

On Classical Cloning and No-cloning

Nicholas J. Teh

Abstract

It is part of information theory folklore that, while quantum theory prohibits the generic (or universal) cloning of states, such cloning is allowed by classical information theory. Indeed, many take the phenomenon of no-cloning to be one of the features that distinguishes quantum mechanics from classical mechanics. In this paper, we argue that *pace* conventional wisdom, in the case where one does *not* include a machine system, there is an analog of the no-cloning theorem for classical systems. However, upon adjoining a non-trivial machine system (or ancilla) one finds that, *pace* the quantum case, the obstruction to cloning disappears for pure states. We begin by discussing some conceptual points and category-theoretic generalities having to do with cloning, and proceed to discuss no-cloning in both the case of (non-statistical) classical mechanics and classical statistical mechanics.

Keywords: philosophy of information theory, cloning, classical-quantum interface, foundations of classical mechanics

Contents

1	Introduction	2
2	Cloning	6
3	Category-theoretic framework	7
4	Categories for classical mechanics	11
4.1	Set and Rel	12
4.2	Lagrangian systems	13
4.3	Symplectic systems	14
4.4	Poisson systems	16
5	No-cloning in classical mechanics	17
5.1	Self-replication	18
5.2	Cloning	23
5.3	Uniform no-cloning and categorical collapse	26
5.4	Quantum cloning and rigidity	28
5.5	Implications for the CBH theorem	31

6 No-cloning in classical statistical mechanics	32
7 Future directions of research	36
Appendix A Cloning in non-interacting classical statistical mechanics	37

1. Introduction

Although the subject of cloning (and no-cloning arguments) is a ubiquitous theme in the philosophy and foundations of quantum theory, it is noteworthy that in the case of classical Hamiltonian mechanics, the subject has received almost no attention.¹ Perhaps this is due to the overwhelming foundational (and pragmatic?) interest in *quantum* information theory, or perhaps it is due to a complacent acceptance of the conventional wisdom that ‘cloning is trivially possible in classical physics, and thus uninteresting’.² Either way, it is not our intent here to speculate on the sociological reasons for this neglect, but instead to redress it by giving a conceptual overview of the subject and proving some elementary but foundational results therein.

We shall thus begin from scratch, i.e. by first reviewing some basic material in classical (symplectic) mechanics, and then building up to our main conclusions. However, our main result is simple (if surprising) and can be stated without further ado:

On one natural construal of ‘classical cloning’, viz. the straightforward analog of ‘quantum cloning’ *without* a machine system, conventional wisdom is actually *wrong* (in both statistical and non-statistical classical mechanics). On the other hand, if one requires that cloning is implemented by a machine system with sufficiently many degrees of freedom,³ then cloning is permitted in the non-statistical case, but not in the statistical case.

An important corollary of this result, which is relevant to recent philosophical discussion (e.g. [40], [24], and references therein), is the following:

The ‘no cloning’ information-theoretic axiom of the Clifton-Halvorson-Bub (CBH) characterization theorem [13] for quantum mechanics does *not* distinguish classical mechanics from quantum mechanics.

This claim is further qualified in §5.5.

It should be emphasized from the outset that we will here treat ‘cloning’ as a problem in dynamical systems, viz. given initial data at some initial time t_0 , can one find a suitable classical dynamical process

¹On the other hand, if by ‘classical’ one merely means ‘non-quantum’ then the subject has received some attention in the form of both toy theories and classical probability; see e.g. [26, 29, 38, 19]. We outline exactly what we mean by ‘classical physics’ below, which is largely in keeping with (the focal sense of) traditional usage.

²For an argument that addresses the practical possibility of classical cloning, see [44], which argues that such operations can be realized up to an arbitrarily high fidelity provided that the ‘resolution’ of the probability distributions is bounded.

³In a sense that we make precise in §5.

that clones this data at some final time t_1 ? This treatment is very much in keeping with the physical literature on quantum cloning ([35]), although that literature is often also concerned with a related albeit distinct issue: viz. what are the implications of prohibitions on cloning (or broadcasting) for the trade-off between ‘gaining information about a state’ and ‘disturbing that state’.⁴ This is not the place to explore either this trade-off or classical measurement theory, although we shall offer some brief remarks in §5 in order to assure the reader that our main result does little to undermine our faith in the reliability of Xerox devices and other such classical cloning machines.⁵

Let us also be more specific about what we mean by ‘classical physics’: we mean Hamiltonian mechanical systems, both statistical and non-statistical, i.e. those mechanical systems described by symplectic geometry and appropriate dynamical functions, with statistics given by a probability measure on the ‘phase space/symplectic manifold’. In particular, this excludes non-Hamiltonian Lagrangian systems, as well as general classical field theories (the latter mostly for ‘technological’ reasons, since it is not entirely clear how one should define a ‘symplectic manifold’ for field theories). This is not an enormous restriction, since a large class of well-behaved theories are both Lagrangian and Hamiltonian, and it may in any case be possible to extend our results to field theories. It will also be convenient to introduce a slight regimentation of our language at this stage: we shall always take classical mechanics to mean *non*-statistical classical mechanics. This will cause no confusion as we shall always include the adjective ‘statistical’ when we wish to refer to classical statistical mechanics.

Our discussion comprises three themes, which when taken together provide the narrative structure of this overview:

- Category-theoretic reasoning about no-cloning: In a series of papers on category theory and quantum mechanics [5] [6], Baez has advocated category theory as a perspicuous framework for reasoning about the basic structural features of physical theories, and in particular the issue of cloning. This is a laudable line of thought, especially since categories can in some sense be seen as laying bare the ‘grammar’ or ‘logical form’ of physical theories (in a sense that we shall explain later). Nonetheless, we wish to issue two caveats here: first, the general caveat that the relevant categories for physics are not ‘God-given’ and need to be carefully chosen if they are to capture the features of interest; and second, a more specific warning against an invalid inference that people sometimes make when they discuss (admittedly in a heuristic way) categories and cloning.

Our category-theoretic discussion will also serve as a convenient setting in which to provide the neces-

⁴Note that this trade-off is not about the ‘uncertainty principle’, although there has been much confusion in the physical literature about this issue. See e.g. [20] for a masterly discussion of the real issues underlying the ‘information gain v. state disturbance’ trade-off.

⁵Nonetheless, it is an interesting question whether the deep structure of classical mechanics (i.e. symplectic topology) has implications for both topics and we hope to explore this elsewhere.

sary background on symplectic mechanics, and to take note of the obstacles that arise when one tries to straightforwardly translate category-theoretic discussions of quantum cloning to the classical case. Furthermore, when one restricts classical mechanics to its linear form and formulates it in category-theoretic terms, one sees that the structure that prohibits ‘uniform cloning’ (i.e. a category’s being ‘compact closed’) is precisely the same as that identified by Abramsky [1] as prohibiting quantum cloning.

- No-cloning in classical (non-statistical) mechanics: Our results in this area concern both cloning and a closely related operation that we call ‘self-replication’. We shall prove a no-self-replication result for classical mechanics in a way that extends and complements an earlier example of Baez-Fenyes [18]. We shall argue that while Baez-Fenyes initially took themselves to have proven a ‘no-cloning’ result (however, see Fenyes [17] for a reversal of this position), this claim is misleading, on both a semantic, and more deeply, a conceptual level. As we shall see, once one understands cloning in an *operational* way, i.e. as being implemented by machine system, it is in general possible to clone states in classical mechanics.
- No-cloning in classical statistical mechanics: Current understanding of cloning in the statistical mechanics case is actually rather more advanced than in the non-statistical case: a no-cloning result for additive states is folklore among experts, and a rigorous and general no-cloning argument was proven in [14] (henceforth DPP). However, DPP’s proof seems to be unnecessarily circuitous and introduces superfluous machinery such as a ‘distance measure’ on the space of probability distributions. Here we connect the literature on non-statistical no-cloning to statistical no-cloning by showing how our techniques furnish us with a no-self-replication argument in the statistical case (the full no-cloning argument is similar, but rather more complicated). We furthermore discuss the relationship between the statistical case and the non-statistical case on general grounds.

Clearly, these themes and results are not exhaustive, but we hope that they provide philosophical clarification concerning the possibility of classical cloning. We now proceed to describe the detailed plan of the paper.

§2-3 discuss general concepts that apply to both classical and quantum physics. In §2, we explain what we take to be an appropriate definition of cloning, and also the closely related concept of self-replication. §3 discusses category theory as a general framework for thinking about self-replication. We review Baez’s conception of self-replication in the quantum case, as well as the standard arguments for why there is no self-replication map. We also caution against a certain invalid inference that is sometimes made when people discuss this issue informally, viz. the inference from ‘no Cartesian product’ to ‘no-cloning’. This section lays the ground for our attempt to find an analog of no-cloning for classical physics. However, for reasons that

we shall explain, we are unable in this paper to give a fully category-theoretic account of either cloning or replication in the classical case.⁶ Thus, our use of category theory in this paper is primarily as a tool to organize the relevant concepts.

In §4, we specialize to those classical categories that are relevant to the discussion of mechanics. We discuss the *naive* Lagrangian and symplectic (Hamiltonian) categories in order to emphasize that one’s choice of category is crucial for disclosing the structural aspects in which one is interested, and also so as to fix notation and the (minimal) technology that we will later need for symplectic mechanics. For the sake of completeness we also list the Poisson category; although all the results below translate straightforwardly from the symplectic category to the (non-degenerate) Poisson sub-category and so we will not re-prove them. (The Poisson geometric approach will also be useful for understanding how our work relates to the CBH [13] theorem.) It will emerge from our discussion that these naive classical categories are somewhat unsatisfactory for our purposes, since they do not offer the resources for reasoning about cloning in a purely category-theoretic way: this problem is in part remedied by the strategy adopted in §5.3.

§5 discusses the status of self-replication and cloning within classical mechanics. More specifically, in §5.1 we present the classical no-self-replication argument from several points of view, beginning with the tangent space geometry of a symplectic manifold, and proceeding to consider this result from progressively more global points of view, i.e. Lagrangian submanifolds and symplectic invariants. §5.2 discusses the case of classical cloning, and shows that there is in general no obstruction to this process, although there is a certain constraint, viz. a certain submatrix of the cloning map’s Jacobian must be orientation-reversing. We then explain how this constraint is related to the fact that cloning is impossible when the machine system’s state space has fewer dimensions than those of the source system, but not otherwise. The next two sections connect the above results to foundational work in *quantum* cloning: first, §5.3 shows that Abramsky’s abstract no-go theorem for quantum self-replication can also be applied to the case of classical self-replication; and second, §5.4 uses the geometrical formulation of quantum mechanics [36, 22, 8] to show that the difference between classical cloning and quantum no-cloning can be understood in terms of the ‘rigidity’ of their respective geometries. Finally, §5.5 discusses the relevance of our result to the CBH theorem.

§6 extends our results to the case of classical statistical mechanics. In particular, we show that the self-replication case of [14]’s result on no-cloning can be recovered by means of the simple argument that we have developed in §5. §7 concludes with a list of open problems and indicates directions for future research.

⁶With the exception of the linear case, which we discuss at the end of §5.

2. Cloning

We take cloning to be a process through which the information in a source state a is copied to a target (or ‘blank’) state b – thereby overwriting it with the information contained in a – by a cloning machine.⁷ That is to say, a cloning process involves three systems: a source system M_s , a target system M_t , and a machine system M_m . The source and target systems are required to be copies of the same state space, although a priori, there is no such requirement on the machine system. Later, in §5.2, we shall see that the size of the machine system is relevant to the issue of whether classical cloning is possible.

A state of the composite system can thus be written as $(x, y, z) \in M_s \odot M_t \odot M_m$, where \odot represents the relevant composition operation. For instance, in quantum theory on flat spacetime, the component systems are Hilbert spaces, and the composition operation is given by the tensor product.⁸ In the case of Hamiltonian classical mechanics, the component systems are symplectic manifolds. The composition operation is correspondingly more complicated than that of mere smooth manifolds (whose composition is given by the direct product): since symplectic manifolds also carry an algebraic structure (i.e. a symplectic form), we shall presently see that we need a notion of composition for this structure in addition to taking the direct product of the subsystem manifolds.

A cloning process can then be represented by the following map ϕ ,

$$\phi \equiv (\phi_1(x, y, z), \phi_2(x, y, z), \phi_3(x, y, z)) : M_s \odot M_t \odot M_m \rightarrow M_s \odot M_t \odot M_m \quad (1)$$

$$\phi : (x, b, z) \mapsto (x, x, z') \quad (2)$$

for which there exists an initial machine state z and an initial blank state b such that, for any initial source state x (‘any’ with a restriction of the quantifier if necessary!), the map overwrites the blank state b with x .

This conception of cloning as requiring a cloning machine makes sense on both the semantic and the operational level. First, the semantic level: notice that, in the ordinary sense of ‘cloning’, it would be incoherent to say that one had literally cloned oneself; to step into a cloning machine and thereby be cloned is not to literally clone oneself; it is to bring about one’s being cloned by the machine. But more important from the physical perspective is that cloning could not be realized except by some machine, and so we need to introduce a machine system and a machine state into the composite system if we are to discuss cloning in a physically realistic way.

⁷Indeed, this is the standard conception of cloning used by careful treatments, e.g. [35], although informally (and especially in the context of quantum cloning), this is often abbreviated to what we shall call ‘no-self-replication’ below. In this paper, we shall always take care to distinguish the two cases, although it will be useful for the reader to remember that both cases are often called ‘cloning’ in the literature.

⁸It is clear that a different operation, i.e. the Connes tensor product, will have to be used for QFT on more general manifolds. For an exposition of this point, see Segal’s unpublished MS [37], where he shows that the standard tensor product won’t work for Type III factors. See also [45], where the Connes tensor product is deployed in the context of conformal field theory.

We can contrast cloning with a closely related process we will call *self-replication*, which involves only M_s and M_t , and for which there is a self-replication map, i.e. there exists an initial blank state b such that for all initial source states x , the map takes an initial state $(x, b) \in M_s \odot M_t$ to a final state (x, x) . From the purely formal point of view, self-replication is a special case of cloning, viz. cloning in the case where M_m is trivial (i.e. a point). However, in what follows, we shall take cloning to involve a *non-trivial* machine system, and thus reflect a substantive difference between cloning and self-replication in our terminology (with the exception of §5.3, where we defer to the established usage of ‘uniform cloning’ to mean ‘uniform self-replication’).

In the case of quantum mechanics, it is possible to prove a no-self-replication theorem and show that this extends to a no-cloning theorem. We shall see that in the case of classical mechanics, one can also prove a no-self-replication theorem, but this does *not* in general extend to a no-cloning theorem. We also note that there is a variant of self-replication that is often used in category-theoretic discussions of cloning, viz. the diagonal map

$$\Delta : M_s \rightarrow M_s \odot M_t, \tag{3}$$

$$\Delta : x \mapsto (x, x). \tag{4}$$

Under ‘good conditions’ (i.e. if we are working in a monoidal category), this definition leads to the standard definition of self-replication given above. However, from the physical point of view, the standard definition is conceptually prior, as the cloning process should begin with a target system in some blank state instead of just creating the target system out of thin air.

Finally, we shall understand the question of whether cloning is ‘generically possible’ as the question of whether one can build a cloning machine that clones some arbitrary state x to a specific blank state b . In other words, our cloning machine should not be ‘state-dependent’. (Such cloning machines are often also called *universal* cloning machines; see e.g. [35].) When people speak of ‘no-cloning theorems’, they usually have in mind results which show that cloning is not *generically* possible in the above sense.⁹ After all, even in the case of quantum mechanics (i.e. the case in which we expect the strongest prohibition on cloning processes), it is clear that non-generic cloning processes are possible: a trivial example is furnished by the fact that for any initial and final states, there is a unitary map connecting them.

3. Category-theoretic framework

In [5] and [6], Baez introduces the idea that the quantum no-self-replication theorem can be traced to the structure of the category used to model quantum mechanics.¹⁰ (This insight was then refined by Abramsky

⁹The ‘genericity’ of a cloning machine should not be confused with another property that is usually assumed of cloning machines, viz. ‘perfection’ – the latter property is said to hold if the source state is cloned with perfect fidelity.

¹⁰See for instance p. 19 of [5] and p. 15 of [6].


in [1], which shows that any compact category with cloning ‘collapses’.) This is a very attractive and natural idea, as category theory has often been touted as a framework for capturing the ‘grammar’ or ‘logical form’ of a physical theory. However, as we shall see, one needs to be careful about exactly what physical inferences one is licensed to make from the presence or absence of certain structures in a category.

The outline of Baez’s argument in [6] is as follows:

- P1 Physical processes are to be identified with the morphisms of some relevant category.
- P2 A self-replication process is possible only if the relevant category has Cartesian products.¹¹
- P3 Quantum categories, i.e. categories appropriate for quantum theory, do not have Cartesian products (although they do have symmetric monoidal products).
- C Thus there is no self-replication process for a quantum category.

Let us consider the motivation for these premises in turn. First, P1. In order to make such an identification, we need to find a category such that (i) the objects represent physical systems, viz. as their space of states, and (ii) the points of an object (morphisms from the terminal element into the object) represent states of the system. We also need to be able to identify the morphisms of the category with processes that take the system from some initial state to some final state. Finally, we should be able to identify a physical composition operation with some mathematical notion of composition that exists in the category (e.g. pullback, product, symmetric monoidal functor, and so on). It is quite plausible that such categories can be found for a large class of physical systems, and at least in the case of quantum systems this has been shown by [5], [6], and [2]. On the other hand, we must also bear in mind that there are categories describing physical systems which do not have such an interpretation: for instance, neither of the ‘naive’ classical categories that we consider below can be interpreted in this way.

Next, P2. Let us first recall that a (Cartesian) *product* $A \times B$ of two objects A and B in a category is defined by the following diagram,



where π_1 and π_2 are projections onto A and B respectively. That is to say, $A \times B$ is a product iff for any object Q and morphisms $f : Q \rightarrow A$ and $g : Q \rightarrow B$, there is a unique morphism $(f, g) : Q \rightarrow A \times B$ which

¹¹For instance, Baez writes that ‘The Wootters-Zurek argument that ‘one cannot clone a quantum state’ is also based on the fact that the tensor product of Hilbert spaces is not cartesian.’ (p. 15, [6])

makes the diagram commute.

Now if a category has products, it is clear that we can define a unique diagonal morphism

$$\Delta : A \rightarrow A \times A \tag{5}$$

by setting $Q = B = A$ and f and g above to the identity morphism $1_A : A \rightarrow A$. The diagonal morphism Δ is a perfectly reasonable candidate for a ‘self-replication morphism’ in this context, and thus the converse of P2 is uncontroversial. However, there is a more dubious assumption made in P2, viz. the idea that *all* self-replication processes arise in this way from a product structure. As we will see shortly, it is possible to challenge this assumption: in particular, it is plausible that other conceptions of cloning are possible in classical mechanics. The existence of rival conceptions implies that while Baez’s categorical schema may be good as a test for whether self-replication processes of a certain sort exist, it will not give us a general recipe for producing no-self-replication theorems. Furthermore, even if we restrict ourselves to the discussion of self-replication that is induced by a product structure, it is simply false that the non-existence of Cartesian products implies the non-existence of cloning: even though cloning morphisms exist, a Cartesian product could fail to exist because suitable projections π_1, π_2 do not exist.

Finally, P3. First, what is a suitable quantum category for our purposes? Based on P1, this should be (at least) a category with finite-dimensional Hilbert spaces as objects and linear operators as morphisms. And based on our usual notion of composition in quantum mechanics, a composite state is the tensor product of states, and a composite system is the tensor product of systems. It is easy to see that the tensor product is not a Cartesian product. For first, there are no linear operators that could play the role of the projections in the above diagram. And second, by a standard argument that we will rehearse shortly, there can be no self-replication morphism $\Delta : A \rightarrow A \otimes A$ in this category. Then, by an elementary proposition of category theory, viz. ‘a monoidal product – in this case the ordinary tensor product – induces a product structure iff there exist natural diagonals and projections’, it follows that the tensor product does not induce a Cartesian product.

The standard argument for the non-existence of a no-self-replication morphism runs as follows: assume that such a self-replication morphism Δ exists, and let us for simplicity work with a system whose basis states are $|0\rangle$ and $|1\rangle$. Then $\Delta|0\rangle = |00\rangle$ and $\Delta|1\rangle = |11\rangle$. Now consider a general state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Applying Δ , assumed linear, to both sides of the identity, we find:

$$\Delta|\psi\rangle = \alpha|00\rangle + \beta|11\rangle, \tag{6}$$

while

$$\Delta|\psi\rangle = |\psi\psi\rangle = \alpha^2|00\rangle + 2\alpha\beta|01\rangle + \beta^2|11\rangle \tag{7}$$

Thus, a generic (i.e. basis-independent) self-replication morphism does not exist. It is also clear that this result immediately extends to the case of no-cloning, i.e. with a non-trivial machine system appended to

the source and target systems. This proof is entirely due to the *linear* structure of quantum states and operators, and it explains why quantum categories do not have a product structure.¹² On the other hand, a category might fail to have a product structure for other reasons – in the next section we shall see that classical categories do not in general have a product structure; but this is due to complications concerning how to compose bilinear forms on respective subsystems and *not* due to the linear structure of states.

So much for the quantum case, where it seems to us that Baez’s schema has not only been successful, but also a fruitful source of interactions between quantum mechanics and category theory.¹³ On to the case of classical physics: here it is common to see the claim that one can clone a classical state, especially when made in contrast with the no-cloning theorem of quantum information theory. Witness Vedral in his recent textbook [42]:

No cloning implies that there is a limited amount of information which we can learn about an unknown quantum state. This is in stark contrast to classical physics: we can always learn everything about a classical system even if we have only one copy of it, simply because classical measurements are noninvasive and do not destroy the state of the measured system. For example, when we photocopy a piece of paper, we do not destroy the content of the writing on it. (p. 32 [42])

Indeed the idea seems not only to be pervasive, but necessary for life as we know it. For instance, Gisin et al [35] note that biological reproduction is possible because the encoding in DNA is a classical process (a point about which Wigner was ignorant; hence his worry that ‘biological reproduction appears to be mysterious from the point of view of the physicist’ [47]).

In early work on this topic [6], Baez himself characterized classical physics as taking place in the category *Set* (or categories modeled on *Set*) which is of course a Cartesian category, i.e. a category with products, and thus comes equipped with a cloning morphism. And indeed this seems to be a fair characterization of generic classical information theory, in which it is uncontroversial that we can clone classical bits (for an exposition of this point, see e.g. Section 3 of [9]).

Nonetheless, it is important to notice that there is a conceptual gap between the processes of classical information theory and those of some classical physical theory (treated as a theory of information-processing): while the information processes of the latter are constrained by additional structure and parameters (dynamical laws, symmetries, etc.), the former are not. And thus it makes sense to ask whether the tasks that one can perform in the former (such as cloning) straightforwardly carry over to the latter; it may be that

¹²This argument was first given by Ghirardi in an anonymous referee’s report for Foundations of Physics, although it was first published in [48], followed closely by [21]. A slightly different proof, which uses the property of unitarity, is given in [34].

¹³Indeed, it is one of a family of projects which try to understand and illuminate quantum mechanics in terms of category theory. For some other approaches to categorical quantum mechanics, see [2], [43], [16], and [28].

the additional structure presents an obstacle to these tasks.

The observation that the additional structure of classical physics (over above abstract classical information theory) might constitute an obstacle to cloning is not new: Baez first conjectured this in [7] and Fenyes gave a proof of a special case in [18]. Baez-Fenyes work with the symplectic (analogously: Poisson) formulation of Hamiltonian mechanics and locate the obstacle in the preservation of the symplectic form (analogously: Poisson bracket). The next three sections can thus be seen as an extension of their work; in particular, we shall recover Fenyes’s result as a special case of our no-cloning argument for classical mechanics. §4 surveys some of the categories that are relevant to classical mechanics. §5 and §6 are a departure from category theory: here we prove a no-cloning theorem by working directly with the objects themselves, and direct the reader to [27] for analogous category-theoretic arguments.

4. Categories for classical mechanics

What are the categories whose objects describe classical mechanical systems and whose morphisms thus describe operations on those systems? Below we shall see that the category of interest depends on our chosen formulation of classical mechanics. For the Lagrangian formulation, we have a category of Lagrangian systems (see below for the definition), whereas for the Hamiltonian formulation, we have a category of symplectic manifolds that describe phase spaces, or more generally, a category of Poisson manifolds that perform the same function. It is natural to think of symplectic manifolds as state spaces (some of) whose morphisms represent time evolution of the system, whereas the morphisms of a Lagrangian system are most naturally thought of as symmetries of the system. We shall thus focus on the Hamiltonian perspective in our arguments. However, it is useful for completeness to present the Lagrangian perspective, especially because it is rarely discussed from a categorical point of view.

We should at this point stress that the ‘naive’ categories discussed below do not in fact give us the resources for reasoning about cloning in a category-theoretic way. As we shall see, neither the Lagrangian nor the symplectic category below comes equipped with a category-theoretic notion of ‘point’ or ‘state’: in a Lagrangian system, a state is given by the position and velocity of constituent particles at some point in time, and in a Hamiltonian system, a state is given by the (generalized) position and momentum of constituent particles at some point in time, but in neither case can we express a state as a morphism in the relevant category. This problem can be partially overcome by adopting the strategy of §5.3. But in the rest of this paper we shall merely treat categories as a useful tool for organizing information about objects whose properties we then examine in a more mundane way.

Before we proceed to describe the more ‘physical’ Lagrangian and symplectic classical categories, we briefly review a ‘proto-physical’ classical category that many have claimed provides a universe of mathematical discourse for classical information theory, viz. the category of sets. This example serves to illustrate

two morals, viz. that (i) such categories need to be supplemented with much richer structures if they are to be useful for information theory, and (ii) relaxing the structure of morphisms in a category can often lead to a situation in which the original product structure is no longer a product. The second moral hints at the kinds of moves that we are forced to make in [27].

4.1. *Set and Rel*

Abstracting away from the laws of physics, the most ‘naive’ example of a classical category that one might come up with is probably **Set**, viz. the category whose

- Objects are sets.
- Morphisms are functions between sets.

And if one conceives of cloning as being given by the diagonal morphism $\Delta : A \rightarrow A \times A$, one can then use the fact that **Set** comes equipped with a categorical product \times (i.e. the Cartesian product) to show that cloning is possible in this category. Note that the existence of cloning in this sense doesn’t mean that one will be able to make any interesting information-theoretic inferences without adding further structure to the category: in particular, in information theory, we usually want to treat the morphisms (or information processes) as objects, and this will mean at least adding some kind of ‘higher structure’ to the category.¹⁴

However, functions are not the only morphisms between sets that we can study. More generally, we can also define the category **Rel** whose

- Objects are sets.
- The set of morphisms $\mathbf{Rel}(X, Y)$, between two sets X and Y , is the set of all subsets of $X \times Y$. In other words, the morphisms in **Rel** are relations between sets.

Notice that in **Rel**, the Cartesian product \times is no longer a categorical product! This can be easily seen by considering the ‘empty Cartesian product’, i.e. a singleton set $\{*\}$, and observing that there is no unique map from some arbitrary object A to $\{*\}$. In other words, $\{*\}$, which was a terminal element in **Set**, is no longer a terminal element in **Rel**. Thus, if one thinks that system composition is given by the Cartesian product, it turns out that this product no longer induces cloning in **Rel**. On the other hand, the disjoint union \sqcup of sets is a categorical product in **Rel** - indeed, a biproduct - but this notion of composition seems too weak to be of any relevance for information systems. The moral of this example is clear: in order to argue that cloning is possible for some class of systems, one needs (i) to be careful about one’s choice of category for those systems; and (ii) to be able to justify why the morphisms of that system should be restricted or relaxed in the chosen way.

¹⁴For a first stab in this direction, see [3].

4.2. Lagrangian systems

One of the most basic and powerful ways of conceptualizing classical mechanics is via the Lagrangian formalism. This is both more and less general than the Hamiltonian formalism. It is more general in the sense that a Lagrangian analysis can be given without assuming that the configuration space manifold M admits of a space-time splitting (whereas the Hamiltonian formalism requires a space-time splitting). On the other hand, there are many Hamiltonian systems that do not admit of a Lagrangian formulation, because, e.g. the inverse Legendre transform does not exist.¹⁵

A Lagrangian system is a pair (L, M) , where the manifold M is the configuration space (i.e. the space of positions of the constituent particles) and the Lagrangian $L : TM \times \mathbb{R} \rightarrow \mathbb{R}$ is a smooth real-valued function, where TM is the tangent bundle of M , and the second argument of L , i.e. $t \in \mathbb{R}$, allows for time-dependent Lagrangians. In order to obtain the dynamical behavior of the objects described by (L, M) , we adopt the following recipe: fix an initial configuration $q_0 \in M$ for time $t = t_0$ and a final configuration $q_1 \in M$ for time $t = t_1$ and consider the space of paths $P(q_0, q_1)$ between these points. The action functional $S : P(q_0, q_1) \rightarrow \mathbb{R}$ is defined by $S(\gamma) := \int_{t_0}^{t_1} L(q, \dot{q}, t) dt$, where γ is a path in $P(q_0, q_1)$ along which the integration takes place. The extremals (i.e. the critical points) of S as the paths vary are the solutions to the equations of motion for the system described by (L, M) – in common parlance, they are the ‘on-shell’ dynamical trajectories for objects traveling from q_0 to q_1 in the time interval $[t_0, t_1]$. The equations for the extremals, i.e. the Euler-Lagrange equations, can be obtained directly from the Lagrangian by considering the kernel of the Euler-Lagrange operator D :¹⁶

$$DL = 0, \quad \text{where} \quad D := \frac{d}{dt} \frac{\partial}{\partial \dot{q}} - \frac{\partial}{\partial q}. \quad (8)$$

Each Lagrangian system gives us all the information we need in order to derive the dynamics of the system.

We might thus try to construct a category of Lagrangian systems and see if we can give any sensible interpretation to its morphisms. If a manifold is equipped with some extra structure, we often require that a map from it to some other similarly structured manifold preserve that structure: following this line of thought, let us for simplicity restrict our attention to time-independent Lagrangians, so that $L : TM \rightarrow \mathbb{R}$, and let a morphism from (L, M) to (L', M') be a diffeomorphism $f : M \rightarrow M'$ such that L' pulls back to L , i.e. $f^*L' = L$.¹⁷ A moment’s thought shows that there is indeed a sensible interpretation: these morphisms capture our intuitive notion of a symmetry of a Lagrangian system.

¹⁵There are also of course classical systems with nonholonomic constraints, which are strictly speaking neither Lagrangian nor Hamiltonian.

¹⁶Strictly speaking, the Euler-Lagrange operator should be defined as a coboundary operator acting on the variational bicomplex (where the Lagrangian L is written as a form), but we have suppressed this technicality in favor of the elementary formula above. See e.g. [15] for a thorough exposition of the geometric approach to Lagrangian mechanics.

¹⁷This definition will strike many as being too strong, as it implies that there are no morphisms in this category that are not isomorphisms. One could try to relax these morphisms to relations of some sort, as we do in §5.3.

The category of Lagrangian systems gives us quite a lot of information about classical mechanical systems. However, within our framework, its morphisms are most naturally thought of as (general) symmetries and not as providing time-evolution from one state to another; (recall that the states of a Lagrangian system (L, M) are given by the positions and velocities (q^i, \dot{q}^i) of constituent particles at some point in time, i.e. coordinates of a bundle chart TU on TM , where U is a coordinate chart on M). The Hamiltonian formalism, on the other hand, offers a formalism in which time evolution transparently appears as a specific automorphism (viz. the Hamiltonian flow) of a symplectic manifold.¹⁸

We shall now describe how to transit from the Lagrangian notion of state to the Hamiltonian notion of state. Consider the 1-form $\theta_L := \frac{\partial L}{\partial \dot{q}^i} dq^i$ on TM and its analog on T^*M , viz. the Liouville or canonical 1-form $\theta := p_i dq^i$.¹⁹ The transition is accomplished by a fiber-wise mapping $\tau_L : TM \rightarrow T^*M$ such that $\theta_L = \tau_L^*(\theta)$. This mapping is called the Legendre transform associated with L . It is a local diffeomorphism iff L is nondegenerate.²⁰ Assuming non-degeneracy, we can then define the Hamiltonian $H : T^*M \rightarrow \mathbb{R}$ by $H \circ \tau_L := \dot{q}^i \frac{\partial L}{\partial \dot{q}^i} - L$ and thus begin to discuss the Hamiltonian formulation of classical mechanics.

4.3. Symplectic systems

When the Legendre transform $\tau_L : TM \rightarrow T^*M$ is a diffeomorphism, the Euler-Lagrange equations on TM can be recast as Hamilton's equations on T^*M , viz.,

$$\dot{q}^i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q^i}. \quad (9)$$

From these equations, we observe that a Hamiltonian H defines a corresponding vector field $X_H = \frac{\partial H}{\partial p_i} \frac{\partial}{\partial q^i} - \frac{\partial H}{\partial q^i} \frac{\partial}{\partial p_i}$ whose integral curves take earlier states to later states.²¹ Thus, the one-parameter group $\{g_t\}_{t \in \mathbb{R}}$ (parameterized by the time t) of diffeomorphisms of T^*M generated by X_H is known as 'the Hamiltonian flow'.

It turns out that the Hamiltonian flow on T^*M preserves the closed, nondegenerate 2-form ω on T^*M , i.e. the symplectic form, which is the exterior derivative of θ :

$$\omega := d\theta = dp_i \wedge dq^i. \quad (10)$$

By definition, the Hamiltonian flow satisfies Hamilton's equations. More generally, any smooth automorphism of T^*M that preserves the symplectic form (a 'canonical transformation') preserves Hamilton's equations.

¹⁸This contrast is of course somewhat superficial and formal, as surely both (q, \dot{q}) and (q, p) equally represent the instantaneous state of the system.

¹⁹Here we are using standard coordinates on a chart of T^*M . Let (U, ϕ) be a chart on M ; we adopt the coordinates $(q^1, \dots, q^n, p_1, \dots, p_n)$ of the chart T^*U on T^*M , where the second n coordinates parametrize the fiber of T^*U .

²⁰Recall that a Lagrangian is said to be nondegenerate iff its Hessian $\frac{\partial^2 L}{\partial \dot{q}^i \partial \dot{q}^j}$ is invertible. For an exposition of this point, and the transition from Lagrangian to Hamiltonian mechanics in general, see e.g. [32, 4, 11, 12].

²¹Indeed, X_H is obtained from H by 'raising' the index of dH with the symplectic form, which yields a vector field. The vector field is in 1-1 correspondence with the form field because the symplectic form is non-degenerate.

All this generalizes quite nicely within the framework of symplectic geometry. The state space T^*M equipped with the form $dp_i \wedge dq^i$ is the prototypical example of a symplectic manifold; more generally, a symplectic manifold is a pair (M, ω) , where M is a manifold and ω is a symplectic form.²² We can construct a category whose objects are symplectic manifolds: in this category a morphism from (M_1, ω_1) to (M_2, ω_2) is given by a ‘symplectomorphism’, i.e. a diffeomorphism $f : M_1 \rightarrow M_2$ such that $f^*\omega_2 = \omega_1$. Evidently, symplectomorphisms are a generalization of the notion of a canonical transformation, of which the Hamiltonian flow is an instance. We should however bear in mind that it is the Hamiltonian flow that describes the dynamics of a classical mechanical system.

Notice that just as in the case of our naive ‘Lagrangian category’, all the morphisms in this symplectic category are *isomorphisms*. Since the category-theoretic notion of a ‘state’ of some object A (representing a state space of the system) is usually taken to be a point of A , i.e. a morphism from the terminal object to A , it follows that in these categories, we do not have access to the notion of a ‘state’ that is purely categorical. This obstacle is overcome in [27]; at any rate it will not stop us from reasoning about states in this paper: we shall simply discuss states in a conventional, i.e. non-categorical way.

Despite the limitations of this symplectic category, it does come equipped with an operation that is crucial for our discussion, viz. a notion of composition for objects: given two objects (M_1, ω_1) and (M_2, ω_2) , we can form a composite object $(M_1 \times M_2, \Omega) =: M_1 \otimes M_2$. But there is no unique way to construct the composite symplectic form Ω out of ω_1 and ω_2 . In fact, any choice of the following form will do:

$$\Omega = \lambda_1 \pi_1^* \omega_1 + \lambda_2 \pi_2^* \omega_2, \quad \lambda_1, \lambda_2 \in \mathbb{R}/\{0\}, \quad (11)$$

where π_1 and π_2 are the projections of $M_1 \times M_2$ onto the first and second factors of the Cartesian product respectively. Two common choices are

$$\Omega = \pi_1^* \omega_1 + \pi_2^* \omega_2, \quad (12)$$

and the twisted composite form

$$\Omega = \pi_1^* \omega_1 - \pi_2^* \omega_2. \quad (13)$$

It is easy to be impressed by the fact that while the Cartesian product is a categorical product in the category of smooth manifolds (thus the possibility of forming $M_1 \times M_2$), there is no such product in the category of symplectic manifolds. This is due in the first instance to the fact that one cannot define a symplectomorphism between manifolds of different dimension, and so obviously one cannot define anything like a projection here.²³ From the absence of such projections, and thus of a Cartesian product, one might be tempted to conclude that there is no-cloning in symplectic mechanics; however, as we will shortly argue, the

²²Other ‘prototypes’ of symplectic manifolds in physics include coadjoint orbits and Kahler manifolds.

²³Note however that a generalization of the notion of symplectomorphism, viz. a Lagrangian correspondence, can be defined between symplectic manifolds of different dimension, and this may offer a way of recovering the notion of a projection.

no-cloning argument in classical mechanics in fact turns on a more central feature of symplectomorphisms, viz. that they are required to preserve the symplectic form. Before we do so, we describe a yet more general category for classical mechanics, viz. the category of Poisson manifolds.

4.4. Poisson systems

A Poisson manifold is a manifold M whose algebra $\mathcal{A} = C^\infty(M)$ of classical observables is a Poisson algebra. A Poisson algebra, in turn, is a Lie algebra whose Lie bracket (called the Poisson bracket, and denoted $\{\cdot, \cdot\}$) is a bi-derivation²⁴ with respect to pointwise multiplication in \mathcal{A} . The derivation property of the Poisson bracket implies that, in coordinates, it takes the more familiar form:

$$\{f, g\}(x) = \eta^{ij}(x) \partial_i f(x) \partial_j g(x). \quad (14)$$

The tensor η^{ij} is called the Poisson tensor, and defines a global section of the bundle $TM \wedge TM \rightarrow M$.²⁵

Poisson manifolds form a category, in which

- An object is a Poisson manifold $(M, \{\cdot, \cdot\})$.
- A morphism between $(M_1, \{\cdot, \cdot\}_1)$ and $(M_2, \{\cdot, \cdot\}_2)$ is a smooth map $\phi : M_1 \rightarrow M_2$ such that $\{f\phi, g\phi\}_1 = \{f, g\}_2 \phi$ for all $f, g \in C^\infty(M_2)$.

Just as with the category of symplectic manifolds, the category of Poisson manifolds comes equipped with a notion of composition for two objects for which there are analogues of (11), (12) and (13).²⁶

It should be clear from these definitions that Poisson manifolds are very closely related to symplectic manifolds. In order to say exactly how they are related, one needs the notion of a nondegenerate Poisson manifold, which is a Poisson manifold whose Poisson tensor η is non-degenerate everywhere on M . We can then state the relevant theorem relating symplectic manifolds to Poisson manifolds, viz. that the category of symplectic manifolds is anti-isomorphic to the category of non-degenerate Poisson manifolds.²⁷ Indeed, the direction that says ‘every symplectic manifold is a Poisson manifold’ is obvious and familiar from how, in the geometric formulation of Hamiltonian mechanics, one defines the Poisson bracket in terms of the symplectic form, i.e. one defines $\{f, g\} := \omega(X_f, X_g)$, where $f, g \in C^\infty(M)$, ω is the symplectic form, and X_f and X_g are the vector fields associated to f and g by the symplectic form. It is easy to verify that this bracket is indeed a Poisson bracket.

In sum, all the symplectic constructions that we use for our no-cloning result below can be reformulated in terms of Poisson geometry. Although we will not need the added generality of Poisson structures in our

²⁴That is to say, the Poisson bracket satisfies Leibniz’s law in both arguments.

²⁵Here we use the ‘wedge’ notation to indicate that each fiber of the bundle is the antisymmetrized tensor product of two tangent spaces of M .

²⁶We shall not need this notion of composition here, and so we refer the reader to [39] for details.

²⁷See e.g. p. 53 of [39].

no-cloning argument, Poisson geometry is helpful when we consider questions related to quantization (see, e.g. [30]).²⁸ For our purposes, the notion of a Poisson algebra of observables will be useful in order to see how our results relate to the CBH characterization theorem [13], which we briefly discuss at the end of the next section.

5. No-cloning in classical mechanics

In this section, we present several perspectives on self-replication and cloning in classical mechanics. We begin with the tangent space geometry of a symplectic manifold, and proceeding to consider our subject from progressively more global points of view, i.e. Lagrangian submanifolds and symplectic invariants. We then sketch the connections between our topic and the more widely discussed case of quantum no-cloning.

Before proceeding to discuss cloning, we first prove an elementary no-self-replication result in classical mechanics with symplectic manifolds. Let us remind the reader of what we mean here by ‘self-replication’: a source system M_s self-replicates itself to a target system M_t just in case there exists a blank state b such that for all initial source states x , there is a self-replication map that sends (x, b) to $(x, x) \in M_s \otimes M_t$. (Recall that, as in §4.3, here ‘ \otimes ’ does not denote the tensor product, but rather the composite symplectic manifold constructed from the direct product of the component system manifolds and equipped with the composite symplectic form.) We will show that assuming the existence of such a map leads to a contradiction.

Since we are working within the symplectic category, our systems are symplectic manifolds, and our morphisms are symplectomorphisms. Furthermore, the source system and the target system must be copies of the same symplectic manifold (M, w) , because otherwise s would by definition be unable to self-replicate into two qualitatively identical copies of itself. (On the other hand, when considering cloning, the machine system M_m need not have the same dimension as M_s ; §5.2 argues that cloning is only possible when the dimension of M_m is at least that of M_s .) However, one need not assume that M_s and M_t come with the same Hamiltonian function (which prescribes the dynamics) or even that they individually have a Hamiltonian function at all; indeed our argument will not place any restriction on the Hamiltonian functions of either the subsystems or the composite system.

Even before giving any rigorous arguments, it is intuitively clear from our setup that a self-replication map cannot exist. Consider that one of the key geometric features of symplectomorphisms is that, if a region of space U evolves according to such a map, the area of its ‘symplectic shadows’ (i.e. the projection of U onto the symplectic planes²⁹ in the manifold) will be preserved. However, given that a self-replication map sends (x, b) to (x, x) for all x , it is intuitively obvious that such a map will not preserve the symplectic

²⁸However, also see [25] for why ‘deformation quantization’ is a somewhat anemic form of what we usually mean by quantization.

²⁹A symplectic plane is a plane made up of a pair of conjugate coordinates (q, p) .

shadows of a region $U \times \{b\} \subset M_s \times M_t$. Indeed it will duplicate the total area occupied by the symplectic shadow of U at the initial time. We shall go over this in more detail below; suffice it for now to remark that the principle prohibiting self-replication is naturally expressed in terms of global (i.e. non-infinitesimal) symplectic geometry. It is gratifying to see this connection, as it is precisely global methods that have led to the remarkable progress in symplectic geometry in the last two decades, as epitomized by Gromov’s non-squeezing theorem and the famous ‘symplectic camel’ (just as it is impossible for a camel to pass through the eye of a needle, a unit ball cannot be ‘symplectically squeezed’ through a hole of radius less than one!³⁰).

It may be helpful at this juncture to allay a potential worry that the reader might have: Surely we are familiar, within the world of classical physics, with the possibility of gathering information about a classical state without disturbing it. But if self-replication is impossible, then we seem to have found a situation in which one can only gain information about a state by disturbing it. Two remarks will serve to clarify this situation: First, the self-replication case is not generic, and indeed it is not the operationally important case of §3.2 (i.e. cloning) since it does not include a non-trivial machine system. Including a machine system gives us enough ‘space’ to clone, thus protecting our cherished intuitions about the classical world, albeit in a subtle way. Second, the processes that we discuss abstract away from the details of classical measurement theory, which are necessary for the notion of ‘information gathering’ to be realized. On the latter topic, one finds varying claims in the literature, some of which are in tension with each other: for instance, [44] claims that even if one cannot strictly speaking clone (in the statistical case), one can always do so up to an arbitrarily high fidelity; on the other hand, [33] (p. 372) argues that Gromov’s non-squeezing theorem places fairly strong constraints on just how precisely one can measure a state in (both statistical and non-statistical) classical mechanics. We cannot enter this interpretive debate here, but instead leave it to a further investigation.

5.1. Self-replication

We now embark upon our first no-cloning argument for classical mechanics. Recall from (11) that the symplectic form Ω on the composite system is defined via pullback of the symplectic forms ω of the component systems. For simplicity, we shall take the symplectic form on $M_s \otimes M_t$ to be given by $\lambda_1 = \lambda_2 = 1$, i.e. by (12), so that:

$$\Omega_{(x,y)}[(u, q), (v, r)] = \omega_x(u, v) + \omega_y(q, r) \quad u, v \in T_x M_s \quad q, r \in T_y M_t, \quad (15)$$

where (x, y) is a point of $M_s \times M_t$. In order to be absolutely clear here, it may help to illustrate what we mean by writing out the above symplectic forms in coordinates. Let a component system M be of dimension $2n$. In the canonical basis (viz. the basis in which the coordinates are $(q^1, \dots, q^n, p_1 \dots p_n)$), the component

³⁰See [33] for a precise statement of this and similar theorems, as well as the different techniques that can be used to prove them, e.g. J -holomorphic curves, generating functions, etc.

system's symplectic form ω takes the *canonical* form:

$$\omega \equiv J = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}, \quad (16)$$

where I_n is the n by n identity matrix. On the other hand, because we have decomposed the tangent space of $M_s \otimes M_t$ as the direct sum of its component systems' tangent spaces, the composite symplectic form Ω is written in a basis in which it takes the $4n$ by $4n$ *block diagonal* form:

$$\Omega = \begin{bmatrix} J & 0 \\ 0 & J \end{bmatrix}. \quad (17)$$

Since any symplectomorphism $\phi : M_s \otimes M_t \rightarrow M_s \otimes M_t$ preserves Ω , we obtain the condition

$$\phi^* \Omega[(u, q), (v, r)] = \Omega[(u, q), (v, r)] \quad (18)$$

on any such map, which can also be expressed in matrix form (and without the use of tangent vectors) as:

$$\phi_*^T \Omega \phi_* = \Omega. \quad (19)$$

From the definition of a 'pullback' ϕ^* , we have $\phi^* \Omega[(u, q), (v, r)] = \Omega[\phi_*(u, q), \phi_*(v, r)]$, and we can use this along with (18) and (15) to obtain the constraint:

$$\Omega[\phi_*(u, q), \phi_*(v, r)] = \omega(u, q) + \omega(s, r). \quad (20)$$

Note that it is by no means obvious that (20) can be met for any choice of ϕ . In fact, we shall show the self-replication maps have push-forwards ϕ_* that violate (20).

Let us begin by considering some generic smooth (not necessarily self-replicating!) map

$$\phi(x, y) = (\phi_1(x, y), \phi_2(x, y)), \quad (x, y) \in N := M_s \otimes M_t \quad (21)$$

whose Jacobian matrix (or, in geometric terminology, 'push-forward') is

$$\begin{aligned} \phi_* : T_{(x,y)}N &\equiv T_x M_s \oplus T_y M_t \longrightarrow T_{(\phi_1, \phi_2)}N \equiv T_{\phi_1} M_s \oplus T_{\phi_2} M_t \\ \phi_* &= \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \\ \phi_* : (u, q) &\mapsto (Au + Bq, Cu + Dq) \end{aligned} \quad (22)$$

$$(23)$$

where $A = \partial_1 \phi_1$, $B = \partial_2 \phi_1$, $C = \partial_1 \phi_2$ and $D = \partial_2 \phi_2$, and all these partial derivative matrices are evaluated at $(x, y) \in N$. (In order to reduce notational clutter, we have abbreviated the basepoints of the tangent spaces from e.g. $\phi_1(x, y)$ to ϕ_1 ; this will cause no confusion in our context.)

Having expressed ϕ_* in matrix form, we can now write (20) more explicitly as:

$$\omega(Au + Bs, Av + Bt) + \omega(Cu + Ds, Cv + Dt) = \omega(u, v) + \omega(s, t). \quad (24)$$

It will be convenient for us to further sharpen this condition by studying how it appears on the copy of $T_x M_s$ that passes through the origin of $T_{(x,y)}N$: that is to say, by evaluating it on a pair of (non-zero, non-identical vectors) $(u, 0), (v, 0) \in T_{(x,y)}N$. In that case, we see that consistency amounts to

$$\omega(Au, Av) + \omega(Cu, Cv) = \omega(u, v). \quad (25)$$

Similarly if we evaluate (24) on $(0, s), (0, t) \in T_{(x,y)}N$, we see that consistency amounts to

$$\omega(Bs, Bt) + \omega(Ds, Dt) = \omega(s, t). \quad (26)$$

We thus see that the conjunction of (i) the composition rule for the symplectic structures of the component systems (15); and (ii) the preservation of the symplectic structure of the composite system (18), places strong constraints on the partial derivatives of ϕ at $(x, y) \in N$. We will now show that self-replication violates these constraints.

The most general notion of a self-replication map is a map ϕ as before, but with the added requirement that

$$\phi(x, b) = (x, x), \quad (27)$$

where $b \in M_t$ is our special ‘blank state’. Since $\partial_1 \phi_i = 1$ at $(x, b) \in N$, we find that $\phi_*(x, b)$ is given by:

$$\begin{bmatrix} 1 & B \\ 1 & D \end{bmatrix}. \quad (28)$$

The first column of ϕ_* now violates the constraint that we derived above, for if we evaluate (24) on $(u, 0), (v, 0) \in T_{(x,b)}N$ for $A = C = 1$, we find that

$$2\omega(u, v) = \omega(u, v) \quad \text{i.e.} \quad \omega(u, v) = 0 \quad \forall u, v \in T_x M_s, \quad (29)$$

i.e. the symplectic form on M_s is completely degenerate! Evidently, this result is fairly robust: it doesn’t matter what the values of B and D are.

Let us briefly review what we have done in the foregoing: we have formalized the intuitive notion of a cloning-map in classical mechanics, ‘lifted’ the notion of cloning to a condition on maps between tangent spaces, and observed that at the tangent space level, there is a powerful obstruction to the existence of such maps. It is thus natural to ask if this non-existence result can be formulated at the level of symplectic manifolds instead of at the level of their tangent spaces. The answer to this question is ‘yes’, and it follows immediately by thinking of (the graph of) a symplectomorphism as a Lagrangian submanifold of some larger manifold. We are thus able to formulate a second version of the no-cloning argument from this perspective.

In order to do so, we shall need to introduce some terminology. By a Lagrangian submanifold L of an n -dimensional symplectic manifold M , we mean a submanifold of dimension $\frac{n}{2}$ on which the symplectic form of M restricts to zero. Now, it is an (easy) fact that the graph of the symplectomorphism $\phi : X \rightarrow Y$ is a Lagrangian submanifold of the product symplectic manifold $X \times \overline{Y}$, where we use the bar notation to denote the manifold equipped with its twisted (i.e. minus) symplectic form. Indeed, this fact can be used to generalize the morphisms of a symplectic category to Lagrangian submanifolds that are not graphs of symplectomorphisms (see e.g. [46] for an exposition of this point). However, for our purposes it will suffice to consider Lagrangian submanifolds that are graphs of symplectomorphisms; we shall call these ‘graph’ Lagrangian submanifolds.

Let $N := M_s \otimes M_t$. We now translate the claim that ‘a self-replication morphism $\phi : N \rightarrow N$ (i.e. ϕ such that $\phi(x, b) = (x, x)$ for some b and all x) does not exist’ into the following claim about Lagrangian submanifolds:

Proposition 5.1.1. *There does not exist a ‘graph’ Lagrangian submanifold in $N \otimes \overline{N}$ that contains the surface $\{(x, b, x, x) \mid x \in M\}$, where the first two coordinates denote a point in N and the last two coordinates denote a point in \overline{N} .*

Proof. Let us first note the general form of such a Lagrangian submanifold: it consists of the points $L \equiv \{(x, y, \phi_1(x, y), \phi_2(x, y)) \mid (x, y) \in N\}$. We are interested to find out if L contains the surface $\Lambda \equiv \{(x, b, x, x) \mid x \in M\}$.

From the invariance of the symplectic form under ϕ , we know that $\phi_*^T \Omega \phi = \Omega$, where T denotes the matrix transpose. In other words,

$$\begin{bmatrix} A^T & C^T \\ B^T & D^T \end{bmatrix} \begin{bmatrix} J & 0 \\ 0 & J \end{bmatrix} \begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} J & 0 \\ 0 & J \end{bmatrix}, \quad (30)$$

where we have used the notation of (17) and (22). This gives us the equations

$$A^T J A + C^T J C = J \quad (31)$$

$$B^T J B + D^T J D = J \quad (32)$$

$$B^T J B + D^T J C = 0 \quad (33)$$

$$A^T J B + C^T J D = 0 \quad (34)$$

But since by the reasoning used to obtain (28) L can only contain Λ if $A = C = I_n$, clearly it cannot do so. \square

There is another way in which one might try to ‘globalize’ the self-replication argument above, and it takes as its starting-point the deep insight that the condition of symplectic invariance (i.e. $\phi^* \Omega = \Omega$) does not really involve derivatives at all. Instead, it can be expressed in terms of the preservation of a topological invariant called a ‘symplectic capacity’, as demonstrated by the following result:

Proposition 5.1.2. *Let $\psi : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$ be a diffeomorphism, c a symplectic capacity on \mathbb{R}^{2n} , and E an arbitrary ellipsoid in \mathbb{R}^{2n} . Then the statement that ψ is either a symplectomorphism or an anti-symplectomorphism³¹ is equivalent to the statement that, for all c and for all ellipsoids E , $c(\psi(E)) = c(E)$.*

³¹An anti-symplectomorphism ψ is a map such that $\psi^* \omega = -\omega$.

Proof. See Prop 12.10 and Lemma 12.11 on pp. 378-9 of [33]. □

We now explain what a capacity is and use this concept to give a no-self-replication argument. A symplectic capacity c is a functor that assigns to every subset $U \subset \mathbb{R}^{2n}$ a number $c(U) \geq 0$ or $+\infty$ and that has the following properties:

- c is invariant under symplectomorphisms, i.e. $c(\psi(U)) = c(U)$.
- c is monotone with respect to set inclusion.
- c scales like an area under dilations.
- c satisfies the normalization conditions $c(B^{2n}(R)) = \pi R^2 = c(Z_j^{2n}(R))$, where $B^{2n}(R)$ is the ball of radius R , and $Z_j^{2n}(R)$ is the cylinder (infinite in coordinates $i \neq j$) whose circular cross-section lies in the x_j, p_j -plane and is defined by $x_j^2 + p_j^2 \leq R^2$.

The existence of symplectic capacities is highly non-trivial and we shall not discuss the matter here.³² We only wish to show how no-self-replication is related to this invariant that captures the deepest aspects of symplectic geometry. In fact there exist many symplectic capacities, and they are all bounded from below by the minimum symplectic capacity c_0 , which is defined as follows:

$$c_0(U) := \sup_f \{ \pi r^2 \mid f(B^{2n}(r)) \subset U \} = \pi R^2,$$

where R is the supremum of the radii of balls that we can embed into U via the symplectomorphisms f .

In order to give a no-self-replication argument using symplectic capacities, let M_s be \mathbb{R}^2 and M_t be \mathbb{R}^2 , both equipped with the standard symplectic structure. Let us consider a ball $B^4(R)$ of radius R centered at the origin of the symplectic manifold $\mathbb{R}^2 \times \mathbb{R}^2$, which is contained in a cube (also centered at the origin) whose sides have length $2R$. By definition, $c_0(B^4(R)) = \pi R^2$. Under the action of the cloning map ϕ , the ball gets stretched into a region whose ends along the diagonal touch the vertices of the cube (this is true, because setting $b = 0$,³³ we have $\phi(x, 0) = (x, x)$). But ϕ is volume-preserving since it is assumed to be a symplectomorphism; so the middle of the region must be squashed to compensate for the stretch along the diagonal. This means that the largest ball that one could symplectically embed into the ellipsoid has a radius smaller than R , and so c_0 is not preserved.³⁴ It then follows from Prop 5.1.2 that ϕ is not a symplectomorphism, i.e. on pain of contradiction, no such self-replication map can exist.

³²The locus classicus for this notion is Chapter 12 of [33].

³³We can do so without any loss of generality, since a translation is a symplectomorphism.

³⁴Another way of putting this point would be to note that any symplectic embedding into region would involve ‘squeezing’ the projection onto either the M_s or the M_t planes, but this is prohibited by Gromov’s non-squeezing theorem.

A somewhat more heuristic, but slicker, argument involving capacities can be given by using the concept of a Hofer-Zehnder capacity c^{HZ} , which is a capacity such that for a compact convex set U in a symplectic manifold,

$$c^{HZ} = \int_{\gamma} p_i dq^i,$$

where γ is the shortest positively-oriented Hamiltonian periodic orbit on the boundary ∂U of U . Applying this to the setup of the previous paragraph and assuming that the self-replication map ϕ transforms $B^4(R)$ into a convex region, we intuitively see that because of the squashing that compensates for the stretch, the shortest Hamiltonian periodic orbit would now be shorter than a great circle of the sphere $\partial B^4(R)$ (which is the shortest Hamiltonian periodic orbit on that sphere). This contradicts the assumption that ϕ is a symplectomorphism.

5.2. Cloning

So much by way of proving our no-self-replication theorem. We now discuss the case of a general *cloning* map and argue that the no-self-replication argument does *not* extend to this case as well. Indeed, we shall show that the obstruction to cloning that was present in the self-replication case is now no longer an obstruction, but merely a constraint, and we shall give a geometric interpretation of this constraint and of why it can be satisfied.

Recall that in order for cloning to make sense semantically and also operationally, one needs three systems: a cloning machine system M_m , a source system M_s , and a target system M_t . Although *a priori* there need not be any constraint on the dimensionality of M_m (apart from its being even), let us assume for now that $M_m = M_s = M_t = \mathbb{R}^2$; we shall comment on the dimensionality of the machine system later. Let $N := M_s \otimes M_t \otimes M_m$.

The cloning map ϕ is described by:

$$\phi(x, y, z) \equiv (\phi_1(x, y, z), \phi_2(x, y, z), \phi_3(x, y, z)) : N \rightarrow N \quad (35)$$

$$\phi : (x, b, z) \mapsto (x, x, z'), \quad (36)$$

for some b, z , and all x . Its push-forward $\phi_*(x, b, z)$ is given by:

$$\begin{bmatrix} 1 & A & B \\ 1 & C & D \\ E & F & G \end{bmatrix}. \quad (37)$$

As before, if ϕ is to be a symplectomorphism then ϕ_* must satisfy the equation

$$\phi_*^T \Omega \phi_* = \Omega, \quad (38)$$

where now $\Omega = \text{diag}(J, J, J)$. However, this time one can see by explicitly multiplying the matrices that no inconsistency can be generated for the matrix elements of ϕ_* . Indeed, in [17] an explicit example is given of a linear symplectomorphism that takes the form (37).

Nonetheless, one can still obtain some non-trivial information about ϕ_* from symplectic invariance: in particular, it strongly constrains the matrix element obtained by multiplying the row $(1, 1, E)$ of ϕ_*^T with the column $(J, J, JE)^T$ of $\Omega\phi_*$, which results in the matrix equation

$$E^T J E = -J. \quad (39)$$

Indeed, by writing out the matrix elements of E explicitly, we immediately see that this forces the condition that $\det E = -1$. We now explain the meaning of this negative determinant condition and why it follows from ϕ being a symplectomorphism.

We begin our explanation by reminding the reader that the simplest symplectic invariant is actually the relative integral invariant $\int_\gamma p_i dq^i$, where γ is an oriented 1-cycle (or closed curve). In other words, if ψ is a symplectomorphism then we have

$$\int_\gamma p_i dq^i = \int_{\psi(\gamma)} p_i dq^i, \quad (40)$$

which, if the region bounded by γ is simply connected, is just the absolute integral invariant defined by the integral of the symplectic form (a 2-form) over the region bounded by γ .³⁵

Again, set $M_m = M_s = M_t = \mathbb{R}^2$ and let us consider the action of a cloning map ϕ on a closed oriented 1-cycle γ lying in $M_s \times b \times z$. From the definition of a cloning map, we immediately see that $\phi(\gamma)$ has identical projections onto both the M_s and M_t planes, viz. γ_s and γ_t respectively, which are both copies of γ because $\phi_1(x, b, z) = \phi_2(x, b, z) = x$. If we apply the relative integral invariant to this situation, we obtain:

$$\int_\gamma p_i dq^i = \int_{\gamma_s} p_i dq^i + \int_{\gamma_t} p_i dq^i + \int_{\gamma_m} p_i dq^i, \quad (41)$$

where γ_I is the projection of the image $\phi(\gamma)$ onto the M_I plane, for $I = m, s, t$. But since the first two terms on the RHS are both identical to the LHS, it is clear that ϕ is a symplectomorphism only if the integral over γ_m cancels one of the other integrals on the RHS. In order for this to be the case, γ_m must be given the *opposite* orientation to γ_s and γ_t , and this is precisely the meaning of the orientation-reversing condition $\det E = -1$ above. In other words, a cloning symplectomorphism ϕ is only allowed if the positive contribution to the integral invariant (or the symplectic area) gained from cloning is compensated for by a negative contribution, and the latter requires an orientation-reversing (and area-preserving) element of the Jacobian.

Thus far, we have been assuming that M_m has the same dimension as M_s and M_t . Let us consider the simplest situation in which the dimension of M_m is smaller than M_s and M_t , viz. $M_m = \mathbb{R}^2$ and

³⁵This follows from Stokes' theorem for differential forms.

$M_s = M_t = \mathbb{R}^4$. M_s is then covered by the global coordinates (q^1, q^2, p_1, p_2) and can be decomposed into two symplectic planes, viz $P_{s1} \equiv (q^1, p_1)$ and $P_{s2} \equiv (q^2, p_2)$. We can thus consider the above integral invariant taken over the sum of two 1-cycles γ_1 and γ_2 that live in P_{s1} and P_{s2} respectively. In order to see what goes wrong with cloning in this situation, let us consider the case where ϕ is a linear symplectomorphism (i.e. a 10 by 10 matrix) and the two 1-cycles are circles of unit radius (the result will easily generalize to the Jacobians of non-linear symplectomorphisms). Evidently, E is now a 2 by 4 matrix

$$\begin{bmatrix} E_{11} & E_{12} & E_{13} & E_{14} \\ E_{21} & E_{22} & E_{23} & E_{24} \end{bmatrix}. \quad (42)$$

which satisfies

$$E^T J' E = -J, \quad (43)$$

i.e. the analog of (39), except that now J' is a 2 by 2 symplectic matrix and J is a 4 by 4 symplectic matrix. In other words, (43) tells us that

$$\det \begin{bmatrix} E_{11} & E_{13} \\ E_{21} & E_{23} \end{bmatrix} = \det \begin{bmatrix} E_{12} & E_{14} \\ E_{22} & E_{24} \end{bmatrix} = -1, \quad (44)$$

which is precisely the right condition for E to send γ_1 and γ_2 to a reversed-orientation cycle in M_m (indeed, each of these submatrices is area-preserving on γ_1 and γ_2 respectively). This, coupled with the fact that E has fewer rows than columns immediately shows us that (40) will not be satisfied, since the action of the cloning map on γ_1 and γ_2 creates two circles (each with positive area π) in M_t , whereas it creates an ellipse with negative area whose magnitude is greater than 2π in M_m .

For instance, when

$$E = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & -1 & -1 \end{bmatrix}. \quad (45)$$

then E maps the unit circles γ_1 and γ_2 to a reversed-orientation circle of radius 2 in M_m . The point here generalizes for $M_s > 4$: the positive contribution to the symplectic area is given by the sum of areas of n circles of radius R in M_t , whereas the compensating negative contribution is given by the area of a circle whose radius is nR – thus the symplectic area cannot be preserved by such a transformation.³⁶ On the other hand, if the dimension of M_m is larger than that of M_s , then it is always possible for E to contain square blocks that create an equal number of reversed-orientation 1-cycles in M_m to those in M_t .³⁷

Let us take stock. It is striking that cloning, but not self-replication, is allowed by the geometry of classical mechanics, i.e. symplectic geometry. The reason may be summarized thus: Symplectic geometry is

³⁶This heuristic argument is complemented by the algebraic argument given in Prop 5 of [17] for the same conclusion.

³⁷For instance, imagine that E is a 4 by 2 matrix acting on a 2 vector. Then if the second and fourth rows vanish, the analog of (39) tells us that E maps a circle in M_s to a reversed-orientation circle in M_m .

given its distinctive – and often counter-intuitive! – flavor by the fact that in it, one can measure ‘oriented area’ (by which we mean the *symplectic area*, i.e. $\int \omega$) using the symplectic form, just as one can measure lengths in Riemannian geometry using the metric. And just as isometries preserve lengths in Riemannian geometry, symplectomorphisms preserve the oriented area in symplectic geometry. As seen from the action of the self-replication map ϕ on a 1-cycle in M_s , self-replication is a process that *doubles* the original symplectic area, and so such a process cannot be a symplectomorphism. The difference between self-replication and cloning is that, in the latter case, the total system contains additional degrees of freedom that allow the cloning map to include an additional, negative (i.e. reverse-oriented) contribution to the symplectic area that compensates for the doubling: thus a cloning map can be a symplectomorphism. These additional degrees of freedom are precisely what we collectively labeled as ‘the machine system’ above.

Think now of the total system as the universe. It is curious that even allowing arbitrary initial data, one could never have a dynamical process (i.e. a symplectomorphism) that copied one half of the universe to the other half. Indeed one could not even have such a dynamical process within 5/6-ths of the universe! What we have learned from the foregoing is that one must at least have a third of the universe available as ‘machine system’ degrees of freedom of such a copying process is to occur in the rest of the universe.

5.3. Uniform no-cloning and categorical collapse

This section relates the above ideas to the category-theoretic literature on cloning, and may be skipped without interrupting the narrative flow of the main article.

In the category-theoretic literature on ‘cloning’, authors often discuss ‘uniform cloning’ within a monoidal category: by which they mean a co-associative and co-commutative monoidal natural transformation

$$\Delta_A : A \rightarrow A \otimes A, \tag{46}$$

where \otimes is the monoidal product (e.g. the diagonal morphism of (5)). Clearly ‘uniform cloning’ is a category-theoretic generalization of what we have earlier called ‘self-replication’; although in this section we shall continue to say ‘uniform cloning’, as we trust that the reader will not confuse it with what we have earlier called ‘cloning’.

Abramsky has proven the following no-go theorem for uniform cloning, which is meant to capture the ‘categorical essence’ of no-cloning in quantum mechanics [1, 2]:

Theorem 5.3.1. *Let \mathcal{C} be a compact closed category equipped with uniform cloning. Then every endomorphism is a scalar multiple of the identity.*

Proof. See [1]. □

More informally, compact categories can only have uniform cloning on pain of ‘collapsing’ in the sense that their morphisms (i.e. dynamical processes acting on a system) become trivial. Thus the no-go argument against quantum cloning is furnished by the contrapositive of the theorem: since we know that quantum

mechanics has non-trivial dynamical processes and that its category carries a compact structure, it follows that quantum mechanics cannot have uniform cloning.

We shall now give a brief argument (due to Heunen and the author) showing that there is a *classical* category that is compact and admits non-trivial dynamical processes; and so the Abramsky no-go theorem immediately applies to it. This is perhaps not too surprising given the above no-self-replication result for classical mechanics; however it does remind us of the moral that high-level categorical structures such as ‘compactness’ can have very different sources, and are not specific to the quantum world.

In order to exhibit such a classical category, we restrict ourselves to a linear toy model of classical mechanics, viz. the linear symplectic category **Symp**. In order to define this category, we will first need the notion of a linear Lagrangian correspondence (LLC). Recall that a Lagrangian subspace L of a symplectic vector space (V, ω) is a vector subspace such that $\omega|_L = 0$ and $\dim L = \frac{1}{2} \dim V$. Let V_0 and V_1 be symplectic vector spaces. An LLC

$$L_{01} : V_0 \rightarrow V_1$$

is a Lagrangian subspace $L_{01} \subset \overline{V_0} \times V_1$, where the bar $\bar{\cdot}$ notation is used to denote the vector space equipped with its minus symplectic form. We have written it as a morphism above because it is really a generalization of a symplectomorphism, and can be composed. For instance, let $L_{01} : V_0 \rightarrow V_1$ and $L_{12} : V_1 \rightarrow V_2$. Then the composite $L_{01} \circ L_{12}$ is given by

$$L_{02} \equiv L_{01} \circ L_{12} := \{(v_0, v_2) \mid \exists v_1 : (v_0, v_1) \in L_{01}, (v_1, v_2) \in L_{12}\}$$



Thus the linear symplectic category **LinSymp** is defined to have

- Symplectic vector spaces V_0, V_1, \dots as objects.
- LLCs $L_{01} : V_0 \rightarrow V_1$ as morphisms.
- $1_V := \{(v, v) \mid v \in V\} \subset \overline{V} \times V$ as the identity morphism.

It is furthermore easy to see that **LinSymp** is a symmetric monoidal category.

We shall now argue that **LinSymp** carries a compact closed structure. The only non-trivial part of this argument is to identify the right structure to play the role of a ‘dual’ for compactness; it is then trivial to show that the unit and counit maps satisfy the ‘snake axioms’. For a precise statement of the snake axioms, we refer the reader to Definition 6 of [1]; here it will suffice to note that, diagrammatically, they amount to the fact that ‘cup’ and ‘cap’ diagrams

can be ‘straightened out’ as follows:



In order to define the unit and counit, we shall choose $\overline{(\cdot)}$ as our dual $(\cdot)^*$. We then define the unit morphism

$$\eta_V : 0 \rightarrow \overline{V} \times V \tag{47}$$

$$\eta_V : * \mapsto E := \{(v, v) | v \in V\}, \tag{48}$$

and the counit is defined similarly. Evidently this is a Lagrangian subspace since $* \times E$ has half the dimension of $\overline{V} \times V$ and the symplectic form Ω of $\overline{V} \times V$ vanishes when it is restricted to E . We leave it as an (easy) exercise to the reader to show that the unit and the counit satisfy the snake axioms. Thus, **LinSymp** is compact closed and satisfies the assumptions of Abramky's no-uniform-cloning result.

5.4. Quantum cloning and rigidity

In the foregoing sections, we showed that self-replication (or cloning with a trivial machine system) is prohibited in classical mechanics, whereas cloning (with a sufficiently large machine system) is permitted. How do these results compare with their quantum analogs? The answer is well-known, and we summarize it in the table below:³⁸

³⁸Note that in the quantum case we are considering generic states and not merely orthogonal states.

No-cloning results		
Process	Classical	Quantum
Self-replication	Prohibited	Prohibited
Cloning	Permitted	Prohibited

Why is (generic) quantum cloning prohibited whereas classical cloning is permitted? In this subsection we briefly sketch an answer to this question that draws upon the differing geometries of classical and quantum mechanics.

Let us begin by discussing the notion of ‘rigidity’. Consider a smooth manifold, whose structure is preserved by diffeomorphisms – this geometry is very flabby, as the diffeomorphisms can squeeze or stretch regions of space into any shape whatsoever, as long as the transformation is smooth. On the other hand, one might define a volume form on the manifold, and consider the group of volume-preserving diffeomorphisms: while still being relatively flabby, this geometry is more rigid than that of a smooth manifold, as now the morphisms can only transform one region to another of the same volume. A somewhat more rigid structure has mostly concerned us in the foregoing, viz. a symplectic manifold (a smooth manifold that carries a symplectic form) whose structure-preserving maps are symplectomorphisms. This additional rigidity corresponds to the (non-trivial) fact that the group of symplectomorphisms is a proper subgroup of the group of volume-preserving diffeomorphisms.³⁹ However, it is also evident that symplectic geometry still contains a fair bit of flabbiness: this is shown by results such as Darboux’s theorem (i.e. in a small neighborhood of a point, the symplectic form can be expressed as $dq^i \wedge dp_i$ ⁴⁰) and the related fact that a lot of symplectic geometry boils down to topology. The residual flabbiness of symplectic geometry stands out in contradistinction to Riemannian (or pseudo-Riemannian) geometry: the latter is far more rigid than the former, because it requires the preservation of length (i.e. the Riemannian metric). For instance, note that a neighborhood of a point in Riemannian geometry need not be flat, but will in general receive curvature corrections as the metric changes in the vicinity of the point. We will now use this intuitive concept of ‘rigidity’ to highlight a fundamental difference between classical mechanics and quantum mechanics, which furthermore explains the results listed in the table above.

As we showed above, the various cloning and no-self-replication results in classical mechanics are governed by the preservation of symplectic area, which is the distinguishing feature of symplectic geometry. In order to see how the quantum case differs from this, we shall avail ourselves of a geometric formulation of quantum mechanics due to Gibbons, Brody, Ashtekar and Schilling, et al [36, 8, 22]. In this geometric formulation, quantum dynamics is formulated as the evolution of a point on a Kahler manifold that preserves the Kahler

³⁹It is also demonstrated (albeit loosely) by results such as the ‘symplectic camel’ and the Gromov non-squeezing theorem.

⁴⁰This in turn means that there is a sense in which symplectic manifolds are locally (i.e. not just at a point) flat, i.e. one thinks of the symplectic form as the curvature of a connection 1-form $p^i dq_i$.

metric.⁴¹ Since every Kahler manifold is a symplectic manifold (but not conversely), and the (Riemannian) Kahler metric is induced from the symplectic form by means of a complex structure, so the geometric formulation is the perfect framework within which to compare results about dynamics in the classical and quantum case, and in particular, results about cloning (which is standardly formulated as a certain sort of dynamical problem). As we shall explain, these differences are due to the additional rigidity that comes from the (Riemannian) Kahler structure in the quantum case.

First, let us recall that a Kahler manifold (M, ω, \hat{I}) is a $2n$ -dimensional smooth manifold M equipped with a symplectic form ω and an integrable, compatible, almost complex structure \hat{I} . This data is sufficient to ensure that a Kahler manifold is not only a symplectic manifold but also carries a Riemannian metric (called the Kahler metric) that is defined as follows:

$$g(\cdot, \cdot) = \omega(\cdot, \hat{I}(\cdot)). \quad (49)$$

In order to see the connection with quantum mechanics, observe that the (Hermitian) inner product $\langle \cdot, \cdot \rangle$ on the n -dimensional Hilbert space \mathcal{H} , can be decomposed into the sum of a Kahler part and a symplectic part as follows:

$$\langle u, v \rangle \equiv G(u, v) + i\Omega(u, v), \quad (50)$$

where u, v are vectors in \mathcal{H} , Ω is a symplectic form, and $G(u, v) = \Omega(u, Iv)$ is a Kahler metric, where I is a complex structure. The geometrical formulation of quantum mechanics is then obtained by projecting these structures down to the Kahler space $\mathbb{C}P^{n-1}$ (i.e. ‘gauging out’ the phase freedom). It turns out that upon this ‘symplectic reduction’, (i) the physical states of the quantum system are in 1-1 correspondence with the points of $\mathbb{C}P^{n-1}$, and (ii) the time-evolution is given by a flow on $\mathbb{C}P^{n-1}$ that preserves both the symplectic and Kahler structures. In other words, if ϕ is a time-evolution map and g the (reduced) Kahler metric, then we require that $\phi^*g = g$. Or in matrix form, and plugging in the symplectic form ω explicitly, we get:

$$\phi_*^T \omega \hat{I} \phi_* = \omega \hat{I}. \quad (51)$$

The most basic case of a quantum cloning operation, i.e. self-replication, is a map $\Delta : \mathcal{H} \otimes \mathcal{H} \rightarrow \mathcal{H} \otimes \mathcal{H}$ such that $\Delta|\psi\rangle|b\rangle \mapsto |\psi\rangle|\psi\rangle$ where $|\psi\rangle$ is an arbitrary state and $|b\rangle$ is some blank state. Here, both the initial and final composite states are separable composite states of $\mathcal{H} \otimes \mathcal{H}$. In the geometrical formulation, the separable states correspond precisely to pairs of points; one in each factor of the image of the Segre embedding

$$\mathbb{C}P^{n-1} \times \mathbb{C}P^{n-1} \rightarrow \mathbb{C}P^{n^2-1},$$

⁴¹This is all we shall need to know about the geometric formulation of quantum mechanics for our purposes. For subtleties concerning how probabilities are related to distances on the Kahler manifold, as well as concerns about measurement (which lies without the ambit of what we called ‘dynamics’), we refer the reader to the original literature.

where n is the dimension of each factor Hilbert space, and n^2 is the dimension of the composite Hilbert space. In other words, a quantum self-replication map, if it exists, is a Kahler structure-preserving endomorphism on a copy of $\mathbb{C}P^{n-1} \times \mathbb{C}P^{n-1}$ embedded in $\mathbb{C}P^{n^2-1}$. This definition is generalized in the obvious way for cases of cloning with a non-trivial machine system. It will be convenient in what follows to abstract away from the fine structure of $\mathbb{C}P^n$ and merely discuss isometries on generic Kahler manifolds, as this will suffice to prove our results.

We are now in a position to repeat our previous point-wise calculations concerning self-replication/cloning; all we have to do is substitute the symplectic 2 by 2 matrix J with JI , where for the complex structure matrix I we can choose

$$I = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

Notice that $JI = \text{Id}$.

For the self-replication case, it is clear that this substitution will generate an inconsistency in the first equation of (31). Let us thus turn to the case of cloning and try to understand why there is no inconsistency in the classical case, whereas one exists in the quantum case. For simplicity, we restrict ourselves to considering finite-dimensional systems in which the machine system has the same dimensionality as the source and target systems. Upon making the required substitution into (39), we obtain

$$B^T JIB = -JI. \tag{52}$$

But it is clear that this is inconsistent, because the diagonal entries of $B^T B$ must be positive, whereas the diagonal entries of the RHS are negative. Evidently, what was merely a constraint in the classical case has become an obstruction in the quantum case, whence the symplectic structure was ‘rigidified’ into a Kahler (Riemannian) metric.

5.5. Implications for the CBH theorem

Having considered no-cloning in classical mechanics from several different points of view, let us now briefly turn to the broader implications of this result for how we think about the differences between the classical and the quantum world. A much-discussed result in the foundations of quantum mechanics by Clifton, Bub, and Halvorson [13], reconstructs quantum mechanics from three information-theoretic axioms, i.e. no-signalling, no-broadcasting, and no-bit-commitment. More precisely, it reconstructs quantum mechanics by implementing these axioms in a C^* -algebraic setting. In the course of executing this task, CBH also prove several theorems (e.g. theorems 2 and 3 of [13]) that purport to *distinguish* quantum mechanics from classical mechanics. It is this ‘distinguishing’ aspect that we shall be concerned with here. We shall also only be concerned with the no-broadcasting axiom of their theorem, which is a generalization of no-cloning to mixed

states: indeed, for our purposes it will suffice to restrict our attention to pure states and thus to no-cloning proper.⁴²

Recent criticism (see e.g. [40], [38]) of the CBH result turns on the idea that they assume *too much* by requiring that physics be couched in a C^* -algebraic framework. For instance, Timpson [40] objects that (i) the complex structure assumed by C^* -algebras may already be a departure from a neutral starting point, and (ii) the existence of toy models that satisfy the information-theoretic axioms of CBH, but not their C^* -algebraic implementation, indicates that the C^* -algebraic framework smuggles in substantive and possibly tendentious assumptions. In contrast to these criticisms, our results seem to show that the CBH theorem assumes *too little!*

Let us explain. From the perspective of classical mechanics, there is nothing wrong with thinking of the observables on phase space in an algebraic fashion. Indeed, in §4.4 we did precisely this by regarding the set of classical observables as a Poisson algebra. The link between the classical observables and the quantum observables (i.e. the self-adjoint part of a non-Abelian C^* -algebra) is then given by the process of quantization.⁴³ Thus, it is quite natural to think of classical observables as the ‘Abelian limit’ of a C^* -algebra, as CBH do. And as is well known, the states of an Abelian C^* -algebra equipped with the weak-* topology form a compact Hausdorff space that can be thought of as classical ‘state space’. However, what is missing from this notion of ‘state space’ is precisely the notion of a classical symplectic (or Poisson) structure, which was crucial to our no-cloning results above. In particular, Theorem 2 and Theorem 3 of [13] (which roughly say that ‘no-cloning is possible iff the relevant C^* -algebras are Abelian, and thus ‘classical’) do not require the cloning map to preserve a classical Poisson structure. Exploring how the ‘no-cloning’ part of the CBH theorem should be reformulated in order to take into account the classical Poisson structure is an interesting project; however, it lies beyond the scope of our present investigation.

6. No-cloning in classical statistical mechanics

In [14], Daffertshofer, Plastino and Plastino (DPP) present a proof of a no-cloning result in classical probabilistic dynamics, assuming overlapping probability distributions with non-trivial support as their dynamical variables. Their proof proceeds in two steps: first, they assume Liouville dynamics and show that the Kullback-Leibler information distance (a ‘metric’⁴⁴ on the space of probability distributions) is preserved by the Liouville equation. Next, they show that cloning a probability distribution does not

⁴²That is to say, we will here be concerned with the distinguishing of pure state classical mechanics from pure state quantum mechanics. On the other hand, the mixed classical states cannot be broadcast has been obtained by [14], and so the criticism holds more generally in the case of mixed states.

⁴³For a rigorous description of the algebraic approach to quantization, see e.g. Chapters 1 and 2 of [31].

⁴⁴The Kullback-Leibler distance is not strictly speaking a metric in the mathematical sense, as it is not symmetric with respect to the probability distributions.

preserve the Kullback-Leibler metric; and thus by the contrapositive of the first step, it follows that cloning is incompatible with Liouville dynamics.

Actually, the framework that DPP work in is somewhat more general than that of *mechanics* proper, i.e. they only assume a probability density on a smooth manifold M that is governed by a continuity equation. In this continuity equation, the probability density P plays the role of a dynamical variable, and the quantity $\bar{v}P$ describes the flux of P , where \bar{v} is the flow velocity vector field. At this level of generality, the connection has not yet been made with mechanics, since the smooth manifold does not carry any symplectic structure (or in more physical parlance, it is not yet a ‘phase space’). However, if we restrict this setup to the case of statistical *mechanics* (i.e. probability densities on a symplectic manifold), then the continuity equation becomes what is usually known as the Liouville equation

$$\frac{dP}{dt} = -\Omega(X_H, X_P), \quad (53)$$

which makes explicit use of the symplectic structure Ω . In what follows, we shall always assume this restriction as we wish to forge a link with the non-statistical mechanics results of the previous sections.

One virtue of the DPP approach is that it clarifies the relationship between information geometry and no-cloning in classical statistical mechanics. However, the DPP approach also seems to us to suffer from several defects (in the restricted case of mechanics). First, it looks like a rather circuitous way of proving this result, and one is led to wonder if it is possible to give a more direct proof of no-cloning in classical statistical mechanics. In this section, we shall give a direct proof of no-self-replication (i.e. no-cloning with a trivial machine system) by using the symplectic structure, as we did in the non-statistical case above. (The full proof of no-cloning is similar, though slightly more complicated, and we shall treat it, along with no-broadcasting, in a future investigation.) Second, as DPP note, the Kullback-Leibler ‘metric’, which is given by

$$K(P_1, P_2) := \int dx P_1(x) \ln \frac{P_1(x)}{P_2(x)}, \quad (54)$$

(where $P_1(x)$ and $P_2(x)$ are distinct probability distributions) is *not* well-defined when P_1 and P_2 are delta function distributions.⁴⁵ Thus, the DPP approach does not extend to cover the case of non-statistical classical mechanics; nor does it tell us how statistical classical no-cloning is related to the above non-statistical classical no-cloning result. However, as we shall presently see, statistical classical cloning is prohibited for (at least) essentially the same reasons that non-statistical classical cloning is prohibited.

Let us now fix our notation for this section. As in the previous section, we shall take M_s , and M_t to be our source and target systems (i.e. symplectic manifolds) respectively. The coordinates for a point in

⁴⁵The Kullback-Leibler metric is also known as the relative entropy of the two probability distributions. See [41] for the physical interpretations and applications of relative entropy.

the composite system will be written in terms of coordinates for the respective subsystems in the following manner:

$$(x, y) \in M_s \otimes M_t.$$

While the states of classical mechanics are points of phase space (often called ‘pure states’ within the context of statistical mechanics), classical statistical mechanics uses a more general notion of state, viz. ‘a (time-dependent) probability distribution over phase space’. Evidently, this more general notion includes pure states as delta functions localized at some point of phase space. We shall use $P \equiv P(x, y)$ to denote an initial composite probability distribution, and $P_s(x)$ and $P_t(y)$ to denote the initial probability distributions over M_s , and M_t respectively. Similarly, $Q(x, y)$ denotes a final composite probability distribution, and Q_s and Q_t are the final probability distributions over the respective subsystems.

For the sake of simplicity, we shall only give a no-self-replication argument here for the simplest case considered in [14], viz. the case in which the initial composite state P factorizes into its initial subsystem states as $P = P_s P_t$, where P_t is some specially chosen blank state and P_s is an arbitrary source state. A cloning process is a map for which there exists an initial blank state P_t such that for all initial source states P_s , the map takes P to a final composite state Q that factorizes as $Q(x, y) = P_s(x)P_s(y)$. We will now see why such a process is impossible.

Consider that, in the Liouville picture of classical statistical mechanics, the initial state I is given by the pullback ϕ^*F of F by the time-evolution map ϕ . Thus, if the Hamiltonian flow from the initial time to the final time is given by the symplectomorphism $\psi = \phi^{-1}$, this means that

$$P_s(x)P_s(y) = Q(x, y) = P(\phi(x, y)) = P(\phi_1, \phi_2) = P_s(\phi_1)P_t(\phi_2), \quad (55)$$

where we have suppressed the coordinate-dependence on x, y, z in the last two expressions. (Note that no Jacobian factor is necessary here, as ϕ is a symplectomorphism and so its Jacobian matrix is 1.)

We now proceed to give a no-self-replication by differentiating (55) with respect to x and y , evaluating each of these derivatives respectively on the diagonal point (a, a) , and taking their difference. By taking the x -derivative, we obtain

$$\partial_x P_s(x)P_s(y) = P_t(\phi_2)\partial_{\phi_1} P_s(\phi_1)\partial_1 \phi_1 + P_s(\phi_1)\partial_{\phi_2} P_t(\phi_2)\partial_1 \phi_2, \quad (56)$$

and by taking the y -derivative, we obtain

$$\partial_y P_s(x)P_s(y) = P_t(\phi_2)\partial_{\phi_1} P_s(\phi_1)\partial_2 \phi_1 + P_s(\phi_1)\partial_{\phi_2} P_t(\phi_2)\partial_2 \phi_2. \quad (57)$$

Since (56) and (57) agree on any point (a, a) , it follows that if self-replication is possible, the following equation must be satisfied on the diagonal for arbitrary $P_s(\phi_1)$:

$$0 = P_t \partial_{\phi_1} P_s \underbrace{[\partial_1 \phi_1 - \partial_2 \phi_1]}_{\alpha} + P_s \partial_{\phi_2} P_t \underbrace{[\partial_1 \phi_2 - \partial_2 \phi_2]}_{\beta}. \quad (58)$$

We now argue as follows: Consider a P_s that has zero slope at ϕ_1 . Then at that point, either $\partial_{\phi_2} P_t = 0$ or $\beta = 0$. Let us first deal with the non-trivial case, in which $M_s = M_t$ is either non-compact, or is compact but P_t is not a uniform distribution (i.e. does not have zero slope everywhere). In that case, we know that $\partial_{\phi_1} P_t$ cannot vanish at every point on the diagonal, and so we only have to choose a P_s that is zero at that point, in which case β must vanish there. But now that we know β vanishes at that point, we could equally well choose P_s non-zero there, in which case it is clear that α also vanishes.

The fact that there is a point on the diagonal at which $\alpha = \beta = 0$ has immediate consequences when combined with the requirement of symplectic invariance, i.e. (31). In particular, it means that we have

$$A^T J A + C^T J C = J \tag{59}$$

$$A^T J A + C^T J C = 0, \tag{60}$$

which is a contradiction. Thus self-replication is inconsistent with the requirement that ϕ be a symplectomorphism.

In the case where $M_s = M_t$ is compact and P_t is a uniform distribution then the process is trivially inconsistent, because we get from (55)

$$P_s(x)P_s(x) = (\text{const}) \cdot P_s(\phi_1(x)), \tag{61}$$

and we can always choose a P_s such that at some point x , $(P_s(x))^2$ is different from $(\text{const})P_s(\phi_1)$. Alternatively, by differentiating (61) one could generate an inconsistency by simply choosing a P_s such that at some point a , $P_s(a) \neq 0$, but $\partial_{\phi_1} P_s(\phi_1(a)) = 0$.

How is our argument related to the argument of DPP in [14], at least within our restricted setting of classical statistical *mechanics*? Recall that DPP show indirectly, i.e. via the Kullback-Leibler metric, that a cloning process is inconsistent with the Liouville equation

$$\frac{dP}{dt} = -\Omega(X_H, X_P), \tag{62}$$

for the evolution of statistical states (where X_H and X_P are the vector fields associated to the Hamiltonian function H and the probability distribution P respectively).⁴⁶ On the other hand, we have shown that a self-replication process (i.e. a cloning process with trivial machine system) is inconsistent with the preservation of the symplectic form. But since the Liouville equation assumes the preservation of the symplectic form, we have also shown that a cloning process is inconsistent with the Liouville equation. In other words, we have managed to reproduce a special case of DPP's results from a rather more basic principle, viz. the preservation of the symplectic form. Is it possible to use the Liouville equation to argue *directly* against

⁴⁶The association of vector fields to functions (i.e. 0-forms) is possible because of the non-degeneracy of the symplectic form Ω . For instance, in order to construct X_H , we just use the inverse symplectic form to raise the index of dH .

the possibility of cloning? At least in the case of *non-interacting* subsystems, this is certainly possible; we relegate the proof to Appendix A.

7. Future directions of research

In the foregoing sections, we have taken the reader on a guided tour of classical cloning: we began with some conceptual analysis and category-theoretic generalities, moved on to self-replication and cloning in the case of classical mechanics, and ended by considering self-replication in the case of classical statistical mechanics. However, as with any field as young and unexplored as this one, there is much work that remains to be done! We thus conclude this paper by discussing some unresolved issues and indicating future directions of investigation into the theme of classical no-cloning.

- Does classical self-replication admit of a category-theoretic treatment in a manner similar to quantum self-replication? In [27], we discuss this issue for the ‘toy model’ case of symplectic vector spaces, instead of symplectic manifolds, and generalize the notion of ‘state’ to cover Lagrangian submanifolds (these are in fact *semi-classical* states). However, it would be nice to have a categorical treatment that covered generic classical states, as well as more general symplectic manifolds.
- What is the relationship between the above no-self-replication argument and information-theoretic concepts such as entropy?⁴⁷ For instance, can entropy constraints be used to *derive* a restriction on classical no-self-replication? And if so, is there a category-theoretic method of formalizing information-theoretic axioms such that, when we specialize to the case of a ‘symplectic category’, we automatically obtain these restrictions on no-self-replication?
- Another line of investigation is to answer the following question: Even if perfect self-replication isn’t permitted in classical mechanics, how much non-perfect self-replication is permitted? For instance, in the case of quantum mechanics, Buzek and Hillery [10] were able to find non-perfect cloning machines that still clone states with a reasonably high degree of fidelity. Analogously, one might try to find bounds on classical self-replication with non-perfect fidelity. Furthermore, in the case of quantum mechanics, it is known that there is a tight conceptual relationship between ‘optimal (non-perfect) cloning’ and the physical principle known as ‘no-signalling’; it would thus be interesting to investigate whether an analogous correspondence holds in the case of classical mechanics.⁴⁸
- As we point out in §5, the (no-cloning component of the) CBH theorem [13] does not take into the account the Poisson/symplectic structure of classical phase space. It may be illuminating to reformulate

⁴⁷A partial answer to this question is given by the connection with relative entropy in the work of DPP [14].

⁴⁸See Section E1 of [35] for the details of the relationship between optimal cloning and no-signalling.

a version of the CBH theorem that takes this into account, and thus recover the above results from a C^* -algebraic point of view.

- Our results in this paper only apply to the case of (statistical and non-statistical) classical *mechanics*, but many classical systems are not ‘mechanical’ in this narrow sense. Thus, in order to establish classical no-cloning in full generality, we would have to extend our results to the case of classical field theory. One potential obstacle to this direction of generalization is that the notion of a symplectic manifold does not quite make sense for general field theories; instead, one has to consider what is known as a multi-symplectic manifold (for an accessible presentation of multisymplectic geometry in the context of field theory, see [23]). Our desired generalization would thus consist in redeploying our argument in the context of multisymplectic geometry.⁴⁹

Appendix A. No-cloning in non-interacting classical statistical mechanics

Following DPP, we shall for simplicity assume that the joint probability distribution of the initial state P of the composite system factorizes as $P = P_s P_t$, where P_s and P_t are the initial probability distributions of the source and target system respectively. A self-replicating system in classical statistical mechanics is a map that takes the initial state P to the final state Q , where Q is assumed to factorize as $Q = P_s P_s$.

The Liouville equation for the evolution of an arbitrary probability distribution F is:

$$\frac{d\rho}{dt} = -\Omega(X_H, X_\rho), \quad (\text{A.1})$$

where X_H and X_ρ are the vector fields associated to the Hamiltonian function H and the distribution ρ respectively (note that in this section, as in the last, we will not actually need to specify H).

Applying (A.1) to our system and integrating from the initial time t_0 to the final time t_1 , we obtain

$$Q - P = P_s(P_s - P_t) = \int_{t_0}^{t_1} -\Omega(X_H, X_F) dt. \quad (\text{A.2})$$

But then, by our earlier method of decomposing Ω into the pullbacks of its constituent symplectic forms ω (notice that X_H can really be written in terms of its components in TM_s and TM_t) and reapplying the Liouville equation to the result, we immediately see that the RHS of (A.2) is:

$$\int_{t_0}^{t_1} -\Omega(X_H, X_\rho) dt = (P_s - P_s) + (P_s - P_t). \quad (\text{A.3})$$

Comparing this with the LHS of (A.2) shows us that either $P_s = P_t$ or $P_s = 1$. Thus self-replication is either trivial (the identity map acting on two copies of the same distribution) or absurd.

⁴⁹One complication in doing so may be the fact that multi-symplectic manifolds do not in general have a Darboux Theorem – this topic requires further investigation.

As we stressed earlier, true cloning is more constrained than mere self-replication, because one needs to include a machine subsystem M_m as part of the composite system. We now consider whether there is a cloning process that takes the composite initial state $P = P_m P_s P_t$ to the composite final state $Q = Q_m P_s P_s$, where P_m is the initial state of the machine and Q_m is the final state of the machine.

Arguing in exactly the same way as before, we obtain the following consistency condition from the Liouville equation:

$$\begin{aligned} Q_m P_s P_s - P_m P_s P_t + (-P_m P_s P_s + P_m P_s P_s) \\ = (Q_m - P_m) + (P_s - P_t), \end{aligned}$$

which in turn gives:

$$(Q_m - P_m)P_s^2 + (P_s - P_t)P_m P_s = (Q_m - P_m) + (P_s - P_t). \quad (\text{A.4})$$

This simple equation gives us a surprising amount of insight into the possibility of statistical cloning. Let us first consider the case of trivial cloning, in which $P_s = P_t$. Here (A.4) implies that either $Q_m = P_m$ or $P_s = 1$. We can interpret this result as saying that either the machine system has trivial evolution, in which case one can also have non-trivial source/target states that undergo trivial evolution; or the machine state has non-trivial evolution, in which case the source/target states must be trivial. This makes good physical sense: in a system with overlapping distributions (of the subsystems), the non-trivial evolution of one distribution would of necessity interact with the others unless the others were trivial, i.e. identically equal to one.

Let us now consider the case where $P_s \neq P_t$. Here we see that since P_s^2 and $P_m P_s$ are both positive real-valued functions, (A.4) implies that $P_s^2 = 1$ and $P_m P_s = 1$. Thus, we have the result that $P_s = P_m = 1$. This tells us any substantive notion of cloning is impossible. Indeed, we can sharpen our statement slightly: we have shown that the only ‘cloning’ process compatible with the consistency conditions imposed by Liouville dynamics is the following:

$$1 \cdot 1 \cdot P_t \mapsto Q_m \cdot 1 \cdot 1. \quad (\text{A.5})$$

The fact that this process requires the initial source state to take a particular – in fact trivial – value is sufficient to rule it out as a substantive cloning process. It is a further curiosity that such processes constrain the initial machine state to be 1.

- [1] Abramsky, *No-cloning in categorical quantum mechanics*, arXiv:0910.2401v1 [quant-ph] (2009).
- [2] Abramsky and Coecke, *A categorical semantics of quantum protocols*, arXiv:quant-ph/0402130v5 (2007).
- [3] Benjamin Allen, *The category-theoretic arithmetic of information*, arXiv:0803.3608v3 [math.CT] (2008).

- [4] V.I. Arnold, *Mathematical methods of classical mechanics*, Springer, 1989.
- [5] J. Baez and M. Stay, *Physics, topology, logic and computation: A rosetta stone*, arXiv:0903.0340v3 (2009).
- [6] John Baez, *Quantum quandaries: A category-theoretic perspective*, Structural Foundations of Quantum Gravity (Steven French, Dean Rickles, and Juha Saatsi, eds.), OUP, 2006.
- [7] ———, *Classical mechanics notes, lecture 19*, <http://math.ucr.edu/home/baez/classical/classical19.pdf> (2008).
- [8] D. C. Brody and L. P. Hughston, *Geometric Quantum Mechanics*, J. Geom. Phys. **38** (2001), 19–53.
- [9] J. Bub, *Quantum entanglement and information*, The Stanford Encyclopedia of Philosophy.
- [10] V. Bužek and M. Hillery, *Quantum copying: Beyond the no-cloning theorem*, Phys. Rev. A **54** (1996), no. 3, 1844–1852.
- [11] Jeremy Butterfield, *On symmetry and conserved quantities in classical mechanics*, <http://philsci-archive.pitt.edu/id/eprint/2362> (2005).
- [12] ———, *Handbook of the philosophy of physics*, ch. On Symplectic Reduction in Classical Mechanics, North-Holland, 2007.
- [13] R. Clifton, H. Halvorson, and J. Bub, *Characterizing quantum theory in terms of information-theoretic constraints.*, Found. Phys. **33(11)** (2003), 1561.
- [14] Daffertshofer, Plastino, and Plastino, *Classical no-cloning theorem*, Physical Review Letters **88**, no. **21** (2002).
- [15] Deligne and Freed, *Classical field theory*, Quantum Fields and Strings: A Course for Mathematicians, AMS, 1999.
- [16] A. Doering, *Topos theory and neo-realist quantum theory*, arXiv:0712.4003v1 [quant-ph] (2007).
- [17] A. Fenyves, *Cloning in classical mechanics*, arXiv:1010.6103v1 [math-ph] (2010).
- [18] ———, *There’s no cloning in symplectic mechanics*, <http://math.ucr.edu/home/baez/dual/no-cloning.pdf> (2010).
- [19] C. A. Fuchs and C. M. Caves, *Quantum information: How much information is in a state vector?*, arXiv:quant-ph/9601025v1 (1996).

- [20] Christopher A. Fuchs, *Information gain vs. state disturbance in quantum theory*, Fortschritte Der Physik-progress of Physics **46**, 535–565.
- [21] Gian Carlo Ghirardi and T. Weber, Nuovo Cimento 78B **9** (1983).
- [22] G. W. Gibbons, *Typical states and density matrices*, J. Geom. Phys. **8** (1992), 147–162.
- [23] Gotay, Isenberg, Marsden, and Montgomery, *Momentum maps and classical relativistic fields. part i: Covariant field theory*, arXiv:physics/9801019v2 (2004).
- [24] Alexei Grinbaum, *Reconstruction of quantum theory*, Br J Philos Sci **58 (3)** (2007), 387–408.
- [25] Sergei Gukov and Edward Witten, *Branes and quantization.*, Adv. Theor. Math. Phys. **13** (2009), no. 5, 1445–1518 (English).
- [26] L. Hardy, *Disentangling nonlocality and teleportation*, arXiv:quant-ph/9906123v1 (1999).
- [27] C. Heunen and Author, *Classical cloning ii*, MS (2011).
- [28] Chris Heunen, Nicolaas Landsman, and Bas Spitters, *A topos for algebraic quantum theory*, Communications in Mathematical Physics **291** (2009), 63–110, 10.1007/s00220-009-0865-6.
- [29] K. Kirkpatrick, *‘quantal’ behavior in classical probability*, Found Phys Lett **16(3)** (2003), 199–224.
- [30] Maxim Kontsevich, *Deformation quantization of poisson manifolds*, Letters in Mathematical Physics **66** (2003), 157–216, 10.1023/B:MATH.0000027508.00421.bf.
- [31] N. Landsman, *Mathematical topics between classical and quantum mechanics*, Springer, 1998.
- [32] Marsden and Abraham, *Foundations of mechanics*, Benjamin/Cummings, 1978.
- [33] D. McDuff, *Introduction to symplectic topology*, OUP, 1998.
- [34] Peres, *Quantum theory: Concepts and methods*, Kluwer: Dordrecht, 1995.
- [35] Scarani, Iblisdir, Gisin, and Acin, *Quantum cloning*, Review of Modern Physics **77** (2005).
- [36] Schilling and Ashtekar, *Geometrical formulation of quantum mechanics*, arXiv:gr-qc/9706069v1 (1997).
- [37] Graeme Segal, *Qft notes*, Unpublished Manuscript.
- [38] R. Spekkens, *In defense of the epistemic view of quantum states*, Phys. Rev. A **75** (2007).
- [39] Leon Takhtajan, *Quantum mechanics for mathematicians*, AMS, 2008.

- [40] C. Timpson, *Quantum information theory and the foundations of quantum mechanics*, arXiv:quant-ph/0412063v1 (2004).
- [41] V. Vedral, *The role of relative entropy in quantum information theory*, Rev. Mod. Phys. **74** (2002), no. 1, 197–234.
- [42] Vlatko Vedral, *Introduction to quantum information science*, OUP, 2007.
- [43] J. Vicary, *A categorical framework for the quantum harmonic oscillator*, arXiv:0706.0711v2 [quant-ph] (2007).
- [44] Thomas A. Walker and Samuel L. Braunstein, *Classical broadcasting is possible with arbitrarily high fidelity and resolution*, Phys. Rev. Lett. **98** (2007), no. 8, 080501.
- [45] Wassermann, *Operator algebras and conformal field theory iii*, Invent. math. **133** (1998), 467–538.
- [46] Weinstein, *Symplectic categories*, arXiv:0911.4133v1 [math.SG] (2009).
- [47] E.P. Wigner, *The probability of the existence of a self-reproducing unit*, Symmetries and Reflections, Indiana University Press, Bloomington, 1967.
- [48] Wootters and Zurek, *A single quantum cannot be cloned*, Nature **299** (1982), 802.