Lecture 2 Handout: Classical Mechanics II

Bryan W. Roberts

1. Lagrangian mechanics

The Force Mechanics picture of **particles pushed around by forces in the void** might be called the "Ruđer Bošković" picture in honour of the Croatian polymath who advocated it. The picture **faces some deep foundational problems** that troubled natural philosophers throughout the 17th and 18th centuries. For example, it appears to entirely ignore how such particles interact with each other! Is it through some strange "action at a distance"? And, how do we deal with continuum quantities, their boundaries, and how *they* interact?

The Lagrangian approach to mechanics **addresses (or at least brackets) some of these problems**. It extends the reach of what kinds of systems mechanics can discover, while at the same time restricting certain pathological examples in force mechanics. It also introduces several methodological techniques that are of independent interest in both physics and philosophy. The *locus classicus* for these lessons is Butterfield (2004).

In the 18th century Lagrange realised that much of classical mechanics could be derived from a simple principle, that there is a special quantity called "action" associated with the trajectory of a physical system, and that the only possible trajectories are those along which the action is "stationary". We shall build up all these ideas systematically.

1.1. Families of curves in configuration space.

Lagrangian Mechanics takes place on a slightly different "space" than Force Mechanics. We will similarly posit a manifold M with a (Riemannian) metric, called the **configuration space**, and take the state of a system to be represented by a point on that manifold. But we will now view that state more broadly, to represent the state of the entire system as a whole. This means that, to describe a single particle in Euclidean space, we use $M = \mathbb{R}^3$, whereas for two particles we use $M = \mathbb{R}^3 \times \mathbb{R}^3$. The configuration can also be interpreted quite broadly to represent other quantities, such as the angles of a double pendulum, or many more abstract things. It can even capture the configuration of continuous fields, in which case M will be infinite dimensional.

To explain the motion of a system, Lagrangian mechanics asks a slightly different question than we did in Force Mechanics. Instead of beginning with an initial position and velocity in space and asking what trajectory follows, we ask **what curve**

through configuration is travelled to go from one point to another in configuration space. In order to do this, we must first identify the set of all possible configurations of a system, the configuration space — a dramatically different starting point from Force mechanics. We will also need some powerful language for talking about a family of smooth curves between two points in configuration space, and its properties.

We begin by assuming that the curve that a particle will actually travel is one member of a smooth **one-parameter family of curves**, indexed by a parameter $\lambda \in \mathbb{R}$ (Figure 1). Let $\mathbf{x}_\lambda(t)$ denote an arbitrary member of the family of curves. Let us also assume that the function $\lambda \mapsto \mathbf{x}_\lambda(t)$ is smooth.

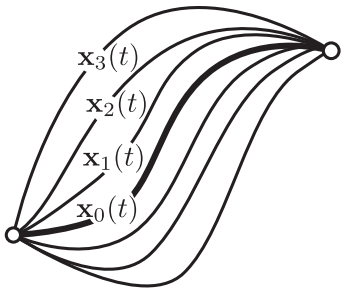

FIGURE 1. A one parameter family of curves between two endpoints.

Let $\mathcal{C}$ be a set of curves with compact support[1]; this is our way of expressing a sense in which each curve has finite length. A *functional* on curves is a function $F : \mathcal{C} \to \mathbb{R}$, which assigns a number to each curve. Because functionals are actually functions of functions, they may seem abstract at first. But **you are already familiar with concrete examples of functionals**. If you think of smooth curves each associated with a curved pieces of string, then the *length* of the string is a functional. This length functional $\Lambda : \mathcal{C} \to [0, \infty]$ assigns a length $\Lambda(\gamma) \in \mathbb{R}$ to each curve γ.

Now consider an arbitrary functional $\Phi : \mathcal{C} \to \mathbb{R}$. Whenever we are considering a family of curves $\mathbf{x}_\lambda(t)$, this family will give rise to a family of real numbers in the range of Φ, namely $\Phi(\mathbf{x}_\lambda(t))$. This language allows us to ask questions like which curve in the family $\mathbf{x}_\lambda(t)$ has the shortest length, i.e. the smallest value of $\Lambda(\mathbf{x}_\lambda(t))$. We can also talk about quantities such as the rate of change $d\Phi(\mathbf{x}_\lambda(t))/d\lambda$ of a function with respect to the parameter λ in a family of curves. This tells us about how quickly the value of Φ is changing from one curve to the next. Given a one-parameter family of curves $\mathbf{x}_\lambda(t)$, a function $\Phi : \mathcal{C} \to \mathbb{R}$ of the set of curves on a manifold is said to

[1]A subset $S \subseteq M$ is called *compact* if every open cover of M has a finite subcover; the compact sets on $\mathbb{R}^3$ are the closed and bounded regions. A curve $\gamma : I \to M$ has compact support iff the set $\{\gamma(t) \mid t \in I\}$ is empty outside of a compact subset of M.

be *stationary* or *extremal*[2] at $\mathbf{x}_{\lambda_0}(t)$ if and only if the the value of Φ is constant at $\mathbf{x}_{\lambda_0}(t)$, in that,

$$\frac{d}{d\lambda}\Phi\left(\mathbf{x}_\lambda(t)\right)\Big|_{\lambda=\lambda_0} = 0. \tag{1}$$

As a shorthand, this condition is sometimes written $\delta\Phi = 0$. Notice that being stationary is not a property of any individual curve $\mathbf{x}(t)$. It is property of a curve with a family of curves, which describes how a particular curve $\mathbf{x}_{\lambda_0}(t)$ is changing with respect to the entire family. **This will be the essential property** that we will shortly use to determine how a system changes.

1.2. **Action, Lagrangians, Lagrangian systems.** Lagrangian mechanics uses the concept of a stationary functional introduced above to determine what the path of a system will be. It does this using the concepts of action and of a Lagrangian.

An *action functional* for a manifold is any functional $\Phi : \mathcal{C} \to \mathbb{R}$ on the set of curves for which there exists a smooth function L such that,

$$\Phi(\mathbf{x}(t)) = \int_{t_1}^{t_2} L\left(\mathbf{x}(t), \dot{\mathbf{x}}(t)\right) dt$$

where $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ are endpoints of a curve $\mathbf{x}(t)$. The function L is then called the *Lagrangian* associated with Φ.

Why are we concerned with this particular kind of functional? One significant reason is that it encodes the empirical assumption that **action is defined by "local" facts about a curve**. Written as an integral in this way, it consists in "adding up" the values of some quantity (the Lagrangian) that is locally defined on the basis of the configuration and velocity variables. In this sense the Lagrangian L is "local": it depends on facts at a point. Thus, although the action is a global property of a curve as a whole, the fact that we express it as the integral of a Lagrangian means it is still the sum of its locally-defined parts.

Another observation is the **space on which the Lagrangian is defined**. The variable $\mathbf{x}$ is a point in the configuration space M. But the velocity $\dot{\mathbf{x}}$ is a vector at that point. The set of all vectors at a point $\mathbf{x}$ is called the *tangent space* $T_{\mathbf{x}}M$ at $\mathbf{x}$. To collect together all the vectors at all the points, we take the disjoint union of all the tangent spaces of M, which is called the *tangent bundle* TM. It is itself a manifold. Formally speaking, the Lagrangian L is a smooth function on the tangent bundle, $L : TM \to \mathbb{R}$.

Just as there was an important, very common class of force in Newtonian force mechanics, so there is an important and common class of Lagrangian. Suppose that we are describing a particle in space, and take the configuration space M to be the

[2]Many texts also use the phrase extremal as a synonym for stationary, and write that a curve $\mathbf{x}_{\lambda_0}(t)$ extremises Φ with respect to the family $\mathbf{x}_\lambda(t)$. This is because a zero derivative on a curve means that Φ achieves a (local) maximum or minimum, i.e. an extremum.

Euclidean manifold $\mathbb{R}^3$. A Lagrangian on such a space is said to have *standard-form* if it can be written as,

$$L(\mathbf{x}, \dot{\mathbf{x}}) = \tfrac{m}{2}\dot{\mathbf{x}} \cdot \dot{\mathbf{x}} - U(\mathbf{x}),$$

where $m > 0$ is a real number and $U(\mathbf{x})$ is a function that depends only on position[3]. The most common Lagrangian mechanical descriptions in classical mechanics have this form. Note that we are also using a dot-product; this is thus a place where the fact that M is equipped with a Riemannian metric (often assumed to be Euclidean) is needed.

1.3. **Stationary action and the Euler-Lagrange equations.** Dynamical motion in a Lagrangian system is underpinned by the following simple principle. Let M be a manifold, let $L : TM \to \mathbb{R}$ be a Lagrangian, and let Φ be the associated action functional. Hamilton's *principle of stationary action*[4] says that the possible motion of a system from one point in M to another is given by a curve $\mathbf{x}_{\lambda_0}(t)$ for which **the action is stationary**, $\delta\Phi = 0$, i.e. Equation (1). Such curves are sometimes said to *extremise* action. Curves with this property are the possible trajectories in Lagrangian mechanics.

To make practical calculation easier, it turns out that a curve extremises action if and only if it satisfies the *Euler-Lagrange equations*, in that,

$$\frac{\partial L(\mathbf{x}, \dot{\mathbf{x}})}{\partial x_i} = \frac{d}{dt}\frac{\partial L(\mathbf{x}, \dot{\mathbf{x}})}{\partial \dot{x}_i}$$

for each $i = 1, 2, 3, ...$ (up to n for an n-dimensional manifold). The Euler-Lagrange equations are the "differential form" of the principle of least action[5]. In practice it is usually easier to find solutions to these, as opposed to seeking direct solutions to the principle of stationary action.

Why do Lagrangians generally take this standard form? They don't always — but one can still seek to explain why so many of them do. One approach to this is to use *Jauch's theorem*, which shows a sense in which Galilei invariance guarantees that the Hamiltonian a standard form. One can then use the relationship between Lagrangian and Hamiltonian mechanics (given by the Legendre transformation) to argue for the standard form of the Lagrangian on this same basis.

1.4. **Why position and velocity?** Most differential equations representing fundamental quantities in physics or either first order, like the Schrödinger equation and

[3]The first term represents the kinetic energy of motion of a single particle. When considering multiple particles, this term is replaced by the sum of each such energy for each particle.

[4]This is sometimes referred to as Hamilton's principle after William Rowan Hamilton's essentially modern application of it in 1834. However, a similar version was stated in the 18th century by Lagrange, and there are other proto-versions that date back to the 17th century.

[5]Proving the equivalence of these statements is a good exercise. There is only one non-trivial part of the proof that you may assume, which is called the *fundamental lemma of the calculus of variations*; the rest is standard multivariate calculus. The proof may be found in (Arnold; 1989, §12) or (Goldstein et al.; 2002, §2.2).

Hamilton's equations, or second order, like $\mathbf{F} = m\ddot{\mathbf{x}}$ and Lagrange's equations, when the Lagrangian depends only on position and velocity. That's curious: if writing action in terms of a Lagrangian just captures the fact that physics is local, then one might expect L to depend on higher order derivatives as well. So, **why do we only use position and velocity?** And more generally, **Why are the differential equations in physics so low order?**

Some have argued that the answer lies in a complex metaphysical story about the nature of forces and accelerations of point particles (Easwaran; 2013). This account has all the problems associated with the Force-particle picture of reality, and more. I will not go into it here, but you are encouraged to read the debate for yourself.

Alternatively, Swanson (2019) pointed out that one can view low-order differential equations as a consequence of energy being bounded from below. Consider a local coordinate system $(x_1, \ldots, x_n)$ for some region of a manifold M of dimension n. This gives rise to an "induced" local coordinate system on the tangent bundle TM, denoted by $(x_1, \ldots, x_n, \dot{x}_1, \ldots, \dot{x}_n)$. The *local energy function* h associated with a Lagrangian $L : TM \to \mathbb{R}$ is then defined by,

$$h := \sum_{i=1}^{n} \dot{x}_i \frac{\partial L}{\partial \dot{x}_i} - L. \tag{2}$$

In most uses of Lagrangian mechanics to represent a physical system, this function corresponds to the system's local energy. For most physical systems, energy has a lower bound. This captures the empirical fact that **known physical systems have a (lowest-energy) ground state**. If they did not, then one could extract unbounded work from a system, and it would just keep dropping down into lower and lower energy states, which appears not to happen in our world. With this in mind, we can now state:

Fact 1 (Ostrogradski's Theorem)**.** *Let $L(x, \ldots, x^{(n)})$ be a non-degenerate[6] Lagrangian. If the associated local energy function h is bounded from below, then $L = L(x, \dot{x})$ is a function of x and $\dot{x}$.*

Thus, the fact that the Lagrangian has this particular functional dependence can be viewed as arising out of the fact that energy is bounded from below. There are of course some further subtleties to this argument; see Swanson (2019) for further discussion.

1.5. **Lagrangian Mechanics vs Force Mechanics.** The mathematical relationship between Lagrangian and Force mechanics is filled with subtleties, many of which can be found in Butterfield (2004), and most of which I will not go into here. Let me instead restrict attention to one special case:

[6]A Lagrangian is called *non-degenerate* or *regular* if the determinant of the Hessian is non-vanishing, $\det[\partial^2 L/\partial \dot{x}^2] \neq 0$. It is equivalent to the condition that the Legendre transformation for L is a local diffeomorphism from TM to T^*M (Marsden and Ratiu; 2010, §7.2).

(1) Every Lagrangian in standard form admits a Force analogue;
(2) *Not* every Force mechanical system admits a Lagrangian analogue in standard form; but,
(3) Every *conservative* Force mechanical system admits a Lagrangian analogue in standard form.

(1): It is geometrically intuitive that this is always possible. Lagrangian mechanics allows one to define a family of curves $\mathbf{x}(t)$ that "cover" the configuration space M. This in particular assigns a value to the second derivative of $\ddot{\mathbf{x}}$ of these curves at each point in M, by defining the force $\mathbf{F}$ to be proportional to $\ddot{\mathbf{x}}$, we arrive at a Newtonian system that predicts the same curves through M as our Lagrangian system does.

One can also see this in a more concrete way by considering a standard-form Lagrangian $L(\mathbf{x}, \dot{\mathbf{x}}) = \frac{m}{2}\dot{\mathbf{x}} \cdot \dot{\mathbf{x}} - U(\mathbf{x})$. Then the right-hand-side of the Euler-Lagrange equations becomes,

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{x}_i} = \frac{d}{dt}\left(\frac{m}{2}\frac{\partial \dot{x}_i^2}{\partial \dot{x}_i}\right) = \frac{d}{dt}(m\dot{x}_i) = m\ddot{x},$$

for each $i = 1, 2, 3, \ldots$. Meanwhile, the left-hand-side of the Euler-Lagrange equations becomes,

$$\frac{\partial L}{\partial x_i} = \frac{\partial}{\partial x_i} U(\mathbf{x}).$$

Thus, by defining $\mathbf{F} := -\nabla \cdot U = \nabla \cdot L$ we obtain a force satisfying $\mathbf{F} = m\ddot{\mathbf{x}}$ for each curve $\mathbf{x}(t)$ in M.

(2): The argument above runs both ways: if $\mathbf{F}$ is not conservative, i.e. not expressible as the gradient of a potential U, then a standard form Lagrangian cannot be used to capture it.

(3): Suppose the force is conservative, i.e. $\mathbf{F} = \nabla \cdot U(\mathbf{x})$. Then there exists a Lagrangian system that describes the same possible particle trajectories, in that a curve $\mathbf{x}(t)$ satisfies Newton's equation if and only if it also satisfies the Euler-Lagrange equations. Namely, define L to have its standard form with this U. Then the previous calculation shows that L satisfies the Euler-Lagrange equations. Note that we have also now shown that the conservative-force Newtonian systems correspond precisely to the standard-form Lagrangians and vice versa.

When a Newtonian force system is not conservative, it does not necessarily admit a description in terms of Lagrangian mechanics. These are known as **non-holonomic systems**. There are ways of dealing with them in Lagrangian mechanics, but they escape the scope of these notes; see Butterfield (2004).

2. Hamiltonian Mechanics

When you first learn Hamiltonian mechanics, it's often from an **extremely practical physics perspective:** when you're analysing a tricky system like the

motion of a double-pendulum, it's just easier to model with a different formalism. You begin by describing the kinematics of a system in terms of position $\mathbf{q} = (q_1, q_2, \ldots, q_n)$ and momentum $\mathbf{p} = (p_1, p_2, \ldots, p_n)$. Then you write down a system of differential equations called *Hamilton's equations*, which describe how those quantities change over time:

$$\frac{d}{dt}q_i(t) = \frac{\partial h(\mathbf{q}, \mathbf{p})}{\partial p_i}, \qquad \frac{d}{dt}p_i(t) = \frac{\partial h(\mathbf{q}, \mathbf{p})}{\partial p_i},$$

where h is a real-valued function called the *Hamiltonian*, usually of the form,

$$h(\mathbf{q}, \mathbf{p}) = \sum_{i=1}^{n} \tfrac{1}{2m_i} p_i^2 + V(q_1, \ldots, q_n),$$

for some set of real numbers $m_i > 0$ and some real-valued function V of the q_i variables, which is related to the forces in the system by the relation, $\mathbf{F} = \nabla \cdot V$.

Work by 20th century mathematicians brought a significantly richer perspective on **what this framework is, why it works, and what it says about the world**.

2.1. **Phase space.** Hamiltonian mechanics begins by writing down a space representing **all possible states** that a physical system can occupy in a particular structured way. We assume this space of states has the structure of a smooth manifold M, with each point $p \in M$ representing one possible state of the world. That we assume this space is a manifold encodes the facts that,

- *(coordinates)* The state of the system can be described using a (possibly infinite-dimensional) local coordinate system; and
- *(smoothness)* It is possible to pass smoothly from one state to another.

We next assume that the space of physical states has a **symmetry structure** that is captured by a mathematical object called a 'symplectic form'. A *symplectic form* is a 2-form on M, denoted ω; that is, ω is a bilinear mapping from pairs of vectors on M to $\mathbb{R}$ which skew-symmetric (i.e. $\omega(X, Y) = -\omega(Y, X)$ for vector fields X and Y). It is also closed ($d\omega = \mathbf{0}$) and non-degenerate (i.e. $\omega(X, Y) = 0$ for all Y only if $X = \mathbf{0}$). The pair (M, ω) is called a *symplectic manifold*. When it used to represent mechanical systems, it is often also called a *phase space*.

Why do we introduce some tricky gadget as a symplectic form for representing physical states? One reason is that it allows us to **build in an assumption about symmetries** that is deep in the practice of physics. That assumption, in rough form, is the following:

> (symmetry-observable correspondence) *Every observable admits a continuous group of symmetries along which it is conserved, and vice versa: every continuous group of symmetries generates an observable that it conserves.*

This statement expresses a Noether-like idea. It is **built into the fabric of Hamiltonian mechanics** through the use of the symplectic form.

To make the statement above meaningful, we need to make precise **what we mean by observables and by symmetries**. Suppose we identify observables with a smooth assignment of a real-number values to each state of the world. That is, an observable is a smooth function $f : M \to \mathbb{R}$. This is standard practice: such a function can be used in an obvious way to represent a physical quantity (like total energy, angular momentum, etc.) associated with each possible state of the system.

We take a continuous "symmetry" to be represented by a diffeomorphism $t \mapsto \varphi$ that preserves the symplectic form, in that $\varphi_t^* \omega = \varphi$ for all $t \in \mathbb{R}$, where φ^* is the pushforward of φ. Such a map is called a *symplectomorphism*, and is the basic concept of isomorphism for a symplectic manifold. This allows us to express our Noether-like assumption above in precise form.

Fact 2. *Let (M, ω) be a symplectic manifold. Then for every smooth function h, there exists a smooth vector field X saisfying,*[7]

$$dh = \iota_X \Omega. \tag{3}$$

When this is true, the family of integral curves $t \mapsto \varphi_t$ are a group of symplectomorphisms (symmetries) that preserve h, i.e. $h \circ \varphi_t(p) = h(p)$ for all t. Conversely, for every smooth 1-parameter group of symplectomorphisms $t \mapsto \varphi_t$ there exists a smooth function h (unique up to addition by a constant) for which Equation 3 is true.

The vector field X associated with a smooth function in this way is sometimes called the *symplectic flow* associated with h, and h is called *Hamiltonian generator* of the vector field X.

2.2. **Hamiltonian dynamics.** Although it might not look like it, this statement is actually the familiar Hamilton's equations in disguise. Seeing this requires two observations.

First, note that most fundamental physical systems are assumed to admit a continuous **time-translation symmetry**. This captures the assumption that time is continuous, and that experiments can be repeated day to day: apply the same conditions tomorrow, and you will get the same experimental results. In the context of Hamiltonian mechanics, this means that we can view **the passage of time as a 1-parameter group of symplectomorphisms**. The associated vector field is usually assumed to be **complete**, to capture the assumption that time can be parametrised

[7]Here, d is the exterior derivative for M, and $\iota_X \Omega$ denotes the contraction of the vector field X with the first index of Ω. This latter statement in Penrose notation is expressed $d_b h = X^a \Omega_{ab}$, or equivalently, $X^a := \Omega^{ab} d_b h$, where Ω^{ab} is the inverse of Ω_{ab} (i.e. $\Omega^{ab}\Omega_{bc} = \delta^a_c$). These definitions are discussed in more detail in any differential geometry textbook, although the notation (sadly!) varies dramatically from book to book.

by the entire real line. By the discussion above, this 1-parameter group has a generator; when the group is interpreted as time-translation, the generator is called the *Hamiltonian* or the *Hamiltonian function* $h : M \to \mathbb{R}$ for the physical system.

Second, it turns out that if (M, ω) is a symplectic manifold, then M has **even dimension** $2n$. And, a basic theorem of symplectic mechanics makes it is possible to define a special local coordinate system around every point $p \in M$, using a binary operator $\wedge$ on forms called the *wedge product*.[8]

Fact 3 (Darboux's Theorem)**.** *Let* (M, ω) *be a symplectic manifold of dimension* $2n$*. Every point* $x \in M$ *admits a neighbourhood* U *with a local coordinate chart* $(q_1, q_2, \ldots, q_n, p_1, p_2, \ldots, p_n)$ *centred at* x*, called a* local Darboux coordinate system*, such that,*

$$\omega = (dq_1 \wedge dp_1) + (dq_2 \wedge dp_2) + \cdots + (dq_n \wedge dp_n).$$

In this local Darboux coordinates, the statement of Equation (3) about time translation symmetry is then **equivalent to the usual expression of Hamilton's equations given at the outset of this section**, with t interpreted as time.

Like Lagrangian and Force mechanics, Hamiltonian mechanics has a standard form associated with the Hamiltonian h. In coordinates interpreted as canonical position q and momentum p, that form is,

$$h(q, p) = \tfrac{1}{2m}p^2 + v(q)$$

for some real-valued function v.

There is also a "standard form" of the manifold itself in Hamiltonian mechanics, which is the case that there exists a manifold Q such that $M = T^*Q$, where T^*Q is the *cotangent bundle* over Q. Note that this is *not necessarily the case*. It is only true given an empirical assumption, that it is possible to globally separate the canonical position and momentum variables associated with M so that it can be written as a cotangent bundle. However, it is quite standard to use Hamiltonian mechanics in situations where this is the case. And, when it is the case, there exists an isomorphism (called the *Legendre transformation*) which transforms between Lagrangian and Hamiltonian mechanics. It is thus thus exactly in this case that the two are empirically equivalent.

3. Some philosophical questions

- Compare how the frameworks of Force Mechanics and Lagrangian Mechanics expand or restrict what is physically possible.
- In what sense (if any) does nature always respect a stationary action principle?
- In what sense (if any) should action be expressible in terms of a Lagrangian, and one that depends only on position and velocity?

[8]For simplicity, I'll define the wedge product here only for the relevant case of two 1-forms *alpha* and β: then their *wedge product* is the 2-form defined by $(\alpha \wedge \beta)(X, Y) := \alpha(X)\beta(Y) - \alpha(Y)\beta(X)$. For the general definition, see Marsden and Ratiu (2010, §4.2).

- How can the standard form of the Lagrangian (or the Hamiltonian) be motivated?

References

Arnold, V. I. (1989). *Mathematical methods of classical mechanics*, 2nd edn, Springer-Verlag New York, Inc.

Butterfield, J. (2004). Between laws and models: Some philosophical morals of Lagrangian mechanics. `http://philsci-archive.pitt.edu/1937/`.

Easwaran, K. (2013). Why physics uses second derivatives, *The British Journal for the Philosophy of Science* **65**(4): 845–862.

Goldstein, H., Poole, C. and Safko, J. (2002). *Classical Mechanics*, 3rd edn, Addison Wesley.

Marsden, J. and Ratiu, T. (2010). *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems*, 2nd edn, Springer-Verlag New York, Inc.

Swanson, N. (2019). On the Ostrogradski Instability, or, Why Physics Really Uses Second Derivatives, *The British Journal for the Philosophy of Science* . `http://philsci-archive.pitt.edu/15932/`.

Email address: `b.w.roberts@lse.ac.uk`