On North's "The Structure of Physics"

Noel Swanson and Hans Halvorson*

September 6, 2012

In "The 'structure' of physics: a case study," (*Journal of Philosophy* 106: 57–88) Jill North argues that when faced with with competing physical hypotheses, other things being equal, we should believe the one that imputes *less structure* to the world.[1] Such arguments have a distinguished pedigree; both contemporary and historical interpreters of physics have famously used Ockhamist principles to draw conclusions about the ontology of spacetime. North's paper gives a more generalized account of Ockham's razor inspired by contemporary structural realism, and applies it not to spacetime structure, but *statespace* structure. As a case in point, she argues that the Hamiltonian formulation of classical mechanics requires less statespace structure than the Lagrangian formulation, and that we should therefore believe the former theory rather than the latter.[2]

In this comment, we question North's conclusion about classical mechanics and point out some general challenges involved in comparing amounts of structure postulated by physical theories. North crucially relies on symmetry criteria to quantify and compare amounts of theoretical structure. Our criticism of her argument is twofold. First, the different criteria she employs are neither conceptually nor mathematically precise. We clarify them and discuss some subtleties involved in their application to the case of Hamiltonian and Lagrangian mechanics. Second, in order to apply

*Department of Philosophy, Princeton University

[1]"It is not that, other things being equal, we should go with the fewest entities, but that we should go with the *least structure*." (North p. 64)

[2]This argument makes two crucial assumptions which we will not directly address here. The first is that any additional structure (e.g. dynamical vector fields, spacetime structure) posited by the two theories is roughly equal or similarly weighted in favor of Lagrangian mechanics. The second is that statespace structure actually corresponds to some kind of real, physical structure. Many physicists and philosophers of science view statespace as a formal tool for efficiently encoding certain modal commitments of scientific theories, an aid for making calculations and predictions, but not part of the theory's ontology. North's position presupposes some form of *statespace realism*, although she leaves its exact contours flexible:

> Whether to say that statespace exists as a concrete thing; or that all there is, fundamentally is statespace and its structure; or an altogether less radical view, is left open here. What's not open is that we should take statespace seriously, where this simply means that statespace has a definite, natural structure, which we can be right or wrong about depending on the way the world is, and which we should try to minimize. (North p. 33)

Even if one thinks that statespace has an important explanatory role in physics, and is not merely a formal tool, it is far from clear that statespace structure needs to be minimized. As we will see later on, however, even with this premise in place, great care must be taken to accurately assess the connection between the mathematics and the physics.

such criteria it is essential to correctly identify the symmetries of the models used by competing theories to describe the world. We argue that North's analysis fails to do this. Lacking such a characterization, we conclude that at present none of the proposed (mathematically precise) accounts of symmetry is adequate to support an argument that Hamiltonian mechanics imputes less structure to the classical world.

Does Hamiltonian statespace have less structure than its Lagrangian counterpart (as North contends)? Initially, this seems like a question best left to physicists and mathematicians to settle. But it is here that we find a host of unexpected philosophical challenges. What interpretive principles should we use to ascertain a theory's structural commitments? Can we read these commitments off from the mathematical structure of the theory? And how then do we compare the amounts of structure imputed by different theories?

In regards to this last question, North provides a helpful suggestion — amount of structure is inversely related to size of symmetry groups. She observes the following:

> The geometric structure of a mathematical space is given by quantities that remain intact under changes in coordinates. There is a difference, then, between the features ascribed to a space by the coordinate system being used, and the intrinsic features of the space itself. There is a difference between (genuine) *structure* and (mere) *description* of that structure. (p. 62)

According to a standard interpretation of symmetries which North heavily relies upon, the symmetries of a mathematical space correspond directly to possible coordinate transformations on the space. They represent mere changes in description. Symmetries therefore give us an abstract way of characterizing the intrinsic structure of the space. Intuitively, the more symmetries a space has, the less mathematical structure there is.

The same idea applies to physical theories. Theories describe the world using mathematical models. These models are invariant under certain symmetry transformations, which in turn characterize the intrinsic physical structure of the model. Only those mathematical structures that are invariant under allowed changes in description are interpreted as being physically significant. North proposes that given two rival theories, we can compare the amount of structure they posit by comparing the symmetries of their respective models — one model has less structure than another if it admits a larger group of symmetries.[3] But what is the relevant notion of size here? How are we supposed to count and compare numbers of symmetries?

Some of North's arguments suggest that *dimension* might play this role. For example, she notes that if $\langle V, p\rangle$ is an n-dimensional inner-product space, then its group $\mathrm{Aut}(\langle V, p\rangle)$ of automorphisms has dimension $2n^2 - 2$; in contrast, if $\langle V, \omega\rangle$ is an n-dimensional symplectic vector space, then $\mathrm{Aut}(\langle V, \omega\rangle)$ has dimension $2n^2 + 2$. She concludes from these facts that a symplectic form has less structure than a metric.[4] So, we might rephrase North's symmetry principle as follows:

(SYM) Given models X and Y, if the dimension of $\mathrm{Aut}(X)$ is greater than the dimension of $\mathrm{Aut}(Y)$, then X has less structure than Y.

[3] "[An inner product on a vector space V] is a stronger structure [than a symplectic form on V], in that it admits a smaller group of symmetries." (North p. 87)

[4] ibid, p. 87.

But this criterion has serious defects. First, it only applies to symmetry groups that have a dimension — e.g. symmetry groups that are themselves manifolds — and there are many interesting symmetry groups to which this concept does not apply.[5] Second, this criterion yields obviously wrong verdicts in numerous cases. For example, if V and W are inner product spaces with $\dim V \geq 2$ and $\dim W \geq 2$, then $\dim[\mathrm{Aut}(V \otimes W)] > \dim[\mathrm{Aut}(V)]$. Hence it entails that a system whose state space is $V \otimes W$ has less structure than a system whose state space is V. It would follow, in particular, that a pair of electrons has less structure than one electron!

A more promising alternative employs subset inclusion as a measure of relative size:

(SYM*) If $\mathrm{Aut}(X)$ properly contains $\mathrm{Aut}(Y)$ then X has less structure than Y.

Intuitively, if for every allowable coordinate transformation on Y, there is a corresponding transformation on X, but not vice versa, there must be additional structure on Y that is not preserved by the extra transformations on X. In certain situations, this intuition can be made more precise using the connection between symmetries and definability provided by the Beth-Svenonius definability theorem:

> Let $\mathcal{L} \subseteq \mathcal{L}^+$ be first order languages, with R a relation symbol in $\mathcal{L}^+$. Let T be a complete theory stated in $\mathcal{L}^+$. R is explicitly definable in T in terms of $\mathcal{L}$ iff for every model $M \vDash T$, the subset R^M is invariant under all automorphisms of $M|\mathcal{L}$, the structure obtained from M by omitting all relations and operations not contained in $\mathcal{L}$.[6]

The Beth-Svenonius theorem provides a handy tool for showing when one theory T^+ imputes more structure than another theory T (supposing that the vocabulary of T is a subset of the vocabulary of T^+): just show that there are symmetries of T-models that do not leave all the T^+-imputed structure invariant. Although North never explicitly mentions this theorem, it is plausible to think that she has something like it in mind when she claims that symmetries are inversely related to "levels of structure" (North p. 65). For example, it is something like (SYM*) that seems to underlie her claim that a metric cannot be defined in terms of a symplectic form: "So long as a transformation leaves the symplectic structure intact, it can alter any metric structure" (North p. 74).

Thus, (SYM*) receives justification from the analogy to the first-order case (the Beth-Svenonius theorem). Nonetheless, this criterion is of limited applicability for comparing the state spaces of physical theories: X has more structure than Y only in cases where X and Y are structures on a common base set S, for example $X = \langle S, R_1, \ldots, R_n \rangle$ and $Y = \langle S, R'_1, \ldots, R'_m \rangle$, where the R_i and R'_j are relations on the set S. In the context of comparing models of Hamiltonian and Lagrangian mechanics, (SYM*) can therefore only be employed if there exists a suitable translation between the underlying set of states in each model (a fact which will prove important later on). We will employ (SYM*) as a precisification of North's symmetry criterion throughout the remainder of our comment.

We now turn to North's main argument that Lagrangian mechanics imputes more structure than Hamiltonian mechanics. She begins with the claim that Lagrangian statespace always has the

[5] See, for example, Doplicher and Piacitelli, "Any compact group is a gauge group," *Reviews of Mathematical Physics* 14 (2002): pp. 873–886.

[6] see W. Hodges, *Model Theory*, (Springer, 1993): p. 516.

structure of a Riemannian manifold, whereas Hamiltonian statespace always takes the form of a symplectic manifold. It is well known that a symplectic form defines a volume element, but does not define a metric. Based on these facts, North argues as follows:

> The thing is, as far as we can tell, we need only symplectic structure to do classical mechanics. This structure suffices for the theory; it does not "leave anything out." And there is a clear sense in which a space with a metric structure has more structure than one with just a volume element. Metric structure comes with, or determines, or presupposes, a volume structure, but not the other way around. (In the same way that a metric space comes with, or determines, or presupposes, a topology, and not the other way around.) Intuitively, knowing the distances between the points in a space will give you the volumes of the regions, but the volumes will not determine the distances. Metric structure adds a further level of structure. (North pp. 74–75)

This argument is supposed to show that a metric has more structure than a symplectic form. But a symplectic form is not the same thing as a volume element (a fact which can be quickly seen by noting that an odd-dimensional manifold can have a volume measure, but can have no symplectic form). In general, both a metric and a symplectic form define corresponding volume measures; but a volume measure defines neither a metric, nor a symplectic form.[7]

Although this argument is a red herring, North's subsequent comments partially close the gap. She notes that given a manifold with both a symplectic form and a metric, there are symmetries of the symplectic form that do not preserve the metric; symplectic structure does not define metric structure. This indicates that if $\mathrm{Aut}(M, g) \subseteq \mathrm{Aut}(M, \omega)$, it will be a proper subset; however, North never provides an argument for the antecedent. In fact, it is false. Given any Riemannian manifold (M, g) that admits a symplectic form ω, there is a symmetry J of (M, g) that does not preserve ω.[8] In particular, isometries can reverse orientation while symplectomorphisms must preserve orientation. Thus, metric structure does not define symplectic structure either. Symplectic manifolds and Riemannian manifolds each possess a piece of structure the other lacks, hence it does not follow that Lagrangian mechanics imputes more statespace structure than Hamiltonian mechanics. But unless these physical theories are identical to the theories of Riemannian and symplectic geometry, we should be wary of drawing any definite conclusions from these considerations. Indeed, there are two glaring problems with the preceding analysis.

First, North's initial assumption about the structure of Lagrangian statespace is suspect. Models of both Lagrangian *and* Hamiltonian mechanics can be found which posit a metric on statespace. The crucial difference, according to North is that the Lagrangian dynamics require such a metric whereas the Hamiltonian dynamics do not. But this is false. In what is arguably its best, most general formulation, Lagrangian mechanics does not make use of a metric at all.[9] Even when a metric is

[7] What is more, given a manifold with both a symplectic form and a metric, the volume measures defined by these two different structures do not necessarily agree. If they coincide, the space is said to possess Kähler structure.

[8] Let $\{e_1, e_2\}$ be an orthonormal basis, and define $J : V \to V$ by setting $Je_1 = e_2$ and $Je_2 = e_1$ and extending linearly. Clearly J is an isometry, but J does not preserve the symplectic form since $\omega(Je_1, Je_2) = \omega(e_2, e_1) \neq \omega(e_1, e_2)$.

[9] Instead it employs the naturally defined *almost-tangent structure* of Lagrangian statespace to frame the Euler Lagrange equations. See de Leon and Rodrigues, *Methods of Differential Geometry in Analytical Mechanics*, (North Holland, 1989): Ch. 2 & 7 for details. North discusses this formulation briefly in her appendices (pp. 83–88), arguing that

present in the Lagrangian (or the Hamiltonian) framework, its physicality is dubious. Generally, the statespace metric is induced by a metric on the underlying configuration space representing the total kinetic energy of the system; dynamically possible solutions are geodesics, paths of extremal kinetic energy. But neither the total kinetic energy nor the metric are invariant under constant velocity boosts, so-called *Galilean transformations*, which are among the symmetries of classical spacetime. It appears that in many cases a metric is simply a useful calculational tool, precisely the kind of thing that North thinks should not have deep ontological significance.

Despite this unresolved issue, there is a second, deeper problem which takes precedence. A specific Hamiltonian system involves both a statespace (M, ω) as well as a Hamiltonian function $H : M \to \mathbb{R}$, and symmetries of *that* system will have to preserve some data specific to the Hamiltonian function. Similarly, a specific Lagrangian system includes not only a statespace TQ (which is required to be the tangent bundle of an underlying configuration space manifold, Q) but also a Lagrangian function $L : TQ \to \mathbb{R}$. The addition of a Hamiltonian or Lagrangian function not only specifies the dynamics, but also allows for the definition of a number of auxiliary geometric structures.[10] When asking questions about what kinds of structure can be defined on statespace we must avail ourselves of the resources of the entire physical theory, not just mathematical subtheories. It might still be possible to use symmetry criteria to compare the structure of Lagrangian and Hamiltonian mechanics, but more care must be taken to identify the models and symmetries of the two theories.

We recall that, according to standard accounts, the symmetries of a Hamiltonian model (M, ω, H) are the "canonical transformations," which is just a synonym for "symplectomorphisms" of (M, ω).[11] In contrast, the symmetries of a Lagrangian model (TQ, L) are usually required to be "point transformations." A point transformation $\varphi : TQ \to TQ$ is a diffeomorphism of the form Tf, where $f : Q \to Q$ is a diffeomorphism.[12] (Here Tf is the pushforward map of tangent vectors.) Canonical transformations preserve the form of Hamilton's equations of motion whereas point transformations preserve the form of the Euler-Lagrange equations of motion (although neither are required to preserve the exact form of the Hamiltonian or Lagrangian function). At the end of her second appendix, North makes the following observation "whereas all point transformations are canonical transformations — point transformations form a subgroup of the set of all canonical transformations" (p. 88), citing Abraham and Marsden.[13] Deploying this fact, we should be able to use (SYM*) to conclude that Hamiltonian mechanics has less structure than Lagrangian mechanics.

But something is not quite right with this argument: a Lagrangian symmetry is an automorphism of a tangent bundle TQ, and a canonical transformation is an automorphism of a symplectic manifold (M, ω). Without a suitable mapping of the states in TQ to the states in M, comparing

---

even in this more general case there will be a difference in structure between the two theories. Much of the force of this argument rests on her identification of canonical and point transformations as the symmetries of Hamiltonian and Lagrangian mechanics. As we will go on to see, this assumption is problematic.

[10] For example, once a Lagrangian function is specified, it is possible to define a symplectic form on TQ along with an almost-complex and an almost-Kähler structure.

[11] see Defn. 3.2.5 in R. Abraham and J.E. Marsden, *Foundations of Mechanics: 2nd Edition*, (American Mathematical Society, 2008): p. 177.

[12] ibid, p. 181.

[13] "the point transformations clearly form a subgroup of the set of all canonical transformations," ibid, p. 181.

the symmetry group of a Lagrangian system to the symmetry group of a Hamiltonian system is like comparing apples to oranges.

The key to resolving this impasse is to note that Abraham and Marsden's claim is about point and canonical transformations on the cotangent manifold T^*Q, with its canonical symplectic form ω. The uncontroversial fact which they point out is that any pullback diffeomorphism $T^*f : T^*Q \to T^*Q$ preserves the canonical symplectic form ω. If we consider a Hamiltonian statespace T^*Q as isomorphic (via the Legendre transformation) to the Lagrangian statespace TQ, then we can interpret point transformations on T^*Q as corresponding to point transformations on TQ. But if we consider T^*Q as a bare symplectic manifold — ignoring its relation to TQ — then we have a larger group of symmetries, namely all symplectomorphisms of (T^*Q, ω). Here the Legendre transform plays the role of a suitable translation between the states of the two systems, allowing for the application of (SYM*). It is important to note, however, that the Legendre transform will only be a statespace isomorphism if the two systems are *hyperregular*.[14] We only claim that this is a sufficient condition for comparing two models using (SYM*). In other cases, weaker mappings might serve as a suitable translation. This question requires further study, however it is clear that any comparison requires some translation scheme, a point not made clear by North's explanation.

Thus there is a precise sense in which for certain well-behaved systems, the passage from a Lagrangian description to a Hamiltonian one involves expanding a group of symmetries. Following North, we could take this as evidence for a corresponding reduction in the amount of postulated structure. But are these the *right* symmetries? Recall the original motivation for (SYM*): symmetries are identified with coordinate transformations (changes in description) of a given theoretical model. The structure of the world, according to the theory, corresponds to the symmetry-invariant structure of the model. When we turn our attentions to particular models of Hamiltonian and Lagrangian mechanics, we find that canonical and point transformations do preserve important physical structures (e.g. the equations of motion), but we also find exceptions. There are models of Hamiltonian mechanics with canonical transformations that do not represent mere changes in description, but rather map physically distinct solutions of the equations of motion onto each other. In Lagrangian mechanics, there are many models with additional symmetries that cannot be classified as point transformations. These facts dramatically call into question North's assumption that point and canonical transformations accurately characterize the structural commitments of each theory.

We illustrate this point with two examples. The first comes from the classical Kepler problem. Consider two bodies in 3-dimensional space interacting via a central force that is proportional to the inverse square of the distance between them (e.g. Newtonian Gravity). The Hamiltonian equations of motion governing the time evolution of the system are invariant under spatial translations, spatial rotations, and time translations, leading to conservation of momentum, angular momentum, and energy via Noether's Theorem. Interestingly, there is a fourth, somewhat mysterious conserved quantity called the Laplace-Runge-Lenz vector.[15] In the system's statespace (which has

[14] see Abraham and Marsden, pp. 218-223.

[15] The expression for the LRL vector is given by:

$$A \equiv p \times L - mk\frac{r}{|r|}, \tag{1}$$

where p is the system's total momentum, L the total angular momentum, m the reduced mass, r the position vector

6-dimensions), the canonical transformation that generates the LRL vector continuously maps solutions of the equations of motion representing orbits with the same total energy but different eccentricity onto one another. Such solutions correspond to great circles constrained to lie on a 4-dimensional sphere in statespace; the LRL transformation rotates these circles around the sphere. While these solutions are, in a sense, geometrically isomorphic in statespace, they project down onto physically inequivalent orbits in spacetime. To see this, note that the eccentricity vector measures how elliptical a particular orbit is. The orbits linked by the LRL symmetry have the same energy (hence the same semi-major axis), but different eccentricities, and thus have different intrinsic geometries in spacetime (as well as different angular momenta). If we retain a realist view towards spacetime or spatio-temporal relations, then the LRL transformation cannot be symmetry of the full Hamiltonian model of the 2-body system even though it is a canonical transformation.[16]

The second example is one that we have already encountered. Consider any Lagrangian system whose equations of motion fit the form of Newton's second law. These equations of motion are invariant under Galilean boosts. Although such boosts are widely viewed as symmetries of classical spacetime, they are not point transformations. By definition point transformations do not alter the velocity coordinates on TQ, therefore any velocity-dependent spacetime symmetry cannot be represented as a point transformation in the Lagrangian framework. This suggests that in a wide range of interesting cases, point transformations do not exhaust the symmetries of Lagrangian systems. Thus neither canonical transformations nor point transformations appear to accurately represent the symmetries of Hamiltonian and Lagrangian systems — the former are too permissive, the latter too restrictive.

A skeptical reader will have noticed a potentially troubling circularity in the style of reasoning that has been employed thus far: the symmetries of a physical model determine what its intrinsic structure is, but it is precisely this structure that determines which transformations count as symmetries in the first place. This is not a vicious circle, but it does mean that care must be exercised in order to fruitfully employ (SYM*) and other symmetry criteria. In some instances we might have an antecedent understanding of the physically significant structure in both models. The corresponding symmetry groups then serve as a convenient means of abstractly characterizing and comparing this physical structure. In other cases, we might have an independent procedure for determining a model's symmetries, allowing us to settle open interpretational questions about the model's description of the world. Of course, the majority of cases will likely lie between these two poles, requiring the interpreter to bring general theoretical constraints and guiding symmetry principles into reflective equilibrium with physical intuition.

North's arguments eschew analysis of individual models entirely. She believes that in general we do have an independent procedure for determining the symmetries of a theory. The method she invokes exploits symmetries of the theory's *laws*, and has been widely applied in debates about the

between the two bodies, and k a force-specific constant.

[16]There are additional examples of canonical transformations that are not full symmetries of Hamiltonian models, although, since such transformations are system-specific and their behavior subtle, they have received little attention among philosophers and even physicists. Another interesting example is given by symmetry associated with the quadrupole moment of the harmonic oscillator. In "Symmetry and Equivalence" forthcoming in *The Oxford Handbook of Philosophy of Physics*, Gordon Belot raises a number of similar concerns regarding symmetry arguments in classical physics including a discussion of the LRL vector and the harmonic oscillator.

structure of spacetime. One begins by looking for transformations of the independent and dependent variables of the theory that leave the mathematical form of the fundamental laws unchanged. Such transformations map dynamical solutions of these equations onto other dynamical solutions which are nomologically indistinguishable. Ontological parsimony dictates that, ceteris paribus, one only has reason to infer the minimum amount of physical structure required to formulate the fundamental dynamical laws. Hence, symmetry-linked solutions are interpreted as merely different descriptions of a single solution. In other words, symmetries in the laws (mappings of solutions of dynamical equations onto solutions) are taken to indicate symmetries in the complete theoretical model (mappings of the model onto itself).

North suggests that we take this same procedure and use it to help interpret the structure of statespace. She asserts that "the Lagrangian equations of motion are invariant under the set of point transformations; the Hamiltonian, under the canonical transformations" (pp. 87–88), and therefore, these transformations indicate "which coordinate changes the [theories respectively take] to be mere arbitrary changes in description" (p. 73). But as we have seen, this inference leads to trouble in specific cases. There are a number of possible reasons why.

First, the antecedent of North's claim, while true, is misleading; canonical and point transformations do not perfectly classify the symmetries of Hamiltonian and Lagrangian equations of motion. There are symmetries of Hamilton's equations that are not canonical, and there are even more symmetries of the Euler-Lagrange equations that are not point transformations. Interestingly, many of these additional Lagrangian symmetries correspond to canonical transformations when the models are translated into the Hamiltonian framework. For example, Galilean boosts always preserve the symplectic form on Hamiltonian statespace (for Newtonian systems), and are thus examples of canonical transformations. Similarly, the canonical LRL transformation can be represented in the Lagrangian framework as a particular velocity-dependent transformation that preserves the Euler-Lagrange equations for the Kepler problem. Even if North is right, and Galilean boosts and LRL transformations are symmetries of certain Hamiltonian models, but not of their Lagrangian counterparts, it cannot be for the reason North claims. The equations of motion are invariant under such transformations in both theories.

Second, it is not clear that the statespace equations of motion always play the role of truly *fundamental* laws. In the spacetime cases cited by North as exemplars, interpreters look for symmetries of Newton's second law, its special relativistic cousin, or the Einstein field equations. These laws have a wide modal scope, applying to a range of different physically possible worlds. As a result, symmetries of these laws are stable under changes to physically contingent aspects of the systems being described. This suggests that these nomological symmetries reflect facts about the fundamental structure of spacetime shared by models of the theory. In contrast, statespace equations of motion frequently apply only to a single, precisely defined system. As a result, symmetries in these equations can be sensitive to slight perturbations in the system. For example, in the Kepler problem, if a third interacting body is added or the central force otherwise deviates from its $1/r^2$ form, the LRL symmetry fails to hold. Assuming this perturbation is small, the LRL vector will continue to be approximately conserved; however the perfect symmetry is broken — solutions corresponding to orbits with different eccentricity will be distinguishable solely based on their intrinsic shape in statespace. Because of this sensitivity, physicists frequently refer to the LRL transformation as an "accidental" symmetry, one that holds due to contingent facts about the arrangement and motion of matter in a particular physical system, not necessary facts about the fundamental kinematic structure of the classical world. An interpreter who uncritically projects symmetries of the statespace equations of motion onto the world runs the risk of misclassifying such accidental symmetries as true symmetries.

Third, Hamiltonian and Lagrangian models do not consist solely of a statespace with an accompanying dynamical vector field defined on it. There are also mappings from this structure into spacetime which, as illustrated by the Kepler 2-body example, can encode important physical information. In ignoring this fact, North fails to realize that symmetries of particular models usually must preserve certain relations between statespace and spacetime. This serves as a constraint on which symmetries of the equations of motion should be interpreted as symmetries of the full model, altering many of our initial judgments about which statespace structures are physically significant. For example, in many cases (including the Kepler problem) Hamiltonian statespace is not simply a symplectic manifold, but the cotangent bundle, T^*Q, of a configuration space manifold, Q. North interprets this structure as physically insignificant because it is not preserved under canonical transformations. But if we have reasons to doubt that canonical transformations are always symmetries of Hamiltonian models, then we cannot simply dismiss this structure as North does.

Part of the difficulty in making broad structural comparisons the way North does, is that neither Hamiltonian nor Lagrangian mechanics purports to be a complete description of the world. A complete physical model is generated only once particular forces and dynamical laws are specified, along with how the statespace description of these structures maps onto spacetime. By focusing her attention on the two theories abstractly, and not on particular, fully detailed models, North's analysis only captures the features of a generic model, the structure that every model of Hamiltonian or Lagrangian mechanics has in common. But most models have more structure, and much of this structure appears, on a natural interpretation, to be physically important. The interpretive question which motivates the structuralist razor is not, "which of these two theories posits less *common* structure between its models," but rather "which of these two theories posits less *total* structure on the world." It doesn't help to know that the generic Hamiltonian dynamics only require statespace to be a symplectic manifold if additional structure must be posited on a model-by-model basis in order to illuminate how this statespace description relates to structures in spacetime.[17]

[17] One might resist the pull of these considerations by adopting a more radical form of statespace realism, such as the *statespace monist* view espoused by David Albert in his paper "Elementary Quantum Metaphysics" (in *Bohmian Mechanics and Quantum Theory: An Appraisal*, ed by A. Fine et al., (Kluwer, 1996): pp. 277-284). According to the view, statespace itself is taken to be the fundamental arena in which physics unfolds. A single world-particle moves through this space governed by a Hamiltonian. All additional physical facts supervene on the position of the world-particle in statespace. For instance, facts about the number of ordinary particles and their spatio-temporal positions will only obtain if the Hamiltonian contains suitably detailed "interaction" terms.

Whether or not this is a viable interpretation of classical mechanics is an interesting question that needs to be explored. Seeking a solution to the measurement problem, Albert originally proposed the view in order to explicate the ontology of quantum mechanics. The same motivating factors do not apply in the classical case. Furthermore, there are a number of interpretational oddities that plague the classical version of the view that do not arise in the quantum case. These questions aside, North's arguments were not supposed to presuppose any substantive form of statespace realism. Moreover, adopting such a radical view might undermine the reasons for adhering to a form of statespace realism in the first place. Her arguments for statespace realism crucially rely on analogies with spacetime. Summarizing them, she

Finally, it should be noted that there is an additional source of flexibility hidden in the Lagrangian framework that has been overlooked thus far. There exist symmetries of the equations of motion which are generated by subtle transformations of the Lagrangian function itself, and which (in general) cannot be represented as simple changes of coordinates on statespace. Any Lagrangian, L, can be replaced by $L' = L + \hat{\theta}$, where $\hat{\theta}$ represents the function on TQ defined by an arbitrary closed 1-form, θ, on Q.[18] The resulting variational problem produces the same dynamical vector field on TQ. (In contrast, the Hamiltonian is fixed up to an additive constant.) How these extra Lagrangian symmetries should be interpreted is an intriguing open question, especially in light of the considerations raised above. Nonetheless, their existence points to yet another way in which there is greater descriptive freedom in the Lagrangian framework than North presupposes.[19]

Our arguments do not preclude the possibility that a more refined set of symmetry considerations might allow for the implementation of (SYM*), although they do show that one must tread carefully in this domain. The successful application of symmetry arguments will hinge on how we interpret particular models of Hamiltonian and Lagrangian mechanics. This project turns on a number of interesting subtleties, including those discussed here. At present, there does not appear to be an exhaustive, independent criterion for identifying the symmetries of classical Hamiltonian and Lagrangian systems.[20] As such, we must exercise caution in using general symmetry arguments to support structural Ockhamist claims.

Acknowledgments: We thank Gordon Belot, Jeremy Butterfield, Erik Curiel, Shamik Dasgupta, Joe Rachiele, Jim Weatherall, and the participants at the 2011 UC Irvine Philosophy of Physics workshop for helpful comments and discussion.

---

concludes "the claim is just this: that phase space is as much a part of the representational content of classical mechanics as the theory's spacetime is" (p. 29).

[18] See de Leon and Rodrigues, p. 395. A 1-form on configuration space defines a function on the tangent bundle in a natural way: $\hat{\theta}(q, \dot{q}) \equiv \theta_q(\dot{q})$. That is, the value of the function $\hat{\theta}$ at point $(q, \dot{q}) \in TQ$, is given by the contravection of the 1-form with the tangent vector $\dot{q}$ at $q \in Q$.

[19] We thank Erik Curiel for bringing this point to our attention.

[20] An opinion shared by Belot in his forthcoming work, op. cit.