

Excess Structure in the Constrained Hamiltonian Formalism

Clara Bradley

Department of Philosophy, University College London

clara.bradley@ucl.ac.uk

Abstract

Gauge theories are often characterized as possessing ‘redundancy’ or ‘excess structure’. This, in turn, motivates *reducing* the gauge symmetries, commonly through ‘symplectic reduction’ in the Hamiltonian framework. However, there are multiple ways to formulate a Hamiltonian gauge theory. This paper examines the relationship between the formulation of a Hamiltonian gauge theory and the attribution of excess structure. I argue that one can formulate a Hamiltonian gauge theory such that symplectic reduction does not remove structure, and therefore that the role of symplectic reduction cannot be purely to remove ‘excess structure’. I discuss in what sense symplectic reduction is thereby motivated.

1 Introduction

There are several cases in physics where the mathematical formalism of a theory appears to have redundancy: there are objects in the theory that don’t seem to play any role in capturing the empirical content of the theory. There have been attempts to formalize the idea of a theory having redundancy through the notion of “excess structure”¹: a theory has excess structure when it has distinct, non-isomorphic models that are taken to represent the same physical situation. And what it is to remove excess structure is to come up with a new theory with *less* structure—in particular, without the structure distinguishing the problematic models—that captures the same empirical content. This notion has been fruitfully applied to several cases to understand the sense in which a theory does, or does not, have redundancy: from formulations of classical mechanics (North (2009), Barrett

¹Sometimes the term “surplus structure” is used, following Redhead (1975).

(2015*a*)), to formulations of spacetime theories (Barrett (2015*b*), Weatherall (2018)), to formulations of electromagnetism (Weatherall (2016*b*)).

One context where the term ‘redundancy’ is used by both physicists and philosophers is in the context of gauge theories; indeed, sometimes a ‘gauge theory’ is used to mean just a theory with redundancy.² And this is not just an observation; taking gauge theories to have redundancy motivates the procedure of ‘reduction’, which involves reformulating a gauge theory by effectively removing the ‘excess’ quantities from the formalism.

On the other hand, Weatherall (2016*b*) argues that some theories that are called ‘gauge theories’—namely, (the standard formulations of) Yang-Mills theory and General Relativity—are not theories with excess structure, in the sense defined above. The reason is that the ‘gauge symmetries’ act as *isomorphisms* in the theory, and so they do not give rise to distinct models of the theory. It follows that whether a gauge theory is a candidate for excess structure depends on the particular formulation of the theory, and therefore, whether reduction is motivated by considerations of excess structure depends on how the gauge theory that one reduces is formulated.

A standard procedure for reduction, known as ‘symplectic reduction’, takes place in the Hamiltonian formulation of gauge theories. The Hamiltonian formalism provides a useful framework for formulating gauge theories because, as originally shown by Dirac (1964), one can connect gauge symmetry to constraints on the Hamiltonian variables.³ There is, therefore, a natural question to ask: is it the case that the Hamiltonian formulation of gauge theories has excess structure? While many discussions on the Hamiltonian formulation of gauge theories suggest that this is the case⁴, there has not been an explicit attempt to formalize this claim. The aim of this paper is to rectify that.

The question is complicated by the fact that there are disputes about how to characterize a Hamiltonian gauge theory—in particular, which transformations constitute the gauge transformations—that comes down to a disagreement about the equivalence class of Hamiltonians that generate the dynamics. The standard view is that the ‘Extended Hamiltonian’ is the right equivalence class of Hamiltonians and that gauge transformations are generated by arbitrary combinations of first-class constraints. The alternative view—argued for, most prominently, by Pitts (2014*b*)⁵—is that the ‘Total Hamiltonian’ is the right equivalence class of Hamiltonians and that gauge transformations are generated by a

²For example, this is suggested by Earman (2004), Rovelli (2014), Ismael & van Fraassen (2003).

³This is not to say that the Hamiltonian formalism provides the best or most general framework for formulating gauge theories. However, it is important because symplectic reduction provides the basis for standard quantization techniques.

⁴For example, see Thébault (2012), Earman (2002*a*, 2004), Belot (2003).

⁵See also Pitts (2014*a*, 2022, 2024) as well as Pons (2005).

particular combination of first-class constraints. In attempting to answer the question above, we will therefore consider whether different characterizations of a Hamiltonian gauge theory lead to different answers to this question.

In fact, I will show that the formulation of a Hamiltonian gauge theory that takes the Extended Hamiltonian to be the equivalence class of Hamiltonians has *less structure* than the formulation that takes the Total Hamiltonian to be the equivalence class of Hamiltonians. This suggests that the theory formulated in terms of the Total Hamiltonian is one with excess structure, and that the theory formulated in terms of the Extended Hamiltonian is precisely a way of removing this excess structure. This is interesting in itself; it suggests that an argument for the standard view can be made in terms of considerations of excess structure. Moreover, it opposes a recent argument made by Pitts (2022, 2024) that the Extended Hamiltonian is (at best) a “trivial reformulation” of the Total Hamiltonian.

With the theory in terms of the Extended Hamiltonian as a starting point, we then ask whether this is a theory with excess structure by considering whether symplectic reduction can be said to remove structure from this theory. I show that symplectic reduction does not, in fact, result in a theory with less structure than the theory characterized by the Extended Hamiltonian. This might appear surprising, given that the theory characterized by the Extended Hamiltonian has symmetries of precisely the kind that are removed via symplectic reduction. However, I will argue that this appearance results from a misunderstanding of the relationship between the number of gauge symmetries a theory has, the amount of structure it has, and whether it has redundancy.

More carefully, I will argue that there is a way to characterize the kind of redundancy that a gauge theory has compared to the symplectic reduced theory, but that it is conceptually distinct to the kind of redundancy at issue in the context of excess structure. This kind of redundancy is rather one that Bradley & Weatherall (2020) call “representational redundancy”: the redundancy comes from the fact that the theory has the freedom to represent a particular physical situation in lots of different ways. Applying the notion of representational redundancy in the context of Hamiltonian gauge theories helps to clarify the motivations for moving to the symplectic reduced theory, and how these motivations differ from the motivations for moving to a theory without excess structure.

The rest of the paper will go as follows. In Section 2, I discuss the notion of ‘excess structure’ and how the tools of category theory can be usefully applied to formalize this notion. In Section 3, I give an overview of the constrained Hamiltonian formalism for describing gauge theories. In Section 4, I argue that the theory corresponding to the Extended Hamiltonian removes excess structure from

the theory corresponding to the Total Hamiltonian, using the tools described in Section 2. In Section 5, I turn to the relationship between the Extended Hamiltonian theory and the theory obtained via symplectic reduction. I show that the latter does not remove excess structure from the former, and I discuss the sense in which symplectic reduction is nonetheless motivated. In Section 6, I conclude with some broader lessons.

2 Defining Excess Structure

As a starting point, we will take a theory to have ‘excess structure’ whenever it admits non-isomorphic models that are taken to represent the same physical situation. This allows for various refinements.

First, one might argue that if a theory contains distinct models that are taken to represent the same physical situation, yet no plausible alternative theory captures the same empirical content *without* such distinct models, then one shouldn’t regard the theory as having excess structure.⁶ The underlying thought is that structure counts as ‘excess’ only if it can, in principle, be dispensed with. On this view, a theory has excess structure only when there exists another theory with *less* structure that nonetheless captures the same empirical content. In what follows, we will at least presume that presenting such an alternative theory is a way of determining that the first theory has excess structure.

Second, there are different ways that one might spell out what it is for two models to represent the same physical situation, or what it is for two different theories to have the same empirical content. Some philosophers have tried to narrow down the sense in which two models represent the same physical situation through an account of ‘symmetries’ between models of a theory.⁷ However, there are challenges with providing a general account of this kind,⁸ and for our purposes here, it suffices that there is a way to spell out the sense in which two models or theories are empirically equivalent on a case-by-case basis. For example, Barrett (2019) discusses how empirical equivalence between (hyperregular) Lagrangian and Hamiltonian formulations of a theory can be captured in terms of agreement about the collection of integral curves that represent solutions to the equations of motion, under the map that takes a model of Lagrangian mechanics to one of Hamiltonian mechanics and vice versa. Later, when we discuss Hamiltonian gauge theories, we will similarly specify the relevant sense in which the theories in question share the same ‘empirical content’. This leaves open that

⁶This is sometimes expressed as a ‘motivational’ as opposed to ‘interpretational’ view of symmetries, following Møller-Nielsen (2017).

⁷For example, see Baker (2010), Ismael & van Fraassen (2003), Earman (2002*b*), Dewar (2019).

⁸For discussion of some of these challenges, see Belot (2013).

there might be other ways of comparing and interpreting the theories such that they are not regarded as empirically equivalent. Indeed, we will return to the importance of the choice of comparison in determining whether two theories are empirically equivalent in Section 4.

Finally, to say whether one theory has excess structure over another theory, we need a way to compare the structure of two theories. First, we can ask what it means for one *model*—understood as a single mathematical object—to possess more structure than another model. A criterion for comparing the amount of structure of two mathematical objects has recently been proposed by Barrett (2015*a,b*):

(**SYM***): A mathematical object X has at least as much structure as an object Y if
 $Aut(X) \subseteq Aut(Y)$

The intuitive idea behind this criterion is that the amount of structure a mathematical object possesses is related to its *isomorphisms*, since an isomorphism is, by definition, a structure-preserving map. In particular, the *fewer* the isomorphisms that a mathematical object has, the *more* structure it has. And vice versa—more isomorphisms indicate less structure.

However, there are some limitations of this account. First, it only applies when one can compare the set of automorphisms of X and Y , which itself relies on X and Y being similar in a certain way, such as having the same underlying domain.⁹ Second, in order to say whether one theory has excess structure over another, what we care about is not just whether a single model of one theory has more structure than a single model of the other, but whether one theory as a whole—understood as a collection of models with structure-preserving maps between these models—has more structure than the other.

To extend the above criterion to theories more generally, we can make use of the tools of category theory.¹⁰ To think about a physical theory categorically, we take the objects of the category to be the models of the theory, and the arrows between objects to be the isomorphisms between models of the theory. Relations between theories are described by functors between the categories representing those theories, where a functor is a structure-preserving map between categories that takes objects to objects and arrows to arrows.

A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ from the category $\mathcal{C}$ to the category $\mathcal{D}$ is said to be *full* if for every pair of objects A, B of $\mathcal{C}$ the map $F : \text{hom}(A, B) \rightarrow \text{hom}(F(A), F(B))$ induced by F is surjective, where

⁹For further discussion, see Barrett et al. (2023).

¹⁰See, for example, Weatherall (2016*b*), Nguyen et al. (2020), Bradley & Weatherall (2020). For a defense of using category theory to represent theories more generally, see Halvorson (2012, 2016), Halvorson & Tsementzis (2017), Weatherall (2016*a*, 2017).

$\text{hom}(A, B)$ is the collection of arrows from A to B . Similarly, F is said to be *faithful* if for every pair of objects the induced map on arrows is injective. Finally, F is said to be *essentially surjective* if for every object X of $\mathcal{D}$, there is some object A of $\mathcal{C}$ such that $F(A)$ is isomorphic to X . Using this terminology, we say (following Weatherall (2016b)) that a theory represented by category $\mathcal{C}$ has more structure than a theory represented by category $\mathcal{D}$ if a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is not full (but is faithful and essentially surjective). In this case, we say that F *forgets (only) structure*.

We therefore say that a theory, represented by a category $\mathcal{C}$, has ‘excess structure’ relative to another theory $\mathcal{D}$, when there is a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ that forgets (only) structure while preserving empirical content. The intuitive reason that such a formal notion captures what it is for one theory to have excess structure over another is that a functor failing to be full means that there are some isomorphisms in the second theory that are not mapped to by any isomorphisms in the first theory. In other words, the isomorphisms of the second theory are ‘wider’ than the isomorphisms of the first. If we take on board the principle that *more* isomorphisms indicates *less* structure, this suggests that a functor failing to be full captures precisely the idea that we have moved to a theory with overall less structure.

However, there are other interesting relationships between theories, represented as categories. For example, if the functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is not faithful (but full and essentially surjective), we say that F *forgets (only) stuff*. To draw out the interpretation of a theory having more ‘stuff’, it is helpful to consider an example.¹¹ Take a category **Vect** whose objects are two-dimensional vector spaces and whose arrows are linear transformations. Similarly, take a category **OBVect** whose objects are two-dimensional vector spaces with a fixed, ordered basis, with basis-preserving maps as arrows. Then, there is a functor from **Vect** to **OBVect** that takes all objects in **Vect** to a single object in **OBVect**, and all arrows in **Vect** to the identity map on the chosen object in **OBVect**. This functor is full and essentially surjective, but not faithful, and so it forgets *stuff*. The reason is that the objects in **Vect** have multiple automorphisms that correspond, via the functor, to the identity map on the corresponding model in **OBVect**. We can understand this difference as characterizing the fact there are multiple ways that one can ‘use’ the objects of **Vect** to represent some ordered basis (such as a choice of ‘North’ and ‘East’) via its non-trivial automorphisms, whereas there is no corresponding freedom using the objects of **OBVect**, because one ordered basis is privileged from the rest.

Does this correspond to a kind of redundancy? One relevant observation is that there also exists

¹¹The following example is taken from Bradley & Weatherall (2020).

a functor in the opposite direction, from **OBVect** to **Vect**, that forgets only *structure*. Indeed, this captures the fact that the objects of **OBVect** are just those of **Vect** plus the structure of a fixed ordered basis. So in this case, the theory that has more ‘stuff’ is really the one with less structure. We therefore might see the additional ‘stuff’ that **Vect** has as being a feature of the fact that it is a theory *without* redundancy.¹²

On the other hand, one might argue that there *is* a kind of redundancy associated with the fact that there are multiple automorphisms of the models of **Vect**: it indicates that there is nothing privileged about one vector over any linear transformation of that vector, and yet such vectors correspond to distinct vectors in the vector space. Bradley & Weatherall (2020) label the idea that a model can represent some object or physical situation in multiple ways via its non-trivial automorphisms “representational redundancy”. One might argue that removing this kind of redundancy can be achieved by moving to **OBVect** or, depending on one’s representational concerns, by moving to a category whose objects are equivalence classes of vectors under the linear transformations, where the only automorphisms are the identity map.

This suggests that there are different ways one could interpret the fact that one category has more ‘stuff’ than another. It might be that the first is regarded as a theory without excess structure, or it might be that the first is regarded as a theory with redundancy, as captured by ‘representational redundancy’. Which is regarded as the ‘right’ interpretation might depend on the particular representational concerns one has in some context. As we will discuss in Section 5, this is precisely the situation in the context of *symplectic reduction*. However, to make the point in this context, we first need to discuss how one should formulate (the structure of) a Hamiltonian gauge theory. To do this, we will make use of the “constrained Hamiltonian formalism”, to which we now turn.

3 The Constrained Hamiltonian Formalism

The standard geometric way of formulating a Hamiltonian theory is to take its state space to be a *symplectic manifold*.¹³ A symplectic manifold consists of a pair (M, ω) where M is a smooth manifold and ω is a *symplectic form*: it is a two-form (a smooth, anti-symmetric tensor field of rank (0,2)) that is closed and non-degenerate. A Hamiltonian theory is a theory whose state space is given by the cotangent bundle of configuration space T^*Q i.e. the points of the manifold are $\{(q_i, p_i), i = 1, \dots, N\}$

¹²Indeed, this is exactly what is argued for in Bradley & Weatherall (2020).

¹³For details, see Abraham & Marsden (1987).

where q_i are the generalized positions and p_i the generalized (canonical) momenta. T^*Q comes naturally equipped with a one-form, the *Poincaré one-form*, given by $\theta = p_i dq^i$. The corresponding symplectic two-form is given by $\omega = \mathbf{d}\theta = dp_i \wedge dq^i$.

Given a function $f : T^*Q \rightarrow \mathbb{R}$, one can uniquely define a smooth tangent vector field X_f through:

$$\omega(X_f, \cdot) = \mathbf{d}f \quad (1)$$

where $\{\cdot\}$ represents any vector field tangent to T^*Q . An important function is the *Hamiltonian* $H(q_i, p_i)$, where the unique vector field associated with H via the above equation is interpreted as the generator of the solutions to the equations of motion i.e. the equations of motion can be expressed as:

$$\omega(X_H, \cdot) = \mathbf{d}H \quad (2)$$

The “constrained Hamiltonian formalism” concerns a collection of Hamiltonian theories where there are constraints of the form $\phi_m(q_i, p_i) \approx 0$ for $m = 1, \dots, M$ where M is the number of constraints and the meaning of $\approx$ is that one can substitute the left-hand side for the right only on the subspace where the equation holds.¹⁴ Constraints that arise from the definition of the canonical momenta $p_i = \frac{\partial L}{\partial \dot{q}_i}$ for Lagrangian function L are called the *primary* constraints. We can think of the primary constraints as defining a submanifold of T^*Q , which we call the *primary constraint surface*, Σ_p .

The *Total Hamiltonian* is the equivalence class of Hamiltonians defined as $H_T = H + u^m \phi_m$ (using the Einstein summation convention) where $u^m(q_i, p_i)$ are arbitrary functions, which is equivalent to H when one restricts to the points of Σ_p . Taking the Total Hamiltonian to generate the dynamics on T^*Q means that there are multiple possible solutions given some initial state. In other words, it leads to the failure of a well-posed initial value problem. This apparent indeterminism is the root of gauge symmetry in the Hamiltonian context: gauge symmetries are understood to relate states or solutions that are equivalent from the perspective of restoring unique evolution.

We can alternatively see the underdetermination in the dynamics by formulating the theory intrinsically on Σ_p . To do this, we define an induced two-form $\tilde{\omega}_p = i^* \omega$, where $i : \Sigma_p \rightarrow T^*Q$ is the inclusion map. We say that a (primary) constraint is *first-class* when its associated vector field is tangent to Σ_p (and *second-class* when it is not tangent to Σ_p). Whenever there are first-class primary constraints, $\tilde{\omega}_p$ is *presymplectic*, which means that it is degenerate: it has null vector fields—vector

¹⁴This formalism originates with Dirac (1964) and Bergmann (1961). For details, see Henneaux & Teitelboim (1994).

fields X that satisfy $\tilde{\omega}_p(X, \cdot) = 0$ —that correspond to the primary first-class constraints. This means that the equations of motion, written intrinsically as $\tilde{\omega}_p(X_H, \cdot) = dH|_{\Sigma_p}$, have non-unique solutions: they are unique up to arbitrary combinations of vector fields associated with the primary first-class constraints.

However, the theory formulated intrinsically on Σ_p may be such that the solutions to the equations of motion are not well-defined everywhere i.e. the vector fields X_H that define the solutions to this equation may not be tangent to the primary constraint surface everywhere. For the solutions to be tangent to the constraint surface, it must be that $\tilde{\omega}(X_H, X_{\phi_a}) = dH(X_{\phi_a}) = 0$ for vector fields X_{ϕ_a} associated with the primary constraints. This may define a further collection of constraints called *secondary constraints*, and we can think of these additional constraints as leading to the specification of a further submanifold. Continuing this process of requiring that the solutions to the equations of motion be tangent to the constraint surface terminates in a *final constraint surface*, $(\Sigma_f, \tilde{\omega}_f)$, with dynamics given in terms of $H|_{\Sigma_f}$ via $\tilde{\omega}_f(X_H, \cdot) = dH|_{\Sigma_f}$, which has solutions everywhere that are unique up to an arbitrary combination of the vector fields associated with the primary and secondary first-class constraints. The integral curves of the null vector fields of $\tilde{\omega}_f$ are sometimes called the *gauge orbits*.

We can alternatively think of the dynamics on the final constraint surface in terms of solutions on T^*Q where, instead of taking the Total Hamiltonian to generate the dynamics, one takes the *Extended Hamiltonian* to generate the dynamics, defined as $H_E = H_T + w^j \varphi_j$ where φ_j are the first-class secondary constraints and $w^j(q_n, p_n)$ are arbitrary functions. Indeed, Dirac (1964) argued that the Extended Hamiltonian is the ‘right’ equivalence class of Hamiltonians, and correspondingly that the gauge transformations are generated by arbitrary combinations of first-class constraints, which connect solutions generated by the Extended Hamiltonian. This is now the standard view.

However, some have found this position unfounded; in particular, Pitts (2014b) argues that the Extended Hamiltonian is unmotivated because he takes the *Total* Hamiltonian to correctly generate the dynamics, and so he takes the gauge transformations between solutions to be more restrictive than the transformations relating solutions to the Extended Hamiltonian. One can take Pitts to be arguing that a Hamiltonian theory should be formulated intrinsically on the primary constraint surface, and not on the final constraint surface.

There are two reactions that one could have to the disagreement between the standard view and that of Pitts. First, one could maintain that the question of whether arbitrary combinations of first-

class constraints generate gauge transformations or not comes down to whether the Total Hamiltonian or Extended Hamiltonian is considered the right equivalence class of Hamiltonians, and this needs to be independently settled. Second, one could maintain that what the debate shows is that there is no Hamiltonian-independent way to characterize the gauge transformations, but that we can think of these different forms of the Hamiltonian as equivalent, and so there is no conflict.¹⁵ We will presently argue that the first option is right: there is a way to settle the question of which equivalence class of Hamiltonians should be regarded as the right one through considerations of excess structure.

4 The Extended Hamiltonian Removes Excess Structure

We have already seen that there is a sense in which the Total Hamiltonian and the Extended Hamiltonian naturally correspond to the theories formulated intrinsically on the primary constraint surface and the final constraint surface, respectively. However, given the inconsistencies that arise with the intrinsic dynamics on Σ_p , the corresponding theory will need to be such that it essentially takes the dynamics to be defined only at those points that satisfy the secondary constraints. We will therefore think of the theory characterized by the Total Hamiltonian as being defined geometrically in terms of the structure $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ where H is the Hamiltonian restricted to the primary constraint surface and φ_i are the secondary constraints, and we will take the dynamics of this theory to be given by two equations: $\tilde{\omega}_p(X_H, \cdot) = \mathbf{d}H$ and $\varphi_i = 0$. Similarly, we can think of the theory described by the Extended Hamiltonian as characterized geometrically by the structure $(\Sigma_f, \tilde{\omega}_f, H)$ where H is Hamiltonian restricted to the final constraint surface, and the dynamics are given by the single equation $\tilde{\omega}_f(X_H, \cdot) = \mathbf{d}H$.

Characterizing the theories in this way naturally provides a sense in which the Extended Hamiltonian theory regards more solutions as equivalent compared to the Total Hamiltonian theory. Consider the vector fields of the form $X_H + a^j X_{\varphi_j}$ where X_{φ_j} are the vector fields associated with the secondary first-class constraints φ_j . These vector fields correspond to equivalent solutions according to the Extended Hamiltonian theory, since the solutions to $\tilde{\omega}_f(X_H, \cdot) = \mathbf{d}H$ are unique up to arbitrary combinations of vector fields associated with the first-class constraints. On the other hand, they are distinct solutions according to the Total Hamiltonian theory, in the sense that they are solutions to distinct Hamiltonian functions on Σ_p . The reason is that X_{φ_j} are not null vector fields of $\tilde{\omega}_p$, while they are of $\tilde{\omega}_f$.

¹⁵I take this to be one way of interpreting the position maintained by Pitts (2022, 2024).

However, in another sense, these theories are *empirically equivalent*. The reason is that the solutions that Total Hamiltonian theory distinguishes between that the Extended Hamiltonian theory does not are solutions that correspond to distinct Hamiltonian functions, according to the Total Hamiltonian theory, only at points *off* the final constraint surface. Since the points that do not lie on the final constraint surface are regarded as empirically inaccessible according to both theories, this suggests that by the lights of *both* theories, there is no way to empirically distinguish such solutions—they are generated by Hamiltonians that are equivalent at all points that are empirically accessible.¹⁶

We can therefore say that the theories agree about empirical content but disagree about which curves on the state space correspond to (physically) equivalent solutions. This suggests that we are precisely in a situation of ‘excess structure’: the Total Hamiltonian theory draws distinctions that one does not need to make in order to capture the empirical content of the theory.

Let us spell this out precisely, starting with the idea that a Total Hamiltonian theory has fewer symmetries, or isomorphisms, than the associated Extended Hamiltonian theory. Take the symmetries of the Total Hamiltonian theory to be given by diffeomorphisms of Σ_p , $f : \Sigma_p \rightarrow \Sigma_p$ that preserve the two form $\tilde{\omega}_p$ in the sense that $f^*(\tilde{\omega}_p) = \tilde{\omega}_p$ and preserves the Hamiltonian H and secondary constraints φ_i in the sense that $f^*(H) = H$ and $f^*(\varphi_i) = \varphi_i$. Similarly, take the symmetries of the Extended Hamiltonian theory to be given by diffeomorphisms of Σ_f , $g : \Sigma_f \rightarrow \Sigma_f$ such that $g^*(\tilde{\omega}_f) = \tilde{\omega}_f$ and $g^*(H) = H$. In other words, the symmetries are just the automorphisms of the structures $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ and $(\Sigma_f, \tilde{\omega}_f, H)$ respectively. Notice that these symmetries preserve solutions to the equations of motion since the solutions are determined by the two-form and the Hamiltonian. Then:

Proposition 1 *The automorphisms of $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$, restricted to their action on Σ_f , are a subset of the automorphisms of $(\Sigma_f, \tilde{\omega}_f, H)$.*¹⁷

The core part of the proof of this proposition is contained in the fact that the null vector fields of $\tilde{\omega}_p$ are a subset of the null vector fields of $\tilde{\omega}_f$ (when we consider the action of these vector fields on the final constraint surface) and arbitrary transformations along the null vector fields are automorphisms of the respective structure. This shows that there are ‘fewer’ isomorphisms of the Total Hamiltonian theory compared to the Extended Hamiltonian theory. As we pointed out in Section 2, having *more* isomorphisms is associated with having *less* structure. Therefore, the above proposition suggests

¹⁶For further discussion of whether the Extended Hamiltonian can be empirically distinguished from the Total Hamiltonian, see Pitts (2014b) and the response by Pooley & Wallace (2022).

¹⁷See A.1 for proof.

that the Total Hamiltonian theory should be thought to have more structure than the Extended Hamiltonian theory.¹⁸

We can spell out the comparison of structure more precisely using the tools from category theory:

Define the category **TotHam** (the category of Total Hamiltonian theories) as having as objects the models $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ and as arrows between objects $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$, $(\Sigma_p, \tilde{\omega}'_p, H', \varphi'_i)$ the diffeomorphisms $f : \Sigma_p \rightarrow \Sigma_p$ such that $f^*(\tilde{\omega}'_p) = \tilde{\omega}_p$, $f^*(H') = H$ and $f^*(\varphi'_i) = \varphi_i$ i.e. the symmetries are taken to be the symplectomorphisms that preserve the Hamiltonian and the secondary constraints.

Similarly, define the category **ExtHam** (the category of Extended Hamiltonian theories) as having as objects the models $(\Sigma_f, \tilde{\omega}_f, H)$ and as arrows between objects $(\Sigma_f, \tilde{\omega}_f, H)$, $(\Sigma_f, \tilde{\omega}'_f, H')$ the diffeomorphisms $g : \Sigma_f \rightarrow \Sigma_f$ such that $g^*(\tilde{\omega}'_f) = \tilde{\omega}_f$ and $g^*(H') = H$.

Consider the functor $F : \mathbf{TotHam} \rightarrow \mathbf{ExtHam}$ that takes each model $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ to its restriction to the points that satisfy the constraints $\varphi_i = 0$, i.e. the associated model $(\Sigma_f, \tilde{\omega}_f, H)$, and that takes the arrow f to its action on Σ_f (since f preserves the secondary constraints, f preserves Σ_f , and so this is well-defined). Then:

Proposition 2 $F : \mathbf{TotHam} \rightarrow \mathbf{ExtHam}$ forgets (only) structure.¹⁹

Inasmuch as forgetting structure in category-theoretic terms captures what it means for one theory to have less structure than another, this proposition captures a precise sense in which the Extended Hamiltonian formalism removes structure from the Total Hamiltonian formalism. Moreover, F preserves empirical content because it takes the Hamiltonian in the model $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ to its restriction to the points of Σ_f , which are precisely the points where the Hamiltonian is empirically accessible. Therefore, the proposition allows us to capture in precise terms the sense in which characterizing gauge theories in terms of the Total Hamiltonian has ‘excess structure’: there is an alternative way to characterize the empirical content of such theories but with less structure, by characterizing the theories in terms of the Extended Hamiltonian.

An interesting consequence of this argument is that it provides a response to a recent argument by Pitts (2022, 2024) that moving to the Extended Hamiltonian is a “trivial reformulation” of the theory characterized by the Total Hamiltonian. In particular, Pitts argues that we can always add new ‘gauge’ transformations by adding new terms to the Hamiltonian while preserving empirical content.

¹⁸Indeed, this reasoning can be seen as an application of (**SYM***), allowing for the comparison to be made at the points where the domains of the objects coincide. As we show below, this is further justified by the category-theoretic comparison.

¹⁹See A.2 for proof.

To support this, Pitts invites us to consider a simple example. Take a Lagrangian given by:

$$L_1 = \frac{1}{2}\dot{q}^2 \quad (3)$$

This describes a particle moving in a straight line with uniform velocity; the equation of motion is $\ddot{q} = 0$. Now consider the Lagrangian:

$$L_2 = \frac{1}{2}(\dot{q} - \dot{\mu})^2 \quad (4)$$

where μ is either an arbitrary function of time or a dynamical variable. This Lagrangian is invariant under the transformation $q \rightarrow q + \epsilon$, $\mu \rightarrow \mu + \epsilon$ where ϵ is an arbitrary function of time. Since not all choices of ϵ are symmetries of L_1 , Pitts argues that we have added symmetries by moving from L_1 to L_2 . Moreover, L_2 gives rise to the equation of motion $\ddot{q} - \ddot{\mu} = 0$, which says that $q - \mu$ represents a particle moving in a straight line with uniform velocity. Therefore, the argument goes, we have added symmetries while keeping the empirical content the same.

In the Hamiltonian framework, the move from L_1 to L_2 can be seen to add gauge symmetries in the following way: there are no constraints corresponding to L_1 and so there is no gauge freedom. However, treating μ as a dynamical variable, we find that the Hamiltonian formulation of L_2 has a constraint, $p_q + p_\mu$, which is first-class. So, there are new gauge transformations in the Hamiltonian formulation; the transformations $q \rightarrow q + \epsilon$, $\mu \rightarrow \mu + \epsilon$ generated by the first-class constraint.

Pitts calls this process of revising a Lagrangian by adding a new variable (or “splitting one quantity into two”) and thereby adding new symmetries “de-Ockhamization”. Arguably, it is a trivial way of adding symmetries, as demonstrated by the above example. But, Pitts argues, this is exactly what is going on when one moves from the Total to the Extended Hamiltonian: one adds a new term to the Hamiltonian—some arbitrary function of the secondary constraints—and in doing so, one introduces new symmetries while preserving empirical content. Therefore, these additional gauge symmetries that the Extended Hamiltonian theory is associated with should not be regarded as “genuine” gauge transformations, in just the same way that the symmetries of L_2 are not physically interesting symmetries.

However, once we understand the relationship between the Total Hamiltonian and the Extended Hamiltonian in terms of Propositions 1 and 2, we can see that the process of “de-Ockhamization” is *not* what is going on when one moves from the Total Hamiltonian to the Extended Hamiltonian.

The reason is that one cannot set up a proposition analogous to Proposition 1 in the case of de-Ockhamization and simultaneously maintain that the two theories are empirically equivalent. To see this, recall that Proposition 1 stated that the symmetries of the Total Hamiltonian theory are a subset of those of the Extended Hamiltonian theory, when we consider their action on the same state space. This last clause is important: to compare the symmetries of two theories, we need some mapping between the two theories so that we can say which symmetries are being identified.

If we want to claim that L_1 and L_2 are empirically equivalent, then the natural choice of mapping is to take q to represent the position of the particle in L_1 and $q - \mu$ to represent the position of the particle in L_2 .²⁰ Therefore, we ought to consider the relationship between the symmetries under this standard of comparison. But now we run into a problem: the transformations of q that preserve the Euler-Lagrange equations for L_1 are the same as the transformations of $q - \mu$ that preserve the Euler-Lagrange equations for L_2 . Indeed, all we have done is effectively change the label of the variable that represents position. This is clearly a trivial kind of reformulation. However, it does not support the claim that L_2 has additional symmetries, since under the standard of comparison where q is identified with $q - \mu$, the Lagrangians are empirically equivalent and also have the same symmetries.

One might try to respond by saying the following: we have added symmetries by moving to L_2 because the transformations $q \rightarrow q + \epsilon$, $\mu \rightarrow \mu + \epsilon$, where ϵ is an arbitrary function of time, are symmetries of L_2 that are not symmetries of L_1 and that preserve the equations of motion. But under the identification of q with $q - \mu$, these transformations *are* symmetries of L_1 , since they correspond to the identity transformation on q in L_1 . What this response points out is that there are many symmetries of L_2 , differing over the choice of ϵ , that correspond to a single symmetry of L_1 . However, this is not sufficient to set up an analogous proposition to that of Proposition 1; one cannot say that there is a symmetry of L_2 that is not a symmetry of L_1 when we identify q in L_1 with $q - \mu$ in L_2 .

The upshot is that there is no way to maintain simultaneously that L_1 and L_2 are empirically equivalent and that the symmetries of L_1 are a subset of the symmetries of L_2 . Under the mapping that preserves empirical content, every symmetry of the first theory is a symmetry of the second theory and vice versa. Since Propositions 1 and 2 show that the symmetries of the Total Hamiltonian theory are a subset of those of the Extended Hamiltonian theory under the map that preserves empirical content—namely, the map that takes every point on the final constraint surface to itself—this suffices

²⁰If we took q in both Lagrangians to represent the position of the particle, then the Lagrangians would not be empirically equivalent: L_2 would represent a particle whose acceleration is equal to the acceleration of μ , which is not (necessarily) zero.

to show that de-Ockhamization is not what is going on in the move from the Total Hamiltonian to the Extended Hamiltonian. Therefore, whether de-Ockhamization is a way of trivially adding gauge symmetries is orthogonal to the question of the significance of moving from the Total Hamiltonian to the Extended Hamiltonian.

5 Symplectic Reduction Does Not Remove Further Structure

Given that we have convinced ourselves that the theory characterized by the Extended Hamiltonian removes excess structure from the theory characterized by the Total Hamiltonian, a natural question arises: is the theory characterized by the Extended Hamiltonian a theory *without* excess structure, or is there some other theory that has even less structure while preserving empirical content? Indeed, the standard way of removing what is regarded as excess structure from a Hamiltonian gauge theory is to undergo symplectic reduction. We therefore want to consider the relationship between the theory one gets via symplectic reduction and the theory characterized by **ExtHam**.

On the face of it, models of **ExtHam** have a kind of redundancy: states along the gauge orbits on Σ_f cannot be distinguished by the induced two-form, and so the theory doesn't depend upon which point along a gauge orbit one chooses to represent the state of a system. In light of this, it is tempting to make the following interpretive move: one should treat gauge-related states/solutions as being equivalent. Therefore, only gauge-invariant quantities (the “observables”) should be interpreted as physically meaningful. However, if we want to interpret gauge-related states/solutions as being equivalent, then, the argument goes, we should find a new mathematical formalism for describing Hamiltonian gauge theories such that gauge-related states/solutions are identified. The symplectic reduced theory is such a formalism.

For example, Th  bault (2012) says that one could interpret a theory in the constrained Hamiltonian formalism by “instituting a many-to-one relationship between gauge related sequences of points on the constraint surface and the unique sequences of instantaneous states they represent” but that this “does nothing about removing what would seem like superfluous mathematical structure—to dispense with this surplus structure we need to move to the reduced phase space.” Similarly, Belot (2003) says that “because the points of such orbits are dynamically indifferent, the x_i [the gauge variables] are dynamically irrelevant—any way of setting their value leads to the “same” evolution. This suggests in turn that it may be possible to drop the x_i from our theory altogether”. These quotes suggest

that reduction provides an option for removing redundancy present in a model of **ExtHam** since the points of the reduced space are just the equivalence class of points along the gauge orbits.

In more detail, symplectic reduction works as follows: we take the theory defined on the final constraint surface, and then we ‘quotient out’ the gauge symmetries by effectively equivocating between the points along the gauge orbits.²¹ More precisely, we can define a smooth, differentiable manifold $\bar{\Sigma}$ by taking the quotient of Σ_f by the kernel of $\tilde{\omega}$ i.e. the null vector fields of $\tilde{\omega}$. Next, we can define an open, surjective projection map $\pi : \Sigma_f \rightarrow \bar{\Sigma}$ such that we define the reduced two-form $\bar{\omega}$ via $\tilde{\omega}_f = \pi^*(\bar{\omega})$, which acts according to $\bar{\omega}(\bar{X}, \bar{Y}) = \tilde{\omega}_f(X, Y)$ where $\bar{X} = \pi_*(X)$. One can show that $\bar{\omega}$ is well-defined and is symplectic. We can also define a reduced Hamiltonian $\bar{H}$ as the value of H on the equivalence class of points along the gauge orbits i.e. $H = \pi^*(\bar{H})$. This is well-defined because H is constant along the gauge orbits on the final constraint surface. We can therefore write the equations of motion on the reduced space in terms of the reduced Hamiltonian as $\bar{\omega}(\bar{X}_{\bar{H}}, \cdot) = d\bar{H}$, whose solutions are the projection of the solutions to the equations of motion on Σ_f to $\bar{\Sigma}$. We can think of the dynamics on the reduced space as just the dynamics for the gauge-invariant variables. This provides the sense in which the reduced theory preserves empirical content: the predictions for the ‘physical’ variables are the same.

In order to evaluate the argument that reduction removes excess structure, let us consider the relationship between the reduced theory and the theory defined on the final constraint surface given by **ExtHam**. Unlike in the case of moving from the Total Hamiltonian theory to the Extended Hamiltonian theory, we cannot think about the automorphisms of the Extended Hamiltonian theory and the reduced theory as acting on the same state space, since $\bar{\Sigma}$ is not a submanifold of Σ_f or vice versa. Instead, we think of $\bar{\Sigma}$ as the projection of Σ_f under a map that collapses distinctions between the points along the gauge orbits. Therefore, we cannot say that the automorphisms of one are a subset of the other in the same way. Instead, what we can say is that:

Proposition 3 *Every automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$, under the projection $\pi : \Sigma_f \rightarrow \bar{\Sigma}$, is an automorphism of $(\bar{\Sigma}, \bar{\omega}, \bar{H})$, and every automorphism of $(\bar{\Sigma}, \bar{\omega}, \bar{H})$ corresponds to some automorphism(s) of $(\Sigma_f, \tilde{\omega}_f, H)$ under π .²²*

Since this proposition says that all of the symmetries of the Extended Hamiltonian theory corre-

²¹The procedure we describe here is really a special case of what is sometimes also called “symplectic reduction”, or “Marsden-Weinstein reduction”. For more details on reduction as we describe it here, see Henneaux & Teitelboim (1994), Souriau (1997), Thébault (2012). For discussion of symplectic reduction more generally, see Butterfield (2006), Belot (2006).

²²See A.3 for proof.

spond to symmetries of the reduced theory, we cannot infer that one theory has more structure than the other by drawing an analogy with Proposition 1. So let us see whether the category-theoretic approach can shed light on the structural relationship between these theories.

An issue that we face is that some of the models in **ExtHam** are, essentially, partial reductions of other models in **ExtHam**, since we didn't restrict the dimension of the final constraint surface when defining this category, nor the dimension of the null space. Given that we want to capture the relationship between theories and their reduced counterpart, we will restrict to the case where the models of **ExtHam** have the same dimensions, which excludes the case of objects within the category **ExtHam** being related by (partial) reduction.

We can define the category **HamRed** as having objects $(\bar{\Sigma}, \bar{\omega}, \bar{H})$ and arrows between objects $(\bar{\Sigma}_1, \bar{\omega}_1, \bar{H}_1)$ and $(\bar{\Sigma}_2, \bar{\omega}_2, \bar{H}_2)$ to be given by diffeomorphisms $h : \bar{\Sigma}_1 \rightarrow \bar{\Sigma}_2$ such that $h^*(\bar{\omega}_2) = \bar{\omega}_1$ and $h^*\bar{H}_2 = \bar{H}_1$. Take the functor G that takes the object $(\Sigma_f, \tilde{\omega}_f, H)$ to $(\bar{\Sigma}, \bar{\omega}, \bar{H})$ and that takes an arrow $g : \Sigma_f \rightarrow \Sigma_f$ to $\bar{g} : \bar{\Sigma} \rightarrow \bar{\Sigma}$ where $\bar{g}(\bar{x}) = g(\bar{x})$ i.e. $\bar{g}$ is the action of g on the gauge orbits, such that whenever g acts by moving points of Σ_f along the gauge orbits, $\bar{g}$ acts as the identity on $G(\Sigma_f)$. This is well-defined because every arrow $g : \Sigma_f \rightarrow \Sigma_f$ preserves the gauge orbits. Then, the following is true:

Proposition 4 $G : \mathbf{ExtHam} \rightarrow \mathbf{HamRed}$ is full and essentially surjective but not faithful i.e. G forgets only stuff.²³

This shows us that symplectic reduction is not a way of removing ‘excess structure’ from **ExtHam**. In a way, this was to be expected: we already established that gauge symmetries are isomorphisms of the models of **ExtHam**, and so this theory already captures the idea that gauge symmetries relate physically equivalent situations. Nevertheless, the fact that the functor G forgets ‘stuff’ remains significant. As we suggested in Section 2, one can think about a theory with more stuff as exhibiting “representational redundancy”: the models of that theory can be used to represent some object or physical situation in lots of different ways through its non-trivial automorphisms. This seems to correctly capture the role of the gauge transformations in **ExtHam**: they reflect the fact that any point along the gauge orbit is equally well-suited to represent some physical state.²⁴ Moreover, removing this ‘redundancy’ seems well-motivated from the perspective of pursuing a well-posed initial value problem. Therefore, while symplectic reduction does not remove excess *structure*, Proposition

²³See A.4 for proof.

²⁴One might notice the analogy between this reasoning and the one we considered in the vector space example at the end of Section 2.

4 might nonetheless indicate that it removes a relevant form of redundancy from the theory on the final constraint surface—one associated with the failure of a well-posed initial value problem.

However, as we also discussed in Section 2, having *more* ‘stuff’ can indicate that the theory *less* ‘structure’, and therefore that the two notions of redundancy can pull in different directions. Indeed, notice that a different way of removing the ‘representational redundancy’ that **ExtHam** has would be to pick out a representative for each gauge orbit—that is, to *fix a gauge*. By analogy, one might regard symplectic reduction as *adding* structure insofar as it is effectively like fixing a representative for each gauge orbit.

Of course, much more could be said about the similarities and differences between symplectic reduction and gauge fixing. But at the very least, the foregoing discussion shows that the notions of *structure* and *stuff* must be handled with care in the context of gauge theories, and that the corresponding types of redundancy they track should not be conflated. This observation may have important implications, for instance, when considering the motivations behind standard quantization procedures that proceed through—or are equivalent to proceeding through—the reduced theory.

6 Conclusion

We have examined two central questions. First, what is the relationship between the structure of the Total Hamiltonian theory and that of the Extended Hamiltonian theory? Second, what is the relationship between the structure of the Extended Hamiltonian theory and the symplectic reduced theory? We argued that the first relation is best understood in terms of *excess structure*, whereas the second concerns *representational redundancy*. We used this to argue that the motivations for symplectic reduction differ, conceptually, from the motivations for moving from the Total Hamiltonian to the Extended Hamiltonian.

While our focus has been on the context of Hamiltonian gauge theories, there are broader take-aways. First, our intuitions about whether one theory is simpler or more complicated than another can diverge from whether it possesses more or less *structure*. In particular, counting terms in the Hamiltonian or dimensions of the state space is not sufficient to capture the amount of structure that the theories commit to. In moving from the Total to the Extended Hamiltonian, additional terms appear in the Hamiltonian, but no new structure is thereby introduced. Conversely, in symplectic reduction, while the theory is formulated on a lower-dimensional space with the same empirical content,

this does not entail that it has *less* structure.

Second, to say that one theory has ‘excess structure’ relative to another, one needs to be able to say two things simultaneously. First, one theory has more structure than the other. Second, the two theories have the same empirical content. While one might be able to say both things relative to two *different* maps between the theories, being able to say both things relative to a *single* map is importantly different. In Section 4, we discussed how this differentiates Pitts’ notion of ‘de-Ockhamization’ from the move from the Total Hamiltonian to the Extended Hamiltonian. However, it is also relevant in the context of symplectic reduction.²⁵ Indeed, Proposition 3 highlights that there is a sense in which the reduced theory and the Extended Hamiltonian theory have the ‘same’ symmetries, under the map that preserves empirical content: when points on the reduced space correspond to the collection of points along the corresponding gauge orbit, there is no symmetry of one theory that is not a symmetry of the other. Therefore, even though **ExtHam** might be said to have less ‘structure’ than **HamRed**, relative to a functor between the categories, it seems inappropriate to regard **HamRed** as having *excess* structure, at least in the same way that, for example, the Total Hamiltonian theory has compared to the Extended Hamiltonian theory.

Finally, the discussion here emphasizes the conceptual differences between a theory having ‘representational redundancy’ and a theory having ‘excess structure’. While both are related to the amount of structure that a theory has, as represented as a category, the interpretational significance is importantly different. In particular, whether ‘representational redundancy’ is regarded as a genuine sort of redundancy depends on what one’s representational aims are. As we have seen through our exploration of Hamiltonian gauge theories, the distinction between the two notions is especially important for understanding the relationship between theories where the dimension of the state spaces is different; in particular, when the state space of one can be understood as the quotient, under some symmetry group, of the state space of the other.

²⁵In fact, ‘de-Ockhamization’ can be thought of as a kind of ‘un-reducing’, in the following way: Imagine that we start with the Lagrangian $L_2 = \frac{1}{2}(\dot{q} - \dot{\mu})^2$ where q and μ are two dynamical variables. In the Hamiltonian framework, there is a single first-class constraint, $p_q + p_\mu = 0$. This generates a gauge symmetry that preserves the quantity $q - \mu$. Reducing this theory gives rise to a new theory that is equivalent to starting with the Lagrangian $L = \frac{1}{2}\dot{\alpha}^2$ where $\alpha = q - \mu$. Viewed this way, the relationship between $L_1 = \frac{1}{2}\dot{q}^2$ and its “de-Ockhamized” counterpart $L_2 = \frac{1}{2}(\dot{q} - \dot{\mu})^2$ is precisely that between a reduced and an unreduced theory.

A Proofs of Propositions

A.1 Proposition 1

We need to show that every automorphism of $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$, restricted to its action on Σ_f , is an automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$, but that there is some automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$ that is not an automorphism of $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ on Σ_f . The first follows from the fact that every automorphism of $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ preserves the points where $\varphi_i = 0$ (since it preserves φ_i) and so preserves Σ_f , and that $i^*(\tilde{\omega}_p) = \tilde{\omega}_f$ where $i : \Sigma_f \rightarrow \Sigma_p$ is the inclusion map.

For the second, consider the diffeomorphism $g : \Sigma_f \rightarrow \Sigma_f$ that takes each point on Σ_f to another arbitrary point along the gauge orbit associated with the secondary first-class constraints φ_j at that point. We can represent g as the flow of the vector field associated with $\alpha^j d\varphi_j$ where α^j are arbitrary functions. Since $d\varphi_j = 0$ on Σ_f , $\alpha^j d\varphi_j = 0$. This means that $\alpha^j d\varphi_j$ is closed i.e. $d(\alpha^j d\varphi_j) = 0$. But this means that one can (locally) associate a vector field Y with $\alpha^j d\varphi_j$ via $\tilde{\omega}_f(Y, \cdot) = \alpha^j d\varphi_j$. It follows that the flow of Y on Σ_f consists of maps that preserve $\tilde{\omega}_f$.²⁶ So g is a diffeomorphism that takes $\tilde{\omega}_f$ to itself. Moreover, $g^*H = H$ because H is gauge-invariant on the final constraint surface. So g is an automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$.

Now consider the diffeomorphism $f : \Sigma_p \rightarrow \Sigma_p$ that takes each point on Σ_p to another point along the vector field associated with the secondary first-class constraints at that point. We can similarly consider f to be the flow of the vector field associated with $\alpha^j d\varphi_j$. In order to associate a vector field X with $\alpha^j d\varphi_j$ via $\tilde{\omega}_p(X, \cdot) = \alpha^j d\varphi_j$, it must be that $d(\tilde{\omega}_p(X, \cdot)) = d(\alpha^j d\varphi_j) = 0$. But $d_b(\alpha^j d_a \varphi_j) = \alpha^j d_b d_a \varphi_j + d_{[a} \varphi_j d_{b]} \alpha^j$. The first term vanishes since $d(d\varphi_j) = 0$ by Poincaré's Lemma. However, the second term does not necessarily vanish, since α^j is an arbitrary function of the canonical coordinates (so $d\alpha^j$ is not necessarily zero). In such cases, one cannot associate with f a vector field via $\tilde{\omega}_p(X, \cdot) = \alpha^j d\varphi_j$, and so the flow of the vector field associated with $\alpha^j d\varphi_j$ does not take $\tilde{\omega}_p$ to itself. Therefore, there are automorphisms of $(\Sigma_f, \tilde{\omega}_f, H)$ that do not correspond to an automorphism of $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$.

A.2 Proposition 2

We want to show that F is not full i.e. that it fails to be surjective on arrows, but that it is faithful and essentially surjective. That F is not full follows from Proposition 1: the arrows $g : (\Sigma_f, \tilde{\omega}_f, H) \rightarrow$

²⁶This follows from Abraham & Marsden (1987) Proposition 3.3.6. (when we extend the proposition to presymplectic manifolds).

$(\Sigma_f, \tilde{\omega}'_f, H')$ in **ExtHam** such that $g \neq F(f)$ for any arrow f in **TotHam** are transformations along the null vector fields associated with secondary first-class constraints. That F is essentially surjective follows from the fact that every object of **ExtHam** is the restriction of some object $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ to the surface defined by $\varphi_i = 0$. Finally, to show that F is faithful, we need to show that if two arrows f, g between objects $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i), (\Sigma_p, \tilde{\omega}'_p, H', \varphi'_i)$ of **TotHam** are distinct, then their action on Σ_f is distinct. In other words, we want to show that if $f|_{\Sigma_f} = g|_{\Sigma_f}$, then $f = g$.

Suppose that $f|_{\Sigma_f} = g|_{\Sigma_f}$. If f, g are arbitrary gauge transformations, then the only way that f and g could differ is if they move points off of Σ_f by differing amounts along the vector fields associated with the primary first-class constraints. But H is not constant along the vector fields associated with the primary first-class constraints off of Σ_f . Since f, g must preserve H by definition, f must be equal to g . If f, g are symplectomorphisms that are flows along vector fields other than the vector fields associated with the primary first-class constraints, then the only way that f and g could differ is if at least one changes the secondary constraints. But since f, g must preserve φ_i by definition, f must be equal to g . So F is faithful.

A.3 Proposition 3

That every automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$, under the projection $\pi : \Sigma_f \rightarrow \bar{\Sigma}$, is an automorphism of $(\bar{\Sigma}, \bar{\omega}, \bar{H})$ follows from the fact that $\tilde{\omega}_f = \pi^*(\bar{\omega})$ and $H = \pi^*(\bar{H})$ and so every transformation from Σ_f to itself that preserves $\tilde{\omega}_f$ and H must preserve $\bar{\omega}$ and $\bar{H}$ under π . That every automorphism of $(\bar{\Sigma}, \bar{\omega}, \bar{H})$ corresponds to some automorphism(s) of $(\Sigma_f, \tilde{\omega}_f, H)$ under π follows from the fact that every transformation on Σ_f that preserves the equivalence class of points along the gauge orbits is an automorphism of $(\Sigma_f, \tilde{\omega}_f, H)$ by the proof of Proposition 1 (since these transformations correspond to moving along the gauge orbits at each point).

A.4 Proposition 4

That G is essentially surjective follows from the fact that every model of the reduced theory is the reduction of some model on the final constraint surface by definition. To show that G is full, notice that if there is an arrow h between $G(\Sigma_f, \tilde{\omega}_f, H), G(\Sigma'_f, \tilde{\omega}'_f, H')$, then there is a transformation between $(\Sigma_f, \tilde{\omega}_f, H), (\Sigma'_f, \tilde{\omega}'_f, H')$ that preserves the equivalence class of points along the integral curves of the null vector fields of $\tilde{\omega}_f, \tilde{\omega}'_f$. Since Σ_f, Σ'_f are of the same dimension (by assumption), this implies that there exists a diffeomorphism $g : \Sigma_f \rightarrow \Sigma'_f$ such that $g^*(\tilde{\omega}'_f) = \tilde{\omega}_f$. Moreover, since H, H' are invariant

along the gauge orbits, and give rise to isomorphic reduced Hamiltonians, it must be that $g^*(H') = H$. So g is an arrow between $(\Sigma_f, \tilde{\omega}_f, H), (\Sigma'_f, \tilde{\omega}'_f, H')$ and $h = G(g)$. Therefore, every arrow between $G(\Sigma_f, \tilde{\omega}_f, H), G(\Sigma'_f, \tilde{\omega}'_f, H')$ is mapped to by at least one arrow between $(\Sigma_f, \tilde{\omega}_f, H), (\Sigma'_f, \tilde{\omega}'_f, H')$, and so G is full.

G is not faithful if there are two distinct arrows in **ExtHam** that map to the same arrow in **HamRed**. Consider some model $(\Sigma_f, \tilde{\omega}_f, H)$ where $\tilde{\omega}_f$ has at least one null vector field and consider two arrows g_1, g_2 from $(\Sigma_f, \tilde{\omega}_f, H)$ to itself corresponding to distinct gauge transformations i.e. two different ways of moving the points of Σ_f to other points along the gauge orbits (as long as $\tilde{\omega}_f$ has at least one null vector field, one can find such distinct g_1, g_2). Then $G(g_1) = G(g_2) = Id$ since g_1 and g_2 preserve the gauge orbits. So G is not faithful.

Acknowledgements

I am especially grateful to Jim Weatherall for helpful guidance on the proofs of the propositions in the paper and for comments on a previous version. I also thank Karim Thébault and Sean Gryb for several conversations that have helped me better understand the constrained Hamiltonian formalism. Finally, I am grateful to the audience at PSA 2024 for their feedback on the material in Section 4, and to Brian Pitts for correspondence.

References

- Abraham, R. & Marsden, J. E. (1987), *Foundations of Mechanics, Second Edition*, Addison-Wesley Publishing Company, Inc.
- Baker, D. J. (2010), ‘Symmetry and the metaphysics of physics’, *Philosophy Compass* **5**(12), 1157–1166.
- Barrett, T. W. (2015a), ‘On the structure of classical mechanics’, *The British Journal for the Philosophy of Science* **66**(4), 801–828.
- Barrett, T. W. (2015b), ‘Spacetime structure’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **51**, 37–43.

- Barrett, T. W. (2019), ‘Equivalent and inequivalent formulations of classical mechanics’, *The British Journal for the Philosophy of Science* **70**(4), 1167–1199.
- Barrett, T. W., Manchak, J. & Weatherall, J. O. (2023), ‘On automorphism criteria for comparing amounts of mathematical structure’, *Synthese* **201**(6), 191.
- Belot, G. (2003), ‘Symmetry and gauge freedom’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **34**(2), 189–225.
- Belot, G. (2006), The representation of time and change in mechanics., in J. Butterfield & J. Earman, eds, ‘Handbook of the Philosophy of Physics’, pp. 133–227.
- Belot, G. (2013), Symmetry and equivalence., in R. Batterman, ed., ‘The Oxford Handbook of Philosophy of Physics’, Oxford University Press, pp. 318–339.
- Bergmann, P. G. (1961), ‘Observables in general relativity’, *Reviews of Modern Physics* **33**(4), 510.
- Bradley, C. & Weatherall, J. O. (2020), ‘On representational redundancy, surplus structure, and the hole argument’, *Foundations of Physics* **50**(4), 270–293.
- Butterfield, J. (2006), On symplectic reduction in classical mechanics., in J. Butterfield & J. Earman, eds, ‘Handbook of the Philosophy of Physics’, pp. 1–131.
- Dewar, N. (2019), ‘Sophistication about symmetries’, *The British Journal for the Philosophy of Science* **70**(2), 485–521.
- Dirac, P. A. M. (1964), *Lectures on Quantum Mechanics*, Dover.
- Earman, J. (2002a), ‘Thoroughly modern mctaggart: or, what mctaggart would have said if he had read the general theory of relativity’, *Philosophers’ Imprint* **2**(3), 1–28.
- Earman, J. (2002b), Tracking down gauge: An ode to the constrained hamiltonian formalism., in K. Brading & E. Castellani, eds, ‘Symmetries in Physics: Philosophical Reflections’, Cambridge University Press, pp. 140–62.
- Earman, J. (2004), ‘Laws, symmetry, and symmetry breaking: Invariance, conservation principles, and objectivity’, *Philosophy of Science* **71**(5), 1227–1241.
- Halvorson, H. (2012), ‘What scientific theories could not be.’, *Philosophy of Science* **79**(2), 183–206.

- Halvorson, H. (2016), Scientific theories., in P. Humphreys, ed., ‘The Oxford Handbook of Philosophy of Science’, Oxford University Press.
- Halvorson, H. & Tsementzis, D. (2017), Categories of scientific theories., in ‘Categories for the Working Philosopher’, Oxford University Press.
- Henneaux, M. & Teitelboim, C. (1994), *Quantization of Gauge Systems*, Princeton University Press.
- Ismael, J. & van Fraassen, B. C. (2003), Symmetry as a guide to superfluous theoretical structure., in K. Brading & E. Castellani, eds, ‘Symmetries in Physics: Philosophical Reflections’, Cambridge University Press, pp. 371–92.
- Møller-Nielsen, T. (2017), ‘Invariance, interpretation, and motivation’, *Philosophy of Science* **84**(5), 1253–1264.
- Nguyen, J., Teh, N. J. & Wells, L. (2020), ‘Why surplus structure is not superfluous’, *The British Journal for the Philosophy of Science* **71**(2), 665–695.
- North, J. (2009), ‘The “structure” of physics: A case study’, *The Journal of Philosophy* **106**(2), 57–88.
- Pitts, J. B. (2014a), ‘Change in hamiltonian general relativity from the lack of a time-like killing vector field’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **47**, 68–89.
- Pitts, J. B. (2014b), ‘A first class constraint generates not a gauge transformation, but a bad physical change: The case of electromagnetism’, *Annals of Physics* **351**, 382–406.
- Pitts, J. B. (2022), ‘First-class constraints, gauge transformations, de-ockhamization, and triviality: Replies to critics, or, how (not) to get a gauge transformation from a second-class primary constraint’, *arXiv: 2212.02944* .
- Pitts, J. B. (2024), ‘Does a second-class primary constraint generate a gauge transformation? electromagnetisms and gravities, massless and massive.’, *Annals of Physics* **462**, 169621.
- Pons, J. M. (2005), ‘On dirac’s incomplete analysis of gauge transformations’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **36**(3), 491–518.
- Pooley, O. & Wallace, D. (2022), ‘First-class constraints generate gauge transformations in electromagnetism (reply to pitts)’, *arXiv:2210.09063* .

- Redhead, M. L. (1975), ‘Symmetry in intertheory relations’, *Synthese* **32**, 77–112.
- Rovelli, C. (2014), ‘Why gauge’, *Foundations of Physics* **44**, 91–104.
- Souriau, J.-M. (1997), *Structure of dynamical systems: a symplectic view of physics*, Vol. 149, Springer Science & Business Media.
- Thébault, K. P. (2012), ‘Symplectic reduction and the problem of time in nonrelativistic mechanics’, *The British Journal for the Philosophy of Science* **63**(4), 789–824.
- Weatherall, J. O. (2016a), ‘Are newtonian gravitation and geometrized newtonian gravitation theoretically equivalent?’, *Erkenntnis* **81**, 1073–1091.
- Weatherall, J. O. (2016b), ‘Understanding gauge’, *Philosophy of Science* **83**(5), 1039–1049.
- Weatherall, J. O. (2017), Category theory and the foundations of classical space–time theories., in E. Landry, ed., ‘Categories for the Working Philosopher’, Oxford University Press.
- Weatherall, J. O. (2018), ‘Regarding the ‘hole argument’’, *The British Journal for the Philosophy of Science* **69**(2), 329–350.