The Extended Hamiltonian is Not Trivial*

Clara Bradley

Department of Logic and Philosophy of Science,

University of California, Irvine

Abstract

This paper critically analyzes an argument made by Pitts (2022, 2024) that extending the form of the Hamiltonian constitutes a trivial reformulation of a theory and therefore doesn't provide insight into the gauge transformations. I argue that a trivial reformulation cannot be used to add new gauge transformations to a theory, and I show that the sense in which extending the form of the Hamiltonian is nontrivial is that it removes structure.

1 Introduction

One way of characterizing the "gauge symmetries" or "gauge transformations" of a physical theory is through the Hamiltonian formalism: gauge symmetries are defined as being transformations generated by *first-class constraints*. However, different ways of writing down the Hamiltonian, which generates the dynamics, lead to different combinations of first-class constraints being the generator of the gauge symmetries. Recently, there has been a debate about the correct characterization of the gauge transformations that rests on different views about the correct form of the Hamiltonian. On the one hand, Pitts (2014) argues that the standard characterization of the gauge transformations as the

* *Draft of August 27, 2024.*

transformations generated by arbitrary combinations of first-class constraints is wrong by showing that one form of the Hamiltonian has solutions whose symmetries are not given by these transformations. On the other hand, Pooley & Wallace (2022) argue that the standard characterization of the gauge transformations is right by showing that these transformations are symmetries under a different but empirically equivalent form of the Hamiltonian.

There is a natural question to ask in reaction to this recent debate: if both sides agree about the characterization of the gauge transformations *given a particular form of the Hamiltonian*, what is at stake in using one form of the Hamiltonian or another? In more recent papers, Pitts (2022, 2024) responds to Pooley & Wallace (2022) by arguing that their use of a different form of the Hamiltonian is a trivial kind of reformulation of a theory that artificially adds new gauge transformations and, therefore, that it does not provide a physically interesting alternative. More generally, Pitts argues, we can add new symmetries by changing the form of the Hamiltonian in lots of different ways, but we should not think that this leads to a new understanding of the gauge transformations.

In this paper, I critically analyze the arguments made by Pitts (2022, 2024). I argue that a trivial reformulation of a theory in the way suggested by Pitts cannot be used to add new gauge symmetries to a theory, and I show that the move to a new form of the Hamiltonian is therefore not a trivial reformulation. In doing so, I clarify what the consequences are of adopting different forms of the Hamiltonian, and I argue that there is a natural reason to think that in fact, the form of the Hamiltonian used by Pooley & Wallace (2022) is the correct one – it is associated with a formulation of the theory with *less structure*.

2 The Constrained Hamiltonian Formalism

To set up the debate between Pitts (2014) and Pooley & Wallace (2022), let us begin with a brief overview of the constrained Hamiltonian formalism and the standard approach to the connection between gauge symmetries and the form of the Hamiltonian. Following Dirac (1964), we start with a Lagrangian $L = L(q_n, \dot{q}_n)$ with corresponding action $I = \int L(q_n, \dot{q}_n)dt$, from which one derives the equations of motion, the Euler-Lagrange (E-L) equations:

$$\frac{d}{dt}\frac{\partial L(q_n, \dot{q}_n)}{\partial \dot{q}_n} = \frac{\partial L(q_n, \dot{q}_n)}{\partial q_n}$$

Moving to the Hamiltonian framework, we introduce *canonical momenta* variables $p_n = \frac{\partial L}{\partial \dot{q}_n}$. When these momenta are not independent of each other, there are constraints of the form $\phi_m(q_n, p_n) \approx 0$ for $m = 1, ..., M$ where M is the number of constraints and the meaning of $\approx$ is that the relationship holds *weakly*: one can substitute the left-hand side for the right only on the subspace where the equation holds. Constraints of this kind are called the *primary constraints*.

The *Total Hamiltonian* is defined as $H_T = H + u^m \phi_m$ where $H = p^n q_n - L$ is the Hamiltonian and $u^m(q_n, p_n)$ are arbitrary functions.[1] The Total Hamiltonian is weakly equal to the Hamiltonian H, so one can think of it as an equivalence class of Hamiltonians, differing over the choices of u^m. From the variation in H_T, one can derive Hamilton's equations of motions with constraints: For any dynamical variable g, $\dot{g} = \{g, H_T\}$ where $\{\}$ is the Poisson bracket.

In order for the solutions of the equations of motion to be consistent with the primary constraints, in the sense that the primary constraints hold at all times along a solution, it ought to be the case that $\dot{\phi}_m = \{\phi_m, H_T\} \approx 0$. When this reduces to an equation independent of the u^m's, $\chi_k(q_n, p_n) \approx 0$, we say that $\chi_k(q_n, p_n)$ are *secondary constraints*. If we have a secondary constraint, then we get new consistency conditions by requiring $\dot{\chi}_k \approx 0$, which again can give rise to further secondary constraints.

A dynamical variable $R(q_n, p_n)$ is said to be *first-class* if $\{R, \phi_j\} \approx 0$ for $j = 0, ..., M + K$ where M is the number of primary constraints and K is the number of secondary constraints. If it is not first-class, it is called *second-class*. One can show that the Total

[1] The raised and lowered indices indicate a sum over n and m.

Hamiltonian can be written as a sum of a first-class Hamiltonian and a linear combination of primary, first-class constraints.

Using this formalism, Dirac (1964) argues that the gauge transformations – understood as transformations relating any two states that are possible evolutions from an initial state under the dynamics generated by the Total Hamiltonian – are generated by arbitrary combinations of the first-class constraints.

However, this definition of the gauge transformations means that the dynamics is given by the Total Hamiltonian, which includes the arbitrariness associated with the primary first-class constraints, but there is also arbitrariness associated with the secondary first-class constraints. This mismatch between the dynamics and the arbitrariness led Dirac to suggest that one should also add the first-class secondary constraints to the Total Hamiltonian, giving rise to the *Extended Hamiltonian*, $H_E = H_T + w^b \chi_b$ where χ_b are the first-class secondary constraints and $w^b(q_n, p_n)$ are arbitrary functions.

3 Pitts' Counterexample

Pitts (2014) argues that classical Electromagnetism is a counterexample to the standard view that arbitrary combinations of first-class constraints generate gauge transformations. The reasoning is as follows. The Lagrangian for classical Electromagnetism can be written in observer-dependent form as

$$\mathcal{L}(\vec{A}, V; \dot{\vec{A}}, \dot{V}) = \int \frac{1}{2}(\dot{\vec{A}} - \nabla V)^2 - \frac{1}{2}(\nabla \times \vec{A})^2 - (V\rho + \vec{A} \cdot \vec{J})$$

where $\vec{A}$ and V are time-dependent functions on $\mathbb{R}^3$ and the integral is over $\mathbb{R}^3$. The conjugate momenta are $p_{\vec{A}} = \frac{\delta L}{\delta \dot{\vec{A}}} = \dot{\vec{A}} - \nabla V$ and $p_V = \frac{\delta L}{\delta \dot{V}} = 0$. This means that there is one primary constraint, $\phi_0 = p_V$. The Total Hamiltonian (under appropriate boundary conditions) can be written as:

$$H_T = \int \frac{1}{2}(p_{\vec{A}}^2 + \vec{B}^2) + \vec{A} \cdot \vec{J} + \lambda p_V - V(\nabla \cdot p_{\vec{A}} - \rho) \tag{1}$$

where the integral is over $\mathbb{R}^3$ and λ is an arbitrary function.

The evolution of the primary constraint is $\{p_V, H_T\} = \nabla \cdot p_{\vec{A}} - \rho$, so there is a secondary constraint given by $\phi_1 = \nabla \cdot p_{\vec{A}} - \rho$. The evolution of the secondary constraint is zero, so there are two constraints in total, and both constraints are first-class.

The equations of motion for $\vec{A}$ and V are given by:

$$\begin{aligned} \frac{\partial \vec{A}}{\partial t} &= p_{\vec{A}} + \nabla V \\ \frac{\partial V}{\partial t} &= \lambda \end{aligned} \tag{2}$$

What Pitts (2014) shows is that if $(\vec{A}(t), V(t); p_{\vec{A}}(t), p_V(t))$ satisfies these equations of motion, then transforming this solution by an arbitrary combination of the first-class constraints, $\int \alpha\phi_0 + \beta\phi_1$, where α and β are arbitrary functions of the canonical variables and time, does not in general satisfy the equations of motion

To see why this is the case, notice that $\{\vec{A}, \int \alpha\phi_0 + \beta\phi_1\} = \nabla\beta$ and $\{V, \int \alpha\phi_0 + \beta\phi_1\} = \alpha$, and so the transformed quantities are $A' = A + \nabla\beta$ and $V' = V + \alpha$. The conjugate momenta are unchanged. We can therefore write the transformed equations of motion for $\vec{A}$ and V as:

$$\begin{aligned} \frac{\partial \vec{A}'}{\partial t} &= \frac{\partial \vec{A}}{\partial t} + \frac{\partial \nabla\beta}{\partial t} = p_{\vec{A}} + \nabla(V + \alpha) \\ \frac{\partial V'}{\partial t} &= \frac{\partial V}{\partial t} + \frac{\partial \alpha}{\partial t} = \lambda \end{aligned} \tag{3}$$

Since we assumed that $\frac{\partial \vec{A}}{\partial t} = p_{\vec{A}} + \nabla V$, the first equation is satisfied only when $\frac{\partial \nabla\beta}{\partial t} - \nabla\alpha = 0$. On the basis of this argument, Pitts (2014) concludes that arbitrary combinations of first-class constraints do not generate gauge transformations for the dynamics generated by the Total Hamiltonian; rather, only a *specific* combination of them do.

4 Pooley and Wallace's Response

In response to Pitts (2014), Pooley & Wallace (2022) show that if one starts with the Extended Hamiltonian for classical Electromagnetism rather than the Total Hamiltonian, then arbitrary combinations of first-class constraints *do* generate gauge transformations. To see this, consider the Extended Hamiltonian for classical Electromagnetism, where we add to the Total Hamiltonian the secondary constraint multiplied by an arbitrary function μ:

$$H_E = \int \frac{1}{2}(p_{\vec{A}}^2 + \vec{B}^2) + \vec{A} \cdot \vec{J} + \lambda p_V - (V + \mu)(\nabla \cdot p_{\vec{A}} - \rho) \tag{4}$$

With this Hamiltonian, the equations of motion become:

$$\begin{aligned} \frac{\partial \vec{A}}{\partial t} &= p_{\vec{A}} + \nabla(V + \mu) \\ \frac{\partial V}{\partial t} &= \lambda \end{aligned} \tag{5}$$

When we now consider the transformation generated by an arbitrary combination of the first-class constraints, $\int \alpha\phi_0 + \beta\phi_1$, we find:

$$\begin{aligned} \frac{\partial \vec{A}'}{\partial t} &= \frac{\partial \vec{A}}{\partial t} + \frac{\partial \nabla\beta}{\partial t} = p_{\vec{A}} + \nabla(V + \mu + \alpha) \\ \frac{\partial V'}{\partial t} &= \frac{\partial V}{\partial t} + \frac{\partial \alpha}{\partial t} = \lambda \end{aligned} \tag{6}$$

Notice that since μ, α and β are all arbitrary functions, we can rewrite the first equation as

$$\frac{\partial \vec{A}'}{\partial t} = \frac{\partial \vec{A}}{\partial t} = p_{\vec{A}} + \nabla(V + \mu')$$

where μ' is arbitrary. This is just the untransformed equation of motion, with μ' in place of μ. Therefore, arbitrary combinations of first-class constraints generate gauge transformations on solutions, for the dynamics generated by the Extended Hamiltonian.

Moreover, Pooley & Wallace (2022) argue that the dynamics generated by the Extended Hamiltonian is *empirically equivalent* to the dynamics generated by the Total Hamiltonian. In particular, what they notice is that the difference between the solutions of the Total and Extended Hamiltonian lies only in the quantity that plays the role of the electric field: when the Total Hamiltonian is used to generate the dynamics, it is $\dot{\vec{A}} - \nabla V$ that plays the role of the electric field, but when the Extended Hamiltonian is used, it is $p_{\vec{A}}$. And so, given that it is the electric field that one measures directly through its interaction with charges (and not the quantities $\dot{\vec{A}}$ and V), there is no empirical difference between these choices of Hamiltonian.

5 Triviality Argument

Given Pooley and Wallace's response to Pitts (2014), there seem to be two possible reactions. 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 or Extended Hamiltonian is considered the right equivalence class of Hamiltonians. 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. This second reaction is the approach that Pitts (2022, 2024) takes: he argues that the Extended Hamiltonian can be seen as a trivial kind of reformulation of the theory. Moreover, he argues that we can construct similar kinds of reformulation that allow one to conclude that quantities other than the first-class constraints generate gauge transformations, which is not maintained by either side of the debate. So, we haven't gained insight into the gauge transformations by making this move; rather, we have stated the same thing in a different (more complicated) form.

To see what Pitts means by a trivial reformulation, let us consider a simple example presented in Pitts (2022). Take a Lagrangian given by:

$$L = \frac{1}{2}\dot{q}^2 \tag{7}$$

This describes a particle moving in a straight line with uniform velocity; the equation of motion is $\frac{d^2q}{dt^2} = 0$. The symmetries of this equation of motion are spatial translations and boosts. Now consider the Lagrangian:

$$L = \frac{1}{2}(\dot{q} - \dot{\mu})^2 \tag{8}$$

where μ is either an arbitrary function of time or a dynamical variable. This Lagrangian is invariant under the transformation $q \to q + \epsilon$, $\mu \to \mu + \epsilon$ where ϵ is an arbitrary function of time. So, Pitts argues, we have 'added' a symmetry. Moreover, this Lagrangian gives rise to the equation of motion $\frac{d^2(q-\mu)}{dt^2} = 0$, which says that $q - \mu$ represents a particle moving in a straight line with uniform velocity. Therefore, the new Lagrangian has the same physical content as the previous Lagrangian.

More precisely, we can see in the Hamiltonian formulation that the first Lagrangian does not have any gauge freedom since there are no constraints. However, treating μ as a dynamical variable, we find that the Hamiltonian formulation of the second Lagrangian has a constraint, namely $p_q + p_\mu$, which is first-class. So, there is a new gauge transformation in the Hamiltonian formulation: it is the transformation $q \to q + \epsilon$, $\mu \to \mu + \epsilon$ 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". In the above sense, Pitts argues it is trivial: it doesn't change the physical content of the theory, and therefore it is just a more complex redefinition of the original theory.

To further push this point, Pitts shows that we can do the same thing to reach the conclusion that second-class constraints generate gauge transformations, which is arguably a *reductio ad absurdum*. Take the Lagrangian for Electromagnetism, but add a photon mass term $-\frac{1}{2}m^2(\vec{A}^2 - V^2)$. This is called "Proca Electromagnetism". The primary constraint is the same, but the secondary constraint has an additional term of m^2V. This has the consequence that both constraints are second-class: $\{p_V, \nabla \cdot p_{\vec{A}} + m^2V - \rho\} = m^2$.

The time derivative of the secondary constraint fixes the value of λ: one gets $\lambda = \nabla\cdot\vec{A}$. Therefore, one can remove the arbitrariness in the Total Hamiltonian. In particular, the primary second-class constraint generates a transformation that takes $V \to V + \alpha$ where α is an arbitrary function of time, as in ordinary Electromagnetism, but it does not generate a gauge transformation. Consider the equations of motion for p_v:

$$\frac{\partial p_V}{\partial t} = \nabla \cdot p_{\vec{A}} - \rho + m^2 V \tag{9}$$

The right-hand side is just the secondary constraint and so is equal to 0. But if we transform $V \to V + \alpha$, the right-hand side is equal to $m^2\alpha \neq 0$.

However, let's consider "de-Ockhamizing" this theory by replacing V with $V + \mu$ in the equations of motion where μ is arbitrary. Now the above equation of motion is satisfied when we transform $V \to V + \alpha$, since the right-hand side is just equal to an arbitrary function. Moreover, we can think of the de-Ockhamized equations of motion as resulting from an "extended Hamiltonian", where V is replaced by $V + \mu$.

Although extending the Hamiltonian in this way isn't to add a linear combination of constraints, Pitts suggests that it is analogous to what Pooley & Wallace (2022) do in order to recover the claim that arbitrary combinations of first-class constraints generate gauge transformations in Electromagnetism: one redefines a quantity by adding a new variable. In doing so, one introduces new symmetries, but these should not be regarded as "genuine" gauge transformations; if they were, then one would have to conclude that second-class constraints generate genuine gauge transformations as well. Therefore, Pitts concludes, we should not think that the Extended Hamiltonian supports the claim that gauge transformations are generated by arbitrary combinations of first-class constraints in a non-trivial sense.

6 Response to the Triviality Argument

Let us look more closely at the first example presented by Pitts. Recall that the claim is that the Lagrangian $L = \frac{1}{2}(\dot{q} - \dot{\mu})^2$ has *more* symmetries than the Lagrangian $L = \frac{1}{2}\dot{q}^2$ even though they have the same empirical content. The sense in which it has more symmetries is supposed to be that if we just consider the variable q, then for the original Lagrangian, we can only transform q by spatial translations and boosts and preserve the equations of motion, while for the "de-Ockhamized" Lagrangian, we can transform q by an arbitrary function of time and preserve the equations of motion (with a corresponding change to μ).

There are two kinds of comparison here: First, there is a comparison regarding *empirical content*. Second, there is a comparison regarding *symmetries*. In order to maintain simultaneously that the Lagrangians are empirically equivalent and that they have different symmetries, it must be that we are comparing the Lagrangians in the same way when we make this claim. So, let us consider under what standard of comparison one can make these claims.

Starting with the claim that the two Lagrangians are empirically equivalent, this seems to rely on taking q to represent the position of the particle in the first Lagrangian and $q - \mu$ to represent the position of the particle in the second Lagrangian, since these quantities satisfy the same equations of motion, namely that the second derivative is equal to 0.

However, if we identify q in the first Lagrangian with $q - \mu$ in the second Lagrangian, then we also should compare the transformations of q that preserve the E-L equations for the first Lagrangian with the transformations of $q - \mu$ that preserve the E-L equations for second Lagrangian. But these transformations are the same: the only transformations of $q - \mu$ that preserve the equations of motion for the second Lagrangian are spatial translations and boosts. 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 Pitts' position that the second Lagrangian has additional gauge 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: the transformation $q \to q + \epsilon$, $\mu \to \mu + \epsilon$, where ϵ is an arbitrary function of time, is a symmetry of the second Lagrangian that is not a symmetry of the first Lagrangian and that preserves the same form of the equations of motion. But under the identification of q with $q - \mu$, this transformation *is* a symmetry of the first Lagrangian – it is the identity transformation on q. If instead one said that the transformation $q - \mu \to q - \mu + \epsilon$ is a symmetry of the second Lagrangian by taking μ to be arbitrary, then this would be an 'additional' symmetry, but it would also mean that the Lagrangians are not empirically equivalent via an identification of q with $q - \mu$; one describes a particle moving in a straight line and the other describes a particle whose dynamics is arbitrary. Either way, one cannot simultaneously claim that one has added a new symmetry and preserved the empirical content under the identification of q in the first Lagrangian with $q - \mu$ in the second Lagrangian.

Another way to compare the Lagrangians is to identify the quantity q as representing the position of the particle in both Lagrangians and take μ in the second Lagrangian to represent an additional (perhaps arbitrary) variable. On this standard of comparison, there is a sense in which one has added a symmetry of q by moving to the second Lagrangian: we can transform q by an arbitrary function of time and preserve the equations of motion. But now, we have that the equation of motion for q is: $\frac{d^2q}{dt^2} = \frac{d^2\mu}{dt^2}$. This is a different equation of motion for q compared to the original Lagrangian since it describes a situation where the position of the particle is either an arbitrary function of time, when μ is arbitrary, or where the particle moves in the same way as μ, when μ is a dynamical variable. Therefore, under this standard of comparison, the two Lagrangians are not empirically equivalent, and so it is not a trivial reformulation.

The upshot is that one cannot simultaneously maintain that the two Lagrangians are empirically equivalent and that one has more symmetries than the other. The same is true of the second example in Section 5 of Proca Electromagnetism. For $V \rightarrow V + \alpha$ to be a gauge transformation, V must be arbitrary in the equations of motion. But V is not arbitrary in the original equations of motion; the equations of motion for V are:

$$\frac{\partial V}{\partial t} = \lambda$$

where λ satisfies $\lambda = \nabla \cdot \vec{A}$, and so is not arbitrary. Therefore, if the de-Ockhamization involves replacing V with an arbitrary function of time $\mu' = V + \mu$, the two equations of motion are not empirically equivalent.[2] If instead one wants to maintain that $V + \mu$ plays the same role as V in the original equation of motion, then the equations would be empirically equivalent but they would also have the same symmetries: the transformation $V + \mu \rightarrow V + \mu + \alpha$ would not be a symmetry if $V + \mu$ is understood to be a relabelling of V.

Consequently, there is a kind of trivial reformulation that one can invoke in the examples that Pitts provides, but it isn't the kind where we add new symmetries. Indeed, we should not be surprised that changing the symmetries of the theory in general has new empirical consequences: symmetries tell us which physical situations are equivalent according to a theory, and so theories with different symmetries will disagree about the physical possibilities.

[2]Indeed, if we replace V everywhere in the Total Hamiltonian with an arbitrary function μ', then we would not have any secondary constraints and there would be one first-class constraint that generates a gauge transformation that shifts V by an arbitrary function of time, as one would expect if V is decoupled from the equations of motion. This would be a different theory from Proca Electromagnetism.

7 Is the Extended Hamiltonian Trivial?

We have established that there is a kind of reformulation of a theory that is trivial but does not correspond to adding new symmetries to a theory. On the other hand, there is a kind of reformulation that does change the symmetries of a theory, but in the examples that Pitts gives, this reformulation leads to a Lagrangian that is physically distinct. So the natural question is: What kind of reformulation is going on in the case where one moves from the Total Hamiltonian to the Extended Hamiltonian?

Let us consider this question in the context of classical Electromagnetism. Again, we need to consider what the standard of comparison is supposed to be. If the replacement of V with $V + \mu$ corresponds to a mere relabeling, then the move to the Extended Hamiltonian appears trivial. But this would mean that we understand V in the Total Hamiltonian and $V + \mu$ in the Extended Hamiltonian to have the same symmetries. This does not seem to be what is going on; the Extended Hamiltonian is supposed to come with the addition of new symmetries, namely, the transformations generated by an arbitrary combination of the first-class constraints (that go past the transformations generated by the specific combination of first-class constraints). This suggests that the move to the Extended Hamiltonian corresponds to the second kind of reformulation: one that changes the physical content of the theory. But then, how is this compatible with the claim of Pooley & Wallace (2022) that the Extended Hamiltonian doesn't come with a change in empirical content?

To see what is going on here, I think it is helpful to consider an alternative way of formulating the constrained Hamiltonian formalism as an extension to the geometric way of formulating a Hamiltonian theory on a symplectic manifold. 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 $\{(q_i, p_i), i = 1, ..., N\}$. T^*Q comes equipped with a one-form, the *Poincaré one-form*, given by $\theta = p_i dq^i$. The corresponding two-form is given by $\omega = \mathbf{d}\theta = dp_i \wedge dq^i$, which is symplectic.

Given a function f, one can uniquely define a smooth tangent vector field X_f through:

$$\omega(X_f, \cdot) = \mathbf{d}f \tag{10}$$

where $\{\cdot\}$ represents any vector field tangent to T^*Q. In particular, one can uniquely define a vector field corresponding to the Hamiltonian H through $\omega(X_H, \cdot) = dH$. This provides an alternative way to write Hamilton's equations, since $\{f, H\} = \omega(X_f, X_H) = df(X_H) = \mathcal{L}_{X_H}(f)$. If we define the flow parameter of X_H to be time, then this says that $\{f, H\} = \frac{df}{dt}$, which is Hamilton's equation.

We can understand the constraints $\gamma_j(q, p) = 0$ as giving rise to a smooth, embedded sub-manifold of phase space, Σ, which we call the *constraint surface*. The first-class constraints are those constraints whose associated vector field is tangent to Σ, while the second-class constraints are those constraints whose associated vector field is not tangent to Σ.

We can define an induced two-form on the constraint surface $\tilde{\omega}$ as the pullback along the embedding $i : \Sigma \to T^*Q$ of ω. This induced two-form is in general *degenerate* i.e. it is not invertible. In particular, the null vector fields of $\tilde{\omega}$ – the vector fields X_j for which $\tilde{\omega}(X_j, \cdot) = 0$ – are the vector fields associated with the first-class constraints. This means that one cannot associate a unique vector field with any smooth function on the constraint surface through the equation $\tilde{\omega}(X_f, \cdot) = \mathbf{d}f$, since if X_f satisfies this equation, so does $X_f + X_{\phi_j}$ where X_{ϕ_j} are the vector fields associated with the first-class constraints. We call the geometry of such a surface *presymplectic*. The integral curves of the null vector fields are called the *gauge orbits*. Equivalently, the gauge orbits consist of the set of points that can be joined by a curve with null tangent vectors.

Recall that the Total Hamiltonian is the equivalence class of Hamiltonians defined up to arbitrary combinations of primary (first-class) constraints. These Hamiltonians are equivocated when we move to the submanifold of T^*Q defined by the satisfaction of the primary constraints. We call this submanifold the *primary constraint surface*. Therefore, we can think of the theory described by the Total Hamiltonian geometrically as the object $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$ where Σ_p is the primary constraint surface, $\tilde{\omega}_p$ is the presymplectic two-form defined intrinsically on the primary constraint surface, H is the Hamiltonian restricted to the primary constraint surface, and φ_i are the secondary constraints. The dynamics is given by two equations, $\tilde{\omega}_p(X_H, \cdot) = \mathbf{d}H$ and $\varphi_i = 0$. Notice that on the primary constraint surface, the solution to $\tilde{\omega}_p(X_H, \cdot) = \mathbf{d}H$ is unique only up to arbitrary combinations of vector fields associated with the primary first-class constraints (the null vector fields of $\tilde{\omega}_p$), which characterizes the sense in which there are multiple, equivalent solutions to the equation of motion using the Total Hamiltonian.

Similarly, the Extended Hamiltonian is the equivalence class of Hamiltonians defined up to arbitrary combinations of all the first-class constraints, which are equivocated on the surface defined by the satisfaction of the primary and secondary constraints. We call this the *final constraint surface*. Therefore, we can think of the theory described by the Extended Hamiltonian as the object $(\Sigma_f, \tilde{\omega}_f, H)$ where Σ_f is the final constraint surface, $\tilde{\omega}_f$ is the presymplectic two-form defined intrinsically on the final constraint surface, and H is Hamiltonian restricted to the final constraint surface. The solution to $\tilde{\omega}_f(X_H, \cdot) = \mathbf{d}H$ on the final constraint surface is unique only up to arbitrary combinations of vector fields associated with the first-class constraints.[3]

This characterization naturally provides a sense in which the Extended Hamiltonian theory regards more solutions as equivalent compared to the Total Hamiltonian theory: there are solutions that the Total Hamiltonian theory distinguishes between that the Extended Hamiltonian theory does not distinguish between (when we consider these solutions on the final constraint surface). More generally, we can now give a precise characterization of the fact that the Extended Hamiltonian theory has more symmetries than the Total Hamiltonian theory: the null vector fields of the two-form on the primary constraint surface are a subset of the null vector fields of the two-form on the final

[3]For further detail on why the Extended Hamiltonian formalism is naturally represented on the final constraint surface, see Bradley (2024*a*).

constraint surface, and so the gauge transformations – the transformations along the gauge orbits – of the Total Hamiltonian theory are a subset of the gauge transformations of the Extended Hamiltonian theory (when considered on the final constraint surface).

However, restriction to the final constraint surface does not come with a change in empirical content. One way to put the reason for this is that the solutions to the equations of motion on the final constraint surface correspond precisely to the solutions to the equations of motion on the unconstrained symplectic manifold that satisfy the constraints (equivalently, the solutions to the equations of motion on the primary constraint surface that satisfy the secondary constraints); they are the solutions projected to the final constraint surface. And the symmetries of *these* solutions, i.e. the solutions defined on the final constraint surface, are just the symmetries given by the Extended Hamiltonian theory. Therefore, as long as constraints are considered a physical requirement, the Extended Hamiltonian theory is empirically equivalent to the Total Hamiltonian theory.[4]

Indeed, I think that the geometric framing helps to see exactly what the move to the Extended Hamiltonian formalism corresponds to: it corresponds to moving to a theory with *less structure*, since it is a theory with more symmetries, that nonetheless has the same empirical content. In other words, it shows that the theory defined by the Total Hamiltonian has a kind of 'excess structure': there are points and solutions distinguished by the theory formulated on the primary constraint surface that are not distinguished by a theory that maintains the same physical content and yet has less structure. In particular, the points that lie along the integral curves of the vector fields associated with the secondary first-class constraints are symmetry related in the Extended Hamiltonian formalism, but not in the Total Hamiltonian formalism. We have reason to think that they should be symmetry related precisely because the differences between these points do not seem to be playing any role in the empirical content of the theory.

[4] Gryb & Thébault (2023, ch.8) argue that the secondary constraints can be derived directly from Noether's Second Theorem, demonstrating that the secondary constraints are not just required for consistency on the Hamiltonian side; they are also required on the Lagrangian side. This provides further support for the claim that the content of the theory is given by the solutions on the final constraint surface.

To spell out the sense in which the Extended Hamiltonian formalism has less structure than the Total Hamiltonian formalism more precisely, let us define the theories in category theoretic terms.[5] Take the category **TotHam** to have as objects the models $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$, and let us take the arrows between objects $(\Sigma_p, \tilde{\omega}_p, H, \varphi_i)$, $(\Sigma_p, \tilde{\omega}'_p, H', \varphi'_i)$ to be the diffeomorphisms $f : \Sigma_p \to \Sigma_p$ 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, let us take the category **ExtHam** to have 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 $g : \Sigma_f \to \Sigma_f$ such that $g^*(\tilde{\omega}'_f) = \tilde{\omega}_f$ and $g^*(H') = H$.

Relations between theories are described by functors between the categories representing those theories. A functor $F : \mathcal{C} \to \mathcal{D}$ 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 $F : \mathrm{hom}(A, B) \to \mathrm{hom}(F(A), F(B))$ induced by F is surjective, where $\mathrm{hom}(A, B)$ 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 (2016*b*)) that a theory represented by category $\mathcal{C}$ has more structure than a theory represented by category $\mathcal{D}$ if a functor $F : \mathcal{C} \to \mathcal{D}$ is not full (but is faithful and essentially surjective). In this case, we say that *F forgets (only) structure.*

Consider the functor $F : \mathbf{TotHam} \to \mathbf{ExtHam}$ 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:

[5] Category theory has been used in several places to give a precise sense in which one theory has less structure than another. See, 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).

Proposition 1: $F : \mathbf{TotHam} \to \mathbf{ExtHam}$ forgets (only) structure.[6]

Inasmuch as forgetting structure in category-theoretic terms captures what it is for one theory to have less structure than another, this proposition gives a precise characterization of the sense in which the Extended Hamiltonian formalism has less structure than the Total Hamiltonian formalism. Therefore, contra Pitts, the move to the Extended Hamiltonian should not be thought of as simply redefining a theory in terms of new quantities. Rather, it defines a new theory that removes structure from the Total Hamiltonian theory. This means that there is a genuine disagreement about the gauge transformations between the two theories. But even further, it suggests that the standard view that arbitrary combinations of first-class constraints is correct, by Pitts' own lights: moving to the Extended Hamiltonian is the opposite of "de-Ockhamization", in the sense that it is a simpler theory (in terms of structure) than the Total Hamiltonian theory. So Pitts was right that simplicity considerations matter in the debate between the Total and Extended Hamiltonian, but rather than these considerations pushing one towards to the Total Hamiltonian, they push one towards the theory that captures all of the symmetries that the theory is naturally understood as having: the theory characterized by the Extended Hamiltonian.[7]

8 Are There Different Notions of Gauge Transformations?

A different response to Pitts' triviality argument is found in Mozota Frauca (2024), where it is argued that we should think of the symmetries of the Extended Hamiltonian theory as being of a different kind to either the symmetries of the Total Hamiltonian or the trivial kind argued for by Pitts. In more detail, Mozota Frauca (2024) suggests that

[6]See Appendix for proof.

[7]In Bradley (2024*a*), I argue that the theory characterized by the Extended Hamiltonian is the natural one from the perspective of (the geometric formulation of) the constrained Hamiltonian formalism. We can see the arguments here as complementary to this view: the theory formulated on the final constraint surface is not only the *natural* one, it is also the *simpler* one.

there are three notions of a gauge transformation:

1. Trivial gauge transformation: A transformation that preserves the empirical content of the theory, even if one needs to introduce compensating fields and redefine what counts as physical.

2. Genuine or original gauge transformation: A gauge transformation associated with symmetries of the original Lagrangian action.

3. Extended gauge transformations: A transformation that preserves the empirical content of the theory, with no need to redefine physical fields, even if one needs to introduce compensating fields and redefine accessory fields.

Mozota Frauca (2024) argues that the symmetries of the Extended Hamiltonian should be thought of as the third kind of gauge transformation: in order for them to be symmetries, one needs to introduce compensating fields and redefine accessory fields in the Total Hamiltonian – which is just to move to the Extended Hamiltonian – but because one doesn't have to redefine the 'physical' variables, they should be thought of as closer to genuine gauge transformations than the trivial gauge transformations.

I have argued against (1) as a way of adding symmetries to a theory while preserving empirical content: one cannot claim that any phase space function can generate 'new' gauge transformations by reformulating the theory. On the other hand, Mozota Frauca (2024) seems to agree with Pitts that one can add symmetries in this way, but argues that the way one adds gauge transformations through the Extended Hamiltonian is a special version of this process that can be interpreted differently.

In order to evaluate the arguments in Mozota Frauca (2024), let us consider their motivating example. Consider the following modification of the Lagrangian $L = \frac{1}{2}m\dot{q}^2$:

$$L = \frac{1}{2}m\dot{q}^2 + (\dot{\mu}q + \mu\dot{q}) \tag{11}$$

From this modified Lagrangian, the canonical momentum becomes $p = m\dot{q} + \mu$ and the Hamiltonian $H = \frac{(p-\mu)^2}{2m} - q\dot{\mu}$. These give rise to the equations of motion $\dot{q} = \frac{p-\mu}{m}$

and $\dot{p} = \dot{\mu}$, which together give $m\ddot{q} = 0$. Therefore, the equations of motion for q remain unchanged with this modification.

Mozota Frauca (2024) argues that this reformulation "adds a new symmetry" – the transformation $p \to p+\epsilon, \mu \to \mu+\epsilon$ – since the equation of motion for p is now dependent on the arbitrary function μ. Therefore, Mozota Frauca (2024) argues, this is similar to the examples that Pitts gives in that we have added a new symmetry transformation and preserved the empirical content of the theory. On the other hand, the difference with the examples that Pitts gives is that to say that the empirical content is the same, we only have to point to the fact that the equation of motion for the physical variable, namely q, remains unchanged. We do not need to take some other mathematical quantity to represent position, and it is only the equation of motion for the 'accessory' variable p that changes: in the original theory the equation of motion is $\dot{p} = 0$, while in the second the equation for p is $\dot{p} = \dot{\mu}$. Since p is 'accessory', we don't have to worry about this affecting the physical content of the theory.

But notice that this modified Hamiltonian does not have any new "genuine" gauge symmetry, since there are no constraints. So in what sense has one added a new symmetry? Indeed, if we take the original Lagrangian $L = \frac{1}{2}m\dot{q}^2$ with equation of motion $m\ddot{q} = 0$ and define some quantity $p = m\dot{q} + \mu$ where μ is an arbitrary variable, then we can rewrite the equations of motion as $\dot{p} - \dot{\mu} = 0$. But the symmetries of this equation of motion are just $p \to p + \epsilon, \mu \to \mu + \epsilon$. So there is a sense in which this "new" symmetry is already contained in the original Lagrangian.

What appears to be the puzzle is that from the Hamiltonian perspective, if we identify the quantity represented by "p" in the original and modified Hamiltonian, then the two theories disagree about the equations of motion for this quantity: the original theory says that p is constant, while the modified theory says that p is arbitrary. Therefore, the symmetries of the equations of motion for p are different. We might want to conclude that the two theories must be empirically distinct, since we have added a new symmetry. But, Mozota Frauca (2024) argues, because p is accessory we need not be led to this conclusion.

The issue with this line of reasoning is that if we are already identifying the variable "q" across the two theories – which is implicitly done in order to claim that the equations of motion for q remain the same – then we cannot also claim that "p" refers to the same quantity across the theories, since $p = m\dot{q}$ in the original theory and $p = m\dot{q} + \mu$ in the modified theory. In other words, in identifying the equation of motion of the original theory $\ddot{q} = \frac{\dot{p}}{m} = 0$ and the modified theory $\ddot{q} = \frac{\dot{p}-\dot{\mu}}{m} = 0$, we are already implicitly identifying p with $p - \mu$. But the symmetries of p in the original theory and $p - \mu$ in the modified theory are exactly the same. Therefore, there doesn't seem to be any sense in this example that one can claim that the modified theory adds a new symmetry while keeping the empirical content the same.

Mozota Frauca (2024) goes on to argue that moving to the Extended Hamiltonian – particularly in the case of classical Electromagnetism – is analogous to this example, in the sense that the the Extended Hamiltonian changes the equations of motion for the 'accessory' variables but not the physical variables. I take my earlier arguments to show why moving to the Extended Hamiltonian is not like this example: unlike in the above example, moving to the Extended Hamiltonian is not a matter of redefining quantities; one moves to a wider equivalence class of Hamiltonians with correspondingly more symmetries of the equations of motion for the same variables.

Therefore, we shouldn't think of the symmetries of the Extended Hamiltonian as different in kind from the symmetries of the Total Hamiltonian; rather, they are both "genuine" symmetries of theories formulated on different presymplectic manifolds.

9 Conclusion

In this paper, I argued that a trivial reformulation of a theory in the way suggested by Pitts (2022, 2024) cannot be used to add new gauge symmetries to a theory, and I showed that the move to the Extended Hamiltonian is therefore not a trivial reformulation.

Moreover, I argued that contrary to the suggested terminology used by Pitts of "de-Ockhamization", the move to the Extended Hamiltonian should rather be viewed as a kind of "Ockhamization": it is a move to a theory with less structure that nonetheless has the same empirical content.

This argument is significant for several reasons. First, it highlights that the debate about the correct characterization of the gauge transformations is not simply a matter of terminology: although there is an agreement about the gauge transformations given a particular form of the Hamiltonian, choosing a particular form of the Hamiltonian is to choose a different formulation of the constrained Hamiltonian formalism. Although these formulations arguably have the same empirical content, they are not structurally equivalent, and so inasmuch as one takes structurally inequivalent theories to be inequivalent, the debate must be about which formulation is the right one.

Second, this argument vindicates the orthodox view according to which arbitrary combinations of first-class constraints generate gauge transformations in the Hamiltonian formalism: they are the symmetries of the Extended Hamiltonian theory, and we have reason to think that this theory is preferred over the Total Hamiltonian theory.

However, there is still an open question suggested by Pitts (2014) about how to reconcile this argument with the fact that the symmetries of the Lagrangian seem to match the symmetries of the Total Hamiltonian. That is, if the Extended Hamiltonian theory is a genuinely different theory, and is the preferred one, what should one make of the apparent inequivalence between the Lagrangian and Hamiltonian perspectives? Indeed, I think that one can resolve this puzzle[8], but as a more general remark, I think that it is this question that should be a stronger focus in the debate regarding the definition of the gauge transformations.

[8] In Bradley (2024*b*), I argue that one can formulate Lagrangian mechanics in a such way that renders the Lagrangian framework for gauge theories equivalent to the Extended Hamiltonian formalism.

10 Appendix

Proof of Proposition 1

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.

To show that F is not full, we need to show that there is an arrow $g : (\Sigma_f, \tilde{\omega}_f, H) \rightarrow (\Sigma_f, \tilde{\omega}'_f, H')$ such that $g \neq F(f)$ for some arrow f in **TotHam**. To do this, we will show that there are transformations along the the vector fields associated with the secondary first-class constraints that are not arrows in **TotHam**, but their restriction to Σ_f are arrows in **ExtHam**. 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. Recall that $\tilde{\omega}_p(X_{\varphi_j}, \cdot) = d\varphi_j \neq 0$ for the secondary first-class constraints φ_j, which tells one the change of a function along the vector fields associated with the secondary first-class constraints. So let us consider f to be the flow of the vector field associated with $\alpha^j d\varphi_j$ where each α^j is an arbitrary function of the canonical coordinates. 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 $\alpha^j d\varphi_j$ is exact i.e. can be represented as $d\gamma$ for some scalar field γ. Therefore, 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 $\tilde{\omega}_p(X, \cdot) = \alpha^j d\varphi_j$. This means that in these cases the flow of the vector field associated with $\alpha^j d\varphi_j$ does not consist of symplectomorphisms[9], and so $f^*(\tilde{\omega}_p) \neq \tilde{\omega}_p$. Therefore, f is not an arrow in **TotHam** (for every choice of α^j).

However, let us now consider the transformations along the vector fields associated with the secondary first-class constraints on Σ_f. That is, consider the (gauge) transfor-

[9]See Abraham & Marsden (1987) Proposition 3.3.6.

mation $g : \Sigma_f \to \Sigma_f$ 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 similarly represent this 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 the final constraint surface, $\alpha^j d\varphi_j = 0$. But this means that one *can* associate a vector field Y with $\alpha^j d\varphi_j$ via $\tilde{\omega}_f(Y, \cdot) = \alpha^j d\varphi_j$ since $d(\alpha^j d\varphi_j) = 0$ on Σ_f . Therefore, the flow of Y on Σ_f consists of symplectomorphisms. Moreover, $g^* H = H$ because H is gauge-invariant on the final constraint surface. Therefore, g is an arrow in **ExtHam** (for all choices of α^j). This implies that there are arrows g of **ExtHam** such that $g \neq F(f)$ for some arrow f in **TotHam**. Therefore, we can conclude that F is not full.

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 $f = g$. So suppose that $f|_{\Sigma_f} = g|_{\Sigma_f}$. If f, g are global symplectomorphisms, 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. If f, g are local symplectomorphisms (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. So F is faithful.

References

Abraham, R. & Marsden, J. E. (1987), *Foundations of Mechanics, Second Edition*, Addison-Wesley Publishing Company, Inc.

Bradley, C. (2024*a*), 'Do first-class constraints generate gauge transformations? a geometric resolution'. Unpublished Manuscript.

Bradley, C. (2024*b*), 'The relationship between lagrangian and hamiltonian mechanics: The irregular case'. Unpublished Manuscript.

Bradley, C. & Weatherall, J. O. (2020), 'On representational redundancy, surplus structure, and the hole argument', *Foundations of Physics* **50**(4), 270–293.

Dirac, P. A. M. (1964), *Lectures on Quantum Mechanics*, Dover.

Gryb, S. & Thébault, K. (2023), *Time Regained: Volume 1: Symmetry and Evolution in Classical Mechanics*, Oxford University Press.

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.

Mozota Frauca, Á. (2024), 'In which sense can we say that first-class constraints generate gauge transformations?', *Philosophy of Physics* **2**(1), 10.

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.

Pitts, J. B. (2014), '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', *https://philsci-archive.pitt.edu/21423/* .

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.

Pooley, O. & Wallace, D. (2022), 'First-class constraints generate gauge transformations in electromagnetism (reply to pitts)', *arXiv preprint arXiv:2210.09063* .

Weatherall, J. O. (2016*a*), 'Are newtonian gravitation and geometrized newtonian gravitation theoretically equivalent?', *Erkenntnis* **81**, 1073–1091.

Weatherall, J. O. (2016*b*), '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.