple example, for instance, it is quite natural to define a system of levels that adds a new level to L_1 and L_2, whose worlds are the Cartesian product of the worlds of each of those two levels, and a pair of supervenience maps for the two components' projection maps.[4] This could then be generalized to any system of ontological levels that has a set-sized number of "lowest levels," and might be given a category theoretic expression using enriched categories. But as systems of ontological levels are currently defined, they includes systems with a (proper) class of lowest levels, whose worlds cannot be so combined into a set-sized Cartesian product.

The second specific kind of system of ontological levels that List considers is a *system of descriptive levels*—roughly, a system of ontological levels in which each level is the set of worlds describable in a specific language. More precisely, a language $\mathbf{L}$ is defined as a set of elements—formal expressions called *sentences*—that are equipped with a negation operator $\neg : \mathbf{L} \to \mathbf{L}$ and a bifurcation of the power set $\mathscr{P}(\mathbf{L})$ into two, labeled "consistent" and "inconsistent." It is required in particular that:

---

[3] In more detail, an inverse limit in a category $\mathfrak{C}$ can be characterized as the limit of a functor from a partially ordered set, considered as a small category, to $\mathfrak{C}$, and such limits exist when $\mathfrak{C}$ has small products and equalizers [9, Theorem 2.8.1]. Posetal categories trivially always have equalizers, but their products are just greatest lower bounds.

[4] Thanks to Christian List for this suggestion. Note that the Cartesian product is the product in the category of sets, not in a posetal category—see footnote 3.

1. any set containing a sentence and its negation is inconsistent;
2. inconsistency is preserved by taking supersets;
3. $\varnothing$ is consistent; and
4. every consistent set is contained in a maximal consistent set (i.e., a consistent set containing, for every sentence $\phi \in \mathbf{L}$, either ϕ or $\neg\phi$).

The *ontology* for $\mathbf{L}$, denoted $\Omega_{\mathbf{L}}$, is defined as the set of all maximal consistent subsets of $\mathbf{L}$. Each such subset (i.e., each element of $\Omega_{\mathbf{L}}$) is a *world*. A sentence $\phi \in \mathbf{L}$ is said to be *true at* a world $w \in \Omega_{\mathbf{L}}$ if $\phi \in w$, and the *propositional content* of ϕ, denoted $[[\phi]]$, is defined as the set of all worlds at which ϕ is true.

A system of descriptive levels is thus a system of ontological levels, each of which is the ontology of some language. Within such a system, a higher-level sentence $\phi' \in \mathbf{L}'$ is defined (by List) to be *reducible* to a lower-level sentence $\phi \in \mathbf{L}$ if and only if the propositional content of ϕ is the inverse image of ϕ' under σ, the supervenience map from $\Omega_{\mathbf{L}}$ to $\Omega_{\mathbf{L}'}$. And he defines the higher level of description L' to be reducible to the lower level of description L if every sentence of the higher level's associated language $\mathbf{L}'$ is reducible to some sentence of the lower level's associated language $\mathbf{L}$.

Note that, so defined, not every system of descriptive levels will be one in which the higher levels reduce to the lower levels: as List observes, this provides a sense in which supervenience does not entail reduction. For a concrete example [13], let $\mathbf{L}'$ be the propositional language for level L' whose only sentence-letter is F, and let $\mathbf{L}$ be the propositional language for level L with sentence-letters $\{P_0, P_1, \dots\}$, each equipped with the standard notion of consistency. $\Omega_{\mathbf{L}'}$ only contains two worlds: ω_F, which contains F, and $\omega_{\neg F}$, which contains $\neg F$. Let $\omega \in \Omega_{\mathbf{L}}$ be the world containing every P_i. Define $\sigma : \Omega_{\mathbf{L}} \to \Omega_{\mathbf{L}'}$ as follows: for any $\widetilde{\omega} \in \Omega_{\mathbf{L}}$,

$$\sigma(\widetilde{\omega}) := \begin{cases} \omega_F & \text{if } \widetilde{\omega} = \omega, \\ \omega_{\neg F} & \text{otherwise.} \end{cases}$$

Then the levels $\Omega_{\mathbf{L}}$ and $\Omega_{\mathbf{L}'}$ equipped with the maps σ, $\mathrm{Id}_{\Omega_{\mathbf{L}}}$, and $\mathrm{Id}_{\Omega_{\mathbf{L}'}}$ constitute a system of descriptive levels.

Now observe that $\sigma^{-1}([[F]]) = \{\omega\}$. But $\{\omega\}$ is not a definable subset of $\Omega_{\mathbf{L}}$ since there is no sentence $\phi \in \mathbf{L}$ such that $[[\phi]] = \{\omega\}$.[5] So the sentence $F \in \mathbf{L}'$ is not reducible to any sentence in $\mathbf{L}$; thus, L' is not reducible to L. We will discuss in section 2, however, what sort of conditions we could place on supervenience that would associate it with reduction.

[5] *Proof.* Suppose for reductio that $[[\phi]] = \{\omega\}$. Since ϕ is a finite sentence, not every sentence letter can occur in it. So suppose P_i does not occur in ϕ. Then since $\omega \models \phi$, it must be the case that $\omega' \models \phi$, where ω' is just like ω save that $P_i \notin \omega'$ (as the truth value of a sentence in propositional logic is dependent only on the truth values of the sentence letters occurring in it). But then $\omega' \in [[\phi]]$, although $\omega' \neq \omega$, so we have a contradiction. $\square$

2 Supervenience and reduction

As discussed in section 1, for List, supervenience need not entail reduction, in the sense that there can be systems of levels of description in which the levels are not reducible to one another. In this section, we look at how, by imposing certain assumptions on the system of descriptive levels, we can recover this entailment.

Before we begin, one preliminary observation is in order. For List, the worlds in a system of levels of description are identified as maximally consistent sets of sentences of a formal language, where the notion of a "language" is left very abstract. In this section's analysis, we will assume that the languages we are working with are first-order languages; this will play an important role when we invoke Beth's theorem. However, the languages may be many-sorted. We will also identify the worlds with models of the language, rather than maximal consistent sets of sentences—note that maximal consistent sets of sentences correspond to equivalence classes of elementarily equivalent models. This assumption is more for convenience than anything else; we do not believe that anything of great significance hangs on it.

Now, consider two levels of description L and L' associated with languages $\mathbf{L}$ and $\mathbf{L}'$, respectively. Let σ be a supervenience map from L to L'. We assume two conditions on the relationship between these levels. First, we assume that they are *compatible* with respect to σ, in the sense that if the vocabularies of $\mathbf{L}$ and $\mathbf{L}'$ intersect, then for any $\omega \in \Omega_{\mathbf{L}}$, $\omega|_{\mathbf{L}\cap\mathbf{L}'} = \sigma(\omega)|_{\mathbf{L}\cap\mathbf{L}'}$. This has the consequence that for any $\omega \in \Omega_{\mathbf{L}}$, we can define a unique expansion to a structure of signature $\mathbf{L}\cup\mathbf{L}'$: let $\omega + \sigma(\omega)$ be the structure such that for any symbol S in the vocabulary of $\mathbf{L}\cup\mathbf{L}'$ (including sort symbols),

$$S^{\omega+\sigma(\omega)} = \begin{cases} S^{\omega} & \text{if } S \in \mathbf{L}, \\ S^{\sigma(\omega)} & \text{if } S \in \mathbf{L}'. \end{cases}$$

Let $\Omega_{\mathbf{L}} + \Omega_{\mathbf{L}'} := \{\omega + \sigma(\omega) : \omega \in \Omega_{\mathbf{L}}\}$. This can be viewed as the union of the two levels relative to the supervenience map σ. Note that when its antecedent is satisfied, this assumption entails that ω and $\sigma(\omega)$ share a domain: or, in other words, that the lower-level theory already asserts the existence of higher-level objects (but without explicitly saying what they are like). Such an assumption is reasonably plausible if (for instance) we take the lower-level theory to include some kind of mereological theory that asserts the existence of mereological sums or fusions. Alternatively, one could seek to weaken this assumption by invoking methods for defining new sorts;[6] although such a project would be interesting, we do not undertake it here.

Our second assumption is that this class is *characterizable*, in the sense that there is a set of sentences T such that $\Omega_{\mathbf{L}} + \Omega_{\mathbf{L}'} = \mathrm{Mod}(T)$. This means that the union of the two levels comprises worlds that can be characterized as those in which certain sentences in the union of their languages are true. This assumption is reasonably strong: in the context of psychology and physics, for example, it amounts to the assertion that there is some joint psycho-physical theory such that a distribution of

[6] See, in particular, [2] and [6].

psychological and physical facts is possible if and only if it is in accord with the joint theory. However, it is still nontrivially distinct from directly assuming reducibility. One way to obtain such a joint theory would be to conjoin our psychological and physical theories with a set of bridge laws connecting them; but the assumption of characterizability does not presume that the joint theory takes this "pre-reduced" form. Note further that, in general, assuming characterizability will mean that the antecedent of the first assumption is satisfied: that is, that there is at least one sort of object that both levels describe (and so, per the first assumption, about which they agree). This is plausible for realistic cases, in which we expect the two levels of description to share at least some vocabulary (e.g., empirical or observational terms).

Given these assumptions, it follows that the vocabulary of the higher level $\mathbf{L}'$ is *implicitly defined* by T in terms of the lower-level vocabulary $\mathbf{L}$: that is,

Proposition 1. *For any models $\widetilde{\omega}_1$ and $\widetilde{\omega}_2$ of T, if $\widetilde{\omega}_1|_\mathbf{L} = \widetilde{\omega}_2|_\mathbf{L}$ then $\widetilde{\omega}_1 = \widetilde{\omega}_2$.*

Proof. Suppose that $\widetilde{\omega}_1|_\mathbf{L} = \widetilde{\omega}_2|_\mathbf{L}$. By the assumption of characterizability, there are $\omega_1, \omega_2 \in \Omega_\mathbf{L}$ such that $\widetilde{\omega}_1 = \omega_1 + \sigma(\omega_1)$ and $\widetilde{\omega}_2 = \omega_2 + \sigma(\omega_2)$. It follows that $\widetilde{\omega}_i|_\mathbf{L} = \omega_i$ and $\widetilde{\omega}_i|_{\mathbf{L}'} = \sigma(\omega_i)$ (for $i = 1, 2$). Hence,

$$\widetilde{\omega}_2|_{\mathbf{L}'} = \sigma(\omega_2) = \sigma(\widetilde{\omega}_2|_\mathbf{L}) = \sigma(\widetilde{\omega}_1|_\mathbf{L}) = \sigma(\omega_1) = \widetilde{\omega}_1|_{\mathbf{L}'}.$$

Thus, $\widetilde{\omega}_1 = \widetilde{\omega}_2$. $\square$

Next, by Beth's theorem, it follows that L' is *explicitly defined* by T in terms of $\mathbf{L}$.[7] In the case of single-sorted languages, this means that for any (n-place) relation symbol R of $\mathbf{L}'$, there is an $\mathbf{L}$-formula τ_R such that $T \vdash \forall x_1 \ldots \forall x_n (Rx_1 \ldots x_n \leftrightarrow \tau_R(x_1, \ldots, x_n))$ and similarly for other kinds of symbols. The case of many-sorted languages is a bit more complex.[8] But it follows in either case that for every $\mathbf{L}'$-sentence ϕ', there is an $\mathbf{L}$-sentence ϕ such that $T \vdash (\phi' \leftrightarrow \phi)$. So ϕ' reduces to ϕ. Hence, it follows that L' reduces to L.

Thus, although supervenience maps between different levels (in general) are not associated with reduction, any supervenience map which fulfills our assumptions of compatibility and characterizability will be so associated.

An alternative way of establishing a relationship between supervenience and reduction proceeds not by imposing constraints on the supervenience map between levels, but rather by treating supervenience as a feature of sets of properties rather than worlds. Note that this is more in line with how supervenience is often defined in the literature: for example, as the Stanford Encyclopedia of Philosophy's entry on Supervenience [26] begins, "A set of properties A supervenes upon another set B just in case no two things can differ with respect to A-properties without also differing with respect to their B-properties." In other words, the A-properties supervene upon the B-properties if any two possible worlds which have the same distribution of B-properties also have the same distribution of A-properties.

[7] For discussion of Beth's theorem, see [17].

[8] For a generalisation of Beth's theorem to many-sorted logics, see [2].

To make things more precise, let us focus on the notion of *strong global supervenience* [26, §4.3.2]: the A-properties strongly globally supervene upon the B-properties if and only if for any worlds w_1 and w_2, any B-preserving isomorphism between w_1 and w_2 is an A-preserving isomorphism between them.[9] Let us suppose that the A-properties are those expressed by a higher-level language $\mathbf{L}'$, and the B-properties are those expressed by a lower-level language $\mathbf{L}$. And suppose that our worlds (structures of signature $\mathbf{L} \cup \mathbf{L}'$) are characterizable, where this means the same thing as before: there is a theory T, such that the worlds are the models of the theory. It then follows that

Proposition 2. *The A-properties strongly globally supervene on the B-properties if and only if T implicitly defines $\mathbf{L}'$ in terms of $\mathbf{L}$.*

Proof. From left to right, suppose that the A-properties supervene upon the B-properties: that is, that for any worlds (models of T) ω_1 and ω_2, if f is an isomorphism from $\omega_1|_{\mathbf{L}}$ to $\omega_2|_{\mathbf{L}}$, then f is an isomorphism from ω_1 to ω_2. Now suppose further that we have two models ω_1 and ω_2 such that $\omega_1|_{\mathbf{L}} = \omega_2|_{\mathbf{L}}$. Clearly, then, the identity is an isomorphism between $\omega_1|_{\mathbf{L}}$ and $\omega_2|_{\mathbf{L}}$, and hence must be an isomorphism between ω_1 and ω_2; it follows that $\omega_1 = \omega_2$. So T implicitly defines $\mathbf{L}'$ in terms of $\mathbf{L}$.

From right to left, suppose that T implicitly defines $\mathbf{L}'$ in terms of $\mathbf{L}$, and let ω_1 and ω_2 be models of T such that f is an isomorphism from $\omega_1|_{\mathbf{L}}$ to $\omega_2|_{\mathbf{L}}$. For *reductio*, suppose that f is not an isomorphism from ω_1 to ω_2, so for some $R \in \mathbf{L}'$, $f[R^{\omega_1}] \neq R^{\omega_2}$. But now define ω_1' as follows: it has the same domain as ω_2, and for every $P \in \mathbf{L} \cup \mathbf{L}'$, $P^{\omega_1'} = f[P^{\omega_1}]$. By construction, f is an isomorphism from ω_1 to ω_1'.[10] So ω_1' is a model of T. And since f is an isomorphism from $\omega_1|_{\mathbf{L}}$ to $\omega_2|_{\mathbf{L}}$, it follows that $\omega_1'|_{\mathbf{L}} = \omega_2|_{\mathbf{L}}$, and so that $\omega_1' = \omega_2$. But then it follows that $R^{\omega_1'} = f[R^{\omega_1}] = R^{\omega_2}$, so we have obtained a contradiction. Hence, f must be an isomorphism from ω_1 to ω_2. □

From here, the analysis goes as before (i.e. via Beth's theorem): hence, if the higher-level properties supervene upon the lower-level properties, then (at least for one natural way of formalizing what this means) it will be accompanied by a reduction from the higher level to the lower level.

3 Extensions: partial and non-surjective maps

The first extension to List's framework we consider will be to *partial* supervenience maps. To motivate it, consider the case of a system of ontological levels as described in section 1. These levels consist of possible worlds, each of which provides a full

[9] As McLaughlin and Bennett [26] discuss, there are other notions of global supervenience one can define. For a compelling case that strong global supervenience is the most appropriate precisification of the intuitive notion of global supervenience, see Shagrir [32].

[10] This is an instance of what Button and Walsh [11] call the "push-through" construction.

specification of the facts particular to that level. List takes the maps between levels, representing supervenience relations, to be *functions*, explaining that

> Supervenience means that each lower-level world determines 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. [24, p. 7]

However, the glosses before and after the colon are not equivalent: that the chemical facts, say, may supervene on the physical facts just means that there cannot be any difference between two chemical worlds without a difference between the physical worlds in the chemical worlds' supervenience bases—the preimages of the supervenience map on the two chemical worlds. But this is entirely compatible with there being some physical worlds that do not determine *any* non-trivial chemical worlds—worlds for which the empty set is not a logical model of any of their descriptions. Indeed, we expect there not to be any non-trivial chemical facts at all determined by those of a roiling quantum vacuum, or any non-trivial biological facts at all determined by the chemical facts of the atmosphere of Venus. Insofar as each of these is respectively a way a physical and chemical world could be, there can be physical worlds on which no non-trivial chemical world supervenes, chemical worlds on which no non-trivial biological world supervenes, etc.

List [24, p. 9] already cites approvingly a similar remark by Kim[11] regarding the difficulties of entity-based (rather than world-based) conceptions of levels, so this small extension seems welcome. Indeed, one can easily integrate into the definition of a system of ontological levels its consequence that the supervenience maps need only be partial: their domains of definition needn't be *all* the worlds at a level. (This doesn't by itself conflict with any of the formal properties defining a system of levels as a posetal category, but one must decide exactly how to model a category with partial morphisms: see, e.g., Cocketta and and Lackb [12, §1] for references to a number of options.) Since a system of levels of description is just a system of ontological levels with added structure, the same conclusion applies for them. The procedure for embedding, via a functor, a system of ontological levels into a system of levels of grain also remains the same; the image under the functor of any partial supervenience map will have some equivalence class of worlds in its lower-level domain that will just not have any image in the higher-level codomain within the system of levels of grain.

One way to regain totality for the supervenience maps would be to introduce a *null world* at each level, a world devoid of non-trivial facts. Then one can extend each previously partial supervenience function by mapping all the elements previously outside its domain of definition to the null world in its codomain. For consistency, the image of any null world under a supervenience map would then have to be a null world. This is coherent with the interpretation of the maps: a world devoid of non-trivial physical facts determines a world devoid of non-trivial chemical facts,

[11] Namely, that "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" [20, p. 11].

and there can be no difference in the latter without a difference in the former—i.e., without the introduction of some non-trivial physical facts or other. However, such null worlds should be interpreted cautiously: if one wants the worlds at a level to be models of a theory describing that level, the null world will not be included if the theory has any existentially quantified axioms. So it probably should receive a fixed interpretation as a technical convenience rather than as a genuine possible world.

This second extension is to *non-surjective* supervenience maps. Now, List claims surjectivity just follows from the meaning of supervenience:[12]

> To say that the chemical level supervenes on the physical, or that the biological supervenes on the chemical, is to say that the class $\mathscr{S}$ of supervenience mappings contains one such mapping, $\sigma : \Omega \to \Omega'$, from the relevant lower level to the high one, where σ maps *each* lower-level world $\omega \in \Omega$ to the higher-level world $\omega' \in \Omega'$. We then call ω a *lower-level realizer* of ω'. The surjectivity of σ means that there are no possible worlds at the higher level that lack a possible lower-level realizer. [24, p. 8]

So it does; but why should the *formal apparatus* of a system of levels necessarily commit to the metaphysical thesis that the *whole* higher level does so supervene? Even though we take the supervenience of levels to be plausible, we do not think that *characterizing* systems of levels require assuming it. Indeed, dropping this assumption, one finds the *interpretation* of the levels, the worlds comprising them, and the supervenience maps between them to be hardly different.

The key insight for this is that a supervenience map relates collections of facts at a lower level to facts at a higher level, hence need not relate all worlds at either level at all. If a supervenience map is not surjective, then there will be at least some higher level world that does not supervene on any lower level world in the map's domain. But still the map characterizes exactly which higher-level properties (and worlds) supervene on lower-level ones. That it is not surjective is just to indicate precisely which of these higher-level properties do not so supervene. The map represents a higher *level* supervening on a lower *level* if and only if it is surjective. This is important for applications in philosophy of science: When the worlds of levels are models of empirical theories, the failure of surjectivity indicates that the possibilities that a higher-level theory describes outstrip those that one committed to level-supervenience would expect. This can motivate revising the higher-level theory by restricting what it allows as possible, or generalizing the lower-level theory to allow more possibilities that could provide new supervenience bases for higher-level possibilities.

Furthermore, examples of non-surjective supervenience maps can be found in systems of different geometric levels, each of which is able to describe a wide range of geometric possibilities.[13] Let us focus on the metric and the topological levels. There is a robust sense in which topological properties (such as being an open set or being a continuous function) supervene on metric properties (distances between points). Every metric naturally induces a topology and isometric metric spaces in-

[12] For emphasis we have italicized the word "each" in the passage.

[13] For a philosophical investigation of the notion of geometric possibility and its role in the relationism-substantivalism debate, see Belot [7].

duce homeomorphic topological spaces. So in the jargon of possible worlds, sameness of worlds on the subvening metric level entails sameness of worlds on the supervening topological level. However, the supervenience map from the metric to the topological level (which sends each metric space to the topological space induced by it) is not surjective because not all topological spaces are metrizable. The topological level of description is essentially more general than the metric level, but nonetheless there is a supervenience map between them. Dropping the constraint of surjectivity enables us to capture such cases of supervenience. Thus, the framework gains expressive power without losing any advantages.

One further consequence of dropping surjectivity for the supervenience maps concerns List's argument for the possibility of mutually supervening but non-identical levels in a system of ontological levels. Suppose that in such a system $\sigma : \Omega \to \Omega'$ and $\sigma' : \Omega' \to \Omega$ are supervenience maps. Then by the definition of a category, the maps are closed under associative composition, so $\sigma \circ \sigma' : \Omega \to \Omega$ is a supervenience map. But because the category is posetal, there can be at most map with domain and codomain Ω, namely the identity 1_Ω. Hence $\sigma \circ \sigma' = 1_\Omega$. It follows that σ and σ' must be total *and* surjective.[14] Moreover, if σ and σ' are total *and* surjective then $\sigma \circ \sigma' = 1_\Omega$. This is as expected: two ontological levels are isomorphic in a system of levels if and only if they supervene on the other as a whole. So if a supervenience map is ever identified as non-surjective (or partial), the levels it relates cannot be isomorphic.

Like with the first extension to partial maps, these insights about systems of ontological levels apply equally to systems of levels of description, although the application to systems of levels of grain requires a bit more work—it also requires adopting the first extension. To do this, one must consider the partition at each level to include one special equivalence class, the class of worlds *not included at that level*. Then, if two levels of grain are related by a supervenience map, that map must exclude the equivalence class of worlds not included at the level of its domain within its domain of definition. (Of course, other equivalence classes could fall outside its domain of definition, too.) Moreover, instead of requiring the domain of the map to be at least as fine-grained as the codomain, one only requires that the domain of definition be at least as fine-grained as the image of the map. Outside the image, the elements of the codomain need bear no relationship of refinement to those of the domain.

4 Levels as categories

Another way of extending List's framework is to give a richer account of the internal structure of levels of description. In particular, it can be fruitful to view levels of description as categories of structures.[15] As we will see below, this account is more

[14] If σ were not total, then 1_Ω could not be; if σ' were not surjective, then 1_Ω could not be. The same reasoning applies mutatis mutandis to $\sigma' \circ \sigma = 1_{\Omega'}$. Hence, each of σ and σ' is both total and surjective.

[15] So on this account, systems of levels of description can be viewed as categories of categories.

suitable to capture important scientific levels of description and it establishes a close connection between List's work and current developments in philosophy of science.

More precisely, our proposal is to represent a level of description, L, as a pair $\langle \mathbf{L}, \Omega \rangle$ consisting of a description language, $\mathbf{L}$, and a category, Ω, of $\mathbf{L}$-structures.[16] This account differs from List's own account of levels of description in two ways.

1. The objects in Ω are $\mathbf{L}$-structures rather than maximally consistent sets of $\mathbf{L}$-sentences. This is not a big difference. But, among other things, it allows one to apply model-theoretic notions (such as homomorphism, embedding, isomorphism, etc.) without further ado.
2. Ω is a category rather than a bare class. So to specify a level of description, one does not only specify its structures but also the morphisms (admissible transformations) between these structures. Which morphisms to choose is in general not determined by the language $\mathbf{L}$ alone. The choice of morphisms in Ω is linked to an interpretive choice concerning the description language. It reflects which expressions of $\mathbf{L}$ are taken to be meaningful within the level L. The idea is that an $\mathbf{L}$-expression is only meaningful within L if its extension is invariant under the morphisms in Ω.

A major advantage of this account is that it enables us to deal with levels of description on which there is a non-trivial distinction between meaningful and non-meaningful expressions. Levels of description of this kind are prevalent both in the natural and social sciences because scientific languages often contain auxiliary vocabulary or use numerical descriptions that may be transformed according to certain rules without changing in content.

Let us consider an example from social choice theory: descriptions of individual welfare along the lines of Sen [31]. In contrast to Arrow's ordinal framework [3], Sen's framework rests on a numerical description of the welfare of individuals under given alternatives. We reconstruct a level of description of welfare $L_{\text{wf}} = \langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}} \rangle$ along these lines to illustrate the proposed account.

The language $\mathbf{L}_{\text{wf}}$ comprises the following descriptive symbols: (1) a sort symbol I for individuals, (2) a sort symbol A for alternatives, and (3) a function symbol W of type $I \times A \to \mathbb{R}$. If i and a are terms of the sorts I and A, respectively, then $W_i(a)$ is a term of sort $\mathbb{R}$. It stands for the degree of i's welfare under alternative a. Of course, the language also has mathematical auxiliary symbols such as $\mathbb{R}, +, \cdot, <$, etc. Moreover, we assume that the language has the usual logical vocabulary (connectives, quantifiers, and variables) and formation rules for formulas.

An object in Ω_{wf} is an $\mathbf{L}_{\text{wf}}$-structure that consists essentially of a finite set of individuals, a finite set of alternatives, and a welfare profile for these individuals and alternatives. To put it more rigorously, the objects in Ω_{wf} are those $\mathbf{L}_{\text{wf}}$-structures, ω, such that (a) ω expands the standard model of the auxiliary mathematical part of the vocabulary (including the field of real numbers), and (b) I^{ω} and A^{ω}—the sets associated to I and A in the structure ω—are both finite.

---

[16] We use the term "$\mathbf{L}$-structure" in its usual model-theoretic sense. So an $\mathbf{L}$-structure is an assignment of extensions to the descriptive symbols of $\mathbf{L}$.

Which morphisms to include in Ω_{wf} depends on which expressions of $\mathbf{L}_{\mathrm{wf}}$ one considers to be meaningful. List [23] gives a concise overview of various positions, each of which depends on the extent to which interpersonal comparisons of welfare are considered as meaningful and on which level of measurement (ordinal, interval, or ratio) welfare data are taken to reside. For instance, many economists (including Arrow) consider only ordinal *intra*personal comparisons of welfare as meaningful.[17] According to this view, the formula "$W_i(a) < W_i(a)$" is meaningful. In contrast, the formula "$W_i(a) < W_j(a) \wedge i \neq j$" is not meaningful because it amounts to an *inter*personal comparison of welfare. And "$W_i(a) = 2 \cdot W_i(b)$" is also not meaningful because it goes beyond a comparison of order. On the other end of the spectrum, one finds the view that full interpersonal comparisons are meaningful and welfare can be measured on a ratio scale. According to this view, both "$W_i(a) < W_j(a) \wedge i \neq j$" and "$W_i(a) = 2 \cdot W_i(b)$" are meaningful.

Much depends on which descriptions are taken as meaningful. A level of numerical welfare descriptions where only ordinal intrapersonal comparisons count as meaningful boils down to a purely ordinal level of description of welfare along the lines of Arrow [3]. The meaningful numerical statements of the form "$W_i(a) < W_i(b)$" can be reduced to statements about individual preference orderings of the form "individual i prefers alternative a to alternative b" (or, in short: "$a \prec_i b$"). In contrast, a level of numerical welfare descriptions where interpersonal ratio comparisons of welfare count as meaningful is an essentially more fine-grained level of description. Thus, even if levels of description have the same underlying language and structures, they are not necessarily the same. One also needs to specify what counts as meaningful to identify a level.

As alluded above, different standards of meaningfulness can be captured in terms of choices of morphisms. To discuss such choices in our present example, the notion of a transformation of structures is useful. Given structures $\omega_1, \omega_2 \in \Omega_{\mathrm{wf}}$ with the same set of individuals $\mathrm{Ind} := I^{\omega_1} = I^{\omega_2}$ and the same set of alternatives $\mathrm{Alt} := A^{\omega_1} = A^{\omega_2}$, we call a family $\{f_i\}_{i \in \mathrm{Ind}}$ of functions from $\mathbb{R}$ to $\mathbb{R}$ a *transformation of* ω_1 *into* ω_2 just in case for all $i \in \mathrm{Ind}$,

$$f_i \circ W_i^{\omega_1} = W_i^{\omega_2}.$$

Now the question of which welfare comparisons one considers to be meaningful boils down to the question of which transformations one considers to be admissible in the sense that they do not change any meaningful features of welfare profile structures. The admissible transformation are then taken as morphisms in the category.

To capture the idea that only ordinal intrapersonal comparisons of welfare are meaningful, one takes all positive monotonic transformations as morphisms (cf. Sen [31, Chapter 7*]). More precisely, one takes $\{f_i\}_{i \in \mathrm{Ind}}$ as a morphism from ω_1 to ω_2 just in case it is a transformation of ω_1 into ω_2 such that, for all $i \in \mathrm{Ind}$, f_i is a positive monotonic function. To capture the ideas that full interpersonal and ratio comparisons are meaningful, one takes $\{f_i\}_{i \in \mathrm{Ind}}$ as a morphism from ω_1 to ω_2 just

[17] For a discussion of the thesis that interpersonal comparisons of utility are meaningless and the inference that such comparisons are therefore impossible, see List [22].

in case it is a transformation of ω_1 into ω_2 such that for some $\alpha > 0$, $f_i(x) = \alpha \cdot x$ for all $x \in \mathbb{R}$ and all $i \in \text{Ind}$. Let $\Omega_{\text{wf}}^{\text{ONC}}$ be the category with the former choice of morphisms and $\Omega_{\text{wf}}^{\text{RFC}}$ the category with the latter.[18] Then, according to our proposal, $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{ONC}} \rangle$ and $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{RFC}} \rangle$ are different levels of description. This illustrates how taking morphisms into account allows us to capture differences between levels of description even if they have the same underlying language and structures.

Let us now turn to supervenience in the extended framework. Since levels of description are treated as categories of structures, supervenience relations between them are best viewed as *functors*. Note that any functor maps isomorphic objects to isomorphic objects.[19] This coheres with the idea of global supervenience, that sameness of worlds on the subvening level implies sameness of worlds on the supervening level.

But more importantly, viewing supervenience relations as functors allows us to shed more light on List's requirement of surjectivity. First, it is important to distinguish between surjectivity and essential surjectivity. A functor $F : \Omega \to \widetilde{\Omega}$ is *essentially surjective* if and only if every object $\widetilde{\omega}$ in $\widetilde{\Omega}$ is isomorphic to $F(\omega)$ for some object ω in Ω. From a category-theoretic point of view, essential surjectivity is a fruitful notion. In contrast, surjectivity simpliciter is much too strict. If a level of description is such that every structure has several representationally equivalent variants, there is no good reason to require that a supervenience map from a lower level to this level be surjective. Requiring this would preclude cases of supervenience where a level on which every structure has several representationally equivalent variants supervenes on a level without (an at least equal number of) these variants. A given structure ω at the subvening level would correspond to many representationally equivalent structures at the supervening level. However, a supervenience map can send ω only to one of those. So the others cannot be in the image of the map. Ruling out such cases would be a serious limitation. It would, for instance, make it impossible to capture cases in which non-quantitative descriptions determine corresponding quantitative descriptions up to certain transformations. But many important cases of supervenience between levels of description are of this type, especially those given by theorems in measurement theory that exhibit the existence of numerical representations and their uniqueness up to certain transformations [21, 33, 34].

This suggests that surjectivity has to be given up regardless of the arguments given in section 3. So even if one would like to adhere to the idea that every world on the supervening level must have a corresponding world on the subvening level, one should better explicate this idea in terms of essential surjectivity rather than surjectivity simpliciter. In view of that, it becomes clear that the arguments in section 3 are in fact arguments in favor of dropping even essential surjectivity.

---

[18] "ONC' stands for "ordinal measurability with no interpersonal comparability" and "RFC" is short for "ratio-scale measurability with full interpersonal comparability." These acronyms are due to List [23].

[19] If F is a functor from the category Ω to the category Ω' and ω_1 is isomorphic to ω_2 in Ω, then $F(\omega_1)$ is isomorphic to $F(\omega_2)$ in Ω'.

And it makes sense to go even further. There is no principled reason why a system of levels should only incorporate supervenience functors and no other functors between levels. From the perspective of the proposed extension of List's account, it makes sense to include *all functors* between levels of description as arrows in the category of system of levels. Then, rather than excluding and thereby neglecting many functors from the start, one may explore and classify the entire zoo of functors between levels. For example, one may ask which functors should count as reduction functors and which should be seen merely as supervenience functors. Or by dropping uniqueness of arrows between levels, one can investigate cases of multiple reducibility: that is, different ways of reducing a higher level to a lower level (analogous to different ways of reducing arithmetic to set theory).[20]

Here is an example to illustrate how this could look like. Political scientists describe individual preferences on different levels. On one level, one can describe each individual's preference scores for the alternatives in question. On another level, one can describe which alternatives each individual approves of. The structures of the former level are all logically possible preference profile scores for a set of alternatives and a set of individuals. The structures of the latter level are all possible approval profiles for such sets. Now, one can reduce the level of approval descriptions to the level of preference score descriptions by setting a threshold t such that all alternatives with a preference score above t are taken as approved of. Thus every preference score profile gives rise to an associated approval profile. Relative to this supervenience map, statements of the form "individual i approves of alternative a" reduce to statements of the form "individual i's preference score of alternative a is above t". The crucial point is that there are different choices of thresholds that succeed equally well at reconstructing all approval profiles from preference score profiles. So there are multiple reductions of the level of approval descriptions to the level of preference score descriptions. However, a given approval profile will in general correspond to different preference score profiles under different reductions. So there are substantial differences between such reductions. But it is not clear a priori which reductions are "better" than others and in which respects. In such cases, it makes sense to investigate a variety of possible reductions within a system of levels of description.

To sum up, we propose to generalize List's framework in two ways: (1) we construe levels of description as categories rather than bare classes, and supervenience relations as functors; (2) we allow all sorts of functors to be included in a system of levels of description. An advantage of this radical generalization is that it makes the rich toolbox of category-theoretic concepts available for the analysis of levels of description and their relations. For example, the concepts of natural transformation,

[20] One might seek to model multiple realizability in this way: for example, perhaps one successful reduction translates "pain" by "firing of C-fibres" whilst another translates "pain" by "firing of D-fibres" (where, let us suppose, human brains have C-fibres and Martian brains have D-fibres). However, it is not clear to us what the prospects for this manoeuvre might be; note that neither translation will map the true higher-level claim "both humans and Martians experience pain" to a true lower-level claim.

equivalence, duality, adjunction, and forgetful functor lend themselves well to this endeavour.

This point also plays an essential role in a new strand of research in philosophy of science which uses these category-theoretic concepts to study relations between physical theories (where scientific theories are represented as categories of structures).[21] The proposed generalization of List's framework establishes a close connection to this new work in philosophy of science. Many results from this literature can be viewed as results about how certain types of physical descriptions are related to each other and, thus, how our current system of levels of physical descriptions is structured. But the fruitfulness of the generalized framework is not limited to physics or the natural sciences. As our examples above indicate, the same methods may be applied to levels of description that belong to the social sciences (e.g., economics or political science). We believe that extending and generalizing List's framework makes it better applicable to and, thus, more relevant for philosophy of science in general.

Let us illustrate how the above-mentioned category-theoretic concepts can be used to analyze inter-level relations. As pointed out above, welfare can be described on a quantitative level (in terms of real-valued welfare functions) and on a qualitative level (in terms of mere preference orderings). Let $\langle \mathbf{L}_{\text{ord}}, \Omega_{\text{ord}} \rangle$ be the ordinal level of description. Each object of Ω_{ord} is a structure consisting of a finite set of individuals, a finite set of alternatives and an assignment of a weak ordering over the alternatives to each individual. To reflect that the identities of alternatives and individuals are in general significant, Ω_{ord} contains only identity morphisms. Suppose the ordinal level $\langle \mathbf{L}_{\text{ord}}, \Omega_{\text{ord}} \rangle$ supervenes on the quantitative level $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{RFC}} \rangle$ via a supervenience functor F that maps every welfare profile ω in $\Omega_{\text{wf}}^{\text{RFC}}$ to the induced profile of preference orderings ω' in Ω_{ord} and all morphisms between objects ω_1 and ω_2 in $\Omega_{\text{wf}}^{\text{RFC}}$ to the identity morphism on $F(\omega_1) = F(\omega_2)$ in Ω_{ord}. Then there is a sense in which qualitative descriptions neglect or "forget" some information available on the quantitative level. This informal idea is captured formally by the fact that the supervenience functor F is what is called a *forgetful functor* in category theory. In category theory, there is also a way of making precise *what* a functor forgets in terms of its formal properties.[22] In this case, the supervenience functor is essentially surjective and also full (i.e., surjective on morphisms). But it is neither faithful (injective on morphisms) nor essentially injective on objects. In technical terms, the functor forgets "stuff"—intuitively speaking: the set of possible degrees of welfare.

But note that this analysis is more subtle than the obvious point that $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{RFC}} \rangle$ uses numbers while $\langle \mathbf{L}_{\text{ord}}, \Omega_{\text{ord}} \rangle$ does not. That $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{RFC}} \rangle$ uses numbers is not the point here. The level $\langle \mathbf{L}_{\text{wf}}, \Omega_{\text{wf}}^{\text{ONC}} \rangle$ of numerical welfare descriptions where only ordinal intrapersonal comparisons are taken as meaningful also uses numbers but it should count as equivalent to the purely ordinal level $\langle \mathbf{L}_{\text{ord}}, \Omega_{\text{ord}} \rangle$. And, indeed,

[21] See for example Halvorson and Tsementzis [18] and Hudetz [19].

[22] Baez, Bartels, and Dolan [4] have developed a classification of functors with respect to whether they forget structure, properties, or stuff in terms of whether they are full, faithful, or essentially surjective. For a nice overview and further applications in physics, see Weatherall [36].

their equivalence can be captured in category-theoretic terms using the notion of an equivalence of categories. An *equivalence between categories* C and D is a pair of functors $F : C \rightleftarrows D : G$ that are essentially inverse to each other.[23] That there is an equivalence of categories between $\langle \mathbf{L}_{\mathrm{wf}}, \Omega_{\mathrm{wf}}^{\mathrm{ONC}} \rangle$ and $\langle \mathbf{L}_{\mathrm{ord}}, \Omega_{\mathrm{ord}} \rangle$ can be demonstrated easily by invoking the representation theorem for weak preference orderings. Examples of category-theoretic equivalences abound in the philosophy of physics literature. See, for example, Barrett's work [5] on Hamiltonian and Lagrangian descriptions of classical mechanical systems or Weatherall's work [35, 37] on classical field theories.

Another relation which is naturally captured in category-theoretic terms is that of duality between levels of description. Categories C and D are *dual* to each other if and only if there is an equivalence between C and D^{op}, where D^{op} is the opposite of the category D and it is given by reversing all the morphisms in D. One often find dualities between algebraic and topological/geometric levels of description. For example, Rosenstock, Barrett, and Weatherall [29] show that the usual manifold formulation of general relativity is dual to the algebraic formulation in terms of Einstein algebras in the sense that these theories have dual categories of structures.

Another important category-theoretic concept is that of adjunction. Although this notion is hard to grasp on a pre-theoretic level, it may still be fruitful for analyzing relations between levels of description. In our context, adjunction can be roughly understood as a special type of supervenience where the supervenience functor is accompanied by a second functor in the other direction: its adjoint. As Feintzeig [15] has shown, the concept of adjunction captures the relationship between the level of infinite (limiting) quantum systems and the level of finite quantum systems in quantum statistical mechanics. The basic idea is that the properties of infinite quantum systems are determined in a specific way by the properties of their finite subsystems. This illustrates that the generalized framework allows one to transfer technical concepts to new domains of application where they can be used to capture relations between levels of description for which we would otherwise lack appropriate concepts.

5 Conclusions and Prospects

In section 1, we reviewed List's levels formalism [24], focusing on his account of systems of ontological levels, in particular systems of levels of grain and systems of descriptive levels. One of the results of this review was to qualify his claim that any system of ontological levels can be mapped functorially into some system of levels grain: this is true only for systems of ontological levels that are downward-directed, i.e., for which each pair of levels has a common subvienient level that supervenes on all such common levels. In section 2, we examined the relation between supervenience and reduction in this formalism. We showed that while in general su-

[23] This means that their composition FG is naturally isomorphic to the identity functor on C and GF is naturally isomorphic to the identity functor on D.

pervenience does not imply reduction, it does so imply it when the levels related by the supervenience map are compatible and jointly characterizable. Compatibility requires, roughly speaking, that levels of description related by a supervenience map agree with each other as far as shared vocabulary is concerned. Joint characterizability relative to a supervenience map is a strong condition. It holds when the union of two levels of description relative to a supervenience map admits of a description itself. But in many cases of supervenience between scientific levels of description, this can be expected. So it is quite plausible that in many cases of interest, supervenience and reduction of levels go hand in hand.

After this analysis, in section 3 we proposed two extensions of systems of ontological levels by weakening List's characterization of their supervenience maps as surjective functions. First, we proposed considering merely partial (instead of total) functions, motivated by the idea that not every lower-level world ought to give rise to some higher-level world. The totality of the supervenience maps could be recovered by introducing null worlds at each level, but in many cases such worlds would have to be interpreted with caution, perhaps as just mathematical conveniences. Second, we proposed dropping the requirement of surjectivity. While this requirement is plausible for ontological naturalists, we do not think that it should be encoded into the definition of a system of levels, which, as a formal tool, ought to be propounded as neutrally as is feasible regarding substantive philosophical positions. Moreover, it also fits better with philosophy of science applications, as discussed in section 4, in which it is more fruitful to consider a level not as a bare class (of possible worlds), but as a category itself whose objects are structures of some language. This allows to distinguish levels of description even when they have the same language and the same structures. We illustrated this with an example of two social choice theories whose models were identical but whose isomorphism classes were distinct. Moreover, since one may have isomorphic but distinct models in a level, one might only require that the image of lower-level models under a supervenience map intersect with each isomorphism class of models of the higher level (i.e., the level that is the codomain of the map). These investigations also follow up on a comment made briefly in 1, that List's framework can be well-characterized without the use of category theory, which just provides a compact description and interpretive gloss for it. By contrast, taking levels as categories themselves demands a more robust use of categorial ideas that could also prove to be more fruitful.

It also suggests accordingly several directions for future research, of which we mention two. First, more examples should be formalized within the generalized framework to test its fruitfulness, flexibility, and power. Obvious candidates include levels described by:

- physical theories, such as thermodynamics and statistical mechanics [14], and quantum and classical mechanics [8];
- biological theories, such as classical and molecular genetics, levels of selection, and developmental biology and genetics [10]; and
- theories in political science (and other social sciences) that describe individuals, on the one hand, and groups, on the other [25].

Second, as mentioned in section 4, one can generalize the framework of posetal categories to include many types of functorial relationships between levels besides supervenience. List already mentions reduction as one sort that can be defined for systems of levels of description, but the sort of reduction he has in mind (as alluded in section 2) follows one of the oldest versions of the Nagelian model for reduction [27], understood as deductibility allowing for definitional extension and bridge laws induced through the supervenience maps. Already in 1967 Schaffner [30] (foreshadowed by Nagel [28] himself) suggested that reduction needs to accommodate the way different theories (here describing different levels) are related by approximations. But supervenience maps seem entirely inapt to capture these. Perhaps there is some more expansive functorial relationship that can capture these notions of reduction that seem more central to science, such as those suggested by recent work on topological (and topologically inspired) structures on models of theories [16].

Acknowledgements ND is primarily responsible for sections 1 and 2. SCF is primarily responsible for sections 3 and 5 and for general editing, and secondarily responsible for section 1. LH is primarily responsible for section 4 and for Proposition 1, and secondarily responsible for section 3 and general editing. All authors thank Katie Robertson for many insightful conversations leading to the genesis of this essay, Tomasz Brengos and Christian List for encouraging comments on a previous version, and the audience and organizers of the workshop "New Perspectives on Inter-Theory Reduction" in Salzburg in November, 2017. SCF acknowledges partial support through a Marie Curie International Incoming Fellowship (PIIF-GA-2013-628533).

References

1. Adámek, J., Herrlich, H., Strecker, G.E.: Abstract and Concrete Categories: The Joy of Cats, online edn. (2004). URL http://katmat.math.uni-bremen.de/acc
2. Andréka, H., Madarász, J.X., Németi, I.: Defining new universes in many-sorted logic (2008). URL https://old.renyi.hu/pub/algebraic-logic/kurzus10/amn-defi.pdf. Unpublished Manuscript
3. Arrow, K.: Social Choice and Individual Values. Wiley, New York (1951)
4. Baez, J., Bartel, T., Dolan, J.: Property, structure, and stuff (2004). URL http://math.ucr.edu/home/baez/qg-spring2004/discussion.html
5. Barrett, T.W.: Equivalent and inequivalent formulations of classical mechanics. The British Journal for the Philosophy of Science **forthcoming** (2018). URL https://doi.org/10.1093/bjps/axy017
6. Barrett, T.W., Halvorson, H.: Morita Equivalence. The Review of Symbolic Logic **9**(3), 556–582 (2016)
7. Belot, G.: Geometric Possibility. Oxford University Press, Oxford and New York (2011)
8. Bokulich, A.: Reexamining the Quantum-Classical Relation. Cambridge University Press, Cambridge (2008)
9. Borceux, F.: Handbook of Categorical Algebra, vol. 1: Basic Category Theory. Cambridge University Press, Cambridge (1994)
10. Brigandt, I., Love, A.: Reductionism in biology. In: E.N. Zalta (ed.) The Stanford Encyclopedia of Philosophy, Spring 2017 edn. Metaphysics Research Lab, Stanford University (2017)
11. Button, T., Walsh, S.: Structure and categoricity: Determinacy of reference and truth value in the philosophy of mathematics. Philosophia Mathematica **24**(3), 283–307 (2016)
12. Cocketta, J.R.B., Lackb, S.: Restriction categories I: categories of partial maps. Theoretical Computer Science **270**(1–2), 223–259 (2002)

13. Dewar, N.: Supervenience, reduction, and translation. Philosophy of Science Forthcoming.
14. Emch, G., Liu, C.: The Logic of Thermostatistical Physics. Springer, Berlin (2002)
15. Feintzeig, B.H.: Deduction and definability in infinite statistical systems. Synthese p. forthcoming (2017). URL http://link.springer.com/10.1007/s11229-017-1497-6
16. Fletcher, S.C.: Similarity structure on scientific theories. In: B. Skowron (ed.) Topological Philosophy, p. forthcoming. de Gruyter, Berlin (2019)
17. Halvorson, H.: The Logic in Philosophy of Science (2018). Forthcoming with Cambridge University Press (version of May 20, 2018)
18. Halvorson, H., Tsementzis, D.: Categories of scientific theories. In: E. Landry (ed.) Categories for the Working Philosopher, pp. 402–429. Oxford University Press, Oxford (2017)
19. Hudetz, L.: Definable categorical equivalence. Philosophy of Science **86**(1), 47–75 (2019)
20. Kim, J.: The layered model: Metaphysical considerations. Philosophical Explorations **5**(1), 2–20 (2002)
21. Krantz, D.H., Luce, R.D., Suppes, P., Tversky, A.: Foundations of Measurement, vol. II: Additive and Polynomial Representation. Academic Press, San Diego (1989)
22. List, C.: Are Interpersonal Comparisons of Utility Indeterminate? Erkenntnis **58**(2), 229–260 (2003)
23. List, C.: Social choice theory. In: E.N. Zalta (ed.) The Stanford Encyclopedia of Philosophy, Winter 2013 edn. Metaphysics Research Lab, Stanford University (2013)
24. List, C.: Levels: Descriptive, explanatory, and ontological. Noûs **0**(0), forthcoming (2018). DOI 10.1111/nous.12241
25. List, C., Spiekermann, K.: Methodological individualism and holism in political science: A reconciliation. American Political Science Review **107**(4), 629–643 (2013)
26. McLaughlin, B., Bennett, K.: Supervenience. In: E.N. Zalta (ed.) The Stanford Encyclopedia of Philosophy, Spring 2018 edn. Metaphysics Research Lab, Stanford University (2018)
27. Nagel, E.: The meaning of reduction in the natural sciences. In: R.C. Stouffer (ed.) Science and Civilization, pp. 99–135. University of Wisconsin Press, Madison (1949)
28. Nagel, E.: The Structure of Science: Problems in the Logic of Scientific Explanation. Harcourt, Brace & World, New York (1961)
29. Rosenstock, S., Barrett, T.W., Weatherall, J.O.: On Einstein algebras and relativistic spacetimes. Studies in History and Philosophy of Modern Physics **52**, 309–316 (2015)
30. Schaffner, K.: Approaches to reduction. Philosophy of Science **34**, 137–147 (1967)
31. Sen, A.K.: Collective Choice and Social Welfare. Holden-Day, San Francisco (1970)
32. Shagrir, O.: Concepts of Supervenience Revisited. Erkenntnis **78**(2), 469–485 (2013)
33. Suppes, P.: Representation and Invariance of Scientific Structures. CSLI Publications, Stanford (2002)
34. Suppes, P., Krantz, D.H., Luce, R.D., Tversky, A.: Foundations of Measurement, vol. III: Representation, Axiomatization and Invariance. Academic Press, San Diego (1990)
35. Weatherall, J.O.: Are Newtonian gravitation and Geometrized Newtonian Gravitation theoretically equivalent? Erkenntnis **81**(5), 1073–1091 (2016)
36. Weatherall, J.O.: Understanding gauge. Philosophy of Science **83**(5), 1039–1049 (2016)
37. Weatherall, J.O.: Category theory and the foundations of classical space-time theories. In: E. Landry (ed.) Categories for the Working Philosopher, pp. 329–348. Oxford University Press, Oxford (2017)