

Why surplus structure is not superfluous

Forthcoming in *British Journal for the Philosophy of science*

James Nguyen, Nicholas J. Teh, and Laura Wells

December 5, 2017

Abstract

The idea that gauge theory has ‘surplus’ structure poses a puzzle: in one much discussed sense, this structure is redundant; but on the other hand, it is also widely held to play an essential role in the theory. In this paper, we employ category-theoretic tools to illuminate an aspect of this puzzle. We precisify what is meant by ‘surplus’ structure by means of functorial comparisons with equivalence classes of gauge fields, and then show that such structure is essential for any theory that represents a rich collection of physically relevant fields which are ‘local’ in nature.

Contents

1	Introduction	1
2	Theories as Categories	3
2.1	Relations between models	3
2.2	Relations between theories	4
3	Gauge Theory as a Category	7
3.1	Gauge theory on contractible manifolds	7
3.2	Other candidates for representing $U(1)$ gauge theory	7
3.3	Surplus and inter-theoretical comparisons	10
4	Gauge Theory as a Functor	12
4.1	(Richness) and (Locality)	12
4.2	(Richness) and (Locality) imply (Surplus*)	16
5	Conclusion	18
	Appendix	19

1 Introduction

In recent years, there has been a movement to view scientific theories as categories, and to analyse them using the tools of category theory.¹ One set of motivations for this approach

¹See e.g. [Halvorson, 2012; Halvorson and Tsementzis, 2015; Teh and Tsementzis, 2017; Weatherall, 2016a,b] and references therein.

stems from considerations about how best to represent scientific theories from a philosophical perspective.² For instance, if scientific theories are thought of as—in some sense—collections of models, then we might have reason to consider inter-model relationships to be an important part of the ‘structural’ content of a theory. As such, advocates of the categorical approach recommend it because it not only allows us to clearly and precisely formulate the ‘structure’ of a theory (i.e. of its collection of models), but also operates with a notion of structure that is flexible enough to incorporate inter-model relationships. Furthermore, once such structure has been represented, the framework naturally takes it into account when formalising inter-theoretic comparisons such as ‘theoretical equivalence’ and ‘one theory having more structure than another’.

An important class of theories for testing the fruitfulness of categorical methods is ‘gauge theories’, by which we mean theories of ‘Yang-Mills type’, such as electromagnetism. This is because such theories provide a paradigmatic scenario in which comparisons of structure between two theories (or two formulations of a theory) naturally arise.³ More specifically, it is typically held that gauge theories admit of redundancies in the following sense: locally, a theory whose models use gauge fields is ‘empirically equivalent’ to a corresponding theory whose models use only gauge equivalence classes of such fields. Thus, there is an intuitive sense in which the former theory possesses ‘surplus’ (i.e. structure over and above the gauge equivalence classes) compared to the latter theory. Illuminating and precisifying this sense is a task for which the categorical approach seems ideally suited.

The above notion of ‘surplus’ structure has received much attention within the philosophical literature.⁴ An equally intriguing aspect of such structure, albeit one that has received less attention, is the idea that despite its redundancy in the above sense, it also plays an indispensable role in the theory’s formulation. For instance, it is commonly believed that gauge fields are essential for formulating classical gauge theory as a ‘local’ field theory (in a sense that we will make precise in Section 4).⁵ This leads to an apparent puzzle: How can a ‘redundancy’ be an essential feature of a theory?

Although it is certainly not obvious that the categorical approach has anything to say about this puzzle, we will argue in this paper that it does. In particular, we will show that using categorical tools to precisify ‘surplus’ gives us the resources to understand why such ‘surplus’ plays a representationally significant role in the theory. From a broader perspective, we take our work to be a contribution to both the philosophy of scientific theories literature and the philosophy of gauge theory literature. It contributes to the former by showing that for a large class of scientific theories, namely field theories, it is important to extend what we call the ‘theories as categories’ paradigm to the ‘theories as functors’ paradigm in order to fully accommodate inter-model relationships (in this case: relationships specifying how the local systems represented by the theory can determine global systems) which are relevant for gauge theory. (We say more about why this is important in Sections 4.1.1 and 4.1.2.) It contributes to the latter by showing that our intuitive understanding of ‘surplus’ is ambiguous and its disambiguation rests on the precise standard of sameness that one chooses for a scientific

²Note that we do not claim any originality for this approach on the part of the philosophical literature; the general idea of ‘representing theories as categories’ has been common within theoretical physics for a long time, e.g. in the applications of noncommutative geometry [Connes, 1994] to physics.

³We note that the term ‘gauge’ is used in many different ways within the philosophical literature; for instance, some authors use ‘gauge’ to describe any theory with representationally redundant structure. Our stricter notion of ‘gauge’ is standard within the high-energy theory literature, and is appropriate given that ‘Yang-Mills type’ theories are our concern in the present article.

⁴See for instance [Earman, 2004; Healey, 2007; Redhead, 2003] and the discussion in [Belot et al., 2009].

⁵This view is widely held, but for an explicit and mathematically sophisticated presentation, see [Schreiber, 2016]. For a less formal point of view, see also the discussion of ‘stacks’ in [Witten, 2010].

theory; furthermore, only one of these choices will yield a satisfying resolution of the puzzle.

The structure of the paper is as follows. In Section 2, we discuss the general idea of representing a scientific theory as a category. In Section 3, we then apply the ‘theories as categories’ framework to an idealised version of $U(1)$ gauge theory, where we temporarily (i) restrict our focus to contractible manifolds, and (ii) ignore the local nature of fields. Within this framework, we offer an analysis of ‘surplus’ that captures standard physical intuitions about what constitutes the (locally) redundant structure of a gauge theory. For completeness, we then briefly contrast our analysis with that of Weatherall [2016a,b], which investigates a different concept of ‘surplus’ and thus does not capture the intuitions that we are after.

In Section 4, we remove assumptions (i) and (ii) by representing theories as functors instead of as categories. We then resolve our earlier puzzle by showing that from the perspective of this more realistic notion of a theory, ‘surplus’ (in the sense picked out in Section 3) is necessary in order to accommodate a rich set of physically relevant models while simultaneously representing the ‘locality’ of fields. Finally, Section 5 provides concluding morals for the project of representing scientific theories categorically and points out an intriguing connection between our argument and a highly influential approach to mathematical classification that has its roots in the groundbreaking work of Grothendieck [5960].

2 Theories as Categories

We now describe the general framework for representing a theory as a category.⁶ Section 2.1 discusses an inter-model relationship that will be of interest to us: the ‘representational equivalence’ of two models, meaning that they (accurately) represent the same possible states of affairs. We also discuss several different kinds of categories that one might use to encode representational equivalence. Section 2.2 then reviews how functors can be used to formulate equivalence and inequivalence relations between theories.

2.1 Relations between models

Recall that a model of a theory is a structure that represents a system that is possible according to the theory.⁷ For instance, if we are concerned with a (classical) field theory, then a system will be encoded by some particular field configuration that satisfies the theory’s equations of motion. In the present categorical formalisation of a scientific theory T , such models will be the *objects* of the category \mathcal{C}_T that one is using to represent T .⁸ In order to fully specify \mathcal{C}_T , one also needs to define the *morphisms* (or arrows) of this category, which can be taken to represent various kinds of inter-model relationships. In this section, we will focus on the inter-model relationships represented by *isomorphisms* (i.e. invertible morphisms): Models of \mathcal{C}_T that are related by an isomorphism are taken to represent the same possible systems, despite the fact that the models could be (strictly speaking) distinct objects.⁹ Furthermore, each distinct isomorphism specifies a distinct way of transforming one model into another in a manner that preserves its

⁶In keeping with much of the literature on this subject, we assume that—in this basic case—the relevant categories are not functor categories, which would require us to antecedently define domain and co-domain categories of models. The case of functor categories will arise in what we call the ‘theories as functors’ paradigm of Section 4.

⁷How these structures represent physical systems is a vexing question which we will not address here. See [Frigg and Nguyen, 2016] for an overview of options currently available in the literature.

⁸For convenience, we use 1-categories to represent scientific theories throughout this paper, but our points are easily extended to theories that require an n -categorical formulation.

⁹This is motivated by the fact that isomorphic models in a category \mathcal{C} have the ‘same structure’ (from the perspective of \mathcal{C}), and once this structure has been chosen to capture what, and how, the model represents, then isomorphic models can (albeit not *must*) be taken to be representationally equivalent.

representational content; such information (about the different ways relating representationally equivalent models) will play a vital role in the analysis of Section 4. As we discuss there and in the conclusion, the isomorphisms of a theory—thought of as a category—can also contribute to the representational capacities of a theory, namely by representing the different ways in which the systems represented by the objects (models) are related to one another (e.g. as local subsystems combining into global systems).

An elementary type of category of models, then, is one in which the only inter-model relationship encoded is representational equivalence, i.e. one in which the only morphisms are isomorphisms. Such categories are called *groupoids*. In what follows, we will have reason to consider three kinds of groupoids. First, a groupoid is a *categorical set* \mathcal{S} if and only if its only morphisms are identity morphisms.¹⁰ Each distinct model in the category is thereby taken to represent a different possible system. Second, a groupoid is a *categorical equivalence relation* \mathcal{E} if and only if between any two objects A, B in \mathcal{E} there is at most one morphism in the morphism set $\text{hom}(A, B)$.¹¹ For categorical equivalence relations, clusters of objects connected by morphisms fall into the same representational ‘equivalence class’, and thereby represent the same system. Finally, we will also consider groupoids that are neither categorical sets nor categorical equivalence relations. Like the latter, such groupoids contain disconnected clusters of representationally equivalent objects, but in this most general case, there may be multiple morphisms between (not necessarily distinct) objects within each cluster—these morphisms represent the different ways in which representationally equivalent models can be transformed into one another.

2.2 Relations between theories

Having represented a scientific theory as a category, one can then treat a theory as being itself an object of a category: the category of theories. A simple category that lends itself to this purpose is the 2-category **Cat** whose objects are 1-categories, whose 1-morphisms are functors between categories, and whose 2-morphisms (morphisms between 1-morphisms) are natural transformations between functors. We now explain how to use the morphisms of this category to describe equivalence and inequivalence relations between theories.

2.2.1 Equivalence

Just as we previously used morphisms within a 1-category to describe the standard of sameness for models within a theory, we can use the morphisms of **Cat**, namely functors, to describe a standard of sameness for *theories*. Two options emerge for such a standard: (a) *categorical isomorphism*; or (b) *categorical equivalence*, which includes isomorphism as a special case. Two (categorically represented) scientific theories \mathcal{C}_{T_1} and \mathcal{C}_{T_2} are *categorically isomorphic* if and only if there exist functors $G : \mathcal{C}_{T_1} \rightarrow \mathcal{C}_{T_2}$ and $H : \mathcal{C}_{T_2} \rightarrow \mathcal{C}_{T_1}$ such that $GH = 1_{\mathcal{C}_{T_2}}$ and $HG = 1_{\mathcal{C}_{T_1}}$. And they are *categorically equivalent* if and only if $GH \cong 1_{\mathcal{C}_{T_2}}$ and $HG \cong 1_{\mathcal{C}_{T_1}}$; notice that these ‘quasi-inverse’ functors weaken the equalities in the definition of (a) to natural isomorphisms.¹²

¹⁰The reason for this nomenclature is that each categorical set can be defined from a set and vice versa: if a theory has a set of models M_T , then we can redescribe this data as the categorical set \mathcal{S}_T whose objects correspond to elements of M_T and whose morphisms are all and only identities (and conversely, if we have a categorical set we can ‘forget’ the identity morphisms, leaving us with a bare set of objects). In the literature, this sort of groupoid is also known as a ‘discrete category’.

¹¹Thus, every categorical set is a categorical equivalence relation, but not vice versa.

¹²Where needed we use ‘ \cong ’ to denote isomorphism (both ‘natural isomorphism’ as in the above relations between functors, as well as categorical isomorphism—the meaning will be clear from the context) and ‘ \simeq ’ to denote categorical

We will eventually take up the question of which of these standards best captures equivalence and inequivalence relations between (categorically represented) scientific theories such as gauge theory. For the moment, let us note that mathematical practice recommends (b) for methodological reasons: mathematicians generally take it that we should not limit the expressive resources of our standard of sameness by insisting on categorical isomorphism (thereby assuming that all relevant 2-morphisms are trivial). This is motivated by the fact that there are cases of categorical equivalence which are not categorical isomorphism. For example, in the case of ‘Gelfand duality’ one has a category \mathcal{C}_G of geometric models and a non-isomorphic category \mathcal{C}_A of algebraic models: although the objects of these respective categories look very different, they are mutually interdefinable (albeit only up to isomorphism), and so we would intuitively still want to say that \mathcal{C}_G and \mathcal{C}_A are ‘structurally the same’; the quasi-inverse functors of categorical equivalence (but not the inverse functors of categorical isomorphism) have precisely the weakness that allows us to say this.

In the rest of this paper, we will be concerned with a rather more trivial case of a categorical equivalence that is not an isomorphism. The general form of the case is as follows: (Skel): Given a groupoid \mathcal{C} (and assuming the axiom of choice), one can construct a *skeleton* category, $\text{sk}\mathcal{C}$, by picking an object from each isomorphism class of \mathcal{C} and defining $\text{sk}\mathcal{C}$ to be the full subcategory of \mathcal{C} containing all and only these objects. This skeleton is unique up to categorical isomorphism [Riehl, 2016, p. 43]. It is easy to see that \mathcal{C} is categorically equivalent to $\text{sk}\mathcal{C}$, because the relevant interdefinability relations are almost trivial.¹³

Armed with these conceptual tools, we can now pause to reflect on the question of what is gained by representing theories as groupoids, rather than as sets, of models. Suppose that we represent our theory as a categorical set of models, and that we choose (a)—i.e. categorical isomorphism—as our standard of sameness. Then, since the property of being a categorical set is preserved by (a), and categorical sets and sets are interdefinable, nothing has been gained by passing from sets to categorical sets. On the other hand, choosing (b)—i.e. categorical equivalence—as our standard of sameness reveals a novel feature of the categorical approach, as demonstrated by the following special case of (Skel): Every categorical equivalence relation is equivalent to its skeleton, which is a categorical set [Barnet-Lamb, 2005, p. 15]. This fact shows us that the property of ‘being a categorical set’ is not invariant under (b); thus, passing from sets to categorical sets allows us to take advantage of the representational flexibility afforded by (b).¹⁴ For instance, in Section 3, we will consider a gauge theory example in which formulating the models as a categorical set allows us to pass to a categorically equivalent groupoid (a categorical equivalence relation) of models which is not a categorical set.

The distinction between the standards of sameness (a) and (b) will be crucial for the work that we wish to do in this paper. In Section 3, we will deploy this distinction to demonstrate that the intuitive notion of ‘surplus’ is in fact ambiguous, and the resolution of this ambiguity turns on whether one adopts (a) or (b). Furthermore, Section 4 argues that we have more than merely methodological reasons to adopt (b): Within the context of gauge theory, (b) is part of an explanation of why ‘surplus’ can play the representational role that it does.

Let us stress that all that we have said so far concerns when two (categorically represented) theories are *formally* equivalent to one another. The further question of when these theories are equivalent *tout court* (or ‘theoretically equivalent’) clearly turns on how these formal representations should be *interpreted*. We do not wish to take a stand on ‘how to interpret physical

equivalence.

¹³Note that while the algebraic data of a category and its skeleton are equivalent, geometric properties (such as smoothness) are not in general preserved when one passes to the skeleton.

¹⁴Turning this point on its head, one might also use it to justify categorical equivalence as a standard of sameness: if one sticks to categorical isomorphism—which does preserve the property of being a categorical set—then one loses the distinctive flavor of a categorical formulation.

theories’ or on sufficient conditions for theoretical equivalence here. For our purposes it suffices to note that, in addition to being related by a formal standard of equivalence (e.g. (a) or (b)), two equivalent theories would at least need to be ‘empirically equivalent’, which means that the objects (models) of the formally equivalent categories should represent the same empirical data (i.e. the empirical data that would be accessible if the physical state of affairs represented by the models were actual). In the context of the restricted $U(1)$ gauge theory that we discuss in Section 3 of this paper, we follow the physicists’ standard practice of identifying the empirical content of a field model with its force field F . As we will describe below, the assumption of empirical equivalence places an additional constraint on the functors that relate two categorical representations of $U(1)$ gauge theory.

2.2.2 Inequivalence

Finally, there are cases in which it is fruitful to consider relations between theories that have different amounts of structure; in the context of this paper, such a scenario is provided by our analysis of ‘surplus structure’. By hypothesis, there will be no categorical equivalence between such theories, but one might still try to quantify such differences by means of functorial relationships between theories.

Within the mathematical literature, there is a long tradition of describing such relationships by means of ‘forgetful functors’ between categories. More recently, Baez et al. [2004] have offered a more refined taxonomy of ‘forgetfulness’ in order to facilitate applications to categorical algebraic topology; we will find it convenient to apply this taxonomy to the analysis of ‘surplus’ within our less sophisticated setting.¹⁵ In order to understand Baez et al.’s taxonomy, first recall some textbook terminology concerning a functor $G : \mathcal{C} \rightarrow \mathcal{D}$:

- G is *full* if and only if for all objects A, B in \mathcal{C} , the map (induced by G) from $\text{hom}(A, B)$ to $\text{hom}(G(A), G(B))$ —i.e. $G_{A \rightarrow B} : \text{hom}(A, B) \rightarrow \text{hom}(G(A), G(B))$ —is surjective.
- G is *faithful* if and only if for all $A, B \in \mathcal{C}$, $G_{A \rightarrow B}$ is injective.
- G is *essentially surjective* if and only if for every object X in \mathcal{D} , there is an object A in \mathcal{C} such that $G(A)$ is isomorphic to X .

There exists a $G : \mathcal{C} \rightarrow \mathcal{D}$ which is full, faithful, and essentially surjective if and only if \mathcal{C} and \mathcal{D} are categorically equivalent.¹⁶

Now, á la Baez et al., if G fails to have at least one of these properties then G is a *forgetful* functor. If G fails to be full then it forgets *structure**, if it fails to be faithful then it forgets *stuff*, and if it fails to be essentially surjective then it forgets *properties*.¹⁷ Thus, relative to a choice of functor, when comparing two categories \mathcal{C} and \mathcal{D} we can say that if G fails to be full then \mathcal{C} has more *structure** than \mathcal{D} ; if G fails to be faithful then \mathcal{C} has more *properties* than \mathcal{D} ; and if G fails to be essentially surjective, then \mathcal{C} has more *stuff* than \mathcal{D} .

¹⁵One reason for this convenience is that it will allow us to compare our approach to [Weatherall, 2016b, pp. 1042-3], which also borrows Baez et al.’s taxonomy in order to discuss ‘gauge’.

¹⁶Notice that the left-hand-side of this biconditional is an existential claim. That there exists such a functor does not entail that every functor between the categories is full, faithful, and essentially surjective. Equivalently, that there exists a forgetful functor between the categories does not entail that they are not categorically equivalent.

¹⁷Note that we wish to reserve the term ‘structure’ for the intuitive or pre-theoretical notion of structure, and so we introduce ‘structure*’ as a term of art that replaces Baez et al.’s use of ‘structure’ to refer to the functorial notion. Thus, when we speak of a functor $G : \mathcal{C} \rightarrow \mathcal{D}$ forgetting *structure**, we are not claiming that this captures the intuitive idea that \mathcal{C} , or the scientific theory represented by \mathcal{C} , has ‘more structure’ than the theory represented by \mathcal{D} . We leave open the possibility that a functor that forgets *stuff* or *properties* may also demonstrate that the domain category has more ‘structure’ in some intuitive sense than the codomain category.

3 Gauge Theory as a Category

3.1 Gauge theory on contractible manifolds

When physicists speak of ‘gauge theory’, they typically have in mind Yang-Mills type gauge theory, of which the prototype is $U(1)$ gauge theory (also known as ‘electromagnetism’). We have adopted this nomenclature in our paper, and we shall only address the $U(1)$ case for simplicity. Furthermore, we will treat spacetime as a smooth manifold throughout this paper, since none of our points turn on the presence of spacetime metric structure.

As we mentioned in Section 1, a gauge theory is a field theory, and fields are typically taken to possess ‘locality’, in the sense that field configurations on a manifold will be determined by specifying configurations on a collection of local subregions that ‘cover’ that manifold. However, in this section, we will only consider a mock-up of field theory that makes two simplifying assumptions. First, the theory does not represent any facts about locality, and so a model will simply be a field configuration on the whole of a spacetime. Second, we will only consider models that are set on contractible manifolds. Both of these assumptions will be removed in Section 4.

What are the fields of a $U(1)$ gauge theory? Physicists typically formulate the theory in terms of gauge fields, i.e. 1-forms $A \in \Omega^1(M, \mathfrak{u}(1))$.¹⁸ From a gauge field A , one can define the force field $F := dA$, which is invariant under the gauge transformations:

$$A \mapsto A + g^{-1}dg \tag{1}$$

where $g : M \rightarrow U(1)$.¹⁹

Note that by definition, the force field F satisfies the source-free equation of motion $dF = 0$.²⁰ On a contractible manifold M (and assuming that we have no interest in subsystems on non-contractible subregions of M) the empirical content of the theory can be expressed solely in terms of the gauge-invariant fields F ; whence familiar worries about ‘surplus’ when the theory is formulated in terms of the gauge fields A .

We now turn to the articulation of a $U(1)$ gauge theory on M as a category. Since such a theory represents states of affairs on a particular spacetime M , we will adopt the standard practice of fixing M , and taking a model to be a field configuration on M .²¹ Thus, a category of models will be a category whose objects are such gauge fields, and whose morphisms (if we should choose to include them) are symmetries of the relevant fields.

A specific category of this kind that has been widely used in the mathematical physics literature is the groupoid of gauge fields \mathcal{C}_A . \mathcal{C}_A is a groupoid—that is not a categorical set or equivalence relation—whose objects are the gauge fields $A \in \Omega^1(M, \mathfrak{u}(1))$, and whose morphisms are the gauge transformations given by (1).

3.2 Other candidates for representing $U(1)$ gauge theory

Although \mathcal{C}_A is one of the standard categorical representations of gauge theory within physics, it has scarcely been discussed within the philosophical literature.²² On the other hand, several alternative candidate categories for representing gauge theory have been discussed by Weatherall

¹⁸For an introduction to gauge fields, see [Pokorski, 2000, Chapter 1].

¹⁹Throughout this paper we use the mathematicians’ convention; the typical physicists’ convention would denote the gauge transformation as $A \mapsto A - ig^{-1}dg$.

²⁰There is also a dual source-free equation of motion $d * F = 0$, but we will not use it in this paper.

²¹For more details of the standard modelling practice, see [Benini et al., 2015, Section 2]. In Section 4 we will discuss how to assign field configurations to varying manifolds.

²²For some exceptions, see (Dougherty, forthcoming) and Teh, forthcoming)

[2016a,b], in conjunction with a particular proposal for understanding ‘surplus’. We now consider (simplified versions of) these alternative categories and proceed to develop an analysis of ‘surplus’ which identifies \mathcal{C}_A as possessing the relevant ‘surplus’ structure; we then briefly contrast our proposal with that of Weatherall [2016b].²³

Alternative representations of gauge theory are given by the following groupoids:

- \mathcal{S}_F is the categorical set whose objects are force fields $F \in \Omega^2(M, \mathfrak{u}(1))$ and all of whose morphisms are identities.²⁴
- \mathcal{S}_A is the categorical set whose objects are gauge fields $A \in \Omega^1(M, \mathfrak{u}(1))$, and all of whose morphisms are identities.
- $\mathcal{S}_{[A]}$ is the categorical set whose objects are (gauge) equivalence classes $[A]$ and all of whose morphisms are identities.
- \mathcal{E}_A is the categorical equivalence relation whose objects are gauge fields $A \in \Omega^1(M, \mathfrak{u}(1))$ and whose morphisms are given by the gauge transformations $A \mapsto A + \lambda$, where λ is a closed 1-form, i.e. $d\lambda = 0$ (thus, we know from (1) that $\lambda = g^{-1}dg$ for some g).²⁵ Note that it follows from this definition that the only automorphism of A is the trivial one, i.e. $\lambda = 0$.

We note that while the physically-minded reader may rightly chafe at the definition of \mathcal{E}_A (why arbitrarily omit from (1) all automorphisms but the trivial one?), we will provide a mathematical explanation for this choice in Section 5, albeit one that will be at odds with our (physical) motivations.

In order to give the reader some visual intuition for these definitions, Figure 1 provides an illustration of these groupoids and \mathcal{C}_A . Notice that boxes (a)-(c) and (e) only contain identity automorphisms, whereas box (d) also contains non-trivial automorphisms.²⁶ This last feature explains why box (c) contains exactly two arrows between distinct objects in the same cluster, whereas box (d) contains many such arrows (though not explicitly depicted, this is implied by concatenating non-trivial automorphisms with the depicted arrows).

[Insert Figure 1 about here]

We are now in a position to consider the functorial relationships between these different theories. In light of our discussion of empirical equivalence at the end of Section 2.2.1, we will only consider functors that satisfy the following criterion:

Empirical Equivalence of $U(1)$ Theories: Fix a contractible manifold M . We will say that two categorical representations \mathcal{C} and \mathcal{D} of $U(1)$ gauge theory on M are *empirically equivalent* if and only if for each model $X \in \mathcal{C}$ (which can be a gauge field, equivalence class of gauge fields, or force field depending on the category in

²³This simplification derives from omitting spacetime isomorphisms, since for our purposes we focus exclusively on the field structure of $U(1)$ gauge theory, which, in this section, is distinguished from its metric structure by keeping the background manifold and metric fixed and identifying models with n -forms (or constructions thereof) and the morphisms with the maps between them.

²⁴Recall two points: First, since F is a force field, it is closed, i.e. $dF = 0$. Second, since we are assuming that M is contractible, Poincaré’s Lemma implies that for every F there is a non-unique $A \in \Omega^1(M, \mathfrak{u}(1))$ such that $dA = F$.

²⁵The definition of \mathcal{E}_A comes from the category \mathbf{EM}_2 in [Weatherall, 2016a, pp.1082-3], but where we have omitted spacetime isometries since these will not be relevant for our purposes.

²⁶Explicitly, the non-trivial automorphisms contained in \mathcal{C}_A correspond to shifts by a (non-zero) constant function, which in turn give rise to the ‘rigid’ or ‘global’ $U(1)$ symmetries of electromagnetism, i.e. the (non-trivial) constant maps $g : M \rightarrow U(1)$. It is now clear why physicists typically take \mathcal{C}_A to be the ‘physical category’: when one couples $U(1)$ gauge theory to the matter sector (e.g. scalar or fermion fields), these rigid transformations are used to construct non-trivial conserved charges by means of Noether’s (first) theorem. By contrast (and putting aside boundary terms), only trivial charges can be associated with the non-constant maps.

question) with a force field F (either induced by $F := dA$, or explicitly defined in the model) there is a model $Y \in \mathcal{D}$ with the same F , and vice versa for each $Y \in \mathcal{D}$. We thus say that a pair of functors $G : \mathcal{C} \rightarrow \mathcal{D}$ and $H : \mathcal{D} \rightarrow \mathcal{C}$ *realises an empirical equivalence* if and only if for every model (object) X in \mathcal{C} there is a model Y in the image of G with the same force field F , and vice versa for every model $Y \in \mathcal{D}$ and the image of H .

In other words, such functors must act in one of the following ways on the categories in Figure 1: they take a gauge field A to a force field F such that $F = dA$ (and vice versa); they take an equivalence class of gauge fields $[A]$ to a gauge field A such that $A \in [A]$ (and vice versa); and finally, they take an equivalence class of gauge fields $[A]$ to a force field F such that $F = dA$ for all $A \in [A]$ (and vice versa).

For our purposes, it will suffice to consider the following empirical equivalence-realising functors:²⁷

$G: \mathcal{S}_A \rightarrow \mathcal{S}_{[A]}$ such that for all A in \mathcal{S}_A : $G(A) = [A]$ where $A \in [A]$; identities are sent to identities.

$H: \mathcal{S}_F \rightarrow \mathcal{S}_{[A]}$ such that for all F in \mathcal{S}_F : $H(F) = [A]$ where $F = dA$ for all $A \in [A]$; identities are sent to identities.

$I: \mathcal{S}_{[A]} \rightarrow \mathcal{E}_A$ such that for all $[A]$ in $\mathcal{S}_{[A]}$: $I([A]) =$ a representative A of $[A]$; identities are sent to identities.²⁸

$J: \mathcal{E}_A \rightarrow \mathcal{S}_{[A]}$ such that for all A in \mathcal{E}_A : $J(A) = [A]$ where $A \in [A]$; and for all A, B in \mathcal{E}_A such that $J(A) = J(B)$: $\text{hom}(J(A), J(B)) = \{1_{J(A)}\}$.

$K: \mathcal{S}_{[A]} \rightarrow \mathcal{C}_A$ such that for all $[A]$ in $\mathcal{S}_{[A]}$: $K([A]) =$ some A in \mathcal{C}_A such that $A \in [A]$; identities are sent to identities.

$\tau: \mathcal{C}_A \rightarrow \mathcal{S}_{[A]}$ such that for all A in \mathcal{C}_A : $\tau(A) = [A]$ where $A \in [A]$; and for all A, B in \mathcal{C}_A such that $\tau(A) = \tau(B)$, $\text{hom}(\tau(A), \tau(B)) = \{1_{\tau(A)}\}$.

$M: \mathcal{S}_{[A]} \rightarrow \mathcal{S}_A$ such that for all $[A]$ in $\mathcal{S}_{[A]}$: $M([A]) =$ some A in \mathcal{S}_A such that $A \in [A]$; identities are sent to identities.

Since these functors realise empirical equivalences, the categories in Figure 1 are all empirically equivalent to each other. We now turn to the further inter-theoretic relationships encoded by these functors, as depicted in Figure 2.

[Insert Figure 2 about here]

First, there are (formal) equivalences, as represented in the vertical axis of Figure 2, where \cong denotes categorical isomorphism and \simeq denotes categorical equivalence. It is easy to see that H and its inverse demonstrate that \mathcal{S}_F and $\mathcal{S}_{[A]}$ are isomorphic (cf. [Weatherall, 2016a, p. 1082]). Similarly, I and J are quasi-inverse functors which demonstrate that $\mathcal{S}_{[A]}$ and \mathcal{E}_A are categorically equivalent (but non-isomorphic).

Notice that this equivalence is simply a special case of what we called (Skel) in Section 2.2.1, and in particular, of the fact that a categorical equivalence relation (\mathcal{E}_A) is equivalent to a categorical set ($\mathcal{S}_{[A]}$ and \mathcal{S}_F).

Next, there are inequivalences, as represented in the horizontal axis of Figure 2. The following propositions show that G, K, τ and M are all forgetful functors which encode differences in the structure of theories \mathcal{C}_A , $\mathcal{S}_{[A]} (\cong \mathcal{S}_F \simeq \mathcal{E}_A)$, and \mathcal{S}_A (see the Appendix for the proofs):

²⁷Note that the actions on the morphisms by G, H, I, K and M are determined by the definition of a functor since the only morphisms in the domain categories are identities.

²⁸Note that this functor lacks naturality (in the choice of the representative of the equivalence class). The same applies to K and M .

Proposition 3.2.1. *G forgets (only) structure*. So, with respect to G, \mathcal{S}_A has more structure* than $\mathcal{S}_{[A]}$.*

Proposition 3.2.2. *K forgets (only) structure*. So, with respect to K, $\mathcal{S}_{[A]}$ has more structure* than \mathcal{C}_A .*

Proposition 3.2.3. *τ forgets (only) stuff. So, with respect to τ , \mathcal{C}_A has more stuff than $\mathcal{S}_{[A]}$.*

Proposition 3.2.4. *M forgets (only) properties. So, with respect to M, $\mathcal{S}_{[A]}$ has more properties than \mathcal{S}_A .*

In the next subsection we discuss how these forgetful relationships correspond to the notion of ‘surplus’ structure.

3.3 Surplus and inter-theoretical comparisons

Within philosophy, the terms ‘surplus’ structure or ‘descriptive fluff’ have been used to refer to features of gauge theory that seem to play no representational role in theory (cf. [Earman, 2004; Redhead, 2003]). And although physicists tend not to use either of these terms, it is extremely common for them to claim that ‘gauge symmetry is a *redundancy*, because it relates gauge fields that represent the same physical state’ and that it is the ‘gauge equivalence classes’ that capture the non-redundant physical data (the references for this view are too many to list, but see especially [Tong, 2006, Section 6.1.1], [Alvarez-Gaume and Vazquez-Mozo, 2011, p.64], and [Sundermeyer, 2014, p.497]). Thus, the notion of ‘surplus’ structure or ‘redundancy’ that we are after applies to theories in which (i) gauge fields corresponding to the same F are represented as representationally equivalent, and (ii) the gauge symmetries are included as part of the structure that is ‘surplus’. While in Section 4 we argue that this notion of ‘surplus’ *does* in fact play a representational role when one adopts a sufficiently rich conception of ‘gauge theory’, for now we will only be concerned with the question of how precisely to understand ‘surplus’ within a simplified setting. As a rough first pass, we might say that a theory with ‘surplus’ structure is one that includes the gauge symmetries—understood as isomorphisms—as part of the theory (thus satisfying (i)), and whose structure is in some still-to-be-precisified sense ‘surplus’ with respect to a theory that only includes the gauge equivalence classes (thus satisfying (ii)). In other words, such a theory satisfies

(Surplus): A (gauge) theory has ‘surplus structure’ (which includes the gauge fields and gauge symmetries) over and above its gauge equivalence classes.²⁹

How can this intuitive notion of (Surplus) be precisified by means of categorical tools? As we see it, the ‘theories as categories’ paradigm allows us to provide a four-fold precisification of what it means for a theory to have ‘structure over and above its gauge equivalence classes’.

First, we use categories to represent the two theories that we are comparing: a theory whose field content involves only gauge equivalence classes, and a theory that is a candidate for having (Surplus) structure. We thus take $\mathcal{S}_{[A]}$ to represent the former, whereas the latter can be represented by either \mathcal{C}_A , \mathcal{E}_A , or \mathcal{S}_A .

Second, our target concept of ‘surplus’ applies to theories that use gauge symmetries to encode relations of representational equivalence between gauge fields. Since \mathcal{S}_A does not include these symmetries, it is not the appropriate category for capturing the content of (Surplus). Consequently, our focus will be restricted to comparisons between $\mathcal{S}_{[A]}$ (the theory that does not satisfy (Surplus)) and either \mathcal{C}_A or \mathcal{E}_A (the candidate theories that satisfy (Surplus)).

²⁹Note that this is essentially the formulation given in [Redhead, 2003].

Third, once we have represented theories as categories, we will need to settle on a standard of sameness for theories (thus disambiguating an aspect of what is meant by ‘over and above’ in (Surplus)). If categorical isomorphism is our standard, then we obtain what one might think of as the ‘naive’ reading of (Surplus), on which \mathcal{E}_A counts as having ‘surplus’ over and above the gauge equivalence classes of $\mathcal{S}_{[A]}$. However, as we argued earlier, there are strong methodological reasons for picking categorical equivalence as our standard. We will thus make this choice, and so \mathcal{E}_A will count as ‘the same theory’ as $\mathcal{S}_{[A]}$. Consequently, the only candidate category which satisfies (Surplus) is \mathcal{C}_A .

Fourth, we perform a functorial comparison between \mathcal{C}_A and $\mathcal{S}_{[A]}$ that will clarify the ‘surplus’ structure that \mathcal{C}_A possesses over and above $\mathcal{S}_{[A]}$. Such a comparison is the content of Prop. 3.2.3 (and the left arm of Figure 2), which tells us that the forgetful truncation functor $\tau : \mathcal{C}_A \rightarrow \mathcal{S}_{[A]}$ forgets *stuff* in Baez et al.’s taxonomy of forgetfulness.³⁰

This four-fold precisification of ‘surplus’ structure allows us to define *surplus** structure as the *stuff* that is forgotten by the functor τ . This provides the following characterisation of what it means for a $U(1)$ gauge theory to have such structure:

(Surplus*): A (gauge) theory contains the *stuff* that is forgotten by τ (namely the non-trivial automorphisms of the gauge fields and the result of concatenating them with the morphisms already included in \mathcal{E}_A).³¹

It is worth pausing briefly to contrast our notion of ‘surplus’ with the rather different concept that is being analysed in [Weatherall, 2016b] (albeit also under the heading ‘surplus’), where we will omit spacetime isometries from the latter’s categories since we are only concerned with gauge structure. First, note that both accounts hold that the theory without ‘surplus’ (in each respective sense) is $\mathcal{S}_{[A]}$, and both take categorical equivalence to be the correct standard of sameness for theories. The key difference between these accounts lies in their respective target notions of ‘surplus’. On the one hand, we are interested in a notion of ‘surplus’ that is possessed by theories which take gauge fields to be representationally equivalent (and which represent this by means of gauge transformations between gauge fields); thus, \mathcal{C}_A is our candidate for such a theory and ‘surplus’ is characterised by the *stuff*-forgetting functor $\tau : \mathcal{C}_A \rightarrow \mathcal{S}_{[A]}$ (cf. the left arm of Fig. 2). By contrast, Weatherall’s notion of ‘surplus’ applies to a theory that does not represent gauge fields as representationally equivalent, namely \mathcal{S}_A (which has no gauge transformations between objects and is essentially a set). Thus, his notion of ‘surplus’ is characterised by the structure*-forgetting functor $G : \mathcal{S}_A \rightarrow \mathcal{S}_{[A]}$ (cf. the right arm of Fig. 2), which intuitively ‘forgets the lack of arrows between objects’.³²

We take no issue with Weatherall’s account as an attempt to compare the objects of \mathcal{S}_A (which is essentially a set) with $\mathcal{S}_{[A]}$. However, from the point of view of physical practice, there are at least two good reasons to focus on \mathcal{C}_A and its comparison with \mathcal{S}_A . First, and as we mentioned earlier (cf. fn. 26), \mathcal{C}_A —but not $\mathcal{S}_{[A]}$ or \mathcal{S}_A —contains the global gauge transformations

³⁰Here, one familiar interpretation of *stuff* as ‘entities within the objects that can be equipped with structure’ does not apply because we are not dealing with categories which are structured sets. However, the ‘*structure/stuff/properties*’ distinctions are relevant beyond the case of structured sets, and are standardly used in categorical algebraic topology, for which we refer the reader to [Baez and Shulman, 2007]. For a general account of truncation, see [Lurie, 2009].

³¹To see clearly why it is the non-trivial automorphisms that are important, it is useful to compare the skeleton $\text{sk}\mathcal{C}_A$ —which is categorically equivalent to \mathcal{C}_A —to $\mathcal{S}_{[A]}$; the key difference is precisely the presence of the non-trivial automorphisms.

³²More generally, one way of motivating the focus on \mathcal{S}_A as a candidate category for possessing ‘surplus’ is to adopt a conception of structure that focuses on the structure of *objects* of a category, and uses the morphisms to track how much structure these objects have. An implicit maxim guiding such an approach is something like: the more isomorphisms in a category, the less (relevant) structure its objects have, because isomorphisms are structure-preserving maps; the more there are, the less structure there is to preserve.

(constant maps $g : M \rightarrow U(1)$) which are responsible for (non-trivial) charge conservation when a $U(1)$ gauge theory is coupled to matter fields. Second, and as we will argue in the next section, the surplus* structure of \mathcal{C}_A is required in order to explain the oft-cited conceptual connection between the ‘redundancies’ of a gauge theory, on the one hand, and the ‘locality’ of the theory, on the other.

4 Gauge Theory as a Functor

4.1 (Richness) and (Locality)

In this section, we return to the puzzle that we raised in the introduction: how can ‘surplus’ structure be an essential feature of a theory? Roughly speaking, our strategy for resolving this puzzle will be to argue that in order to represent ‘locality’ (i.e. a property of field theory, as it is standardly conceived) the notion of theory in Section 3 needs to be extended to a more general context; from the perspective of this generalisation, surplus* structure is required in order to represent a rich set of physically relevant models that possess locality.

More carefully now, recall that the formulation of a ‘theory as a category’ in Section 3 rested on two simplifying assumptions: (i) we only considered fields set on contractible regions; and (ii) we ignored ‘locality’ by taking such regions to be entire spacetimes. Recall too that our precisification of ‘surplus’, i.e. surplus*, was predicated on this narrow notion of a ‘theory’. In Sections 4.1.1 and 4.1.2, we will consider how removing (i) and (ii) suggests a move to a broader *functorial* notion of ‘theory’, which is not merely a collection of models, but rather a uniform assignment of collections of models to different spacetimes (including subregions of a spacetime).

The motivation for lifting (i) is:

(Richness): Gauge theories have a rich set of physically relevant ‘topological soliton’ solutions, which are defined on non-contractible manifolds.

The use of the term ‘physically relevant’ here covers the spectrum of empirical and theoretical roles that such solutions play in contemporary physics.³³ On the empirical side: many solitons (instantons in Yang-Mills theory; vortices in condensed matter physics) are used to make predictions about models of field theory. And on the theoretical side, solitons are even more pervasive: for instance, many field-theoretic dualities make essential use of soliton solutions. Thus, any philosophically adequate analysis of gauge theories must take into account the richness of their solitonic models. Since the ‘theories as categories’ of Section 3 were only defined on contractible manifolds, they do not in themselves have the resources to capture (Richness); however, Section 4.1.1 will introduce the machinery of principal G -bundles that allows us to do so. Furthermore, Section 4.1.2 will introduce a notion of ‘gluing’ that forges a link between the ‘theories as categories’ of Section 3 and (Richness).

The motivation for lifting (ii) is the fact that field theories are typically taken to possess:

(Locality): All physically relevant global models (i.e. field configurations on an arbitrary spacetime M) are determined by ‘compatibly gluing’ a set of local models (i.e. field configurations on contractible subregions U_i of M , where $\{U_i\}$ is a cover of M).

In other words, (Locality) is a constraint that relates the physically relevant global models on M to a family of ‘theories as categories’ (whose objects are local models), where the family is parameterised by contractible subregions of M .

³³See [Manton and Sutcliffe, 2004] for an overview of (classical) topological soliton solutions.

By itself, (Locality) does not tell us what the physically relevant global models are. However, (Richness) helps specify this class of models, and as argued below, when (Richness) and (Locality) are conjoined they imply (Surplus*), thus showing how ‘surplus’ structure can be an essential feature of theories that are rich enough to combine three aspects: one can represent global models (i.e. topological solitons); one can represent local models (i.e. models of the sort considered in Section 3); and one can represent the relationship between local and global models (i.e. how the former can be combined to yield the latter).

In order to mount this argument more precisely, we will now proceed to describe the categorical formulation of (Richness) and (Locality). While this is a subject matter about which one can be exceedingly abstract, we will focus on a concrete and elementary example for $U(1)$ gauge theory: the Dirac monopole on the sphere $M = S^2$.

4.1.1 Representing (Richness)

The global ‘topological soliton’ models that exemplify (Richness) are standardly defined by means of the machinery of a principal G -bundle with connection.³⁴ For instance, in the our focal case of $U(1)$ gauge theory, (Richness) requires that we take into account the Dirac monopole model, which is a $U(1)$ -bundle P equipped with a connection 1-form ω over the sphere S^2 . Such global models can be understood as combining two pieces of data: (i) a topological charge that corresponds to the isomorphism class of P ; and (ii) a particular connection ω on P (up to connection-preserving bundle isomorphism). In the Dirac monopole case, it will be particularly useful for us to view the monopole’s global force field $F \in \Omega^2(S^2, \mathfrak{u}(1))$ as a summary of (i) and (ii): F is constructed from the bundle’s curvature $d\omega$, and by a standard result from Chern-Weil theory, the cohomology class of F corresponds to the topological charge of the monopole (and thus to the isomorphism class of P).³⁵

Due to a classical result by Steenrod [1951], we know that isomorphism classes of principal G -bundles over M are in bijection with homotopy classes of maps from M into the classifying space BG of the group G . Thus, in cases where the classifying space BG is well-understood, homotopy classes of maps provide a useful classification of the possible topological charges of a soliton model. Applying this result to this case of the Dirac monopole, we see that the principal $U(1)$ -bundles P over S^2 are in bijection with $[S^2, BU(1) = \mathbb{C}P^\infty] = \pi_2(\mathbb{C}P^\infty) \cong \mathbb{Z}$. Thus, the topological charge of a Dirac monopole model is $n \in \mathbb{Z}$. The charge is also known as the ‘first Chern number’ of the model since it can be computed as $n = \frac{1}{2\pi} \int_{S^2} F$, where F represents the first Chern class of P .

In order to speak of a collection of global models on M , it is standard practice to define the functor $Bun : \mathbf{Man} \rightarrow \mathbf{Grpd}$, which assigns to any M (not necessarily contractible) the groupoid $Bun(M)$ whose objects are principal $U(1)$ -bundles over M with connection (P, ω) , and whose morphisms are connection-preserving bundle isomorphisms over the identity on M . In particular, $Bun(S^2)$ contains all the Dirac monopole models of $U(1)$ gauge theory. We note that while one could say that we are still within the ‘theories as categories’ paradigm since $Bun(M)$ is a category (for a fixed M), this observation misses the point of the functorial perspective on (Richness): our theory assigns global models to arbitrary M (and regions of M) in a uniform manner.³⁶

³⁴A principal G -bundle is a fibre bundle $\pi : P \rightarrow M$ with a continuous right action by the group G on the total space P which preserves fibres.

³⁵Here we note a technical subtlety: The first (integral) Chern class of a $U(1)$ -bundle on M is a complete invariant of the bundle’s isomorphism class, but when we represent it by F , we are omitting possible torsion information in this class. In the case where the connection is flat, such torsion data corresponds precisely to the monodromies that account for the Aharonov-Bohm effect, cf. footnote 45.

³⁶The relevant sense of ‘uniformity’ here includes the preservation of the pattern of morphisms between objects in

What is the relationship between $Bun(M)$ (and thus the Dirac monopole models) and the previously mentioned groupoid of gauge fields \mathcal{C}_A (on M)? Such a relationship is easy enough to specify in the case where M is contractible: $Bun(M)$ and \mathcal{C}_A are then equivalent as categories.³⁷ However, for non-contractible M , the relationship between $Bun(M)$ and \mathcal{C}_A becomes more complicated, because the groupoid of gauge fields \mathcal{C}_A is only defined on contractible manifolds.³⁸ And since the Dirac monopole model requires a non-contractible $M = S^2$, we will need to forge a more subtle relationship between $Bun(M)$ and \mathcal{C}_A if we are to connect the ‘theories as categories’ paradigm—and thus (Surplus*)—to the Dirac monopole models which exemplify (Richness). Representing such a relationship is required by (Locality), which we now proceed to discuss.

4.1.2 Representing (Locality)

Recall that (Locality) requires that every physically relevant global model on a (possibly non-contractible) manifold M is determined by ‘compatibly gluing’ some family of local models parametrised by the contractible cover $\{U_i\}$ of M . For $U(1)$ gauge theory, a local model on U_i is contained in one of the ‘theories as categories’ of Section 3. Thus, from a categorical point of view, the most natural way to conceptualise this ‘family of categories’ parametrised by $\{U_i\}$ is by representing theories as *functors* that assign the categories of Section 3 to contractible subregions $U_i \subseteq M$.³⁹ Functoriality is important here because the ‘theory functors’ capture the relationship between ‘relations between regions of spacetime’ and ‘possible gluing relations between models on those regions’.⁴⁰

We will define such functors by extending the ‘theory as a category’ notation of the previous section in the following way. Suppose we previously used \mathcal{C}_T to denote a groupoid of models for T for some fixed contractible manifold M . Abusing notation, we now use $\mathcal{C}_T : \mathbf{Man}^\odot \rightarrow \mathbf{Grpd}$ to denote a *functor* from the category of contractible manifolds to the category of groupoids, which assigns that groupoid of models to M . Thus, in this new notation, $\mathcal{S}_F : \mathbf{Man}^\odot \rightarrow \mathbf{Grpd}$ takes an arbitrary contractible manifold M to the categorical set $\mathcal{S}_F(M)$; *mutatis mutandis* for the functors \mathcal{E}_A and \mathcal{C}_A .⁴¹ We can now give a preliminary formulation of (Locality) in these terms: it is the requirement that any physically relevant global model in $Bun(M)$ (where M is possibly non-contractible) can be constructed by ‘compatibly gluing’ some collection of local

the functor’s domain. See also the discussion of ‘uniformity’ and its relationship with ‘canonicity’ in Section 3.4 of [Teh and Tsementzis, 2017].

³⁷This equivalence works by mapping $A \in \Omega^1(M, \mathfrak{u}(1))$ to the trivial bundle $M \times U(1)$ with connection $\omega := p_1^* A + p_2^* \mathfrak{m}$, where $\mathfrak{m} \in \Omega^1(U(1), \mathfrak{u}(1))$ is the Maurer-Cartan form, and $p_1 : M \times U(1) \rightarrow M, p_2 : M \times U(1) \rightarrow U(1)$ are projection maps. (Explicitly, the Maurer-Cartan form is obtained by taking for any $g \in U(1), \mathfrak{m}_g(v) := (L_g)_*^{-1}(v)$, where $v \in T_g U(1)$, and $L_g : U(1) \rightarrow U(1)$ is left multiplication by g .) Since every connection on $M \times U(1)$ is of this form, we have the desired bijection between $\Omega^1(M, \mathfrak{u}(1))$ and connection forms on *trivial* principal $U(1)$ -bundles. Furthermore, since *all* principal bundles over contractible M are trivial, we have the desired categorical equivalence between $Bun(M)$ and \mathcal{C}_A .

³⁸One could consider defining $\mathcal{C}_A(M)$ for non-contractible M , but this would depend on the cohomology of M and in general could give only the trivial solution; in other words, in general there is no non-trivial global gauge field defined on non-contractible M .

³⁹Below we will only describe the action of such functors on objects; for the issue of how such functors are defined on morphisms, we refer the reader to [Benini et al., 2015].

⁴⁰We note that one could again claim to be describing this data within the ‘theories as categories’ framework by fixing spacetime regions and adding the relevant relationships by hand, but this would be nothing other than an implicit appeal to functoriality.

⁴¹Notice that this ‘theories as functors’ paradigm includes the ‘theories as categories’ paradigm as a special case: to recover the latter, we simply fix a contractible M on which to evaluate the functors. By highlighting the functorial aspect we want to emphasise the significance of being able to consider varying background spacetimes. It is this feature which is significant when dealing with the locality of a theory.

models $\{A_i \in \mathcal{C}_T(U_i)\}$, where $\{U_i\}$ is a contractible open cover of M .

To go further, we will need to articulate the notion of ‘compatible gluing’ that relates global models to some collection of local models. While there are various abstract and technical ways of precisifying this notion, it will suffice for our purposes to provide an elementary description of gluing for $U(1)$ gauge theory.⁴² Suppose that we have two contractible regions U_1, U_2 with non-trivial intersection $U_{12} := U_1 \cap U_2$, and let $A_1 \in \mathcal{C}_T(U_1)$ and $A_2 \in \mathcal{C}_T(U_2)$ be a pair of local models over each region.⁴³ We then define the groupoid $\overline{\mathcal{C}_T}(U_{12})$ by taking its objects to be models in $\mathcal{C}_T(U_i)$ (where $i = 1, 2$) restricted to U_{12} and its morphisms to be the morphisms in $\mathcal{C}_T(U_i)$ which are similarly restricted.⁴⁴ We say that the local models A_1 and A_2 can be ‘compatibly glued’ if and only if there is at least one morphism in $\overline{\mathcal{C}_T}(U_{12})$ between $A_1|_{U_{12}}$ and $A_2|_{U_{12}}$. A choice of gluing is thus a choice of morphism that relates the local models on U_{12} ; as we will soon see in the case of the Dirac monopole, we will be particularly interested in the homotopy class determined by such a choice.

We can now articulate the notion of $U(1)$ gauge theory for which it makes sense to say that ‘the theory satisfies (Locality)’ with respect to the theories of Section 3. Such an *expanded theory* has three elements: (i) it represents global models by means of a ‘global theory’ $Bun(\cdot)$; (ii) it represents local models by means of a ‘local theory’ $\mathcal{C}_T(\cdot)$; and (iii) it represents global-local relationships such as (Locality)—or lack thereof—by means of $\overline{\mathcal{C}_T}(U_{ij})$ for all non-trivial overlapping regions U_{ij} .

We now have a formulation of (Locality) that applies to a general M . In our subsequent discussion of the Dirac monopole, it will be advantageous to use a less general formulation that makes explicit reference to the monopole’s global force field $F \in \Omega^2(S^2, \mathfrak{u}(1))$.⁴⁵ We will say that an expanded theory satisfies F -(Locality) if and only if the force fields F of any of its physically relevant global models can be constructed by taking some collection of local models—one from each $\mathcal{C}_T(U_i)$ —whose fields gives rise to $F|_{U_i}$, and compatibly gluing these over the non-trivial overlaps $\{U_{ij}\}$.

It should now be evident that the structure of the morphisms in the local theories $\mathcal{S}_A, \mathcal{E}_A, \mathcal{C}_A$ (now viewed as functors) of Section 3 place stringent constraints on the extent to which an expanded theory can satisfy (Locality). Without yet introducing any assumptions about what the physically relevant global solutions of a theory should be, let us consider the implications of taking each of these categories to provide a collection of local models.

First, let us consider local models coming from \mathcal{S}_A . Recall that each $\mathcal{S}_A(U_i)$ is a categorical set whose objects are fields $A \in \Omega^1(U_i, \mathfrak{u}(1))$ and whose only morphisms are the identities. Thus, for any two models $A_1 \in \mathcal{S}_A(U_1), A_2 \in \mathcal{S}_A(U_2)$ such that $F|_{U_{12}} = dA_1|_{U_{12}} = dA_2|_{U_{12}}$, the restrictions $A|_{U_{ij}}, A'|_{U_{ij}} \in \overline{\mathcal{S}_A}(U_{ij})$, can only be compatibly glued by an identity morphism. In other words, if we wish to use local models to construct a global force field F on M , we must be in a situation in which the local gauge fields stem from restricting a global gauge field on M . But on any manifold with non-trivial de Rham cohomology in degree 2 (e.g. $M = S^2$), no global gauge field can give rise to a non-trivial force field F . Thus, an expanded theory that links global models to $\mathcal{S}_A(\cdot)$ can only satisfy F -(Locality) with respect to a trivial global

⁴²For an elegant formulation of gluing in terms of homotopy limits, see equation (3.4) of [Benini et al., 2015]. This more sophisticated perspective further highlights the importance of the ‘theories as functors’ paradigm for field theory: it allows us to describe a global theory functor as a (higher) limit of a local theory functor.

⁴³For simplicity, we will only deal with a pair of contractible subregions, but the extension to an arbitrary number of subregions will be straightforward.

⁴⁴Note that this is not the same thing as applying the functor \mathcal{C}_T to U_{ij} ; the latter will not in general make sense because U_{ij} is not necessarily contractible.

⁴⁵This is less general because it does not suffice to capture Aharonov-Bohm type phenomena, where $F = 0$ and so the relevant cohomology classes for the physics stem from non-trivial 1-cycles in the spacetime. But it will suffice for our purposes, and is at any rate much closer to how physicists intuitively think of the monopole.

model; in particular, monopoles with non-trivial topological charge cannot be represented in this way.

On the other hand, in the case of \mathcal{E}_A (respectively \mathcal{C}_A), there exist non-automorphism morphisms between $A|_{U_{ij}}$ and $A'|_{U_{ij}}$ in $\overline{\mathcal{E}_A}(U_{ij})$ (respectively $\overline{\mathcal{C}_A}(U_{ij})$). Thus, these local theories have the following feature: the structure of their isomorphisms is rich enough to compatibly glue *distinct* local gauge fields which nonetheless give rise to an identical force field on U_{ij} . In the next section, we will consider the extent to which this feature allows expanded theories to represent all monopole models while simultaneously satisfying (Locality).

4.2 (Richness) and (Locality) imply (Surplus*)

Recall that our goal in this section was to explain how ‘surplus’ structure—now precisified as surplus* structure—can be an essential feature of gauge theory, and our strategy was to advert to a notion of ‘theory’ that is more faithful to physical practice—and thus incorporates (Richness) and (Locality)—instead of the simplistic theories of Section 3. In such an ‘expanded theory’, as we earlier called it, (Richness) is a property of the global models, ‘surplus’ (however one conceptualises it) is a property of the local models, and (Locality) is a relation between global models and local models.

Our main claim will be that conjoining (Richness) and (Locality) implies that the relevant local theory needs to satisfy (Surplus*), i.e. it should be \mathcal{C}_A . Hence our slogan that (Richness) and (Locality) imply (Surplus*).

We begin by giving a more abstract argument for the converse of our main claim: an expanded theory that has \mathcal{C}_A as its local theory simultaneously satisfies (Richness) and (Locality). (Richness) tells us that our theory should include as a global model any $(P, \omega) \in Bun(M)$, and so conjoining it with (Locality) yields the requirement that for any (P, ω) , there exists a collection of local models $\{A_i \in \mathcal{C}_T(U_i)\}$ that can be compatibly glued to yield (P, ω) . It will be straightforward to see that this requirement is satisfied for $\mathcal{C}_T = \mathcal{C}_A$ if we note that any (P, ω) has the following ‘local’ description: given some contractible cover $\{U_i\}$ of M , P can be described by patching together the bundle charts $\{U_i \times U(1)\}$ by means of the ‘cocycles’ $h_{ij} : U_{ij} \rightarrow U(1)$ (for any non-trivial overlap U_{ij}); similarly, ω can be described by patching together $\omega_i := \omega|_{\pi^{-1}(U_i)}$, where $\pi : P \rightarrow M$ is the bundle projection map.⁴⁶ Let $\{A_i\}$ be a collection of local models (i.e. gauge fields) along with some choice of gluing (i.e. gauge transformations)

$$A_j = A_i + g_{ij}^* \mathfrak{m} = A_i + g_{ij}^{-1} dg_{ij} \quad (2)$$

for all i, j , where we have emphasized the role of the Maurer-Cartan form \mathfrak{m} . To see that any (P, ω) can be reconstructed from some $\{A_i\}$ along with a choice of gluing, we now make two observations. First, by setting $g_{ij} = h_{ij}$ for all i, j , our choice of gluing $\{g_{ij}\}$ yields precisely the cocycle data that defines P (up to isomorphism); in other words, the gauge transformations in \mathcal{C}_A encode the structure that allows us to reconstruct any P . Second, any connection ω on P can be reconstructed (up to gauge transformations) by choosing $\{A_i\}$ such that $\omega_i = \pi^* A_i + p_2^* \mathfrak{m}$, where p_2 is the projection map $U_i \times U(1) \rightarrow U(1)$.⁴⁷

This line of reasoning also indicates why one should expect our main claim to be true. For suppose that the relevant local theory \mathcal{C}_T does not meet (Surplus*), i.e. it is one of $\mathcal{S}_{[A]}$, \mathcal{S}_F or

⁴⁶More precisely, the principal bundle $\pi : P \rightarrow M$ is given by taking the quotient space $P := (\coprod U_i \times U(1)) / \sim$ where the equivalence relation identifies $(x, g) \in U_i \times U(1)$ with $(x, h_{ij}(x)g) \in U_j \times U(1)$. The maps $\phi : \pi^{-1}(U_i) \rightarrow U_i \times U(1)$ given by $[x_i, g] \mapsto (x_i, g)$ for $x_i \in U_i$ are the local trivializations of $P \rightarrow M$, and the maps h_{ij} are the transition functions between these local trivializations.

⁴⁷For a proof of this reconstruction as well as the general equivalence between (P, ω) and a gluing of local models, see [Sontz, 2015, Thm 14.4].

\mathcal{E}_A , and thus does not contain the additional morphisms (*stuff*, cf. Prop. 3.2.3) contained in \mathcal{C}_A . Then, since the isomorphism class of P is determined by its cocycle data, and since (as we have just argued) the cocycle data is contained in the morphisms of \mathcal{C}_A , it is reasonable to expect that it will not be possible to reconstruct an arbitrary P from a local theory whose morphisms are less rich than those of \mathcal{C}_A . In fact, this is exactly what happens, and we shall now provide a more incisive argument for this conclusion in the case of the Dirac monopole, for which it will be expedient to formulate (Locality) in terms of the monopole's global force field F .

In the Dirac monopole case, $M = S^2$, take the open cover of M to be the contractible regions $U_1 = S^2 \setminus \{s\}$ and $U_2 = S^2 \setminus \{n\}$, i.e. the sphere with the south and north pole removed, respectively. In this context, the conjunction of (Richness) and F -(Locality) amounts to the requirement that the global force field F of any monopole model can be determined by compatibly gluing some pair of local models from $\mathcal{C}_T(U_1)$ and $\mathcal{C}_T(U_2)$. We will now argue that if \mathcal{C}_T does not have surplus* structure, then this requirement cannot be satisfied; furthermore, we will only focus on the case of \mathcal{E}_A since, as we discussed at the end of Section 4.1.2, \mathcal{S}_A only contains enough information to reconstruct the trivial monopole.

We begin by considering how the isomorphism class (of the principal bundle P) of an arbitrary monopole can be reconstructed from \mathcal{C}_A by compatibly gluing a pair of local models $A_1 \in \mathcal{C}_A(U_1)$ and $A_2 \in \mathcal{C}_A(U_2)$, i.e. by means of some gauge transformation $A_1 = A_2 + g_{12}^{-1}dg_{12}$, where $g_{12} : U_{12} \simeq S^1 \rightarrow U(1) \cong S^1$. In order to see how this works, it will be especially useful to recall from Section 4.1.1 that this isomorphism class corresponds to the monopole's topological charge $n \in \mathbb{Z}$, which can in turn be computed from its global force field F by means of the formula $n = \frac{1}{2\pi} \int_{S^2} F$. Then, through a judicious application of Stoke's theorem, it is easy to see that $n = \frac{1}{2\pi} \int_{S^1} (A_1 - A_2) = \frac{1}{2\pi} \int_{S^1} g_{12}^{-1}dg_{12}$. This is nothing other than the classical formula for the winding number (i.e. the number of times g_{12} winds S^1 around S^1), and so we see that the topological charge of a monopole is determined by the 'winding' (or homotopy class) of the particular choice of gauge transformation in $\overline{\mathcal{C}_A}(U_{12})$ that is used to glue A_1 and A_2 . It follows that in order to reconstruct monopole models with an arbitrary topological charge, our local theories need to contain all possible gauge transformations (and thus all possible homotopy classes), which is precisely the surplus* structure of \mathcal{C}_A as compared to $\mathcal{S}_{[A]}$.

By contrast, if we take \mathcal{E}_A as our local theory, then there is exactly one gauge transformation (and thus one possible choice of gluing) in $\overline{\mathcal{E}_A}(U_{12})$ between A_1 and A_2 . Recall from the definition of \mathcal{E}_A that this gauge transformation is described as $A_1 = A_2 + \lambda$, where λ is a closed 1-form. It thus follows that the *only* monopole models that one can construct by gluing local models are ones with the topological charge $n' = \frac{1}{2\pi} \int_{S^1} \lambda$ (where n' is possibly non-zero if λ is non-exact). To rephrase this crucial point: if the local theory does not have surplus* structure, i.e. the full set of gauge transformations, then all the other non-trivial monopole models (i.e. those with topological charge $n \neq n'$) cannot be constructed by gluing local models; this thus concludes our argument for the main claim: '(Richness) and (Locality) imply (Surplus*)'.

It is worth making two quick observations about this argument. First, we note that the argument holds up to categorical equivalence of the relevant local theories. In other words, the conclusion of the argument (it is not possible to simultaneously satisfy (Richness) and (Locality)) would have been the same if we had used \mathcal{S}_F or $\mathcal{S}_{[A]}$ as our local theories instead of \mathcal{E}_A ; and it would have been the same if we had used $\text{sk}\mathcal{C}_A$ instead of \mathcal{C}_A . Second, we note that one could of course advocate taking \mathcal{E}_A as the relevant local theory if one were willing to give up either (Richness) or (Locality). In other words, one could hold that the other monopole models are either physically irrelevant, or that they are non-local in the sense that they cannot be determined by the data of local theories. But by our lights, such a position is not only awkward but difficult to square with both physical practice and our intuitive understanding of fields as 'local' objects.

Let us now summarise our explanation of the claim that ‘surplus* structure is an essential feature of gauge theory’. The apparent tension in this statement stems from the idea that the surplus* structure is redundant structure and thus cannot be essential. Our resolution of this tension turns on our above distinction between two different notions of ‘theory’. With respect to the first and limited notion, which we called a ‘local theory’, surplus* structure is indeed redundant. However, with respect to the second richer and more realistic notion, which we called an ‘expanded theory’, such structure is necessary in order for the theory to perform the representational function jointly required by (Richness) and (Locality). Thus, the tension only arises if one mistakenly tries to shoehorn gauge theory into the mould of the first narrow notion.

5 Conclusion

Recall the motivation for the project of using category theory to represent scientific theories: a theory’s content is not exhausted by its collection of models; it includes information about the relationships *between* models. Category theory is a fruitful tool because it provides a natural framework in which to explicitly represent this information. For the most part, the philosophical literature has only implemented this idea by treating a theory as a category, i.e. its models are represented as objects, and its inter-model relationships are encoded in morphisms between objects. For instance, in the categories of Section 3, the relationship of representational equivalence between models (gauge fields) is encoded in the isomorphisms (gauge transformations) of the category.

The argument of our paper can be seen as stemming from two related observations. First, as we explained in Section 3, one important notion of ‘surplus’ in gauge theory applies to theories in which gauge fields (corresponding to the same F) are taken to be representationally equivalent, and the gauge transformations are themselves included as part of the putative redundancies. We also showed that one can give a categorical analysis of this notion, which we called ‘surplus*’. Second, representational equivalence is only one kind of inter-model relationship; in order to capture other relationships such as global-local relationships, it is natural to move to a ‘theories as functors’ framework, within which the ‘theories as categories’ framework can be embedded. In the context of gauge theory, such a move can be further motivated by the fact that field theories are typically taken to possess ‘locality’, a property that can be precisely represented within the ‘theories as functors’ framework. As we argued above, these two observations are related by the conventional wisdom that there is a tight conceptual link between the ‘surplus’ of gauge theory and representing the ‘locality’ of fields. Thus, an immediate pay-off of the ‘theories as functors’ framework is that it allows us to relate gauge theories on various spacetimes in order to address the following questions: (1) What is the correct analysis of ‘surplus’ for the ‘theories as categories’ of Section 3?, and (2) What is the representational role of this ‘surplus’ in gauge theory more generally? In Section 4, we used this relation to provide an independent check on our answer to (1), namely (Surplus*), thus confirming that our analysis really is responsive to desiderata stemming from physical practice, namely the need to simultaneously satisfy (Richness) and (Locality).

From a broader perspective that concerns the representational content of scientific theories, our argument here provides more evidence for the benefit of thinking about scientific theories in category theoretic terms. As we have seen, although it is the objects of the categories (the models) which directly represent possible systems, the morphisms can also feature in the representational content of $U(1)$ gauge theory. In particular: they play a role in representing how subsystems represented by objects of the theories-as-categories of Section 3 can be composed in the theories-as-functors framework. Consider, for example, the comparison between \mathcal{C}_A and \mathcal{E}_A . The objects in these categories are the same, so it is only the morphisms, in particular the non-

trivial automorphisms (and morphisms generated by these), which distinguish the two categories. And yet as our argument in the previous section showed, because of \mathcal{C}_A 's ability—and \mathcal{E}_A 's inability—to adequately represent how the local subsystems compose to form global systems (in a way that satisfies (Richness) and (Locality)), the theories have inequivalent representational content.

To conclude, we wish to point out a striking parallel between our discussion of how scientific theories should be represented—which involves classifying models/theories up to some standard of sameness—and a far-reaching approach to classification that has been developed within pure mathematics (especially algebraic geometry). The central concept here is that of a ‘moduli space’ \mathcal{M} of objects (cf. our local models), i.e. a space whose points are in 1-1 correspondence with isomorphism classes of such objects (thus classifying them in a naive sense), and which is in addition rich enough to encode the structure of *families* of such objects over a base space B (cf. our families of local models parameterised by a contractible cover of M) in the sense that such families will be in 1-1-correspondence with maps from B to \mathcal{M} .⁴⁸ Despite the fruitfulness of this notion, it was recognised very early on that such moduli spaces may not exist if the objects possess non-trivial automorphisms (cf. the non-trivial automorphisms of objects in \mathcal{C}_A), because the automorphisms will quite generically lead to various non-trivial families of objects over B , which will not be classified by maps from B into a candidate \mathcal{M} . Two standard strategies for addressing this problem are: (i) ‘rigidification’, i.e. simply omitting the non-trivial automorphisms from the objects; and (ii) defining a richer structure called a ‘moduli stack’, in which one does not merely assign (categorical) sets of objects to regions of the base space, but rather a groupoid that includes information about the non-trivial automorphisms of the objects, and thus allows one to keep track of the various non-trivial ways in which they ‘glue’ into families of objects.

It will be clear from this heuristic discussion that taking \mathcal{E}_A as one’s local theory constitutes a rigidification of the models of gauge theory, whereas taking \mathcal{C}_A as one’s local theory is a key step in passing to the richer framework of moduli stacks—thus, from a broader mathematical perspective, one way of phrasing the moral of our philosophical discussion is that, in gauge theory, there are very strong grounds for thinking of a collection of models as a ‘stack’.⁴⁹ For all that, there is still a disanalogy between these pure mathematical strategies and our physically-motivated discussion. In the mathematical case, to rigidify the objects is merely to simplify the problem so that one obtains a well-behaved moduli space. By contrast, in our case, to rigidify the models is to say that one does not care either about representing (Richness), or about representing (Locality).

Appendix

Proposition 3.2.1. *G forgets (only) structure*. So, with respect to G , \mathcal{S}_A has more structure* than $\mathcal{S}_{[A]}$.*

Proof. To see that G is faithful notice that for all A in \mathcal{S}_A , $\text{hom}(A, A) = \{1_A\}$, so the map $G_{A \rightarrow A} : \text{hom}(A, A) \rightarrow \text{hom}(G(A), G(A))$ has to be injective. For all A, A' in \mathcal{S}_A such that $A \neq A'$, $\text{hom}(A, A') = \emptyset$, and thus $G_{A \rightarrow A'}$ is also injective. So G is faithful. To see that G is essentially surjective notice that every $[A]$ in $\mathcal{S}_{[A]}$ is simply an equivalence class of gauge fields, and as such contains at least one A such that A is in \mathcal{S}_A and $G(A) = [A]$. To see that G fails to

⁴⁸For a history of the development of the ‘moduli space’ concept, see [A’Campo et al., 2016], and for an elementary introduction that stresses the notion of ‘classification’, see [Ben-Zvi, 2008].

⁴⁹For a detailed presentation of how to describe a stack of gauge fields, see [Benini et al., 2017].

be full consider an equivalence class of gauge fields $[A]$ in $\mathcal{S}_{[A]}$ that contains distinct gauge fields A, A' (which are in \mathcal{S}_A). Since \mathcal{S}_A contains only the identity morphisms, $\text{hom}(A, A') = \emptyset$. But, $G(A) = G(A') = [A]$, so $\text{hom}(G(A), G(A')) = \{1_{[A]}\}$. So the map $G_{A \rightarrow A'}$ is not surjective.⁵⁰ \square

Proposition 3.2.2. *K forgets (only) structure*. So, with respect to K , $\mathcal{S}_{[A]}$ has more structure* than \mathcal{C}_A .*

Proof. K is faithful for an analogous reason to G : for every object $[A]$ in $\mathcal{S}_{[A]}$, $\text{hom}([A], [A]) = \{1_{[A]}\}$ and every pair of distinct objects $[A]$ and $[A]'$, $\text{hom}([A], [A]') = \emptyset$. Thus both $K_{[A] \rightarrow [A]}$ and $K_{[A] \rightarrow [A]'}$ have to be injective. To see that K is essentially surjective notice that every A in \mathcal{C}_A is in some equivalence class $[A]$ in $\mathcal{S}_{[A]}$ and is isomorphic to every A' which is also in that equivalence class. So either $K([A]) = A$, or $K([A]) = A'$ such that $A \cong A'$ in \mathcal{C}_A . This suffices to show that K is essentially surjective. To see that K is not full consider an arbitrary $[A]$ in $\mathcal{S}_{[A]}$. By construction $\text{hom}([A], [A]) = \{1_{[A]}\}$. But, by the definition of the morphisms in \mathcal{C}_A (given by definition (1)), $\text{hom}(K([A]), K([A]))$ contains automorphisms which are not $1_{K([A])}$. So the map $K_{[A] \rightarrow [A]}$ is not surjective. \square

Proposition 3.2.3. *τ forgets (only) stuff. So, with respect to τ , \mathcal{C}_A has more stuff than $\mathcal{S}_{[A]}$.*

Proof. τ is essentially surjective for the same reason as G . To see that τ is full we consider an arbitrary pair of objects A, A' in \mathcal{C}_A . If $A \not\cong A'$ in \mathcal{C}_A then $\tau(A) \neq \tau(A')$, in which case $\tau_{A \rightarrow A'} : \emptyset \rightarrow \emptyset$. If either $A = A'$, or $A \cong A'$ in \mathcal{C}_A , then $\tau(A) = \tau(A')$, so $\text{hom}(\tau(A), \tau(A')) = \{1_{\tau(A)}\}$. So in each case the map $\tau_{A \rightarrow A'}$ is surjective. To see τ is not faithful consider an arbitrary object A in \mathcal{C}_A . By definition (1), $\text{hom}(A, A)$ contains automorphisms which are not 1_A , and by the definition of $\mathcal{S}_{[A]}$, $\text{hom}(\tau(A), \tau(A)) = \{1_{\tau(A)}\}$. So the map $\tau_{A \rightarrow A}$ is not injective. \square

Proposition 3.2.4. *M forgets (only) properties. So, with respect to M , $\mathcal{S}_{[A]}$ has more properties than \mathcal{S}_A .*

Proof. To see that M is full and faithful note that both the domain and codomain categories contain only identity morphisms. If $[A] \neq [A]'$ then $M_{[A] \rightarrow [A]'} : \emptyset \rightarrow \emptyset$ and if $[A] = [A]'$ then $M_{[A] \rightarrow [A]} : \{1_{[A]}\} \rightarrow \{1_{M([A])}\}$; therefore, for every $[A], [A]'$ in $\mathcal{S}_{[A]}$, $M_{[A] \rightarrow [A]'}$ is a bijection. To see M is not essentially surjective note that by the definition of M , for each $[A]$ in $\mathcal{S}_{[A]}$, $M([A]) = \text{some } A \text{ in } \mathcal{S}_A \text{ such that } A \in [A]$. By construction there is also an $A' \neq A$ in \mathcal{S}_A such that $A' \in [A]$, but $\text{hom}(A, A') = \emptyset$. Furthermore, if one takes a distinct equivalence class $[A^*]$ in $\mathcal{S}_{[A]}$ where $A' \notin [A^*]$, then $M([A^*]) \not\cong A'$ in \mathcal{S}_A . So A' is not isomorphic to anything in the image of M . \square

Acknowledgements

Authors are listed in alphabetical order, the paper is fully collaborative. We are grateful to two anonymous referees for useful comments. We would also like to thank Jeremy Butterfield, Bryan Roberts, Alexander Schenkel, Christopher Schommer-Pries, and Stephan Stolz for helpful discussions and comments on previous drafts. Thanks also to audiences at the following workshops and conferences: *The Physics Old and New: Historical and Critical Perspectives on Physics and the Philosophy of Nature* at the University of Notre Dame, *Symmetries and Asymmetries in Physics* at Leibniz University in Hannover, and the *Ninth European Congress of Analytic Philosophy* in Munich.

⁵⁰Notice that, since $\mathcal{S}_F \cong \mathcal{S}_{[A]}$, G also shows that \mathcal{S}_A has more structure* than \mathcal{S}_F . The proof of Proposition 3.2.1 follows the proof of Proposition 1 in [Weatherall, 2016b, p. 1044-1045]. This should make clear how our $\mathcal{S}_{[A]} \cong \mathcal{S}_F$ and \mathcal{S}_A are analogues of his **EM**₁ and **EM**₂ respectively.

James Nguyen
History and Philosophy of Science
University of Notre Dame
Notre Dame, USA
and
Centre for Philosophy of Natural and Social Science
London School of Economics and Political Science
London, UK
jnguyen4@nd.edu

Nicholas J. Teh
Department of Philosophy
University of Notre Dame
Notre Dame, USA
nteh@nd.edu


Laura Wells
Department of Mathematics
University of Notre Dame
Notre Dame, USA
lwells@nd.edu

References


- A'Campo, N., Ji, L., and Papadopoulos, A. (2016). On the early history of moduli and teichmüller spaces, <https://arxiv.org/abs/1602.07208>.
- Alvarez-Gaume, L. and Vazquez-Mozo, M. A. (2011). *An Invitation to Quantum Field Theory*. Lecture Notes in Physics. Springer, Berlin.
- Baez, J., Bartel, T., and Dolan, J. (2004). Property, structure, and stuff. Quantum Gravity Seminar, University of California, Riverside, <http://math.ucr.edu/home/baez/qg-spring2004/discussion.html>.
- Baez, J. and Shulman, M. (2007). Lectures on n-categories and cohomology, <https://arxiv.org/abs/math/0608420>.
- Barnet-Lamb, T. (2005). Minor thesis: from fibered categories to algebraic stacks, <http://people.brandeis.edu/~tbl/minor-thesis.pdf>.
- Belot, G., Earman, J., Healey, R., Maudlin, T., Nounou, A., and Struyve, W. (2009). Synopsis and discussion: Philosophy of gauge theory, <http://philsci-archive.pitt.edu/4728/>.
- Ben-Zvi, D. (2008). Moduli spaces. In Gowers, T., Barrow-Green, J., and Leader, I., editors, *The Princeton Companion to Mathematics*, pages 408–419. Princeton University Press.
- Benini, M., Schenkel, A., and Schreiber, U. (2017). The stack of Yang-Mills fields on Lorentzian manifolds, <https://arxiv.org/abs/1704.01378>.
- Benini, M., Schenkel, A., and Szabo, R. J. (2015). Homotopy colimits and global observables in abelian gauge theory. *Lett. Math. Phys.*, 105:1193–1222.

- Connes, A. (1994). *Noncommutative Geometry*. Academic Press, San Diego, CA.
- Dougherty, J. (forthcoming). Separability in gauge theories. *Philosophy of Science*.
- Earman, J. (2004). Laws, symmetry, and symmetry breaking: Invariance, conservation principles, and objectivity. *Philosophy of Science*, 71(5):1227–1241.
- Frigg, R. and Nguyen, J. (2016). Scientific representation. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.
- Grothendieck, A. (1959/60). Descent technique and theorems of existence in algebraic geometry. i: G in general. descent by flat morphisms. (technique of descent and existence theorems in algebraic geometry). *Sem. Bourbaki 12 (1959/60), No. 190, 29 p. (1960)*.
- Halvorson, H. (2012). What scientific theories could not be. *Philosophy of Science*, 79(2):183–206.
- Halvorson, H. and Tsementzis, D. (2015). Categories of scientific theories, <http://philsci-archive.pitt.edu/11923/>.
- Healey, R. (2007). *Gauging What's Real*. Oxford University Press.
- Lurie, J. (2009). *Higher topos theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ.
- Manton, N. and Sutcliffe, P. (2004). *Topological Solitons*. Cambridge monographs on mathematical physics. Cambridge University Press, Cambridge.
- Pokorski, S. (2000). *Gauge Field Theories*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2nd edition.
- Redhead, M. (2003). The interpretation of gauge symmetry. In Brading, K. A. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 124–139. Cambridge University Press.
- Riehl, E. (2016). *Category theory in context*. Courier Dover Publications, Mineola, New York.
- Schreiber, U. (2016). Higher prequantum geometry, <https://arxiv.org/abs/1601.05956>.
- Sontz, S. (2015). *Principal Bundles: The Classical Case*. Springer International Publishing, Cham, Heidelberg, New York, Dordrecht, and London.
- Steenrod, N. (1951). *The topology of fibre bundles*. Princeton Mathematical Series 14. Princeton University Press.
- Sundermeyer, K. (2014). *Symmetries in Fundamental Physics*. Fundamental Theories of Physics. Springer International Publishing, Cham, Heidelberg, New York, Dordrecht, and London.
- Teh, N. J. (forthcoming). Gauge theory. In Knox, E. and Wilson, A., editors, *The Routledge Companion to the Philosophy of Physics*. Routledge.
- Teh, N. J. and Tsementzis, D. (2017). Theoretical equivalence in classical mechanics and its relationship to duality. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 59(Supplement C):44–54.


- Tong, D. (2006). Quantum field theory. Lecture notes, University of Cambridge Part III Mathematical Tripos.
- Weatherall, J. O. (2016a). Are Newtonian gravitation and geometrized Newtonian gravitation theoretically equivalent? *Erkenntnis*, 81(5):1073–1091.
- Weatherall, J. O. (2016b). Understanding gauge. *Philosophy of Science*, 83(5):1039–1049.
- Witten, E. (2010). Mirror Symmetry, Hitchin’s Equations, And Langlands Duality. In Freed, D. S., Hopkins, M. J., and Teleman, C., editors, *The Many Facets of Geometry: A Tribute to Nigel Hitchin*, pages 113–128. Oxford University Press.




(a) \mathcal{S}_F




(b) \mathcal{S}_A



(c) \mathcal{E}_A



(d) \mathcal{C}_A



(e) $\mathcal{S}_{[A]}$

Figure 1: Categories of Gauge Theories




Figure 2: Functorial Relationships