

How to Choose a Gauge? The case of Hamiltonian Electromagnetism

Henrique Gomes and Jeremy Butterfield*

Trinity College, University of Cambridge

October 21, 2021

Abstract

We develop some ideas about gauge symmetry in the context of Maxwell’s theory of electromagnetism in the Hamiltonian formalism. One great benefit of this formalism is that it pairs momentum and configurational degrees of freedom, so that a decomposition of one side into subsets can be translated into a decomposition of the other. In the case of electromagnetism, this enables us to pair degrees of freedom of the electric field with degrees of freedom of the vector potential. Another benefit is that the formalism algorithmically identifies subsets of the equations of motion that represent time-dependent symmetries. For electromagnetism, these two benefits allow us to define gauge-fixing in parallel to special decompositions of the electric field. More specifically, we apply the Helmholtz decomposition theorem to split the electric field into its Coulombic and radiative parts, and show how this gives a special role to the Coulomb gauge (i.e. $\text{div}(\mathbf{A}) = 0$). We relate this argument to Maudlin’s (2018) discussion, which advocated the Coulomb gauge.

Contents

1	Introduction	2
1.1	Three morals	3
2	The Hamiltonian framework	6
2.1	The Lagrangian framework and indeterminism	7
2.2	The Hamiltonian formalism and constraints	9
2.2.1	Off the constraint surface	12
3	The case of electromagnetism	14
3.1	A toy example of natural coordinate choices on phase space	14
3.2	Hamiltonian treatment of electromagnetism	15
3.3	Interpreting the Gauss constraint	18
3.4	Radiative and Coulombic parts of \mathbf{E} and \mathbf{A}	19
4	Conclusion and outlook	23
A	Holonomies and non-separability	24

*gomes.ha@gmail.com, jb56@cam.ac.uk

1 Introduction

Recently, Maudlin (2018) has exhorted not only philosophers, but also modern theoretical physicists, to be clearer about their theories' ontology. He writes: 'both the glory and the bane of modern physics is its highly mathematical character. This has provided both for the calculation of stunningly precise predictions and for the endemic unclarity about the physical ontology being postulated' Maudlin (2018, p.5). We agree. Today, more than in yesteryear, a theory's mathematical formalism is often interpreted *en bloc*. No special care is taken to specify: which parts represent ontology, 'what there is' (and within that: what is basic or fundamental, and what derived or composite); and which parts represent 'how it behaves' (which Maudlin (p. 4) calls 'nomology': in particular, dynamics); and which parts represent nothing physical, but instead mathematics (which, though unphysical, can of course be invaluable for calculation).

We also endorse Maudlin's programme to develop presentations of our physical theories that are clear about these distinctions; (though one should of course accept that what we usually consider to be one theory may admit two such meritorious presentations—no uniqueness claim is required.)¹

As a case study of his programme, Maudlin considers a theory that, in the hierarchy of mathematical sophistication of modern physics, sits relatively low, viz. classical electromagnetism. He then applies the results to assess some proposed interpretations of the Aharonov-Bohm effect.

In this paper, we also will consider classical electromagnetism, and with an overall aim similar to Maudlin's—to clarify interpretative issues. But we focus on the Hamiltonian formulation of the theory (Section 3): which we lead up to, by first expounding the Hamiltonian framework for finite-dimensional systems, rather than for field theories (Section 2). This exposition will include comparison with the Lagrangian framework, and the treatment of constraints. Despite our adopting the Hamiltonian framework, some of our conclusions, technical as well as interpretative, will be concordant with Maudlin's: in particular, about the Coulomb gauge having a special role.

It is often remarked that the Hamiltonian framework's use of a time parameter, and-or of a '3+1' split of spacetime, carries the price, for special relativistic theories like electromagnetism, that one loses manifest Lorentz invariance. We of course accept that this is a limitation; though we note that often (including in our discussion below) the Lagrangian framework also uses a time parameter.

But for interpreting the gauge aspects of electromagnetism, the Hamiltonian framework has two countervailing benefits. First: its pairing of momentum and configurational degrees of freedom means that we can pair degrees of freedom of the electric field with degrees of freedom of the vector potential. Second: the Hamiltonian framework illuminates symmetries that are time-dependent of the sort associated (in both it and the Lagrangian framework), with constraints and the failure of determinism. This illumination comes from the way the Hamiltonian framework algorithmically identifies subsets of the equations of motion that represent such symmetries (viz. 'the Dirac algorithm'). Combining these benefits gives an illuminating, and interpretatively clear, splitting of the electric field into its Coulombic and radiative parts. Besides, this splitting gives a special role to the Coulomb gauge, i.e. $\text{div}(\mathbf{A}) = 0$.²

So our plan is as follows. In Section 2, we will present the features of the Hamiltonian framework that will be useful for us. We will first describe how constraints emerge from

¹So this programme is consistent, in particular, with rejecting the logical positivists' aim of presenting physical theories with a once-for-all division of fact and convention: a rejection we share with e.g. Putnam (1975).

²Maudlin also advocates this gauge, but for very different reasons than us. We discuss the differences in Sections 3.3 and 4 (cf. also footnote 4), but in short: he makes a controversial ontological proposal, while we, less contentiously, see the gauge as "merely" natural, because of its reflecting the splitting of the electric field.

the Legendre transform of the Lagrangian, and the implications for (in)determinism. Since constraints will play an important role in the paper, we will spend much of the Section justifying their special status amongst the equations that specify a theory. In particular, we will see that—although constraints imply that, roughly speaking, many velocities \dot{q} in the Lagrangian framework are mapped by the Legendre transformation to one momentum p in the Hamiltonian framework—nevertheless, under reasonable conditions (viz. the constraints being first-class and so generating symmetries), the gauge orbits in the constraint surface within the Hamiltonian framework correspond one-to-one to sub-manifolds of the Lagrangian state space: (cf. the two paragraphs after equation 2.14).

In Section 3 we will apply these ideas to electromagnetism: the Gauss constraint emerges from the Dirac algorithm, and it is the generator of gauge symmetries. Finally, we will split the electric field into the part that is uniquely fixed by the Gauss constraint—and so by the instantaneous distribution of charges—and the remainder. We thereby find a symplectically corresponding decomposition of the gauge potential into a part that is pure gauge and a remainder: call them X and Y , respectively. If we single out the scalar degree of freedom of the electric field that is fixed by the Gauss constraint, by representing it as a gradient, it turns out that, on the other side of the symplectic correspondence, Coulomb gauge is singled out. More precisely: if we demand that in our split of the gauge potential, the part Y of the gauge potential that remains after we extract the pure gauge part X , is dynamically independent from the part of the electric field that is fixed by the Gauss constraint, then this non-gauge part Y will obey the Coulomb gauge equation.

In other, somewhat less technical, words: our main idea is as follows. The Coulombic part of the electric field is its electrostatic-like component, which is determined by the instantaneous distribution of charges. The electric field has the gauge potential as its conjugate. And as to the Coulombic part of the electric field, there is a part of the gauge potential that is not at all conjugate to it (i.e. is symplectically orthogonal to it). That part of the potential satisfies the Coulomb gauge condition.³ Besides, this result is worth expounding, since one usually thinks of a choice of gauge as motivated by calculational convenience, often for a specific problem—and so from a general theoretical perspective, as completely arbitrary: whereas this result shows that the choice is related to a physically natural, and general, splitting of the electric field.

In the short final Section 4, we conclude; and sketch how this work can be set against the background, not just of Maudlin’s discussion, but of other similar enterprises, viz. those pursued in Gomes (2019b, 2021a,b,c,d); Gomes et al. (2021).

We now end this introduction by presenting three overall morals, especially in relation to the Aharonov-Bohm effect.

1.1 Three morals

Elsewhere one of us (Gomes, 2021c, Sec. 3), has urged an interpretation of gauge theories—more specifically: the family of Abelian and non-Abelian classical Yang-Mills theories, which includes electromagnetism—that takes a single physical possibility to correspond to an entire gauge-equivalence class.⁴ In Maudlin’s jargon of ontology (‘what there is’) and nomology (‘how

³So the nomenclature is confusing, since the Coulombic part of the electric field corresponds to, i.e. is conjugate to, the *other* part of the gauge potential than that which satisfies the Coulomb gauge condition. Thus one might facetiously propose that for clarity, we should: (i) associate the adjective ‘Coulombic’ with Coulomb’s historical work on electrostatics, and so apply it only to the electrostatic-like component of the electric field; and (ii) as regards the gauge potential, re-name the gauge condition as the ‘*ortho-Coulomb condition*’.

⁴So this is analogous to what is often called ‘sophisticated substantivalism’ in the philosophy of general relativity; (cf. footnote 7). It is also disanalogous to Maudlin, since in the enterprise of interpreting gauge, he investigates ontology associated with different representatives of each equivalence class. So his reasons for

it behaves’), this interpretation, in brief, is that:

(i) *the ontology* consists of a field over spacetime, that encodes the relations between charges that interact via a particular type of force (e.g. electromagnetism), i.e. the field encodes a relation of sameness of charges across spacetime called ‘parallel transport’;⁵

(ii) *the nomology* describes how this sameness relation is constrained by the distributions of charges in spacetime.

Besides: although this interpretation was articulated in the Lagrangian (covariant) framework, we see no significant obstacle to an appropriate translation to the Hamiltonian framework (though there are some subtleties about how to do this: cf. the end of Section 2.2 and (Belot, 2003, Sec. 11)). Thus one main aim of this paper is to give such a translation.

In either framework, this summary, (i) and (ii), is of course a “high altitude” view of the theory. Nonetheless, we can already see how it accommodates facts that are characterised only non-locally in terms of their spacetime properties. Thus, for example, one can parallel transport an internal quantity around a loop in spacetime, and whether that quantity comes back to its original value or not can depend on facts *outside* the vicinity of the loop (Gomes, 2021d, Sec. 2.3).⁶

In the standard jargon of electromagnetism, this type of non-locality is associated with the Aharonov-Bohm effect and the introduction of *gauge* and *vector potentials*. The most important lesson of the effect is of course that there is physical significance in gauge. This is highlighted already in the opening of the original paper (Aharonov & Bohm, 1959):

In classical electrodynamics, the vector and scalar potentials were first introduced as a convenient mathematical aid for calculating the fields. It is true that in order to obtain a classical canonical formalism, the potentials are needed. Nevertheless, the fundamental equations of motion can always be expressed directly in terms of the fields alone. In the quantum mechanics, however, the canonical formalism is necessary, and as a result, the potentials cannot be eliminated from the basic equations. Nevertheless, these equations, as well as the physical quantities, are all gauge invariant; so that it may seem that even in quantum mechanics, the potentials themselves have no independent significance. In this paper we shall show that the above conclusions are not correct and that a further interpretation of the potentials is needed in quantum mechanics.

But we would like to stress two points of clarification about this quotation. They will lead to three “morals” which we will announce here, but present in detail in Sections 2 and 3.

(1): First: although the potentials appear prominently in the theory, we need not interpret as physically different, two potentials that are related by a certain type of transformation, labeled *gauge transformation*. Thus the ‘independent significance’ of the gauge potentials needs to be taken *modulo* such transformations.

The lack of physical difference between potentials related by gauge transformations is incorporated in the interpretation that we endorse, glimpsed above from a high altitude—and

advocating the Coulomb gauge are very different from our reasons for seeing it as special.

⁵This is argued by Gomes (2021c, Sec. 3) to bear a straightforward analogy to the more familiar parallel transport of vectors on spacetime, which is fixed by the metric.

⁶In more detail: Suppose we take all the *local* gauge-invariant facts to be given/determined. This includes a field called *curvature* in the Abelian theory (and similarly, traces of products of curvature in the non-Abelian theory). Then suppose we have two possible worlds, w and w' , and we can identify a loop in spacetime in w with one in w' . Even if w and w' agree about the distribution of the local gauge-invariant quantities in the vicinity of the loop, they may disagree about whether the original value of the charges will coincide with their values after parallel transport around the loop. This fact thus depends on some local gauge-invariant quantities *outside* the vicinity of the loop. Compare the discussion of holonomies, at the end of this Section and in the Appendix.

usually articulated mathematically in terms of fiber bundles. This yields a precise sense in which it is only the *equivalence class* of potentials under gauge transformations that encodes the parallel transport of internal quantities. Thus in Yang-Mills theories, gauge-related models correspond to the same physical situation.⁷

If one rejects this first point, then instead of articulating an ontology for the equivalence class of gauge potentials, one would try to articulate ontologies for different choices of an element of the class: i.e. for different choices of ‘gauge-fixing’. As we mentioned in footnote 4, this is Maudlin (2018)’s approach; (cf. also (Mulder, 2021, Sections 8-9)). In the end, Maudlin settles on Coulomb gauge as combining a relativistic nomology with an ontology that requires a preferred foliation of spacetime: and is thus friendly to the pilot-wave/Bohmian interpretations of quantum mechanics that he favours.

(2): Our second point of clarification is that, although the experimental set-ups implemented hitherto for observing the Aharonov-Bohm effect have a quantum component—viz. an effect on the phase of a quantum wave-function—the *classical* theory already interprets the situations that give rise to the effect as physically distinct. The reason—in short, and in an elementary 3-vector formalism—lies in the facts that:

- (i) the gauge transformations for the vector potential are $\mathbf{A} \mapsto \mathbf{A} + \nabla\lambda$, for λ any scalar;
- (ii) the magnetic field $\mathbf{B} \equiv \nabla \wedge \mathbf{A}$ is unchanged by the addition of any curl-free vector to \mathbf{A} , i.e. \mathbf{C} such that $\nabla \wedge \mathbf{C} = 0$; while
- (iii) in a non-simply connected space $\nabla \wedge \mathbf{C} = 0$ does *not* imply that $\mathbf{C} = \nabla\lambda$ for some scalar λ . (Point (iii) is often expressed in the language of forms as: a closed form need not be exact.)

So the gauge-equivalence classes of the vector potential \mathbf{A} are narrower than the equivalence classes given by addition of a curl-free vector \mathbf{C} . So the classical theory discriminates as physically distinct two configurations (field distributions) of \mathbf{A} that (a) do not differ by a gradient $\nabla\lambda$ (are not gauge-equivalent), but that (b) *do* differ by a curl-free vector field \mathbf{C} (and so determine the same magnetic field $\mathbf{B} \equiv \nabla \wedge \mathbf{A}$). We agree that hitherto no *classical* way to experimentally manifest this distinction has been found. But nor is there a no-go theorem vetoing such a classical manifestation: that is, a measurement (ultimately, a pointer-reading) that uses some classical probe to register which of two \mathbf{A} -configurations, that are physically distinct since not differing merely by a gradient $\nabla\lambda$, but that determine the same magnetic field, is realized.

Both these two points of clarification (especially of course, the first) are recognised in the literature.⁸ What matters for us is that the second point is related to classical electromagnetism’s exhibiting a type of *non-locality*. This is usually expressed in terms of holonomies; (though a complete treatment without holonomies can be given—see (Gomes & Riello, 2021, Sec. 6), Gomes (2021a)). Thus Myrvold (2010) notices that, for classical electromagnetism *in vacuo*, and for a simply connected region: one can use the composition laws of the holonomies

⁷So this is analogous to saying, about general relativity, that diffeomorphism-related models of the metric correspond to the same physical possibility: an interpretation that we argue (Gomes & Butterfield (2021a)) is recommended, though not mandated, by the mathematics. As argued at length in (Gomes, 2021c,d), there is thus no major interpretative difference between the gauge transformations of Yang-Mills theories and the diffeomorphisms of general relativity.

⁸The second is well expressed by Belot (1998, pp. 532-533, pp. 550-554), who also addresses judiciously the question of what an interpretation of classical electromagnetism can tell us about a quantum world. A related question concerns what we can learn about the world from the sort of idealized descriptions usually given of the Aharonov-Bohm effect: including, for example, excluding from space the region where the local gauge-invariant quantities such as curvature (cf. footnote 6) are distinct, i.e. excluding the solenoid. For recent controversy about such idealizations, cf. Dougherty (2020); Earman (2019); Shech (2018).

to decompose any gauge-invariant function on this region into gauge-invariant functions of its component—i.e. mutually exclusive, jointly exhaustive—sub-regions. But, Myrvold continues: for a region that is not simply connected, there are certain “large” holonomies that cannot be obtained from composition of holonomies confined to the sub-regions. That is, there are certain global gauge-invariant functions that are not *separable*. Myrvold concludes that this type of non-locality only arises for non-simply connected regions (and is exhibited by the Aharonov-Bohm effect).

At first sight, this conclusion is in tension with our own discussion, in Section 3, of the role in the Hamiltonian framework of the *Gauss constraint*, i.e. the equation, in elementary terms, that the divergence of \mathbf{E} is equal to the charge density ρ (and so vanishes in vacuum, which is the case we shall often focus on): $\nabla \cdot \mathbf{E} = \rho$. The starting-point in our discussion of this equation will be the oft-noted fact that it involves a sort of *non-locality*. For it implies that by simultaneously measuring the electric field flux on all of a large surface surrounding a charge distribution, and integrating, we can ascertain the total amount of charge inside the sphere *at the given instant*. So this non-locality is classical, and regardless of whether the volume enclosing the charge distribution is simply connected (the surface surrounding the charge distribution could be topologically a doughnut, not a sphere)—which is apparently, at odds with Myrvold’s conclusion. But here we should recall that Myrvold’s analysis is restricted to electromagnetism *in vacuo*; and fortunately, when we remove the restriction, the tension disappears. That is: when we allow for charges, the same line of argument that Myrvold uses also proves that there are gauge-invariant functions that are not separable—even for simply connected regions. In short: in the presence of charges, we have non-locality (and non-separability) even for simply-connected regions. The Appendix will give more details.

We can sum up this discussion of our relation to Maudlin’s enterprise, and of the Aharonov-Bohm effect, in three morals, as follows:

(1): Even setting aside Aharonov-Bohm phases, *classical* electromagnetism in the Hamiltonian formalism shows a certain kind of (non-signalling) non-locality, namely in the Gauss constraint.

(2): There is no unique physically preferred split of the field’s degrees of freedom into purely local and purely non-local ones. But each such choice of split can be made to correspond to a choice of gauge-fixing.

(3): In particular, the Coulomb gauge corresponds in this way to a natural choice of splitting of the electric field; (which, incidentally, buttresses some of Maudlin (2018)’s arguments). That is: this gauge follows naturally from considering the Gauss constraint to single out a ‘scalar’ part of the electric field that is determined by the instantaneous distribution of charges.

2 The Hamiltonian framework

One of the most distinctive features of the Hamiltonian framework is the fact that the Dirac analysis of constraints provides an algorithm for discovering gauge symmetries. Once the algorithm succeeds, the equations of motion of a theory are divided into ones that come from constraints and those that we normally think of as generating dynamics.

Since the formalism is slightly unfamiliar to philosophers of physics, in Section 2.1 we will first review, for the case of a non-relativistic mechanical system, how under-determination of motion i.e. indeterminism can arise, in the *Lagrangian* framework. Then we will relate this to the constraints in the Hamiltonian framework that are obtained by a Legendre transformation. Then, in Section 2.2 we will introduce the analysis of symmetry using symplectic geometry. In these two Subsections, we will, broadly speaking, keep the relation between the

Lagrangian and Hamiltonian frameworks simple—though sufficient for our points about the case of electromagnetism—by certain restrictions of scope about types of constraint. (These restrictions are summarised in the two paragraphs after equation 2.14.) Thus we will see that both the Lagrangian and Hamiltonian frameworks show the special status of constraints: they are the generators of symmetries and must be imposed prior to any equation of motion, if there is to be a correspondence between the frameworks. We end the Section with some philosophical remarks about the status of states that are not in the constraint surface.

2.1 The Lagrangian framework and indeterminism

In the Lagrangian formulation of mechanics, the focus is on curves in the space of all possible instantaneous configurations of the system. For the mechanics of N point-particles parametrized by $\alpha = 1, \dots, N$, let this configuration space be called \mathcal{Q} . A *Lagrangian* is a map from the tangent bundle on configuration space to the reals: $\mathcal{L} : T\mathcal{Q} \rightarrow \mathbb{R}$. Once integrated along a curve, the Lagrangian yields the *action functional* as a function from curves γ in \mathcal{Q} into the reals:

$$S(\gamma) := \int_{\gamma} dt \mathcal{L}(q^\alpha(t), \dot{q}^\alpha(t)). \quad (2.1)$$

One then obtains, from the least action principle $\delta S = 0$, the Euler–Lagrange equations:

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{q}^\alpha} = \frac{\partial \mathcal{L}}{\partial q^\alpha}. \quad (2.2)$$

If we use the chain rule for the $\frac{d}{dt}$ derivative, we get from (2.2):

$$\ddot{q}^\beta \frac{\partial^2 \mathcal{L}}{\partial \dot{q}^\beta \partial \dot{q}^\alpha} + \dot{q}^\beta \frac{\partial^2 \mathcal{L}}{\partial q^\beta \partial \dot{q}^\alpha} = \frac{\partial \mathcal{L}}{\partial q^\alpha}. \quad (2.3)$$

The accelerations are uniquely determined by the positions and velocities if we can isolate \ddot{q}^β in this equation. A necessary and sufficient condition for this is that the matrix $M_{\alpha\beta} := \frac{\partial^2 \mathcal{L}}{\partial \dot{q}^\beta \partial \dot{q}^\alpha}$ be invertible. If it is not, the accelerations are undetermined, so that the motion is underdetermined by the initial positions and velocities: there is indeterminism at the level of the $2N$ variables, $q^\alpha(t), \dot{q}^\alpha(t)$. Assuming we believe that in physical terms, the motion is indeed determined, this indicates a redundancy in our description of the system. For a philosophical introduction to this indeterminism and redundancy, from a Lagrangian treatment, compare Wallace (2002).

Let us see how this redundancy appears in the Hamiltonian formalism. The idea of the *Legendre transformation* is that at any point $q \in \mathcal{Q}$, the Lagrangian \mathcal{L} determines a map Leg_q from the tangent space $T_q \mathcal{Q}$ at q to its dual space $T_q^* \mathcal{Q}$. Intuitively speaking, this is the transition $\dot{q} \mapsto p$. To be precise, Leg_q is defined by

$$\text{Leg}_q : w = \dot{q}^\alpha \frac{\partial}{\partial q^\alpha} \in T_q \mathcal{Q} \mapsto \frac{\partial \mathcal{L}}{\partial \dot{q}^\alpha} \in T_q^* \mathcal{Q}. \quad (2.4)$$

One easily checks that because the canonical momenta $p_\alpha := \frac{\partial \mathcal{L}}{\partial \dot{q}^\alpha}$ are 1-forms, this definition is coordinate-independent. An equivalent definition, manifestly coordinate-independent and given for all $q \in \mathcal{Q}$, is the Legendre transformation, $\text{Leg} : T\mathcal{Q} \rightarrow T^*\mathcal{Q}$, defined by

$$\forall q \in \mathcal{Q}, \forall v, w \in T_q \mathcal{Q} : \text{Leg}(w)(v) := \left. \frac{d}{dt} \right|_{t=0} \mathcal{L}(w + tv). \quad (2.5)$$

(Here we take w, v to encode the identity of the base-point q , so as to simplify notation, writing $\text{Leg}(w)$ rather than $\text{Leg}(q, w)$ etc.) That is: $\text{Leg}(w)(v)$ is the derivative of \mathcal{L} at w , along the fibre $T_q \mathcal{Q}$ of the fibre bundle $T\mathcal{Q}$, in the direction v . Thus Leg is also called the *fibre derivative*.

Given \mathcal{L} , we define its energy function $E : T\mathcal{Q} \rightarrow \mathbb{R}$ by

$$\forall v \in T\mathcal{Q}, E(v) := \text{Leg}(v)(v) - \mathcal{L}(v); \quad (2.6)$$

or in coordinates,

$$E(q^\alpha, \dot{q}^\alpha) := \frac{\partial \mathcal{L}}{\partial \dot{q}^\alpha} \dot{q}^\alpha - \mathcal{L}(q^\alpha, \dot{q}^\alpha). \quad (2.7)$$

Then one shows that $E \circ (\text{Leg})^{-1}$ is, as one would hope, the familiar Hamiltonian function $H : T^*\mathcal{Q} \rightarrow \mathbb{R}$ given by $(q, p) \equiv (q^\alpha, p_\alpha) \mapsto \dot{q}p - \mathcal{L} \equiv \dot{q}^\alpha p_\alpha - \mathcal{L}$, with each \dot{q}^α being a function of the qs and ps .

Thus the question whether the accelerations are determined by the positions and velocities is translated to the question of whether the momenta $p_\alpha = \frac{\partial \mathcal{L}}{\partial \dot{q}^\alpha}$, are invertible, as functions of the velocities. If $M_{\alpha\beta} = \frac{\partial p_\alpha}{\partial \dot{q}^\beta}$ is not invertible, since α and β run over the same indices (i.e. M is a square matrix), the map from the $\dot{q}s$ to the p_s , at fixed q , is many-to-one. That is: there are constraints among the p_β as functions of the velocities. Dropping reference to the velocities, i.e. writing the constraints as functions on phase space, we therefore write equations such as

$$\Phi^I(q^\alpha, p_\beta) = 0, \quad (2.8)$$

where I parametrizes the constraints.⁹ These conditions are usually called *primary constraints*, to emphasize that no equation of motion was used in their derivation. If the constraints are conserved by the equations of motion, then they correspond, by (the converse of) Noether's theorem, to symmetries of the system: as we will explain in Section 2.2.

Now, the constraint surfaces are submanifolds of $T^*\mathcal{Q}$, and it thus follows from (2.8) that the inverse transformation, from the momenta p to the velocities \dot{q} , must be multi-valued, since the dimension of $T\mathcal{Q}$ is the same as that of $T^*\mathcal{Q}$ (viz. $2N$, and so greater than the dimension of the constraint surfaces). Thus the inverse image in $T\mathcal{Q}$ of the constraint surface (2.8) forms a submanifold.¹⁰ Similarly, if we consider the intersection of the constraint surfaces (2.8), as I varies (which is often called '*the constraint surface*'); and so, the inverse image of that intersection.

Correspondingly, within $T\mathcal{Q}$: a non-trivial kernel for the matrix $M_{\alpha\beta}$ in (2.3) implies that the extrema of the Lagrangian are not isolated: there are 1-parameter families of curves in configuration space that extremize the action functional. In other words, as discussed above: there is indeterminism at the level of the $2N$ variables, $q^\alpha(t), \dot{q}^\alpha(t)$: (again, compare Wallace (2002)).

Here is an **Example** from Henneaux & Teitelboim (1992, Sec. 1.1.1), which we will build on in Section 3.1. Let (q^1, q^2) be coordinates for \mathcal{Q} , and consider $\mathcal{L} = \frac{1}{2}(\dot{q}^1 - \dot{q}^2)^2$. The momenta are $p_1 = \dot{q}^1 - \dot{q}^2$ and $p_2 = \dot{q}^2 - \dot{q}^1$, as is easy to verify. Thus we find the rather simple constraint: $\Phi = p_1 + p_2 = 0$. The Legendre transform maps all of $T\mathcal{Q}$ into this constraint surface in $T^*\mathcal{Q}$. Moreover, the entire line $\dot{q}^2 = \dot{q}^1 + c$ is mapped to $p_1 = -c = -p_2$. This transformation is therefore neither one-to-one nor onto.

To render this transformation invertible, we will need to introduce *Lagrange multipliers*, in Section 2.2, which can be thought of as coordinates on the manifolds that are the inverse values in $T\mathcal{Q}$ (for (2.5)) of a given point on the constraint surface lying in $T^*\mathcal{Q}$.¹¹

⁹Here we assume that the rank of $M_{\alpha\beta}$ is constant, and that the constraints obey regularity conditions. See (Henneaux & Teitelboim, 1992, Ch. 1.1.2).

¹⁰The gradients of Φ^I in the momentum directions along the surface, i.e. the vectors $\frac{\partial \Phi^I}{\partial p_\alpha}$, provide the complete set of null vectors of $M_{\alpha\beta}$, since by 2.8

$$0 = \frac{\partial \Phi^I}{\partial \dot{q}^\beta} = \frac{\partial \Phi^I}{\partial p_\alpha} \frac{\partial p_\alpha}{\partial \dot{q}^\beta} = \frac{\partial \Phi^I}{\partial p_\alpha} M_{\alpha\beta}.$$

In Section 2.2 we will give a coordinate-free version of these statements.

¹¹It is also relatively easy to show that two velocities that lie in the pre-image of the same momentum are

2.2 The Hamiltonian formalism and constraints

To understand the basic features of constraints and the symmetries they generate in a manner that will be helpful in what follows, it pays to introduce the symplectic formalism for Hamiltonian mechanics.

The great advantage of the symplectic formalism is that it treats momentum and configuration variables on a par. By so doing we see phase space \mathcal{P} as a high-dimensional manifold—infinite-dimensional, in field-theory—endowed with a certain geometric structure. In the simple example above, \mathcal{P} would be the $2N$ -dimensional manifold whose geometric structure is a symplectic 2-form, given, in the global coordinates (q^α, p_α) , by:

$$\omega = \sum_{\alpha} dq^\alpha \wedge dp_\alpha. \quad (2.9)$$

Though we have given ω in the specific choice (q, p) of coordinates,¹² it is a coordinate-independent, differential geometric object on \mathcal{P} .

The role of the symplectic form ω is to convert a vector field into a one-form, or vice-versa. And since a scalar function defines a one-form, viz. its gradient, ω converts a scalar function like the Hamiltonian into a vector field, whose integral curves are a flow in phase space. We construe these curves as the dynamical trajectories of the system; (i.e. assuming that the given scalar function encodes the forces operative on and in the system). Thus we take a Hamiltonian function $H : \mathcal{P} \rightarrow \mathbb{R}$ to specify the dynamics by assigning to each dynamical state its total energy. For from any such smooth scalar function, we can obtain a one-form dH , and then use ω to define a vector field X_H , by:

$$\omega(X_H, \bullet) = dH(\bullet). \quad (2.10)$$

This vector field on phase space is then integrated to yield a dynamical trajectory through each point. Compare Figure 1.

Indeed, in the simple mechanical case without constraints, we can plug coordinates into this equation to recover a local description of the dynamics, i.e. the familiar form of Hamilton's equations. In particular, the relation between Poisson brackets, defined as usual by

$$\{f, h\} := \frac{\partial f}{\partial q^\alpha} \frac{\partial h}{\partial p_\alpha} - \frac{\partial f}{\partial p_\alpha} \frac{\partial h}{\partial q^\alpha}; \text{ so that Hamilton's equations with } h \text{ are: } \frac{df}{dt} = \{f, h\}, \quad (2.11)$$

and the symplectic form ω is:

$$\{f, h\} = df(X_h) = \omega(X_f, X_h), \quad (2.12)$$

for $f, h \in C^\infty(\mathcal{P})$.¹³

For consistency, since the exterior derivative squares to zero i.e. $d^2 \equiv 0$, ω must be closed, i.e. $d\omega = 0$. Moreover, if we would like the dynamical trajectory associated to H to be unique, ω must be non-degenerate. That is: $\omega(v, \bullet)$ must be injective, i.e. have only the zero vector in its kernel.

related by a linear combination of the null vectors of $M_{\alpha\beta}$, namely $\frac{\partial \Phi^I}{\partial p_\alpha}$ (cf. footnote 10); the coefficients of this linear combination are the Lagrange multipliers. In Section 2.2 we will see a coordinate-free version of these statements.

¹²These are always available locally, thanks to Darboux's theorem: which states, in modern geometric terms, that a manifold equipped with a symplectic 2-form is locally a cotangent bundle. Cf. e.g. (Arnold, 1989, p. 230-232); or for a philosophical introduction, (Butterfield, 2006, Section 6.6).

¹³Here, as usual, $df(X)$ is the contraction between 1-forms and vectors; and $dH(X)$ is equal to $X(H)$ i.e. the directional derivative of a scalar function f along X .




Figure 1: An illustration of the relation between a scalar function, its gradient, and its associated symplectic vector field, using the canonical symplectic form and the Euclidean metric to identify dH with ∇H , as a vector orthogonal to the level surfaces $H = \text{const}$. Supposing, in the (q, p) -coordinates, that $\nabla H = (\partial_q H, \partial_p H)$, then, writing the inversion map as J_0 , we have: $X_H = -J_0(\partial_q H, \partial_p H) = (-\partial_p H, \partial_q H)$, which is orthogonal to ∇H and thus along $H = \text{const}$.

Although this last condition is always taken to hold on the full phase space \mathcal{P} , it needs to be relaxed in gauge theories, precisely because of constraints. So although we require ω to be non-degenerate in \mathcal{P} , it does not need to be non-degenerate once we restrict it to *the* constraint surface—meaning now the intersection of the constraint surfaces (2.8), as I varies:

$$\Gamma := \{z \in \mathcal{P}, | \Phi^I(z) = 0, \quad \text{for all } I\}. \quad (2.13)$$

As depicted in Figure 1: the Φ^I , as scalar functions on phase space (for each I), have a (differential geometric) gradient, $d\Phi^I$, which are in one-one correspondence with vector fields $X_{\Phi^I} =: X_I$ due to the symplectic structure of phase space: namely, through $\omega(X_I, \bullet) = d\Phi^I$. The key idea is that, just as the flow specified by the Hamiltonian function conserves energy, these vector fields associated to Φ^I are tangential to, and so preserve, the intersection of all the constraint- and energy-surfaces. That is, in a less geometric (and maybe more familiar) language: they not only commute with the Hamiltonian and conserve energy, but also conserve the charges associated with the constraints.

Here we will only consider a certain type of constraints, called *first-class constraints*. These are defined as constraints whose Poisson bracket with every constraint vanishes on the constraint surface Γ (though perhaps not elsewhere). Then the assumption that all the constraints are first-class implies that the flow of each vector field X_I , associated to each constraint, is tangent to Γ . This assumption also means that we need not concern ourselves with the several steps involved in the *Dirac algorithm*.¹⁴ In general, these steps are necessary because the primary constraints 2.8 that emerge from the Legendre transform might fail to be preserved by

¹⁴Without the assumption of first-class, we still have an algorithm for finding whether the constraints generate symmetries. This algorithm can be summarised in the more familiar Poisson bracket notation, as follows. Suppose we are given some initial set of constraints $\Gamma_0 := \{z \in \mathcal{P}, | \Phi^{I_0}(z) = 0, \quad \text{for all } I_0\}$, such that, e.g. for some I_0, J_0 , we have $\{\Phi^{I_0}, \Phi^{J_0}\}_{|\Gamma_0} = \Phi^{K_1} \neq 0$ where clearly Φ^{K_1} is not included among the original constraints, since it does not vanish on Γ^0 . (Here the restriction to Γ_0 serves to emphasize that the vector fields need not commute everywhere on phase space, but only on the surface where the constraints vanish.) We would then add this new constraint to the others, to form a new ‘surface’, Γ_1 , and repeat the test above, until, eventually we get:

$$\{\Phi^{I_n}, H\}_{|\Gamma_n} \equiv \omega(X_{I_n}, X_H)_{|\Gamma_n} \equiv X_H(\Phi^{I_n})_{|\Gamma_n} = 0, \quad \text{for all } I_n,$$

and

$$\{\Phi^{I_n}, \Phi^{J_n}\}_{|\Gamma_n} \equiv \omega(X_{I_n}, X_{J_n})_{|\Gamma_n} = 0, \quad \text{for all } I_n, J_n.$$

(It is possible that the iteration yields only the empty set, in which case the system is dynamically inconsistent).

either the Hamiltonian or by other constraints—but this will be ensured if the constraints are first-class. Thus in general, one must then seek a type of reflective equilibrium: successively imposing further restrictions to submanifolds of phase space, until the corresponding symplectic flows of all the constraints preserve the Hamiltonian and become tangent to the constraint submanifold. Compare footnote 14, and Pons (2005); and for a complete account, (Henneaux & Teitelboim, 1992, Chapter 2).

In the language of symplectic geometry, we define the embedding $\iota : \Gamma \rightarrow \mathcal{P}$ and require that the pullback of the symplectic form $\iota^*\omega$ obey:

$$\iota^*\omega(X_H, \bullet) = d(H|_\Gamma), \quad \iota^*\omega(X_I, \bullet)|_\Gamma = d(\Phi_I|_\Gamma) \equiv 0 \quad (2.14)$$

(since $\Phi_I|_\Gamma \equiv 0$). Thus, on the constraint surface, because we assumed constraints to be first-class, the vector fields X_I generated by the constraints are null directions of the symplectic form restricted to the constraint surface. This generalizes, in a coordinate-free formalism, the content of footnote 11.

These directions are *gauge*. The set of points of phase space that are reached by the X_I 's integral curves, from a given point, is called a *gauge orbit*. And the elements within each such orbit are taken to be physically equivalent. Here, physical equivalence of two points of phase space is understood as a matter of any physical quantity, taken as a phase function, i.e. a real-valued function on phase space, having the same value for the two points. That is: a physical quantity must be *gauge-invariant*: taken as a phase function, it must be constant on each gauge orbit.¹⁵

For simplicity, we have here suppressed a few important qualifications (that are widely recognized). Firstly: in some systems, the requirement that the primary constraints (2.8) be preserved in time implies a new relation between the qs and ps , independent of these constraints. Such a relation is called a *secondary* constraint. Secondary constraints can be first-class: an important example being the Hamiltonian i.e. scalar constraint (governing time-evolution) in canonical general relativity. And some such secondary first-class constraints are *not* gauge generators. But though important, these points do not affect this paper. For details, compare (Henneaux & Teitelboim, 1992, Sections 1.1.5, 1.2, 1.6.3 and 3.3) and Pitts (2014a). A second simplifying assumption—which applies to electromagnetism, and indeed to most familiar physical theories—is that the commutation algebra of the constraints closes irrespectively of the satisfaction of the equations of motion: it forms what is usually called a *closed algebra*. Thirdly, we have also assumed that the constraints are *irreducible*, i.e. that all the constraint equations 2.8 (for both primary and secondary constraints) are independent of each other: (so that roughly speaking, there are no “symmetries among the symmetries” that they generate).

¹⁵ Indeed, the null directions of $i^*(\omega)$ are necessary and sufficient to characterise the generators of gauge symmetry. For suppose that what we know is that a certain class of vector fields X_I is such that $\omega(X_I, \bullet) = 0$. Since the exterior derivative d commutes with pullbacks, if ω is closed, $i^*\omega =: \tilde{\omega}$ is also closed. Thus using the Cartan Magic formula relating Lie derivatives, contractions i and the exterior derivative d :

$$L_{X_I}\tilde{\omega} = (di_{X_I} + i_{X_I}d)\tilde{\omega} = 0;$$

i.e. the first term also vanishes because $\tilde{\omega}(X_I, \bullet) = 0$. So $\tilde{\omega}$ itself is invariant along X_I . Moreover, if we take the commutator of X_I, X_J , i.e. $[X_I, X_J] = L_{X_I}X_J$, contract it with $\tilde{\omega}$, and remember the formula:

$$L_{X_I}(\tilde{\omega}(X_J, \bullet)) = \tilde{\omega}(L_{X_I}X_J, \bullet) + (L_{X_I}\tilde{\omega})(X_J, \bullet),$$

we obtain that, since both $L_{X_I}(\tilde{\omega}(X_J, \bullet)) = 0$ and $L_{X_I}\tilde{\omega} = 0$, it is also the case that $\tilde{\omega}([X_I, X_J], \bullet) = 0$. Thus, by the Frobenius theorem the kernel of the pullback $i^*(\omega)$ forms an integrable distribution which integrates to give the orbits of the symmetry transformation. This means, for the discussion below (in Section 2.2.1) of symplectic reduction, in which a Lie group G acts *ab initio* on the phase space, that: we can define a projection operator $\pi : \Gamma \rightarrow \Gamma/G$; and, ultimately the degeneracy of $i^*\omega$ allows one to define a *reduced symplectic form*, $\bar{\omega}$, on the space of orbits, given by $\pi^*\bar{\omega} = i^*\omega$. See (Marsden, 2007, Ch. 1).

And lastly, we have also assumed that the commutation algebra of the constraints forms a true Lie algebra, i.e. the structure ‘constants’ are true constants, not functions on phase space; (this assumption fails for general relativity).¹⁶

The existence of null vector fields implies the Hamiltonian flow is not unique: if X_H solves (2.14), then so does $X_H + a^I X_I$ for any set of coefficients a^I that are arbitrary functions of time. So, since the dynamics preserves the constraint surface, instead of taking H as the Hamiltonian function generating the dynamics, we may equivalently take the *total Hamiltonian*:

$$H_T = H + a^I \Phi_I. \quad (2.15)$$

Along the constraint surface the dynamics according to H_T will be indistinguishable from that according to H . Here, by ‘indistinguishable’, we mean that two trajectories within Γ that start at a common point in Γ and that are determined, respectively, by the two choices of Hamiltonian, will at any later (indeed: any earlier!) time, lie in the same gauge orbit as each other—and so will at all times agree on the values of all gauge-invariant quantities.

In sum, one of the significant (as well as practical) differences between the Hamiltonian and the Lagrangian frameworks is that in the Hamiltonian framework the symmetries are not ‘guessed’ from the form of the action functional. Instead, they are obtained from the constraints that emerge when the Legendre transformation is applied. That is, the constraints that emerge from the Legendre transformation are associated to vector fields; and, with a few auxiliary assumptions, the flow of each of these vector fields conserves the constraints and the Hamiltonian, and the vector field is thus taken to be the generator of symmetries on Γ . The action of the symmetry on any quantity Q is given by

$$X_I(Q) = dQ(X_I) = \omega(X_I, X_Q) = \{\Phi^I, Q\}. \quad (2.16)$$

Finally let us sum up this discussion of the Lagrangian and Hamiltonian frameworks’ treatments of constraints, by stressing a concordance between them: a concordance despite the many-one mapping from the first to the second, i.e. the fact that at fixed q , many \dot{q} map to a single p . Namely: there is a one-one correspondence between the gauge orbits in the Hamiltonian framework and degeneracy directions of the Lagrangian in the configuration space; (cf. Sections 3.1-2 in [Henneaux & Teitelboim \(1992\)](#)).

2.2.1 Off the constraint surface

Our sketch of the Hamiltonian treatment of constraints would be incomplete without some mention of *symplectic reduction*: a large and important topic (briefly mentioned in footnote 15). However, it is usually pursued, not (as in this paper) by using the the Legendre transformation being many-one to motivate restricting one’s attention to the constraint surface Γ in phase space (as in this paper); but by postulating *ab initio* a smooth action of a Lie group G on phase space, and studying the consequences. So the relation to the Lagrangian framework, and to constraints originating from $\frac{\partial p_\beta}{\partial \dot{q}^\alpha}$ being not invertible, tends to be obscured. But the rest of this paper will not need an account of symplectic reduction: for which, cf. e.g. [Marsden \(2007\)](#) for a complete, but concise, exposition, [Butterfield \(2007\)](#) for a philosophical introduction, and [Gomes & Butterfield \(2021b\)](#) for the relation to the Lagrangian framework.

¹⁶We should also note a controversy. [Pitts \(2014b\)](#) claims that, even under these assumptions, and for electromagnetism, the main (and orthodox) idea above—that points in the same gauge orbit, are physically equivalent—fails. He claims that even a first-class constraint can fail to be a gauge generator, i.e. it can generate instead what he calls a ‘bad physical change’. Our own view is that the main idea holds good. (The dispute turns on the transformation properties of Lagrange multipliers in the canonical Lagrangian; and we think the treatment by e.g. ([Henneaux & Teitelboim, 1992](#), equations 19.11 and 19.13, as clarified and supported by 3.26 and 3.31) answers Pitts’ arguments.)

However, symplectic reduction prompts a philosophical topic we want to address. It is about possibility, i.e. about how we should think of the non-actual i.e. unrealized states in the state-space. We have hitherto said nothing about this, since our discussion has prompted no alteration from how one normally thinks of possibility within the Hamiltonian dynamics of an *unconstrained* system. There, one naturally regards which energy hypersurface the system is actually on (or equally: the actual value of any first integral of the motion) as a matter of initial conditions, mere happenstance. And for all we have said so far, it seems that in general, this attitude applies equally to constrained systems. For in general, it seems that the constraints Φ^I could have taken values other than zero (cf. equation 2.8): the state could have been *off* the actual constraint surface Γ . But the theory of symplectic reduction reveals a wrinkle: indeed, two wrinkles. (Belot, 2003, Section 11) discusses them, as do Gomes & Butterfield (2021b); and we sketch them here.

First, in some cases there is reason to deny that the states lying off the constraint surface Γ are genuinely possible. Belot’s (and our) simplest example is *relationism, a la* Leibniz and Mach, about space. The case can be made for a system of N point-particles in Euclidean space. For this system, while the ‘absolutist’ will take the configuration space \mathcal{Q} to consist of all the ways N particles can be placed in \mathbb{R}^3 , i.e. to be \mathbb{R}^{3N} , the relationist will say that two such placements that differ by a spatial translation and-or a rotation should be regarded as one and the same. That is, the relationist advocates a *relative configuration space*, whose points are sets of relative distances between the particles. This space can be presented as the quotient of \mathbb{R}^{3N} by the obvious action of the Euclidean group; (modulo some technicalities about excising unsuitably symmetric points of \mathbb{R}^{3N}). And when one works through the details of the constraint formalism, it turns out that on the relationist’s view, only states on the relevant constraint surface Γ within $T^*\mathcal{Q}$ are genuinely possible.

Of course, not everyone is a relationist! But also in other cases, there is a similar rationale to ‘endorse the dynamics intrinsic to Γ ’, and reject the states not in Γ . Thus Belot points out that in some field theory cases, theorizing about states off the constraint surface corresponds to treating charges that source the field in question without being affected by it, i.e. treating external sources. For example, in our case-study of electromagnetism: for non-vacuum, the Gauss constraint becomes $\text{div}(E) = \rho$, where ρ is treated as sourcing, but as unaffected by, the electromagnetic field. (Of course, there is a close analogy with Poisson’s equation $\nabla\phi = \rho$ in Newtonian gravity, and its modern descendant, Newton-Cartan gravitation: the mass density ρ sources the potential ϕ but does not self-gravitate.) Since this is an idealization, one has reason to reject the states off the constraint surface as not genuinely possible, and to endorse the dynamics intrinsic to Γ —like the relationist above.

Besides, as Belot goes on to say: faced with this idealization, one should seek theories in which the coupling is “two-way”. Indeed, there are such theories (references in his footnote 74); and—what matters for our present topic—in these theories, one again gets *only one constraint surface*, like in the case of vacuum electromagnetism. That is: once one augments the phase space so as to describe the charges (including: augmenting the Hamiltonian to describe the two-way coupling), one gets just one constraint surface in a higher-dimensional space—*not* a family of surfaces in the original lower-dimensional phase space, indexed by the charge distributions. So again, one has reason to endorse the dynamics intrinsic to the constraint surface. And again: our main theme about gauge structure is illustrated: viz. that null vector fields on the constraint surface are infinitesimal generators of gauge transformations.

Let us summarise this discussion by quoting Belot. He writes (p. 215):—

This [i.e. treating charges as external, i.e. as sourcing the field in question yet without being affected by it] amounts to working off of the constraint surface in order to study the field dynamics in the presence of external sources painted onto spacetime independently of the behavior of the fields. This is, of course, an ad

hoc maneuver—if one wants to study Yang–Mills with sources honestly, one must introduce matter which not only acts upon the field but is also acted upon by it. And when one pursues this upright course, one ends up with a constraint which is a direct analog of the usual Gauss constraint—the null directions of the constraint surface correspond to the infinitesimal generators of gauge transformations. Under this more fundamental approach, there is no physical interpretation for points lying off of the constraint surface—and so we have an excellent reason to prefer an intrinsic reading of the theory.

3 The case of electromagnetism

So much by way of reviewing Hamiltonian constrained dynamics. We turn to our case-study, classical electromagnetism. There are many accounts of its symplectic structure in the literature, both physical and philosophical: some of them of course excellent.¹⁷ But there is a “core” of ideas and results, that is relatively easy to expound and illuminating, without having to plumb the depths of the (elegant) symplectic geometry of the theory. This core is, so far as we know, not articulated in the literature: certainly, it is not stressed.

So as we announced at the start of Section 1, the technical aim of this paper is to expound this core, and show that it sheds light on various ideas, both formal and interpretative. More specifically, we will try to shed light on (1) classical non-locality and (2) the preferred splits of degrees of freedom. (These correspond to the three morals at the end of Section 1.)

In Section 3.1 we will illustrate the ideas that are used in (2) in a simple toy example. We will see how certain considerations of convenience and simplicity can go a long way to selecting gauge-fixing conditions. In Section 3.2 we will apply those ideas to electromagnetism, where they give rise to the Gauss constraint and gauge transformations. In Section 3.3 we will interpret the Gauss constraint as encoding a type of non-locality, *a la* (1). It defines a part of the electric field, viz. the Coulombic field, that is determined by the instantaneous distribution of charges. (Adopting Maudlin (2018)’s terms for a moment: it is not ‘fundamental ontology’, but ‘derivative ontology’, since derived from the charge distribution.) There is here a strong analogy with the elementary Newtonian gravitational potential ϕ , which is sourced by the mass distribution via Poisson’s equation: and which is often said to be “not physically real”, or “a convenient fiction”, since it has no energy or momentum, but only encodes, via its gradient $\nabla\phi$ the infinite battery of counterfactual conditionals about how test-masses located at the spatial point in question would accelerate. Similarly here: the Coulombic field encodes infinitely many counterfactual conditionals about how test-charges would move—if there were also no other part of the field, i.e. no radiative part, enjoying its own dynamics.

Using the decomposition of fields implied by this understanding, we will in Section 3.4 use the symplectic structure of the theory to find the “block-diagonal conjugate structure” of the electric and gauge potential fields, and show how this selects the Coulomb gauge.

3.1 A toy example of natural coordinate choices on phase space

To make our aims more vivid, we return to the simple example we gave at the end of Section 2.1. Consider two identical free particles of mass m on a line, with coordinates q_1 and q_2 . So the phase space \mathcal{P} is 4-dimensional, with coordinates (q_1, q_2, p_1, p_2) . The canonical Hamiltonian

¹⁷Excellent physics expositions include Jackson (1975). Excellent philosophical discussions include: Belot (1998, 2003); Healey (2007). We particularly recommend Belot (2003) for the Hamiltonian formulation of vacuum Yang–Mills theories.

for the system is

$$H = \frac{1}{2m}(p_1^2 + p_2^2). \quad (3.1)$$

Now we add a constraint to the system, namely:

$$\Phi = p_1 + p_2 = 0; \quad (3.2)$$

(which is clearly first-class). So the constraint $\Phi = 0$ requires the total canonical momentum $P := p_1 + p_2$ to vanish, and defines a 3-dimensional constraint surface Γ in \mathcal{P} . The total Hamiltonian is $H_T = H + a\Phi$, where a is an arbitrary function of time. It generates the time evolution

$$\dot{q}_1 = \frac{p_1}{m} + a, \quad \dot{q}_2 = \frac{p_2}{m} + a, \quad \dot{p}_1 = \dot{p}_2 = 0. \quad (3.3)$$

Here we started with the variables that initially seemed natural, centered on each of the particles. But a more natural choice for coordinatizing momentum space (i.e. at each fixed value of (q_1, q_2)) would be $P_+(p_1, p_2) := p_1 + p_2$ and $P_-(p_1, p_2) = p_1 - p_2$. Now at each value of (q_1, q_2) , the constraint surface is given by:

$$\{(P_+, P_-) \mid \Phi(z) = 0\} = \{(0, P_-) \mid P_- \in \mathbb{R}\}. \quad (3.4)$$

In the q -coordinates, the natural conjugate variables to P_- and P_+ are $Q_- = q_1 - q_2$ and $Q_+ = q_1 + q_2$, respectively. The first, Q_- , is the relative distance between the two particles. It is gauge-invariant, $\{q_1 - q_2, \Phi\} = 0$; and its equation of motion contain no arbitrariness,

$$\dot{q}_1 - \dot{q}_2 = \dot{Q}_- = \frac{p_1}{m} - \frac{p_2}{m} = \frac{P_-}{m}. \quad (3.5)$$

On the other hand, from (3.3),

$$\dot{Q}_+ = m\dot{q}_1 + m\dot{q}_2 = 2ma \quad (3.6)$$

is ‘‘pure gauge’’, since a is an arbitrary function of time. Note that (3.5) and (3.6) are given on the full phase space: of course, the symplectic form on the constraint surface would be degenerate, since $P_+ = 0$ there. To sum up: Q_- , Q_+ and P_- are natural coordinates to parametrize the constraint surface.

In what follows, we will try to provide a similarly natural decomposition of electromagnetism’s gauge potential and electric field. Φ will be the Gauss constraint, which we will take to be naturally parametrized by a Coulomb potential. So this potential will play the role of P_+ , and thus the radiative degrees of freedom of the electric field will play the role of P_- . As regards the configuration variables: Q_- will be given by a more complicated function of the original configuration variables, which projects it into Coulomb gauge, but it will be likewise gauge-invariant. And finally, Q_+ will be the pure gauge part of the gauge potential.

3.2 Hamiltonian treatment of electromagnetism

The Maxwell equations are written, in terms of the electromagnetic field tensor, F_{ab} , in four-dimensional Minkowski spacetime M , as:

$$\partial^a F_{ab} = j_b, \quad \text{and} \quad \partial_{[a} F_{cd]} = 0; \quad (3.7)$$

where j is the electric 4-current and square brackets denote anti-symmetrization of abstract spacetime indices (a, b , etc.).

The second equation of (3.7) is called ‘the Bianchi identity’, and it is read as a constraint on the field tensor. A geometric explanation for this constraint is that $F_{ab} = \partial_{[a} A_{b]}$, where

the 1-form A_a is called *the gauge potential*. That is: $F_{ab} = \partial_{[a}A_{b]}$ implies the second equation of (3.7) holds identically, thanks to the commutation of partial derivatives. By employing the gauge potential we can drop the Bianchi identity, and that is what we will do.

Then the equations of motion (3.7) are written as:

$$\partial^a \partial_a A_b - \partial^a \partial_b A_a = j_b. \quad (3.8)$$

These equations (together with the equations for the dynamics of the charges constituting the currents) are obtained from the action functional:

$$S[A] := \int d^4x \left(\partial_{[a}A_{b]} \partial^{[a}A^{b]} + A^a j_a + \mathcal{L}_{\text{matter}} \right), \quad (3.9)$$

where $\mathcal{L}_{\text{matter}}$ is the Lagrangian density for the matter fields; (for illustration, one can take this as the Klein-Gordon Lagrangian or as the Lagrangian for a charged point particle).

Now we choose a spacetime split into spatial and time directions, $M = \Sigma \times \mathbb{R}$. (We recall Section 1's admission that this carries the price of losing manifest Lorentz invariance.) We also assume that the fields have appropriate fall-off conditions at spatial infinity.

Upon such a spacetime decomposition, the components of the electromagnetic tensor recover the familiar electric and magnetic fields: $F_{i0} = E_i$, and $F_{ij} \epsilon_i^{jk} = B_i$ and $j_0 = \rho$ (where we used the three-dimensional totally-antisymmetric tensor, ϵ , or the spatial Hodge star, to obtain a 1-form, and i, j, k are spatial indices, i.e. in Σ), and the first equation of (3.7) becomes the familiar Maxwell equations.¹⁸

Now we perform a Legendre transform. Then the spatial vector fields \mathbf{A} and \mathbf{E} are canonically conjugate. In fact, \mathbf{A} is the configurational variable, and \mathbf{E} is the momentum; since $\mathbf{E} = \partial \mathcal{L} / \partial \dot{\mathbf{A}}$. The Poisson bracket is defined, for F, G two functionals of the fields A and E ;

$$\{F, G\} := \int d^3y \left(\frac{\delta F}{\delta A_i(y)} \frac{\delta G}{\delta E^i(y)} - \frac{\delta F}{\delta E^i(y)} \frac{\delta G}{\delta A_i(y)} \right). \quad (3.10)$$

Now, as described in Section 2.1 and Section 2.2, the Lagrangian has symmetries, which translate into constraints in the Hamiltonian formalism. Analogously to (2.15), we obtain a total Hamiltonian, written as $H_T = H(A, E, j, \lambda) + H_{\text{matter}}$, with:

$$H(A, E, j, \lambda) = \int d^3x \left(\|\mathbf{E}\|^2 + \|\mathbf{B}\|^2 + \lambda(\text{div}(\mathbf{E}) - \rho) + A^i j_i \right), \quad (3.11)$$

where $\partial_{[i}A_{j]} \partial^{[i}A^{j]} =: \|\mathbf{B}\|^2$, $\text{div}(\mathbf{E}) := \partial^i E_i$.

The part of (3.11) that we want to draw attention to is the term $\lambda(\text{div}(\mathbf{E}) - \rho)$. For λ is a scalar function on the spatial surface Σ : λ is the *Lagrange multiplier* we first encountered in Section 2.1. Indeed it is just the a^I of (2.15), but now in the field-theoretic context, when I becomes a continuous index. The constraints corresponding to Φ^I are:

$$G(x) := \text{div}(\mathbf{E})(x) - \rho(x) = 0, \quad \text{for all } x \in \Sigma. \quad (3.12)$$

Accordingly, (3.12) defines not a single constraint, but an infinite set of them: one per spatial point. The values of λ at the various points $x \in \Sigma$ thus give a particular linear combination of the constraints. Hence λ is also called a 'smearing' of the constraints; and it is convenient to define the smeared Gauss constraint: $G(\lambda) := \int_{\Sigma} (\text{div}(\mathbf{E}) - \rho) \lambda$.

To see that we are in the domain of the previous discussion, namely that the constraints are all first class, we can check that the constraints commute, and also commute with the

¹⁸For $\mathcal{L}_{\text{matter}} = mv^a v_a$, for $v^a = (\dot{\gamma})^a$ the 4-velocity of a charged particle whose trajectory is γ , i.e. such that $j^a = ev^a$, we obtain the Lorentz force equation as the equation of motion for the particle.

Hamiltonian constraint. This is easy to verify, since a given linear combination of symplectic flows, which we call a *smear*ed symplectic flow, $X_{G(\lambda)}$, acts on the canonical variables as:

$$X_{G(\lambda)}(A_i(x)) := \{A_i, G(\lambda)\} := \int d^3y \frac{\delta(\lambda(\operatorname{div}(\mathbf{E}) - \rho))}{\delta E^i(x)} = \int d^3y \lambda(y) \partial_i \delta(x, y) = -\partial_i \lambda(x); \quad (3.13)$$

$$X_{G(\lambda)}(E^i(x)) := \{E_i, G(\lambda)\} := - \int d^3y \frac{\delta(\lambda(\operatorname{div}(\mathbf{E}) - \rho))}{\delta A_i(x)} = 0 \quad (3.14)$$

To obtain these results, note that in each line the third term, i.e. the integral of a functional derivative, gets just one term from 3.10; and that the final equation in (3.13) is obtained by integration by parts removing the derivative of a delta-function.

So we see that, as expected, the constraints act on phase space as the familiar gauge transformations: they preserve the value of the electric field (which is gauge-invariant in electromagnetism) and change the value of the gauge potential by a gradient of a scalar function: $A_i \mapsto A_i - \partial_i \lambda$.¹⁹ Since the flow of the constraint does not change the Hamiltonian (and preserves the constraint as well), it generates a symmetry of the system.

From the more geometric viewpoint, the infinite-dimensional phase space whose canonical coordinates are the electric field and the gauge potential has a bona-fide (infinite-dimensional) symplectic geometry (e.g. modeled on Banach manifolds, cf. (Lang, 2012, Chapter 2)). Thus we have a symplectic form, which in vacuum is written as:²⁰

$$\Omega = \int d^3x \delta E^i \lrcorner \delta A_i. \quad (3.15)$$

Here we take δE^i and δA_i to be the fundamental one-forms on phase space, and \lrcorner to be their anti-symmetrization (the exterior product of one-forms in an infinite-dimensional space).

Just as for finite dimension, $\frac{\partial}{\partial x^i}$, $\frac{\partial}{\partial p_i}$ are the vectors tangent to the x^i and p_i -coordinates: so also $\frac{\delta}{\delta E^i(x)}$ and $\frac{\delta}{\delta A^i(x)}$, for each $x \in \Sigma$, are vectors in the infinite-dimensional phase space. And just as for finite dimension, we can find new directions, or vector fields, by the linear sum, $\sum_i a^i \frac{\partial}{\partial x^i}$: here also, we can find new directions by linear sums, i.e. by integrating the fundamental directions smeared with certain coefficients, e.g.

$$\mathfrak{v} = \int d^3y v^i(y) \frac{\delta}{\delta E^i(y)}. \quad (3.16)$$

(where we use a double-struck notation to indicate that these are vector fields on the infinite-dimensional phase space of classical electromagnetism).

Moreover, configuration space, which is the space of smooth gauge potentials $\mathcal{A} = \{\mathbf{A} \in C^\infty(\Sigma, \mathbb{R}^3)\}$ admits an action of the (infinite-dimensional) group of gauge transformations,

¹⁹The Gauss constraint is Lie-algebra valued, even in the Abelian case; that is necessary for it to generate infinitesimal gauge transformations. It just so happens that here the Lie group G is $S^1 = U(1)$, and so the Lie algebra is $\mathfrak{g} \simeq \mathbb{R}$. In the non-Abelian theory, the Gauss constraint is $G^\alpha(x) = D_A E^\alpha - \rho^\alpha$, where, in a given basis, $D_A E^\alpha = \partial^i E_i^\alpha - \epsilon_{\beta\gamma}^\alpha E_i^\beta A^{i\gamma}$, where α, β, γ are Lie-algebra indices and ϵ is the structure constant of \mathfrak{g} in the given basis. Due to the appearance of A in the Gauss constraint, it acts linearly on \mathbf{E} : which is thus only covariant, and not invariant as it is in the Abelian theory.

²⁰It is easy to extend this to the presence of matter. For example, with a Klein-Gordon field, we would add: $\int \sqrt{g} \, d\bar{\psi} \lrcorner \gamma^0 d\psi$. See (Gomes & Riello, 2021, Section 3) for more details on the symplectic geometry of the infinite-dimensional space.

²¹In electromagnetism, the standard basis, namely $(E_i(x), A^i(x))$, for $x \in \Sigma$, is in this language just a basis that takes the coefficient functions to be all multiplied by Dirac deltas, e.g. a subset of smearings $v_x^i = E^i(x)$ defined as $v_x^i(y) = E^i \delta(x, y)$, for each $x \in \Sigma$ and $E^i \in \mathbb{R}$. This is what makes that basis completely local. Here we want to find a particular coordinate system, that is not completely local in this way. But in general, one could use symplectic duality to find the dual of any maximal Lagrangian submanifold in phase space.

$\mathcal{G} = \{\lambda \in C^\infty(\Sigma)\}$, namely: $\mu : \mathcal{G} \times \mathcal{A} \rightarrow \mathcal{A}$; $\mu(\lambda, \mathbf{A}) = \mathbf{A} - \text{grad}(\lambda)$.²² This structure exists also for the non-Abelian Yang-Mills theories. In any case, by using the isomorphism that Ω provides between $T\mathcal{A}$ and $T^*\mathcal{A}$, this action can be lifted to phase space in the usual manner (yielding (3.14), or its appropriate non-Abelian generalization).

3.3 Interpreting the Gauss constraint

We now combine Sections 1.1's and 3.2's discussions of the Gauss constraint with a comment on a proposal of Maudlin's. This will pave the way for our main result in Section 3.4.

As we said at the end of Section 1.1: the Gauss constraint (now in the form of (3.12)) involves a (non-signaling) kind of non-locality. For (by the elementary divergence theorem) the integral of \mathbf{E} over a surface enclosing a spatial region determines the total charge inside the surface *at the given instant*.²³ Like the Aharonov-Bohm effect, this type of non-locality is classical; but unlike the Aharonov-Bohm effect, it does not require an underlying topologically non-trivial spatial domain for the probe systems. (Again, compare Appendix A.)

As mentioned in this Section's preamble, this suggests isolating the part of the electric field that is determined by the instantaneous distribution of charges, and thinking of it as 'derivative' from this distribution. We will shortly pursue that idea. But first we notice that (Maudlin, 2018, p.10) proposes to 'turn this around'. Thus he writes:

Let us propose that this equation [i.e. $\text{div}(\mathbf{E}) = \rho$] represents not a physical law but an ontological analysis: electric charges just are the divergences of electric fields. In this way we reduce both the physical ontology and the nomology, and further gain an explanation of why electric charges cannot exist without electric fields.

So Maudlin proposes that the Gauss constraint be read as a definition of ρ , the electric charge. (He springboards this proposal from the corresponding one (p. 9) about \mathbf{B} : that $\text{div}(\mathbf{B}) = 0$ is completely equivalent to the non-existence of magnetic charges.)

We submit that—perhaps unfortunately—this proposal does not work. There are various problems. The most obvious is that there are particles with the same electric charge but differing in other ways; (they have different masses, and-or different charges for other interactions). That is: particles have characteristics independent of their electric charge. (And if one attempted to define the source of each distinct field in a similar eliminative manner, one would then have to explain why all the different sources and sinks coincide in space and time.) Other problems include: (i) since electromagnetism is a linear theory, must charged sources pass through each other? (ii) how do we explain interactions of the electromagnetic field with (apparent) matter, such as light reflection and refraction?²⁴

²²Due to the presence of stabilizers, \mathcal{A} is not an infinite-dimensional principal \mathcal{G} -bundle, i.e. $\mathcal{G} \curvearrowright \mathcal{A} \rightarrow \mathcal{A}/\mathcal{G}$. But, at least in the Abelian case, this is easy to remedy by considering only the group of pointed gauge transformations (i.e. those that are the identity at a given point). But \mathcal{A} always has a local product structure (i.e. it admits a slice, see Kondracki & Rogulski (1983) and references therein), even in the non-Abelian case.

²³Here, 'determines' can be read as 'gives us knowledge of'. For one can imagine measuring \mathbf{E} precisely throughout the surface, doing the integral, and inferring the total charge. This is, of course, the intuitive basis for quantum theory's charge superselection rule. For using this procedure, one can measure the total charge at an arbitrarily large distance from the system; and this suggests that measuring charge is compatible with measuring any other quantity on the system, so that the charge operator commutes with the operator representing that quantity.

²⁴Notice incidentally that Maudlin is re-inventing the wheel. It is a creditable re-invention, since the wheel has a venerable design: but for all that, the wheel does not roll. That is, without metaphors: Maudlin's proposal is the initial idea of the *electromagnetic world-view* that aimed to reduce mechanics (and all physics) to electromagnetism. In particular, it took the velocity-dependence of particles' masses to suggest that all mass might be of electromagnetic origin. It was advocated in the early twentieth-century by physicists such as Abraham and Mie; for an introduction cf. (Kragh, 1999, Chapter 8). We should also note the ongoing—yet

But the failure of this proposal does not impugn the more modest idea above: that we should isolate the part of the electric field that is determined by the instantaneous charge distribution, and think of it as ‘derivative’ from the distributions. So we think of this part—the Coulombic part—of the field as ‘the price to pay’ for a local formulation of the theory; but it can be replaced by the more fundamental distribution of charges acting at a distance. In other words: there is a certain component of the interaction between charges that does *not* take into account a field that has its own dynamics: all this component needs is the present distribution of the charges themselves. The remaining part of the electric field is then interpreted as ‘fundamental’, in the sense that it has its own dynamics: a dynamics that is not reducible to the dynamics of other components of the theory.

The upshot will be that, since \mathbf{E} and \mathbf{A} are conjugate, and we take part of \mathbf{A} to be ‘pure-gauge’, the above decomposition of \mathbf{E} can be made to correspond, in a particular way, to a certain decomposition of \mathbf{A} . We will do this in parallel to the toy case of §3.1, but here defining the gauge-complementary part of \mathcal{A} to be symplectically orthogonal to a natural parametrization of the Coulombic part of \mathbf{E} , and by defining the Coulombic-complementary part of \mathbf{E} to be symplectically orthogonal to the pure gauge part of \mathbf{A} .

3.4 Radiative and Coulombic parts of \mathbf{E} and \mathbf{A}

Now we come at last to our application of *the Helmholtz theorem*. The theorem states that any vector field \mathbf{Z} on \mathbb{R}^3 can be decomposed (‘split’) as:

$$\mathbf{Z} = \mathbf{X} + \mathbf{Y} \tag{3.17}$$

for a unique pair of curl-free (\mathbf{X}) and divergence-free (\mathbf{Y}) vector fields. So this theorem will be our route to the decomposition of \mathbf{A} and \mathbf{E} . It guarantees that we can mathematically “isolate” the divergence-free components and curl-free components of \mathbf{A} and of \mathbf{E} . In physics jargon, the divergence-free component is called ‘transverse’ (also ‘radiative’), illustrating that it carries two degrees of freedom; and the curl-free component is called ‘longitudinal’, for one degree of freedom.

We stress that—quite apart from this paper’s aims—although the split of the electric field into radiative and Coulombic components is mathematically useful, it is far from being of only mathematical relevance; and from being an arbitrary division of the field’s degrees of freedom. The split is physically meaningful. For instance, the question often arises about whether a process ‘radiates’ electromagnetic waves. For example, the famous “freely-falling electron paradox” asks this question (see e.g. [de Almeida & Saa \(2006\)](#) for a pedagogic review). That is because distinguishing the components of the electric field that are wave-like from those which are ‘Coulombic’ is not straightforward. These are the two characteristics that are expected of any definition of radiation: (i) it is transversal: i.e. it only contains polarizations that are orthogonal to the momenta; (ii) if the charges are confined to a compact region, the Coulombic and radiative parts of the electric field are the coefficients in an inverse radial expansion ($1/R^2$ and $1/R$ terms, respectively). The radiative \mathbf{E}^r that we will specify below has both these features.

Now, E_i and A^i generate the standard coordinates in phase space. These are coordinates in an infinite-dimensional space, which parametrize six degrees of freedom per spatial point.

also very speculative—programme to reduce both matter and radiation to structures in spacetime, sometimes called ‘super-substantivalism’: for which, [Misner & Wheeler \(1957\)](#) is a classic, and [Lehmkuhl \(2018\)](#) is a fine philosophical introduction.

Maudlin’s proposal also connects with the discussion at the end of Section 2.2.1. Namely: *if* the proposal held good so that charge was indeed fully defined by the electric field, then our objection in that discussion, that treatments of charge as sourcing the field but as unaffected by it are an idealization—which should be replaced by a “two-way” coupling—would fall by the wayside.

But if we want to think of the Gauss law as derivative (as we envisaged in Section 3.3), we need to find new coordinates for phase space (defined in terms of the original coordinates) that instead parametrize the radiative and Coulombic parts of the fields, much like what we did in Section 3.1. The Coulombic coordinate will then be uniquely fixed by the charge density distribution, while the radiative part should parametrize the rest of the field. These coordinates will no longer be local in spacetime, but they are nonetheless very useful, as remarked above. And once we have found these new coordinates, we can use the symplectic form Ω (cf. 3.15) to establish the conjugate decompositions of the gauge potential \mathbf{A} . Indeed, it is easiest to expound this re-definition of coordinates in terms of symplectic geometry, since the symplectic form Ω is a coordinate-independent object. Moreover, as in Section 3.1, it should be noted that the symplectic form is only non-degenerate in the total phase space. Therefore, to seek a decomposition through symplectic orthogonality, that decomposition should not be restricted to the constraint surface; rather it selects appropriate coordinates in the total phase space, and the constraint surface fixes some of these coordinates. In short: we are seeking adapted coordinates to parametrize the constraint surface.

Thus, on physical grounds, we would like to decompose the electric field as $E_i = E_i^r + E_i^c$ (with r for ‘radiative’ i.e. divergence-free, and c for Coulombic, i.e. curl-free), where

$$\partial^i E_i^r \equiv 0; \quad \text{and} \quad \partial^i E_i^c = \rho. \quad (3.18)$$

Both components are dynamically independent (the evolution of one is independent of the evolution of the other); and E_i^c is to be completely fixed by the distribution of the charges. This means, in vacuum ($\rho = 0$), that $E_i^c = 0$ and $E_i = E_i^r$. In short: \mathbf{E}^r is the component of the electric field that is not due to the distribution of charges (hence the label ‘radiative’); while \mathbf{E}^c is the component of the field due to the simultaneous distribution of charges. What is the form of this component?

The Gauss constraint reduces the degrees of freedom in \mathbf{E} from three to two. So when we decompose \mathbf{E} , writing: $E_i = E_i^r + E_i^c$, with E_i^c to be completely fixed by the distribution of the charges, as discussed above, we conclude that E_i^r has two degrees of freedom; so that E_i^c has one, i.e. it is ‘secretly’ a scalar. Since one of the terms— E^r —is divergence-free, and we want the degrees of freedom of the other— E^c —to be exhausted by the divergence, i.e. to be ‘secretly’ a scalar, it is convenient to introduce the Helmholtz decomposition: we take E_i^c to be curl-free, so that it is the gradient of a scalar. We thus *choose* to write this vector quantity in the customary way, i.e. as $E_i^c = \partial_i \phi$, for some scalar function ϕ .²⁵ Then, the Gauss constraint $G = 0$ fixes ϕ in terms of its simultaneous distribution of charges, via the Poisson equation:

$$\nabla^2 \phi_c = \rho. \quad (3.19)$$

Now, the purely Coulombic terms of the electric field have their phase space coordinate axis generated by the following vectors on phase space; (recall the discussion around eq. 3.16):

$$\mathbb{E}_c := \int d^3x \partial_i \phi \frac{\delta}{\delta E^i(x)}, \quad \text{for each } \phi \in C^\infty(\Sigma). \quad (3.20)$$

²⁵By the Poincaré lemma, in every simply connected domain, any curl-free vector field—which in the language of differential forms would be written as e.g. $d\mathbf{A} = 0$ —is of a pure gradient form. In the field of tensor algebra, one often refers (confusingly) to a ‘spin-decomposition’. Thus a 2-tensor T_{ab} may have vector and scalar components, of the form $\nabla_a V_b$ and $\nabla_a \nabla_b \phi$, respectively, and a vector V_a may have a scalar component, that is written as $\nabla_a \phi$. The extraction of these components usually is made through something like the Helmholtz decomposition theorem. In fact, the theorem is a special case of the Hodge decomposition theorem for n-forms (Morita, 2001, Chapter 43).

This coordinate can then be fixed by (3.19); with $\phi \stackrel{!}{=} \phi_c$ the electrostatic-like potential.²⁶

And now we can ask what are the subset of degrees of freedom of \mathbf{A} that are symplectically orthogonal to the Coulombic part of the electric field. Let us call these components A_i^r . That is, we seek those vector fields

$$\mathbb{A}^r = \int d^3x A_i^r \frac{\delta}{\delta A_i} \quad (3.21)$$

that are symplectically orthogonal to the Coulombic coordinates of the electric field. That is, the A_i^r are specified by requiring:

$$0 \stackrel{!}{=} \Omega(\mathbb{A}^r, \mathbb{E}_c) = \int d^3x A_i^r \partial^i \phi = - \int d^3x \phi \partial^i A_i^r, \quad \text{for all } \phi \in C^\infty(\Sigma); \quad (3.22)$$

where we applied integration by parts in the third equality and assumed that there is no boundary contribution to the integrals. Since ϕ is an entirely arbitrary test (or coordinate) function we get:

$$\partial^i A_i^r = 0. \quad (3.23)$$

Equation (3.23) is, of course, the *Coulomb gauge* for (the radiative component of) the gauge potential. It is a complete gauge-fixing, i.e. it leaves no more gauge-freedom in the potential. Moreover, just like in Section 3.1, where Q_- was written in terms of the original configuration variables, here we could write:

$$A_i^r(A) := A_i - \partial_i(\nabla^{-2} \partial^j A_j), \quad (3.24)$$

which is also called the *radiative projection*. From this equation, it is clear that for any gauge-transformed $A_i^\lambda := A_i + \partial_i \lambda$, a straightforward computation will show that $A_i^r(A^\lambda) = A_i^r(A)$. That is: A^r is gauge-invariant.

Correspondingly, instead of arguing for (3.18) from physical grounds as above, we can see the radiative part of the electric field as being selected as the component of the electric field that is symplectically orthogonal to the pure gauge part of the gauge potential. Namely, we take the pure gauge part, i.e. the vectors that are along the gauge-orbit, to be given by:

$$\mathbb{A}^c := \int d^3x \partial_i \lambda \frac{\delta}{\delta A_i}, \quad \text{for each } \lambda \in C^\infty(\Sigma). \quad (3.25)$$

Then, in parallel to the calculation (3.22), namely:

$$0 \stackrel{!}{=} \Omega(\mathbb{A}^c, \mathbb{E}_r) = \int d^3x E_i^r \partial^i \lambda = - \int d^3x \lambda \partial^i E_i^r, \quad \text{for all } \lambda \in C^\infty(\Sigma), \quad (3.26)$$

we find the defining equation for the fields that are symplectically orthogonal to the pure-gauge part of \mathbf{A} :

$$\partial^i E_i^r = 0. \quad (3.27)$$

Thus we have found that if we choose to represent the Coulombic degree of freedom by the gradient of a scalar, we obtain the Helmholtz decomposition for the electric field: a unique decomposition in terms of divergence-free and curl-free components. And we obtain a similar decomposition for the gauge potential. That is, we obtain:

$$\mathbf{A} = \mathbf{A}^r + d\lambda \quad \text{and} \quad \mathbf{E} = \mathbf{E}^r + d\phi. \quad (3.28)$$

²⁶One might worry here that not all vector fields define a coordinate system: their vector field commutator must vanish, for that to be the case. For finite dimensions, this is not a worry for $\sum_i a^i \frac{\partial}{\partial x^i}$, since the coefficients are constant, i.e. phase space independent. Happily, the same is true here: for the coefficients of the new directions do not themselves depend on the coordinates.

To sum up: we have shown (by just an integration by parts, in each case) that:

(1) the curl-free (longitudinal/Coulombic) component of \mathbf{E} is symplectically orthogonal to the radiative part of the gauge potential; (cf. equations (3.21) to (3.23)); while on the other hand:

(2) the divergence-free (transverse/radiative) component of \mathbf{E} is symplectically orthogonal to the pure gauge part of gauge field; (cf. equations (3.25) and (3.27)).

Finally, we note as a corollary to these results, that we can similarly orthogonally decompose the symplectic form (3.15) as:

$$\Omega = \int d^3x \left(\delta E_r^i \wedge \delta A_i^r + \delta E_c^i \wedge \delta A_i^c \right) : \quad (3.29)$$

which guarantees that the respective phase space directions are (symplectically) independent in each summand.

This concludes our main results. We end this Section with three comments:—

(i): In the presence of boundaries, the determination of E^r does not require the further stipulation of boundary conditions, *if* the gauge-freedom is taken as unconstrained at the boundary (see Gomes (2019a, 2021a,b)). Namely, we obtain that, at the boundary $E_i^r n^i = 0$, where \mathbf{n} is the vector normal to the boundary.

(ii): What we have just seen is a general feature of the symplectic geometric treatment. Namely: the tangent space to the constraint surface and the gauge orbits are the symplectic orthogonal complements of each other: cf. Lemma 1.2.2 in Marsden (2007). Thus in particular, the radiative part of the electric field is singled out just by the symplectic form and the gauge orbits.²⁷

(iii) : We arrived at the Helmholtz decomposition of the gauge potential by considering the symplectic structure and the Gauss constraint. But the Helmholtz decomposition can also be obtained via an inner product structure on the vector bundles in which the fields we are interested in take values. For instance, we split the vector potential into a component along the gauge orbit and one orthogonal to it—its radiative part—according to a natural metric on the corresponding space of functions. Similarly, according to an inner product on a different vector bundle on which the gauge group acts, we can ask for such a generalized Helmholtz decomposition, that splits any given field on a simply connected space into its pure gauge and ‘radiative’, or gauge-orthogonal components. This was studied for general scalar or vector-valued fields in (Gomes et al., 2019, Section 7), where it was found that the split is only mathematically well-defined where the field is non-vanishing, in which case it selects unitary gauge. (For the point-particle analogue of this split along/orthogonal to the gauge orbit, cf. (Gomes & Gryb, 2021, Section 7).) In those cases, the corresponding fixing of the gauge occurs locally and corresponds to what is labeled in (Berghofer et al., 2021; François, 2019) *an artificial symmetry*: one that can be locally projected out to produce a gauge-invariant “dressed” configuration. The dressed state for a nowhere-vanishing Klein-Gordon field corresponds to Wallace (2014)’s gauge-invariant, local composite field, with which one can give an entirely local account of the Aharonov-Bohm effect. It is likely, but has not been shown, that similarly to how in this paper we obtained Coulomb gauge, we could use the symplectic structure to arrive at the unitary gauge for the Klein-Gordon field.

²⁷The general idea, at least in the first-class, or coisotropic, case, is remarkably simple; and thus merits a quick sketch. From (2.14), vectors v that are tangent to the constraint surface obey $d\Phi^I(v) = 0$ for every I . Thus from (2.16), we obtain $\omega(X_I, v) = 0$.

4 Conclusion and outlook

Using the symplectic or Hamiltonian formalism, we have shown how decompositions of the electric field correspond, through symplectic orthogonality, to decompositions of the gauge potential (which correspond to choices of gauge), and vice versa. Gauge choices thus have a very natural interpretation in terms of choices of decomposition of the electric field.

We have argued that for various reasons a natural decomposition of the electric field takes one part, i.e. component, to be determined by the instantaneous charge distribution. We find that the part of the electric field that remains—i.e. the radiative part, the part that is not involved in the Gauss law and does not “care about” the instantaneous charge distribution—is symplectically *orthogonal* to the pure gauge part of the gauge potential. This orthogonality establishes a firm link between gauge symmetry and locality, as a relation between the Gauss law and the pure gauge part of the potential. But this relation does not, by itself, suffice to select a gauge-fixing.

On the other hand, *if* we define the Coulombic part of the electric field as the gradient of a Coulombic potential, then the gauge *is* fixed: the part of the gauge potential that is symplectically orthogonal to the gradient of the Coulombic potential *is* the gauge potential in Coulomb gauge.

Thus in summary, again: our main idea is that the electric field has the gauge potential as its conjugate, and there is a part of the potential that is not at all conjugate to (i.e. is symplectically orthogonal to) the Coulombic part of the electric field, i.e. the part determined by the instantaneous distribution of charges. *That* part of the potential satisfies the Coulomb gauge condition. And this result is worth expounding, since it shows that a choice of gauge need not be a matter of calculational convenience for some specific problem or class of problems, but can be related to a physically natural, and general, splitting of the electric field.

We also saw how these results prompt a comparison with [Maudlin \(2018\)](#). For Maudlin attempted to shift the Gauss law from, in his terms, the nomology to the derivative ontology: namely, by analysing the charges in terms of the electric field. We have argued that this does not work. In any case, it does not help with fixing a gauge. If instead we analyse the Coulombic part of the electric field away, defining it in terms of the charges by (3.19), then the fundamental ontology (to use Maudlin’s term) can be made to correspond, as we have seen in §3.4, to a configuration space parametrised by the charges and to a choice of the Coulomb gauge for the gauge potential.

But we emphasise that Coulomb gauge is not *mandatory*. It only corresponds, in a well-defined sense, to a particular decomposition that singles out the Coulombic part of the electric field. Nonetheless, in whichever sense that choice of decomposition of the electric field is natural, Coulomb gauge is also natural.

Finally, we offer an *outlook*: we stress that the lessons of this paper go through, with minor modifications, to the non-Abelian domain, and also apply in the presence of boundaries; (see e.g. [Gomes \(2019a\)](#); [Gomes & Riello \(2021\)](#)). In brief, each extension requires one important modification.

First, in the non-Abelian case, the Coulombic split occurs only at the level of *perturbations*. That is, although we can split vectors on phase space, \mathbb{X} , the non-Abelian nature of the theory implies that the split is not integrable (see ([Gomes & Riello, 2021](#), Section 5) and ([Gomes et al., 2019](#), Section 9)). Thus the Coulombic split of a state is always ambiguous.²⁸

In the presence of boundaries, in the Abelian case, the radiative and Coulombic components of the electric field are again defined by symplectic orthogonality with the pure gauge part of the potential, which we deem unconstrained at the boundary (cf. ([Gomes & Riello, 2021](#), Prop.

²⁸In the Abelian case, one can integrate the perturbative split along paths in phase space so as to define a split of any final state; integrability guarantees that the end result is path-independent.

3.3)). Thus we obtain the following modifications, for a bounded region R bounded by ∂R , whose normal is \mathbf{n} :

Instead of (3.23),

$$\begin{cases} \partial^i A_i^r = 0 & \text{in } R, \\ n^i A_i^r = 0 & \text{at } \partial R. \end{cases}; \quad (4.1)$$

instead of (3.27),

$$\begin{cases} \partial^i E_i^r = 0 & \text{in } R, \\ n^i E_i^r = 0 & \text{at } \partial R. \end{cases}; \quad (4.2)$$

and instead of (3.19),

$$\begin{cases} \nabla^2 \phi_c = \rho & \text{in } R, \\ n^i \partial_i \phi = f & \text{at } \partial R. \end{cases}; \quad (4.3)$$

where f is the electric flux through the boundary, $f := n^i E_i$, and here can be seen as an independent variable. Note that: (i) the pure gauge part of the potential is unmodified (since we do not truncate gauge transformations at the boundary ∂R), and (ii) the normal to the electric field at the boundary belongs to the Coulombic part. The radiative part gets no extra degree of freedom at the boundary.²⁹

Acknowledgements

Many thanks to Caspar Jacobs for helpful comments on an earlier version of this paper.

Appendix

A Holonomies and non-separability

The holonomy interpretation of electromagnetism takes as its basic elements assignments of unit complex numbers to loops in spacetime. A loop is the image of a smooth embedding of the oriented circle, $\gamma : S^1 \rightarrow \Sigma$; the image is therefore a closed, oriented, non-intersecting curve. One can form a basis of gauge-invariant quantities for the holonomies (cf. Barrett (1991) and (Healey, 2007, Chapter 4.4) and references therein):³⁰

$$hol(\gamma) := \exp\left(i \int_{\gamma} A\right). \quad (A.1)$$

In more detail: by exponentiation (path-ordered in the non-Abelian case), we can assign a complex number (matrix element in the non-Abelian case) $hol(C)$ to the oriented embedding of the unit interval: $C : [0, 1] \mapsto M$. This makes it easier to see how composition works: if the endpoint of C_1 coincides with the starting point of C_2 , we define the composition $C_1 \circ C_2$ as, again, a map from $[0, 1]$ into M , which takes $[0, 1/2]$ to traverse C_1 and $[1/2, 1]$ to traverse

²⁹Thus, the radiative part is local, according to straightforward definitions of locality (cf. (Wallace, 2019, §2, p.5)). The Coulombic field depends on the distribution of ρ in the region *and* on the boundary flux of the electric field. Thus, even if the charge distribution for two worlds matches inside R , the Coulombic field therein may differ, since in each world it will depend on the independent variable that is the electric boundary flux.

³⁰Of course, any discussion of matter charges and normalization of action functionals would require e and \hbar to appear. However, we are not treating matter, so these questions of choice of unit do not become paramount. As before, if needed, we set our units to $e = \hbar = 1$; as is the standard choice in quantum chromodynamics (or as in the so-called Hartree convention for atomic units).




Figure 2: Two subregions, i.e. R_{\pm} , with the separating surface S . A larger loop γ crosses both regions. But, since γ_1 and γ_2 traverse S along C in opposite directions, $\gamma = \gamma_1 \circ \gamma_2$.

C_2 . The inverse C^{-1} traces out the same curve with the opposite orientation, and therefore $C \circ C^{-1} = C(0)$.³¹ Following this composition law, it is easy to see from (A.1) that

$$\text{hol}(C_1 \circ C_2) = \text{hol}(C_1)\text{hol}(C_2), \quad (\text{A.2})$$

with the right-hand-side understood as complex multiplication in the Abelian case, and as composition of linear transformations, or multiplication of matrices, in the non-Abelian case.

As both Healey (Healey, 2007, Chapter 4.4) and Belot ((Belot, 2003, Section 12) and (Belot, 1998, Section 3)) have pointed out: even classical electromagnetism, in the holonomy interpretation, exhibits a form of non-locality, which at first sight, one might have thought was a hallmark of non-classical physics.

But one naturally asks: does the state of a region nevertheless supervene on assignments of intrinsic properties to the sub-regions of the region (where the sub-regions may be taken to be arbitrarily small)? This is essentially the question whether the theory is *separable*: (see (Healey, 2007, Chapter 2.4), (Belot, 1998, Section 3), (Belot, 2003, Section 12), Healey & Gomes (2021); Myrvold (2010) and (Wallace, 2019, Section 2, p.5)). For this topic, we can focus on Myrvold's definition (Myrvold, 2010, p.427). (It builds on Healey's notion of Weak Separability (Healey, 2007, p. 46) and on Belot's notion of Synchronic Locality (Belot, 1998, p. 540).)

- *Patchy Separability for Simply Connected Regions*. For any simply connected spacetime region R , there are arbitrarily fine open coverings $\mathcal{N} = \{R_i\}$ of R such that the state of R supervenes on an assignment of qualitative intrinsic physical properties to elements of \mathcal{N} .

For electromagnetism in vacuum, it is easy to show that *Patchy Separability for Simply Connected Regions* does hold. In Figure 2, we see a loop γ not contained in either R_+ or R_- . However, we can decompose it as $\gamma = \gamma_+ \circ \gamma_-$, where each regional loop γ_{\pm} does not enter the complementary region (R_{\mp} , respectively), but each traverses the curve c lying within the boundary S each in the opposite direction to the other. Using (A.2), since the contributions to the holonomies of the two paths along c cancel out, the fact that holonomies form a basis of gauge-invariant quantities, then implies that the gauge-invariant state of the whole region

³¹Intuitively, we do not want to consider curves that trace the same path back and forth, i.e. *thin* curves. Thus we define a closed curve as *thin* if it is possible to shrink it down to a point while remaining within its image. Quotienting the space of curves by those that are thin, we obtain the space of *hoops*, and this is the space actually considered in the treatment of holonomies. We will not emphasise this finer point, since it follows from an intuitive understanding of the composition of curves.

supervenes on the gauge-invariant states of its sub-regions. As one might put it: the gauge-invariant state of the universe, according to the theory, supervenes on the gauge-invariant states of its regions. This corresponds to the fact that in the absence of charges, and for a simply connected manifold, we can solve for the Coulombic part of the field simply by setting $E_c = 0$.

But Patchy Separability fails for non-simply connected regions. This is easy to see from Figure 2: if one introduces a hole in the middle of the curve c , say composed of two semi-circles c_1 and c_2 that meet along c , we can no longer write $\gamma = \gamma_+ \circ \gamma_-$, since curves in each region, traversing either c_1 or c_2 , are not aligned and thus their contribution does not cancel out. (Compare (Gomes & Riello, 2021, Section 6.8.2) for how this relates to Coulombic vs. radiative modes, and to the topological non-locality of the Aharonov-Bohm effect.).

It is also easy to see how *Patchy Separability for Simply Connected Regions fails* when charges are present within the regions but absent from the boundary S ; (see in particular (Gomes & Riello, 2021, Section 4.3.2), and footnote 70 in Gomes (2021a)). For, in the presence of charges, we can form *gauge-invariant* functions from a non-closed curve C' that crosses S and has e.g. one positive and one negative charge, $\psi_{\pm}(x_{\pm})$, capping off C' at $x_{\pm} \in R_{\pm}$. That is, the following quantity is a gauge-invariant function:

$$Q(C', \psi_{\pm}) = \psi_{-}(x_{-})\text{hol}(C')\psi_{+}(x_{+})$$

for $C'(0) = x_{-}, C'(1) = x_{+}$. It is easy to check from the transformation properties, $A_i \mapsto A_i + \partial_i \lambda$ and $\psi \mapsto e^{i\lambda} \psi$, that Q is gauge-invariant. Moreover, we cannot break this invariant up into gauge-invariant contributions from the two regions, since we have assumed no charges lie at the boundary.

References

- Aharonov, Y., & Bohm, D. (1959, Aug). Significance of electromagnetic potentials in the quantum theory. *The Physical Review*, *115*, 485–491. Retrieved from <https://link.aps.org/doi/10.1103/PhysRev.115.485> doi: 10.1103/PhysRev.115.485
- Arnold, V. (1989). *Mathematical methods of classical mechanics* (Vol. 60). Springer, Berlin.
- Barrett, J. W. (1991, Sep 01). Holonomy and path structures in general relativity and yang-mills theory. *International Journal of Theoretical Physics*, *30*(9), 1171–1215. Retrieved from <https://doi.org/10.1007/BF00671007> doi: 10.1007/BF00671007
- Belot, G. (1998, 12). Understanding Electromagnetism. *The British Journal for the Philosophy of Science*, *49*(4), 531-555. Retrieved from <https://doi.org/10.1093/bjps/49.4.531> doi: 10.1093/bjps/49.4.531
- 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. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1355219803000042> doi: [https://doi.org/10.1016/S1355-2198\(03\)00004-2](https://doi.org/10.1016/S1355-2198(03)00004-2)
- Berghofer, P., François, J., Friederich, S., Gomes, H., Hetzroni, G., Maas, A., & Sondenheimer, R. (2021). *Gauge symmetries, symmetry breaking, and gauge-invariant approaches*.
- Butterfield, J. (2006). On symmetry and conserved quantities in classical mechanics. In W. Demopoulos & I. Pitowsky (Eds.), *Physical theory and its interpretation: Essays in honor of jeffrey bub* (pp. 43–100). Dordrecht: Springer Netherlands. Retrieved from https://doi.org/10.1007/1-4020-4876-9_3 doi: 10.1007/1-4020-4876-9_3

- Butterfield, J. (2007). On symplectic reduction in classical mechanics. In J. Butterfield & J. Earman (Eds.), *Philosophy of physics* (p. 1 - 131). Amsterdam: North-Holland. Retrieved from <http://www.sciencedirect.com/science/article/pii/B978044451560550004X> doi: <https://doi.org/10.1016/B978-044451560-5/50004-X>
- de Almeida, C., & Saa, A. (2006). The radiation of a uniformly accelerated charge is beyond the horizon: A simple derivation. *American Journal of Physics*, *74*(2), 154-158. Retrieved from <https://doi.org/10.1119/1.2162548> doi: 10.1119/1.2162548
- Dougherty, J. (2020). The Non-Ideal Theory of the Aharonov–Bohm Effect. *Synthese*.
- Earman, J. (2019, May 01). The role of idealizations in the aharonov–bohm effect. *Synthese*, *196*(5), 1991–2019. Retrieved from <https://doi.org/10.1007/s11229-017-1522-9> doi: 10.1007/s11229-017-1522-9
- François, J. (2019). Artificial versus substantial gauge symmetries: A criterion and an application to the electroweak model. *Philosophy of Science*, *86*(3), 472-496. Retrieved from <https://doi.org/10.1086/703571> doi: 10.1086/703571
- Gomes, H. (2019a). Gauging the boundary in field-space. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1355219818302144> doi: <https://doi.org/10.1016/j.shpsb.2019.04.002>
- Gomes, H. (2019b). Noether charges, gauge-invariance, and non-locality. *Comissioned from Cambridge University Press, for special edition celebrating 100 years of Noether's theorems*.
- Gomes, H. (2021a). Holism as the significance of gauge symmetries. *European Journal of Philosophy of Science*, vol 11, 87.
- Gomes, H. (2021b). The role of representational conventions in assessing the empirical significance of symmetries. (to appear in *Studies in History and Philosophy of Modern Physics*).
- Gomes, H. (2021c). Same-diff? Part I: Conceptual similarities (and one difference) between gauge transformations and diffeomorphisms. *Arxiv: 2110.07203. Submitted..*
- Gomes, H. (2021d). Same-diff? Part II: A compendium of similarities between gauge transformations and diffeomorphisms. *Arxiv: 2110.07204. Submitted..*
- Gomes, H., & Butterfield, J. (2021a). Assessing the hole argument. *In preparation*.
- Gomes, H., & Butterfield, J. (2021b). A glimpse of symplectic reduction. *In preparation*.
- Gomes, H., & Gryb, S. (2021). Angular momentum without rotation: Turbocharging relationalism. *Studies in History and Philosophy of Science Part A*, *88*, 138-155. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0039368121000704> doi: <https://doi.org/10.1016/j.shpsa.2021.05.006>
- Gomes, H., Hopfmüller, F., & Riello, A. (2019). A unified geometric framework for boundary charges and dressings: Non-abelian theory and matter. *Nuclear Physics B*, *941*, 249 - 315. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0550321319300483> doi: <https://doi.org/10.1016/j.nuclphysb.2019.02.020>
- Gomes, H., & Riello, A. (2021). The quasilocal degrees of freedom of Yang-Mills theory. *SciPost Phys.*, *10*, 130. Retrieved from <https://scipost.org/10.21468/SciPostPhys.10.6.130> doi: 10.21468/SciPostPhys.10.6.130

- Gomes, H., Roberts, B., & Butterfield, J. (2021). The Gauge Argument: a Noether Reason. *Comissioned from Cambridge University Press, for special edition celebrating 100 years of Noether's theorems.*
- Healey, R. (2007). *Gauging What's Real: The Conceptual Foundations of Gauge Theories.* Oxford University Press.
- Healey, R., & Gomes, H. (2021). Holism and Nonseparability in Physics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2016 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2016/entries/physics-holism/>.
- Henneaux, M., & Teitelboim, C. (1992). *Quantization of gauge systems.* Princeton University Press.
- Jackson, J. D. (1975). *Classical electrodynamics; 2nd ed.* New York, NY: Wiley. Retrieved from <https://cds.cern.ch/record/100964>
- Kondracki, W., & Rogulski, J. (1983). *On the stratification of the orbit space for the action of automorphisms on connections. on conjugacy classes of closed subgroups. on the notion of stratification.* Inst., Acad. Retrieved from <https://books.google.co.uk/books?id=LK0JrgEACAAJ>
- Kragh, H. (1999). *Quantum Generations.* Princeton University Press.
- Lang, S. (2012). *Graduate Texts in Mathematics: Fundamentals of Differential Geometry.* Springer Verlag, New York.
- Lehmkuhl, D. (2018). The Metaphysics of Super-Substantivalism. *Nôus*, 52(1), 24-46. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12163> doi: <https://doi.org/10.1111/nous.12163>
- Marsden, J. (2007). Symplectic Reduction. In *Hamiltonian reduction by stages* (pp. 3–42). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from https://doi.org/10.1007/978-3-540-72470-4_1 doi: 10.1007/978-3-540-72470-4_1
- Maudlin, T. (2018). Ontological clarity via canonical presentation: Electromagnetism and the aharonov–bohm effect. *Entropy*, 20(6). Retrieved from <https://www.mdpi.com/1099-4300/20/6/465> doi: 10.3390/e20060465
- Misner, C., & Wheeler, J. (1957). Classical physics as geometry. *Annals of Physics*, 2(6), 525-603. Retrieved from <https://www.sciencedirect.com/science/article/pii/0003491657900490> doi: [https://doi.org/10.1016/0003-4916\(57\)90049-0](https://doi.org/10.1016/0003-4916(57)90049-0)
- Morita, S. (2001). *Translations of Mathematical Monographs Iwanami Series in Modern Mathematics: Geometry of Differential forms.* American Mathematical Society.
- Mulder, R. (2021). Gauge-Underdetermination and Shades of Locality in the Aharonov–Bohm Effect. *Foundations of Physics.*
- Myrvold, W. C. (2010, 03). Nonseparability, Classical, and Quantum. *The British Journal for the Philosophy of Science*, 62(2), 417-432. Retrieved from <https://doi.org/10.1093/bjps/axq036> doi: 10.1093/bjps/axq036

- 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. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1355219814000586> doi: <https://doi.org/10.1016/j.shpsb.2014.05.007>
- Pitts, J. B. (2014b, Dec). A first class constraint generates not a gauge transformation, but a bad physical change: The case of electromagnetism. *Annals of Physics*, 351, 382–406. Retrieved from <http://dx.doi.org/10.1016/j.aop.2014.08.014> doi: 10.1016/j.aop.2014.08.014
- 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. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1355219805000456> doi: <https://doi.org/10.1016/j.shpsb.2005.04.004>
- Putnam, H. (1975). The Analytic and Synthetic. In *Mind, language and reality: Philosophical papers* (pp. 33–69). Cambridge University Press.
- Shech, E. (2018). Idealizations, Essential Self-Adjointness, and Minimal Model Explanation in the Aharonov–Bohm Effect. *Synthese*, 195(11), 4839–4863. doi: 10.1007/s11229-017-1428-6
- Wallace, D. (2002). Time-Dependent Symmetries: The Link Between Gauge Symmetries and Indeterminism. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: Philosophical reflections* (pp. 163–173). Cambridge University Press.
- Wallace, D. (2014). Deflating the Aharonov-Bohm Effect. *arxiv: 1407.5073*.
- Wallace, D. (2019). Isolated systems and their symmetries, part II: local and global symmetries of field theories. Retrieved from <http://philsci-archive.pitt.edu/16624/>