arXiv:2209.07867v3 [quant-ph] 20 Sep 2026

Time symmetry in quantum theories and beyond

John H. Selby[1], Maria E. Stasinou[2,3], Stefano Gogioso[4,5], and Bob Coecke[6,7,8]

[1]International Centre for Theory of Quantum Technologies, University of Gdańsk, 80-308 Gdańsk, Poland
[2]Research Center 'ATHENA', Patras, Greece
[3]Quantum Neural Technologies S.A.
[4]Spooqy Ltd, London, UK
[5]University of Oxford, Oxford, UK
[6]Relational Intelligence Ltd, Oxford, UK
[7]Wolfson College, University of Oxford, Oxford, UK
[8]Perimeter Institute, Waterloo, Ontario, Canada

September 22, 2026

There is a stark tension among different formulations of quantum theory in that some are fundamentally time symmetric and others are radically time asymmetric. This tension is crisply captured when thinking of physical theories as theories of processes. We review process theories and their diagrammatic representation, and show how quantum theory can be described in this language. The tension between time symmetry and time asymmetry is then captured by the tension between two of the key structures that are used in this framework. On the one hand, the symmetry is captured by a *dagger structure*, which is represented by a reflection of diagrams. On the other hand, the asymmetry is captured by a condition involving *discarding*, which is ultimately responsible for the theory being compatible with relativistic causality. Other formulations of quantum theory go a step further than being time symmetric and are "time-neutral" – that is, they contain no notion of time at all. Such time neutrality can be captured by a third key process-theoretic structure – *cups and caps* – which ensures that there is no essential distinction between inputs and outputs to a process. As with the dagger, the tension between cups and caps and discarding provides a clean process-theoretic characterisation of the tension between time neutrality and time asymmetry.

Next we consider three different ways in which this tension can be resolved. The first of these is closely related to recent work by Lucien Hardy, in which the tension is resolved by adding a time-reversed version of discarding together with a suitable consistency condition. The second, to our knowledge, is a new approach. Here the tension is resolved by adding new systems that propagate backwards in time, and imposing a consistency condition to avoid running into well-known time-travel paradoxes. The final approach that we explore is closely related to work by Oreshkov and Cerf, in which the tension is resolved by removing the constraint associated with discarding. We show two equivalent ways in which this can be done whilst ensuring that the resulting theory still makes sensible operational predictions.

John H. Selby: john.h.selby@gmail.com
Maria E. Stasinou: stasinoumar@gmail.com
Bob Coecke: bob@rel-int.ai

Contents

1 Introduction **3**

2 Process theories **4**
2.1 Example 1: quantum physics . . . 5
2.2 Deterministic and causal process theories . . . 8
2.3 Time reversed process theories . . . 10
2.4 Time symmetric process theories . . . 11
2.5 Time neutral process theories . . . 12
2.6 Example 2: quantum calculations . . . 15
2.7 Process theories with dual systems . . . 16
2.8 Example 3: quantum symmetries . . . 18
2.9 Summary of approaches . . . 20

3 Two approaches to time symmetry **20**
3.1 Causal and retrocausal . . . 20
3.2 Causal or retrocausal . . . 23
3.3 Discussion . . . 26

4 Two equivalent approaches to time neutrality **28**
4.1 Modified composition rule . . . 29
4.2 Modified processes . . . 31
4.3 Equivalence of these approaches . . . 33
4.4 Restoring determinism . . . 34
4.5 Connections to higher-order processes . . . 36
4.6 Beyond quantum theory . . . 38
4.7 Discussion . . . 38

A Diagrams for groups and their representations **45**
A.1 Diagrams for groups . . . 45
A.2 Group representations . . . 47
A.3 Intertwiners . . . 49
A.4 Dual systems and conjugate representations . . . 49

1 Introduction

It is a commonly held belief that quantum theory, or indeed nature itself, is time symmetric. This claim, however, requires qualification – whilst the unitary dynamics of the theory may be time symmetric [1–3], the theory as a whole certainly is not [4–6]. In fact, from this broader perspective, we find that time symmetry is broken in a rather dramatic way. This is exemplified by the fact that while there are many pure states in the theory – corresponding to rays in a Hilbert space – there are no pure effects that can be realised without invoking post-selection, i.e., conditioning on the outcomes when performing measurements. Beyond the pure theory, we find that there is a single effect – discarding (i.e., the (partial) trace) – that can be implemented without post-selection whilst every density matrix corresponds to a preparable state. In process-theoretic terms, the root of this asymmetry is known as causality [7] or terminality [8], and as demonstrated in [9], the time reverse of quantum theory is a remarkably different theory: it has only a single state and describes a theory of eternal noise.

There have been several attempts in the literature to reformulate these time-asymmetric theories in an explicitly time-symmetric way, creating a stark tension with the aforementioned process-theoretic causality. There are many different motivations for doing so. Firstly, for philosophical reasons we may believe that nature should not have a preferred direction of time at a fundamental level, and hence quantum theory as a fundamental description of nature should be time symmetric. If this is the case, then the asymmetry that we observe in our experiments must be an emergent phenomenon, perhaps arising due to particular choices of boundary conditions for the universe [10]. Secondly, in our attempts to reconcile quantum theory and general relativity, there are hints that modifying the role of time in quantum theory will be essential. In particular, it may be the case that we cannot have a predetermined causal structure and so the 'past' and 'future' could be inextricably mixed [11]. If that is the case, formulating quantum theory in a time-neutral way (a stronger condition than time symmetry as we make precise in section 2) may be essential to making progress on the unification of quantum theory and general relativity. Thirdly, if we take the view that quantum theory is a generalisation of Bayesian probability theory, that is, that quantum states are epistemic rather than ontic, then it is natural to want to treat space and time on an equal footing [12]. This requires that we find a time neutral formulation of the theory and that the apparent time asymmetry arises from how we choose to apply the theory to physical scenarios rather than being fundamental.

Without necessarily endorsing any of these motivations, we consider the possible approaches to describing quantum theory in a time-symmetric or time-neutral way. We see that many existing works in this direction [10, 13–17] have a particularly concise process-theoretic description. Moreover, they all correspond to essentially the same process theory, differing only in convention and philosophical perspective rather than in their underlying mathematics. Essentially, these approaches can all be seen as particular quotients of a suitably chosen time-neutral theory.

We also see that there are various other ways to symmetrise quantum theory. We consider two such ways, one of which has close connections to the work in [18, 19] whilst the other is, to our knowledge, a new approach inspired by the view of antiparticles as being particles travelling back in time. Specifically, the former can be viewed as a further restriction of a causal theory to also satisfy a 'retrocausality' condition, while the latter can be viewed as an operational restriction on a time-neutral theory.

It is worth commenting further on the process-theoretic nature of our construction. In particular, it means that whilst we are going to focus on quantum theory as a concrete

example in this paper, these constructions can be applied more generally, for example, to arbitrary generalised probabilistic theories (GPTs) [20, 21], operational probabilistic theories (OPTs) [7, 22, 23], categorical probabilistic theories [24], or even theories that make no reference to probabilities at all (as have often been considered in the categorical quantum mechanics literature [25–28]). One interesting feature to highlight is that the symmetrised version of a theory will often be outside of the framework used to describe the original theory. For example, a time-symmetrised GPT is typically not itself a GPT. This showcases the utility of process theories as a flexible and widely applicable framework for exploring conceivable physical theories.

To summarise, the main contributions of this paper are threefold. First, we show that several existing approaches to time symmetry and time neutrality in quantum theory [10, 13–17, 29, 30] are mathematically equivalent[1], meaning that results and tools developed in any one of them transfer freely to all the others. Second, we introduce two new time-symmetric theories: $\mathbf{QPhys}|_{unital}$, which provides a precise process-theoretic realisation of a theory that is both causal and retrocausal, and $\mathbf{QPart}_G$, which gives a toy process-theoretic realisation of Feynman's intuition that antiparticles are particles travelling backwards in time, with non-signalling between particles and antiparticles playing the role of a consistency condition that prevents causal paradoxes. Third, we construct **QNeut**, a fully time-neutral formulation of quantum theory that contains all process matrices [31], connecting our framework to the study of indefinite causal structure.

The remainder of the paper is structured as follows. In section 2 we introduce the basics of the framework of process theories. In particular, we define the key concepts which underpin this work namely: causality, time symmetry, and time neutrality. We illustrate these abstract definitions with the concrete example of quantum theory. In section 3 we explore the two aforementioned approaches that one could take towards obtaining a time symmetric version of quantum theory. In section 4 we recast various other approaches to time symmetry and neutrality in two equivalent ways. The first is a direct adaptation of the existing literature whilst the second is a more elegant and broadly applicable way to capture the same idea. We prove their equivalence in theorem 4.4 and show that the theory can be made deterministic in theorem 4.8. We end this section by discussing connections to higher-order processes [31–35].

This paper is intended as a broad overview of the possibilities of studying time symmetry and neutrality from a process-theoretic perspective. Each approach deserves further study in its own right. As such, we will try to highlight various open problems and research directions throughout.

2 Process theories

A process theory [36–40] is comprised of a collection of *processes*. For example,

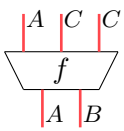

(1)

is a process with *input systems* A and B and *output systems* A, C, and C. Moreover, this collection of processes must be closed under forming *diagrams* by wiring processes

[1] Within their common scope of applicability.

together. For example, the diagram,

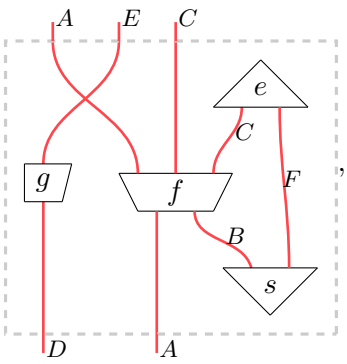

(2)

corresponds to another process in the theory with inputs D and A and outputs A, E, and C. This wiring is subject to the following constraints:

i. outputs cannot be wired to outputs, and inputs cannot be wired to inputs;
ii. the wiring is acyclic;
iii. systems can only be connected if their system types match.

Moreover, we have a notion of diagrammatic equality: two diagrams are equal if they represent the same wiring. For example,

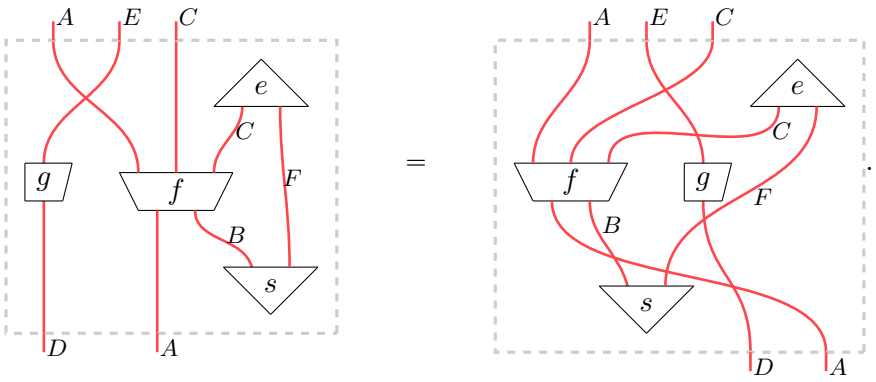

(3)

In other words, the precise layout of processes on the page is not important: all that matters is their *connectivity*.

Remark 2.1. In categorical quantum mechanics [25, 41], every symmetric monoidal category (SMC) is taken to be a process theory – subject to Mac Lane's strictification theorem [42]. Following [43], we also proposed certain operad algebras as process theories [44].

2.1 Example 1: quantum physics

To illustrate this abstract definition, we introduce a process theory that describes finite-dimensional quantum theory, **QPhys**. There are two different kinds of systems in **QPhys**: quantum systems, denoted by purple wires and labelled by finite-dimensional Hilbert spaces $\mathcal{H}$, and classical systems, denoted by grey wires and labelled by finite sets, $\mathbb{A}$. The quantum systems are the fundamental systems of interest within quantum theory, whilst the classical systems represent how we interact with the quantum world. For example, they represent the control variables on experimental devices and the pointers on

measurement apparatus that encode the outcome of the measurement. This is the setting previously considered in [24, 39, 45–47], among other works.

Following these conventions, a general quantum instrument, $\mathcal{E}$, is denoted by

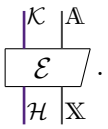

(4)

This has a quantum system, $\mathcal{H}$, together with a classical control variable, $\mathbb{X}$, as inputs and a quantum system, $\mathcal{K}$, together with a classical outcome variable, $\mathbb{A}$, as outputs. Formally, this can be understood as a completely positive trace-preserving (CPTP) map between complex matrix algebras,

$$\mathcal{E} : \mathcal{B}[\mathcal{H}] \otimes \left(\bigoplus_{x \in \mathbb{X}} \mathcal{B}[\mathbb{C}] \right) \to \mathcal{B}[\mathcal{K}] \otimes \left(\bigoplus_{a \in \mathbb{A}} \mathcal{B}[\mathbb{C}] \right). \tag{5}$$

This general definition looks somewhat complicated, but it can be well understood by considering more familiar special cases.

If there is no classical input or output, then we formally treat this as the singleton set $\star := \{*\}$. In this case we have that $\mathcal{B}[\mathcal{H}] \otimes (\bigoplus_{* \in \star} \mathcal{B}[\mathbb{C}]) \cong \mathcal{B}[\mathcal{H}]$, the standard noncommutative matrix algebra associated with a quantum system $\mathcal{H}$. In contrast, if there is no quantum input or output, then we formally treat this as the one-dimensional Hilbert space $\mathbb{C}$. In this case we have that $\mathcal{B}[\mathbb{C}] \otimes (\bigoplus_{x \in \mathbb{X}} \mathcal{B}[\mathbb{C}]) \cong \bigoplus_{x \in \mathbb{X}} \mathcal{B}[\mathbb{C}]$, a commutative matrix algebra commonly associated with a classical system.

A process with no inputs and only a quantum output,

(6)

therefore corresponds to a CPTP map, $\rho : \mathcal{B}[\mathbb{C}] \to \mathcal{B}[\mathcal{H}]$. These are in one-to-one correspondence with positive elements of $\mathcal{B}[\mathcal{H}]$ of unit trace, i.e., quantum states.

A process with a quantum input and a classical output,

(7)

corresponds to a CPTP map, $M : \mathcal{B}[\mathcal{H}] \to \bigoplus_{a \in \mathbb{A}} \mathcal{B}[\mathbb{C}]$. These are in one-to-one correspondence with sets of positive operators in $\mathcal{B}[\mathcal{H}]$ indexed by $a \in \mathbb{A}$, $\{M_a\}_{a \in \mathbb{A}}$, such that $\sum_{a \in \mathbb{A}} M_a = \mathbb{1}_{\mathcal{H}}$. In other words, they are in one-to-one correspondence with destructive POVM measurements.

When we compose a quantum state with a quantum measurement, we obtain a process which has only classical outputs,

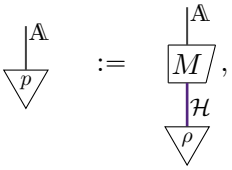

(8)

corresponding to a CPTP map $p : \mathcal{B}[\mathbb{C}] \to \bigoplus_{a\in\mathbb{A}} \mathcal{B}[\mathbb{C}]$. This is simply the sequential composition of the CPTP maps ρ and M, i.e., $p = M \circ \rho$. Processes of this form are in one-to-one correspondence with probability distributions over the set $\mathbb{A}$ and, in particular, with the probability distribution defined by $p(a) = \mathsf{tr}(M_a\rho)$ for all $a \in \mathbb{A}$. Hence, we see that the Born rule is encoded as a special case of how CPTP maps compose in sequence.

More generally, a process with only classical inputs and outputs,

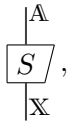

(9)

corresponds to a CPTP map, $S : \bigoplus_{x\in\mathbb{X}} \mathcal{B}[\mathbb{C}] \to \bigoplus_{a\in\mathbb{A}} \mathcal{B}[\mathbb{C}]$. These are in one-to-one correspondence with stochastic maps from $\mathbb{X}$ to $\mathbb{A}$. They therefore map probability distributions over $\mathbb{X}$ to probability distributions over $\mathbb{A}$ via

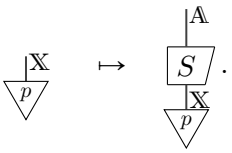

(10)

Some other special cases which are common in the literature are

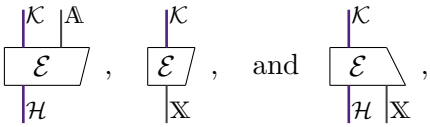

(11)

which describe non-destructive measurements, classically controlled state preparations, and classically controlled CPTP maps, respectively.

Finally, let us consider the set of processes with no outputs,

(12)

They correspond to CPTP maps $E : \mathcal{B}[\mathcal{H}] \otimes (\bigoplus_{x\in\mathbb{X}} \mathcal{B}[\mathbb{C}]) \to \mathcal{B}[\mathbb{C}]$. Note, however, that due to the trace-preservation condition, these processes are unique. On the quantum system they correspond to the (partial) trace whilst on the classical system they correspond to marginalisation. Since they are unique, we introduce a special symbol to denote them:

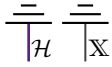

(13)

A special case of the above process is when there is also no input system. Then, there is a unique process with neither input nor output which corresponds to the unique CPTP map from $\mathcal{B}[\mathbb{C}]$ to itself. This can be thought of as the scalar 1 which maps, for example, $\rho \mapsto 1 \cdot \rho = \rho$. We diagrammatically denote the scalar 1 by the empty diagram,

.
(14)

The process-theoretic presentation of quantum theory is a time-asymmetric theory: every system has a large number of states, i.e., density matrices, but only a single effect, i.e., the (partial) trace. Indeed, as demonstrated in an earlier paper, [9], the time reverse of **QPhys** is a remarkably different theory: it has only a single state and describes a theory of eternal noise. We abstractly characterise theories that have a trace-preservation property in section 2.2, and formalise the notion of time symmetry which they (typically) violate in section 2.4.

2.2 Deterministic and causal process theories

The fact that quantum theory has a single scalar, namely the empty diagram, means that **QPhys** is a deterministic process theory:

Definition 2.2 (Determinism)**.** We say that a process theory which has only a single scalar, the empty diagram

 , (15)

is a deterministic process theory.

The reason for this terminology is that such theories describe only processes that can be realised with certainty. That is, they do not capture processes that occur only as one possibility amongst many.

For example, in **QPhys** determinism means that we must consider measurements as channels from a quantum to a classical system, as this is something that we can choose to implement in the lab. One can also think of measurements as being represented by a collection of POVM elements. The individual POVM elements, however, are not part of **QPhys** as they are not things that we can choose to implement – we have no control over which possibility will occur. This should be contrasted with **QCalc**, as defined in section 2.6, which includes (amongst other things) individual POVM elements as effects, leading to scalars in $[0, 1]$.

The trace-preservation condition for quantum processes in **QPhys** can also be elegantly expressed in process-theoretic terms. In short, it says that there is a unique way to go from something to nothing. More formally, for any system A there is a unique process with A as an input with no outputs. These are often known as *discarding* maps as they often correspond to simply throwing away or ignoring a system. As in the case of **QPhys**, we denote these by:

. (16)

Definition 2.3 (Causality [7, 8, 48])**.** We say that any process theory which, for each system A, has a unique process from A to nothing is a *causal* process theory.

Many important conditions immediately follow for any causal process theory:

i. Discarding a composite system is the same as discarding the components:

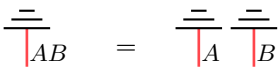 (17)

This means that our theory has a well-defined subsystem structure, because, for example, it provides a unique way to define local states given a global state.

ii. Discarding the outputs of any process is the same as discarding the inputs:

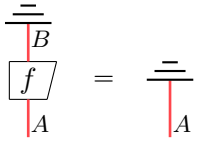 (18)

This is known as *discard preservation* and is the abstract characterisation of trace preservation of quantum theory.

iii. There is a unique scalar, and hence the theory is *deterministic*. This follows by considering the special case of the (unique) discarding effect with a trivial input.

Whilst we have just seen that causality implies determinism, the converse is not necessarily true. For instance, the opposite category $\mathbf{QPhys}^{\mathrm{op}}$[2] has a single scalar 1 but multiple effects for each system. Moreover, in section 4.4 we define a process theory that is deterministic but not causal.

The terminology 'causality' is used because the above condition ensures that the theory is non-signalling between causally separated regions [8] and, more generally, is compatible with relativistic causal structure [49]. For example, suppose we have two spacelike-separated parties who therefore cannot directly signal to one another. They may, however, have a common past and future in which their light cones intersect. This would be represented by a diagram of the shape:

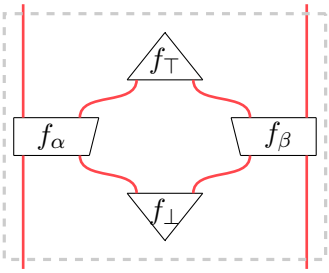

(19)

Then, to describe the local physics from the perspective of the left-hand party, we discard the right-hand output as it is inaccessible to them:

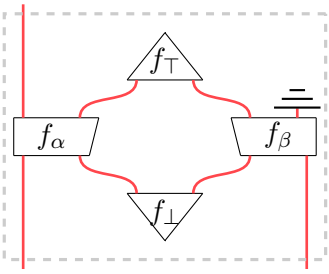

(20)

Now, using the fact that the only effect is the discarding map, and that f_β satisfies Eq. (18), we find that:

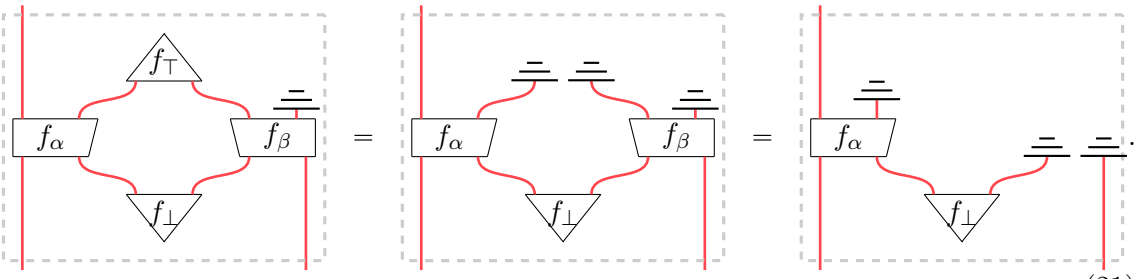

(21)

That is, the right-hand input is simply discarded and is entirely disconnected from the left-hand side. In particular, any control that the right-hand party may have had over the input cannot have any influence over what the left-hand party could observe at their output.

Remark 2.4. In more category-theoretic language, a causal process theory can be thought of as a symmetric monoidal category in which the monoidal unit is terminal [8].

[2]Denoted by $\mathbf{QPhys}_R$ in section 2.3.

2.3 Time reversed process theories

To understand time symmetry we must first consider what it means to time reverse a process theory. That is, we must understand what transformation we are demanding the theory to be invariant under.

For any process theory, **Proc**, there is a very simple way to reverse the arrow of time. We simply read the diagrams in the theory from top to bottom rather than bottom to top. The inputs of a process are then considered outputs and the outputs are considered inputs. It is often, however, more convenient to take an active view of time reversal – that is, we keep the bottom-to-top reading of the diagrams but we flip all of the diagrams upside down. This defines a new process theory, which we denote by $\mathbf{Proc}_R$. This has the same systems and processes as **Proc**, but all of these processes go in the opposite direction. That is,

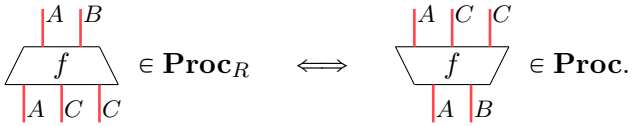

(22)

Moreover, composition in $\mathbf{Proc}_R$ can be defined as the time reverse of the composition of the original processes in **Proc**. This means, for example, that

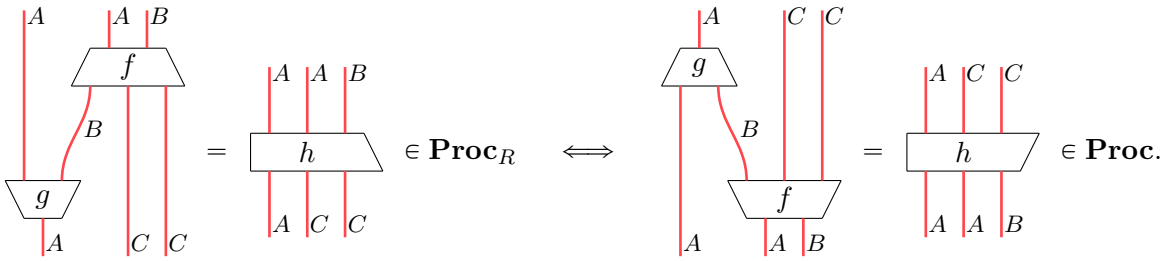

(23)

It is straightforward to see that time-reversing a theory twice leaves it invariant. That is, $\mathbf{Proc}_{RR} = \mathbf{Proc}$.

Remark 2.5. Categorically, time reversal is simply the standard contravariant functor from a category to the opposite category $R : \mathbf{Proc} \to \mathbf{Proc}_R$.

If the process theory, **Proc**, that we start with is a causal process theory, then, as noted in [9], the time-reversed theory will, typically, be remarkably different. It describes a theory with a single state for every system, which is left invariant by every transformation – it is a theory of eternal noise.

Definition 2.6 (Retrocausality)**.** We say that any process theory which, for each system A, has a unique process from nothing to A is a *retrocausal* process theory.

In this somewhat less poetic language, the result of [9] states that the time reverse of a causal theory is a retrocausal theory. In particular, in any retrocausal theory there will be a unique state for each system, which we denote by:

(24)

This, in many theories, can be thought of as a state of uniform noise. Uniqueness of this state implies that they compose as:

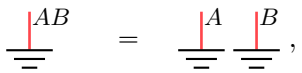

(25)

and, moreover, that every other process will satisfy a *retrocausality* constraint:

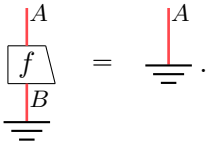

(26)

Finally, note that, like causal process theories, retrocausal process theories are necessarily deterministic (in the sense of Def. 2.2).

2.4 Time symmetric process theories

Refs. [10, 13–17] formulate quantum theory in a time-symmetric way. From a process-theoretic perspective, this means that the process theory is the same as the time-reversed theory. Equivalently, this means that time-reversal is *internal* to the process theory. Formally, we capture this by saying that a process theory is time symmetric if and only if it permits a *dagger* [25, 50, 51].

Definition 2.7 (Dagger)**.** A dagger is a map, †, from the process theory to itself that *reflects diagrams*. Specifically, it acts on processes as

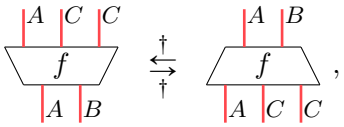

(27)

and, moreover, on diagrams as

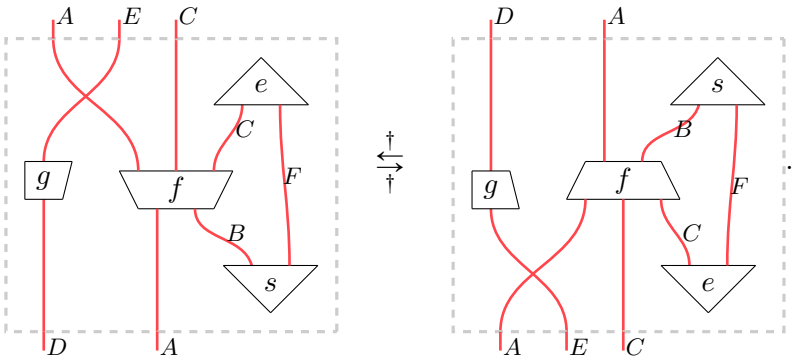

(28)

To understand the connection to time symmetry, note that if a process theory has a dagger, then for every process $f : A \to B$ in the theory there exists a process $f^\dagger : B \to A$ (that is, the symbolic notation for the upside-down f). The process, $f^\dagger$, can be interpreted as the time-reversed version of the process, f. Moreover, it is such that wiring together the time-reversed processes in the reverse order is the same as taking the time reverse of the process representing the diagram. Hence, we say that process theories with daggers are time symmetric.

Note that this is not the same as demanding that every process in the theory is time symmetric[3]. Instead, we are imposing the time symmetry at the level of the theory as a whole. Indeed, it is in this sense that the theories of [10, 13–17, 29, 30] can be considered to be time symmetric.

[3]That would be the case if we additionally required that $f = f^\dagger$ and hence that there were no processes whose inputs differed from their outputs.

Remark 2.8. If the process theory represents an SMC, then any involutive contravariant strong monoidal endofunctor $\dagger : \mathbf{Proc} \to \mathbf{Proc}$, which acts as the identity on objects, provides a dagger for the process theory. This allows us to see that the existence of a dagger implies that **Proc** and $\mathbf{Proc}_R$ are covariantly isomorphic as we can define the covariant isomorphism simply by $R \circ \dagger : \mathbf{Proc} \to \mathbf{Proc}_R$, which is covariant as $\dagger$ and R are each contravariant. Note that **Proc** and $\mathbf{Proc}_R$ are, by definition, contravariantly isomorphic, so the covariance here is the key to having a time symmetric theory.

Proposition 2.9. [9, Thm. 3] A time symmetric theory is causal if and only if it is retrocausal.

This is because there is a unique state for a system given by the dagger of the discarding map:

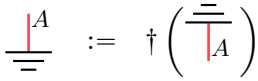

(29)

It is therefore immediately clear, as noted in [9], that **QPhys** is not a time symmetric theory – it is causal but not retrocausal. To see this, it suffices to note that there are multiple distinct states for any (non-trivial) system, e.g., the computational basis states of a qubit:

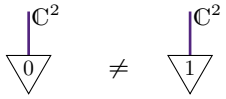

(30)

2.5 Time neutral process theories

In the work of [15], the authors aim to go a step beyond the time symmetry that we have just described and instead want to have a version of quantum theory that is *time neutral*. In process-theoretic terms, this means that we want to forget about the distinction between inputs and outputs entirely. This is possible if the theory has *cups & caps* [25, 52].

Definition 2.10 (Cups & caps)**.** A *cup* is a process, $\bigcup$, in a process theory that can be represented as a bent piece of wire,

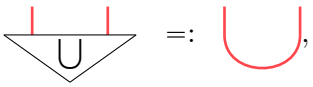

(31)

which allows us to connect inputs to inputs. Similarly, a *cap* is a process, $\bigcap$, which can be represented as a wire bent in the other direction,

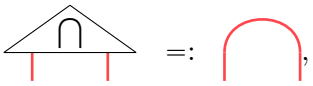

(32)

which allows us to connect outputs to outputs. For our representation of these processes as bent wires to actually make sense, that is, for it to be compatible with the idea that only the connectivity of diagrams matters, these processes must satisfy the following conditions:

(33)

Remark 2.11. There are other conditions that one would expect these two to satisfy such as:

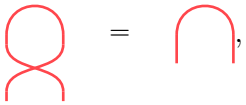

(34)

However, these conditions can be derived from the two above and therefore need not be imposed separately.

Time neutrality goes a step further than time symmetry as it allows us to freely interchange some of the inputs with outputs and vice versa. This is in contrast to time symmetry which requires flipping all of them at once. What this means is that time-neutral theories allow a freer notion of wiring that neglects the input-output structure. For example, the following diagram is permissible:

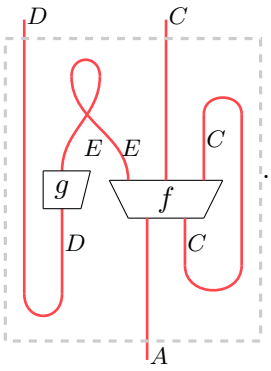

(35)

Theories with cups & caps are therefore said to be time neutral as there is no real distinction between the inputs and outputs of a process[4]. Moreover, there is no meaningful causal order that can be assigned to the processes within general diagrams because this freer notion of wiring allows for cycles. For example, in the diagram,

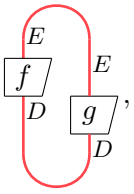

(36)

we have the situation in which f is both in the 'causal future' and the 'causal past' of g.

This process-theoretic notion of a time-neutral theory is stronger than the notion of a time-symmetric theory. That is, any time neutral theory is also necessarily time symmetric, as we can define a dagger using cups & caps,

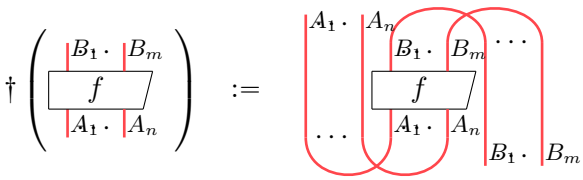

(37)

It can easily be seen that this does indeed satisfy the constraints of Def. 2.7 [38].

[4]Recall from the definition of a process theory in Section 2 that the only distinction between the inputs and the outputs of a process is that outputs cannot be wired to outputs and inputs cannot be wired to inputs, and so cups & caps allow us to circumvent this restriction.

Proposition 2.12. If a time-neutral theory is causal, then there can be only a single process between any two systems.

Proof. As time neutrality implies time symmetry, we immediately have that causality implies retrocausality. Hence, for every system we have a unique state and effect. In particular, this means that,

(38)

Equation (33) then implies that:

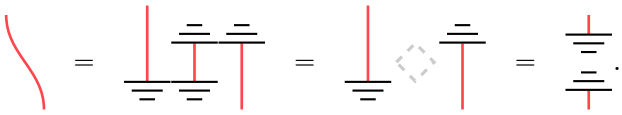

(39)

From the above we find that for any f we have:

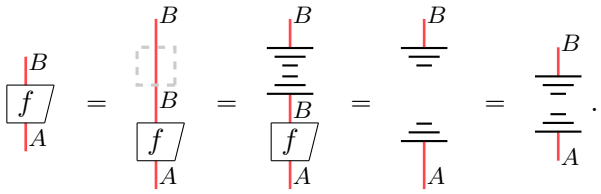

(40)

As this holds for any process f, we have therefore shown that there is a unique process per pair of systems in the theory, namely:

(41)

□

It is therefore clear that the process theory **QPhys** is not a time-neutral theory, since there are multiple distinct processes with the same inputs and outputs. For example:

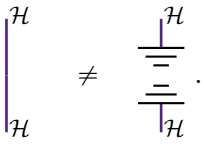

(42)

Remark 2.13. Categorically, cups & caps correspond to the unit and counit in a compact closed category in which the objects are equal to their dual.

One could argue that process theories with cups & caps are not truly time neutral as the individual processes still have a distinction between input and output systems, even if we can now freely interchange them. In [44] we show how to construct fully time-neutral process theories, where there is no distinction between inputs and outputs even at the level of a single process, and then show that these are in fact equivalent to process theories with cups & caps.

2.6 Example 2: quantum calculations

We have seen that our first key example, **QPhys**, is a causal process theory that fails to be either time neutral or time symmetric. There is, however, a closely related theory **QCalc**, which is a supertheory of **QPhys** that is both time symmetric and time neutral. This is the theory in which we often perform calculations about quantum physics – for example, when computing the probability of a measurement outcome.

The systems in **QCalc** are the same as those of **QPhys** and so we will use the same diagrammatic notation for them as we did in **QPhys**. The entire difference between the two theories is a simple modification to the definition of the processes. That is, in **QPhys** processes were defined as CPTP maps, whilst in **QCalc** we drop the trace-preservation condition and allow for arbitrary completely positive maps. In particular, the 'states' of quantum systems, $\mathcal{H}$, are given by arbitrary positive operators and so need not have unit trace, and the 'states' of classical systems, $\mathbb{X}$, are arbitrary $\mathbb{R}^+$-valued functions on $\mathbb{X}$ rather than probability distributions over $\mathbb{X}$.

Notably, this theory is neither a causal nor a deterministic process theory. Indeed, there are many processes which have a system $\mathcal{H}$ as an input and no output,

. (43)

These correspond to CP maps $e : \mathcal{B}[\mathcal{H}] \to \mathcal{B}[\mathbb{C}]$, which can be thought of as positive linear functionals on the cone of positive operators on $\mathcal{H}$. These include, but are not limited to, the quantum effects $\mathsf{tr}(\rho\ \cdot\ )$. Similarly, many processes have a system $\mathbb{X}$ as an input and no output,

. (44)

These correspond to CP maps $r : \bigoplus_{x\in\mathbb{X}} \mathcal{B}[\mathbb{C}] \to \mathcal{B}[\mathbb{C}]$, which can be thought of as positive linear functionals on the cone of $\mathbb{R}^+$-valued functions over $\mathbb{X}$.

We leveraged the causality of **QPhys** to demonstrate that it is neither time neutral nor time symmetric. Because **QCalc** is free of this causality condition, it is both time-symmetric and time-neutral. That is, we will show that it has a dagger as well as cups & caps.

A suitable dagger for **QCalc** is given by the Hermitian adjoint, $\dagger_H$. There are other daggers, such as the transpose; however, the Hermitian adjoint is special because it has the property that, when applied to a state, it defines an effect that tests for that state [39, 53]. That is,

$$\dagger_H :: \rho \mapsto \mathsf{tr}(\rho\ \cdot\ ). \tag{45}$$

Moreover, it inverts reversible dynamics: for any unitary supermap $\mathcal{U}$,

$$\dagger_H :: \mathcal{U} \mapsto \mathcal{U}^{-1}. \tag{46}$$

These two conditions mean that the Hermitian adjoint is well suited to be interpreted as time-reversal for **QCalc**.

We can now turn to cups & caps. In the case of quantum systems, $\mathcal{H}$, these can be expressed as

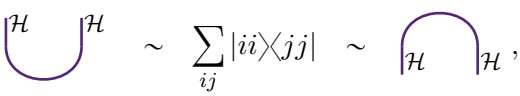

, (47)

for some basis $|i\rangle \in \mathcal{H}$. That is, the cup is a supernormalised version of the Bell state and the cap is a supernormalised Bell effect. For classical systems, $\mathbb{A}$, we express the cup and cap as

(48)

That is, they are the supernormalised perfectly correlated probability distribution and response function, respectively.

Note that the dagger given by the Hermitian adjoint and cups & caps interact in the way that one would expect, namely,

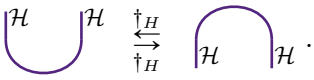

(49)

This implies that we can indeed view the Hermitian adjoint as a reflection of the diagrams.

Many of the processes within this theory have a clear physical interpretation. The trace-preserving processes have the clearest interpretation, since they are exactly the processes in **QPhys**; beyond these, the trace-nonincreasing CP maps can also be interpreted straightforwardly. These can be viewed as processes that occur in some branch of a causal process, that is, processes that only occur as one possibility amongst many.

One may have hoped that one could simply add trace-nonincreasing processes to **QPhys** to obtain a time symmetric theory. Ultimately, however, this is not enough and we must also include some trace-increasing CP maps. For example, the dagger of the discarding effect is a supernormalised maximally mixed state, and hence is a trace-increasing CP map. If we were to try to give these trace-increasing processes a physical interpretation, then we would run into problems. In particular, the theory cannot be given a sensible operational interpretation; for example, it permits 'probabilities' that are greater than 1. A simple example of this comes from composing the cups & caps themselves. In particular, we find that:

(50)

As scalars that are greater than one do not have any physical interpretation, **QCalc** cannot be a good description of nature. Nonetheless, as we mentioned earlier, it is extremely useful as a theory in which we perform calculations relevant to quantum physics. For example, if we want to compute the probability of some measurement outcome given a state, then we can simply compose the associated effect with the state:

(51)

This necessarily results in a sensible probability if the state and effect are both trace-nonincreasing.

2.7 Process theories with dual systems

In the previous subsections, we have assumed that systems are invariant under time reversal. While this is a natural assumption to make, it is also interesting to consider situations

in which this is not the case. For example, in particle physics it is often taken to be the case that time-reversed particles are in fact antiparticles. Indeed, in section 3.2 we define a new kind of time-symmetric quantum theory that captures this idea.

To investigate such situations, we introduce time-directed arrows to our systems,

(52)

and say that these systems are *dual* to one another. We symbolically denote these by $A^{\uparrow}$ and $A^{\downarrow}$, respectively.

Remark 2.14. At this point, categorically, all we are doing is introducing a special kind of relabelling for the systems in the theory. This can be obtained by first defining an involutive function r on the set of atomic systems (i.e., those with no nontrivial decomposition in terms of $\otimes$), using it to pair atomic systems as $(A, r(A))$, and then assigning each pair arbitrarily to $(A^{\uparrow}, A^{\downarrow})$. Note that trivial pairings with $A = r(A)$ are also possible, in which case we would not introduce an arrow. Most choices of involutive function r will not lead to anything particularly meaningful. Meaningful choices are those that allow for a dagger or cups and caps, as we now define.

It is then clear that when we time reverse such a process theory, that is, by reading the diagrams from top to bottom rather than bottom to top, systems will get mapped to their duals. For example,

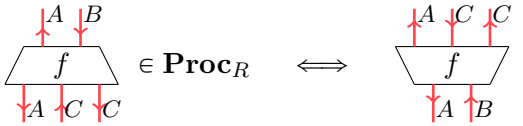

(53)

This means that we must also suitably refine the notion of the dagger, and, hence, the condition for time symmetry for such theories.

Definition 2.15 (Daggers for process theories with duals)**.** A dagger is a map, †, from the process theory to itself that *reflects diagrams*. In particular, it acts on processes as

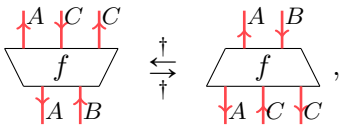

(54)

and, moreover, on diagrams as

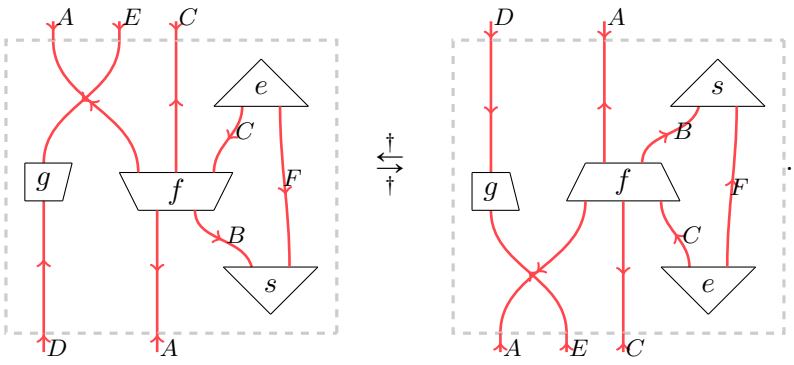

(55)

We can then say that a process theory with duals is time symmetric if and only if it has a dagger.

The definition of cups & caps can also be suitably modified in this context to the bent wires

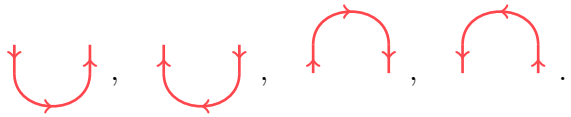

(56)

These bent wires satisfy the obvious diagrammatic equations. A process theory with duals is then said to be time neutral if and only if it has such cups & caps.

2.8 Example 3: quantum symmetries

An example of a category with duals is the category of representations of a group G within quantum theory, denoted by $\mathbf{QRep}_G$. Its objects are representations of G on objects of **QCalc**, and its morphisms are the special processes in **QCalc** known as intertwiners. An extensive discussion of this category in the context of categorical quantum dynamics can be found in [54, 55]. This process theory was previously considered in the context of resource theories of asymmetry as a way to define the free set of processes [56, Example 3.9]. The perspective we take here, however, is that this is a mixed state version of the ideas presented in [57, Penrose (1971)]. These two perspectives are, however, equivalent to one another thanks to a covariant version of Stinespring's dilation theorem [58, 59].

In section 3.2 we motivate $\mathbf{QRep}_G$ physically by creating a toy model for particle physics, in which a particle is identified with a representation of G on a quantum system $\mathcal{B}(\mathcal{H})$. Generally speaking, a particle is defined in the literature as a particular irreducible representation of the Poincaré group, which is infinite-dimensional. In our case, however, we construct a toy model by assuming that we have only finite-dimensional representations of some other group G. However, it is known that the approach also generalises to (certain) infinite groups [55].

This toy model has the potential to shed light on particle physics since it provides a neat description of the interaction of particles with classical systems. For instance, it allows for a clear definition of measurements within particle physics. In addition, it provides a passage from the standard pure-state formulation to a mixed-state formulation of particle physics, for example, via the CPM construction [51].

We now define the category $\mathbf{QRep}_G$ using a diagrammatic notation for groups and their representations that we describe in App. A.

Definition 2.16. Consider a group G. The category $\mathbf{QRep}_G$ consists of the following data:

- Objects are pairs,

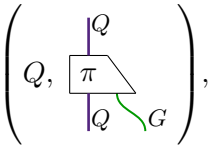

(57)

  where Q is an object in **QCalc** and π_Q is a causal representation (see Eq. (141)) of G on Q.

- The monoidal product of (Q, π_Q) and $(Q', \pi'_{Q'})$ is given by

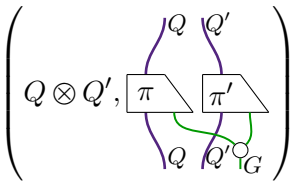

(58)

- The monoidal unit is the pair $(\mathbb{C}, \pi_{\mathbb{C}})$, where $\pi_{\mathbb{C}}$ is the trivial representation:

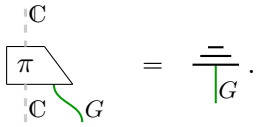

(59)

- Morphisms from (Q, π_Q) to $(Q', \pi'_{Q'})$ are intertwiners in **QCalc**, that is, they are CP maps, $\mathcal{E}$, satisfying the covariance condition:

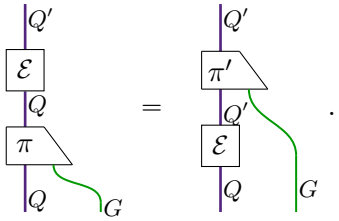

(60)

- Composition is inherited from **QCalc**; the composite of intertwiners is again an intertwiner.
- The identity is the identity process which can also be seen to be an intertwiner.
- Finally, we take each system and its conjugate, (Q, π_Q) and (Q^*, π^*_{Q*}), to form a dual pair, for example,

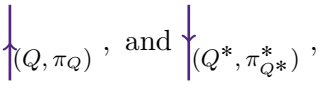

(61)

  where π^* is the conjugate representation defined in Eq. (151). Note that the choice of which is ↑ and which is ↓ is arbitrary as conjugation is involutive.

Note that one can show that cups & caps are intertwiners (Eq. (152)) and hence $\mathbf{QRep}_G$ is a time neutral process theory with duals. Like **QCalc**, however, $\mathbf{QRep}_G$ cannot be directly interpreted as a theory of physics – that is, it does not necessarily make sensible probabilistic predictions. We therefore need to find a physicality condition on the processes in $\mathbf{QRep}_G$ akin to the restriction of **QCalc** to **QPhys** via the causality condition. In section 3.2, we propose an implementation that preserves the time symmetry of $\mathbf{QRep}_G$.

2.9 Summary of approaches

We have now seen that causality in **QPhys** serves as the main obstacle towards time symmetry and time neutrality. In particular, time symmetry together with causality implies that there should be a single state per system, and time neutrality together with causality imply that there should be a single transformation between any pair of systems, neither of which is true within **QPhys**. However, the time symmetric and time neutral theories that we have introduced, **QCalc** and $\mathbf{QRep}_G$, both have the problem that they do not have a coherent operational interpretation, as they have scalars that are greater than 1.

Having identified the root of the time asymmetry within quantum theory, we can then ask what we can do to obtain a symmetric theory. We have identified several possible approaches:

1. We can restrict **QPhys** to a subtheory that additionally satisfies the retrocausality constraint – that is, every system is both causal and retrocausal. This is related to the work in Refs. [19] and [18].

2. We work within $\mathbf{QRep}_G$, a theory built on **QCalc** which is the supertheory of **QPhys**. $\mathbf{QRep}_G$'s systems come in dual pairs, and we restrict the theory to processes satisfying non-signalling between causal and retrocausal systems (Eq. (80)). This gives a theory, $\mathbf{QPart}_G$, in which every system is either causal or retrocausal, serving as a toy model of particle physics. Crucially, this is not a direct extension of **QPhys** but rather a restriction of the time-neutral theory $\mathbf{QRep}_G$.

3. We can start with the supertheory of **QPhys**, **QCalc**, and then, to avoid unphysical predictions, we can:

   (a) Modify the composition rule. This is closely related to the works of Refs. [10, 13–17, 29, 30].

   (b) Modify the processes. This is equivalent to (a) but is a more elegant and adaptable presentation of the theory.

The first two options can lead to a time symmetric theory (but not a time neutral theory) and are formally introduced in the following section, while the third can lead to a time neutral theory which we formally introduce in section 4.

3 Two approaches to time symmetry

3.1 Causal and retrocausal

Given that causality is the main obstacle towards time symmetry within quantum theory, perhaps the most obvious way to symmetrise the theory is to restrict the processes to those that also satisfy a retrocausality condition.

Definition 3.1 (Bicausality)**.** A process theory that is both causal (Def. 2.3) and retrocausal (Def. 2.6) is said to be *bicausal*[5]. In particular, this means that it has a unique effect and a unique state for each system.

[5]Known as double causality in [19].

In such theories, the simple argument against time symmetry of causal theories (namely, that there are more states than effects) breaks down, and so it is at least plausible that bicausal theories can be time symmetric. It is therefore interesting to explore how to construct bicausal theories out of causal theories.

Given any causal process theory, we can construct a subtheory that additionally satisfies a retrocausality constraint. To do so, however, we must pick a particular state for each system that we will then require be the unique state for the system. Typically there will not be a unique way to choose these states, but they nonetheless must satisfy certain consistency conditions for the resulting theory to be well defined. Let us denote these candidate states by:

(62)

These must be chosen such that they satisfy

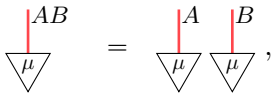

(63)

in order to ensure that the resulting theory is closed under parallel composition. Note that, because they belong to a causal theory, they also automatically satisfy

(64)

We can then restrict the allowed processes within the theory to those that satisfy:

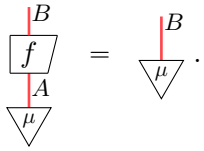

(65)

Having restricted the theory in this way, we can note that the remaining subtheory satisfies the retrocausality constraint by taking

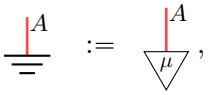

(66)

for all systems A.

For a general causal process theory and a general choice of these states, there is no reason to believe that this will result in a time symmetric theory. That is, even after ensuring that we only consider processes that are both causal and retrocausal, it may still not be possible to define a dagger. In the case of quantum theory, we show shortly that for a suitable choice of states we do end up with a theory with a dagger. However, there are other generalised probabilistic theories (e.g., those that do not satisfy any notion of self-duality [60]) in which this is not possible, at least not without imposing further constraints on the sets of processes.

Returning to the case of **QPhys**, the natural choice to make for the unique state is $\mu_{\mathcal{H}} := \frac{1}{|\mathcal{H}|}\mathbb{1}_{\mathcal{H}}$, that is, the maximally mixed state for the system:

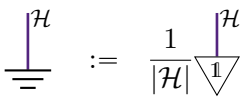

(67)

(Note that the RHS of this equation is actually a diagram living in **QCalc** rather than **QPhys** as it involves a supernormalised state $\mathbb{1}_{\mathcal{H}}$.) It is simple to verify that these indeed satisfy the compositionality condition of Eq. (63). The constraint that is then imposed on the processes of **QPhys** in order to define the subtheory is

(68)

This means that $\mathcal{E}$ maps the maximally mixed state to the maximally mixed state, i.e., that $\mathcal{E}$ is a *unital* CPTP map. In the special case that the inputs and outputs are classical we find that:

(69)

or, in other words, they are bistochastic maps[6]. We will denote the subtheory of unital CPTP maps $\mathbf{QPhys}|_{unital}$.

One may then be tempted to define the dagger using the Hermitian adjoint, $\dagger_H$, as we did for **QCalc**. Note, however, that the Hermitian adjoint maps the discarding map to the supernormalised state, $\mathbb{1}_{\mathcal{H}}$, rather than the unique state that we have defined as, $\frac{1}{|\mathcal{H}|}\mathbb{1}_{\mathcal{H}}$. To take care of these normalisation issues we must therefore define the dagger as

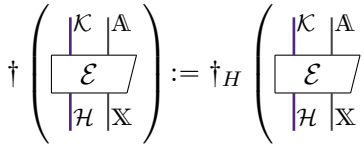

$$\frac{|\mathcal{K}||\mathbb{A}|}{|\mathcal{H}||\mathbb{X}|}. \tag{70}$$

It is straightforward to show that this defines a dagger: in particular, it maps CP maps to CP maps and trace-preserving maps to unital maps, and vice versa. In other words:

Proposition 3.2. $\mathbf{QPhys}|_{unital}$ is a time symmetric process theory.

On the face of it, however, $\mathbf{QPhys}|_{unital}$ does not seem to be a good candidate to describe our world: there is only a single state for every system. The theory is not entirely trivial, however, as it still has interesting transformations. For example, it still contains unitary evolution and, moreover, makes non-trivial classical 'predictions' in the form of the bistochastic matrices. It is not straightforward, however, to see how these bistochastic matrices suffice to explain our everyday experiences. For example, they do not allow us to copy classical information which is, intuitively, an operation that we would expect to be able to do.

Despite these obstacles, this sort of approach is closely related to the approach advocated in Refs. [19] and [18]. The idea is that whilst they may not describe our everyday experiences in the way that we might typically formulate them, they can still be recovered by suitably choosing which classical random variables we condition on. That is, the asymmetry arises as a choice of the direction in which we choose to make inferences, rather than being part of the underlying physics. It is therefore an interesting question to ask how the formalisms of Refs. [19] and [18] (in particular, the idea of choosing which variables to condition on) can be captured within a process-theoretic context. We leave this as an important direction for future research.

[6] Bistochastic maps are typically taken to be square matrices, but this is the natural notion that also applies to the non-square case [19].

3.2 Causal or retrocausal

In contrast to the previous section, in which we imposed both causality and retrocausality on every system, we now formulate a theory in which each system satisfies either causality or retrocausality while the theory remains time-symmetric.

To do so, we work with process theories with duals, where we view a system $A^{\uparrow}$ as a causal system and $A^{\downarrow}$ as its retrocausal counterpart. To enforce this interpretation, we demand that systems $A^{\uparrow}$ have a unique effect and that systems $A^{\downarrow}$ have a unique state. Hence, in particular, processes $A^{\uparrow} \to B^{\uparrow}$ necessarily satisfy the causality condition (Def. 2.3), whereas those from $A^{\downarrow} \to B^{\downarrow}$ satisfy the retrocausality condition (Def. 2.6).

A general process in such a theory, however, has both causal and retrocausal inputs and outputs. Diagrammatically, this is denoted by:

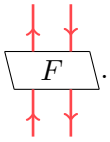

(71)

The important question is: what type of "causality" condition should this process satisfy so that the (retro)causality conditions hold for the systems $A^{\uparrow}$ and $A^{\downarrow}$ individually?

To get to grips with this condition, we work within our example category $\mathbf{QRep}_G$. In this case, systems $A^{\uparrow}$ correspond to pairs (Q, π_Q) and systems $A^{\downarrow}$ to dual pairs $(Q^*, \pi^*_{Q^*})$. This correspondence can be thought of as a toy model for particle physics: particles are defined as the (causal) pairs (Q, π_Q) while antiparticles are defined as (retrocausal) pairs $(Q^*, \pi^*_{Q^*})$. In a sense, this takes seriously Feynman's interpretation of antiparticles as being particles travelling back in time. The flexibility of our construction allows particles to interact with classical systems. Thus, it can be viewed as a way to include the concept of measurement in particle physics. Furthermore, it allows for particles to be in mixed states rather than pure as in the standard formulation. We could also define this directly for **QCalc** using the duality between Q and Q^*. Indeed, this can be seen as the special case in which π_Q is the trivial representation. However, we work with $\mathbf{QRep}_G$ to make these tentative connections to particle physics.

As mentioned above, if we only have causal inputs and outputs then the process should be causal:

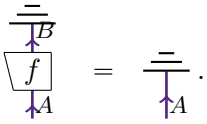

(72)

Similarly, the processes with only retrocausal inputs and outputs should be retrocausal processes:

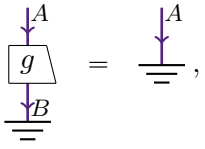

(73)

where, in this case, we take the unique state of the retrocausal system to be the supernormalised maximally mixed state $\mathbb{1}$.

We will now show, however, that imposing these conditions alone does not result in a theory that is closed under composition; we then identify precisely what more is needed to obtain a consistent theory.

To begin, note that Eq. (72) tells us that there is a unique effect for a causal system (the discard) and that Eq. (73) tells us that there is a unique state for a retrocausal

system (the retro-discard). Now, if we have a bipartite effect with both a causal input and a retrocausal input, and we compose it with the retro-discard on the retrocausal system, then we have an effect on the causal system, which, as just argued, must be the discard. Similarly, if we have a bipartite state with both a causal and a retrocausal output, and we compose it with the discard on the causal system, then we are left with a state of the retrocausal system which, as just argued, must be the retro-discard. That is, we find that:

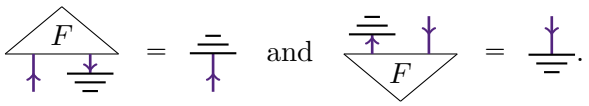

(74)

Cups and caps provide examples of such processes, since they satisfy:

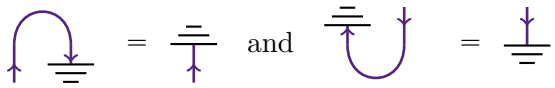

(75)

This example, however, indicates the need for constraints on states and effects beyond those in Eq. (74) as we can compose these cups & caps to achieve violations of Eqs. (72) and (73). Specifically, as noted in Eq. (50), if we compose a cup with a cap then we end up with a scalar other than the empty diagram:

(76)

Composing either of these with a causal process, f, on a causal system will then give us a process that violates the causality condition, i.e.,

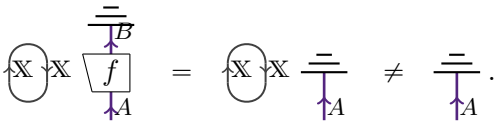

(77)

We have therefore seen that a straightforward application of Eqs. (72) and (73) does not lead to a theory that is closed under composition. This example, however, does give us a hint as to what necessary and sufficient conditions must be imposed on general processes to ensure that Eqs. (72) and (73) always hold.

Specifically, note that the classical cup in this context can be viewed as perfect signalling from a retrocausal to a causal system, whilst the classical cap can be viewed as perfect signalling from a causal to a retrocausal system:

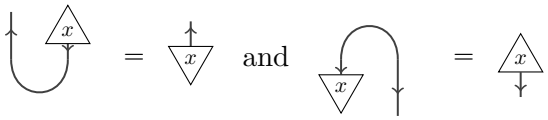

(78)

We can therefore conjecture that it is the signalling between the causal and retrocausal systems that leads us to problems. This is perhaps not surprising as many well-known paradoxes arise from the closed time loops that we could construct via:

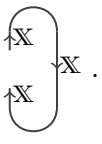

(79)

To prevent such loops from occurring, we must therefore impose non-signalling conditions on our processes. Formally, what this means is that every process F must satisfy both

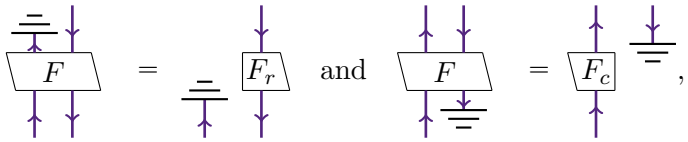

(80)

for some retrocausal process, F_r, and some causal process, F_c. We recover Eqs. (72) and (73) as special cases by considering only causal (or retrocausal) inputs and outputs. Moreover, these non-signalling conditions imply that a general process in $\mathbf{QRep}_G$ must satisfy:

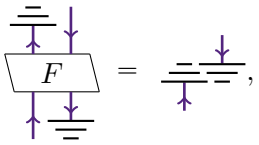

(81)

which rules out interactions that directly send information from a particle to an antiparticle as it implies that:

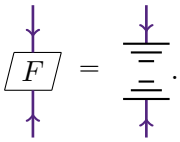

(82)

Note also that only one scalar satisfies this constraint:

(83)

which implies that the theory is deterministic.

Finally, we can be sure that imposing this constraint leads to a consistent theory as it is straightforward to see that this condition is closed under composition. For example, if F and G are both non-signalling, then so is their sequential composite:

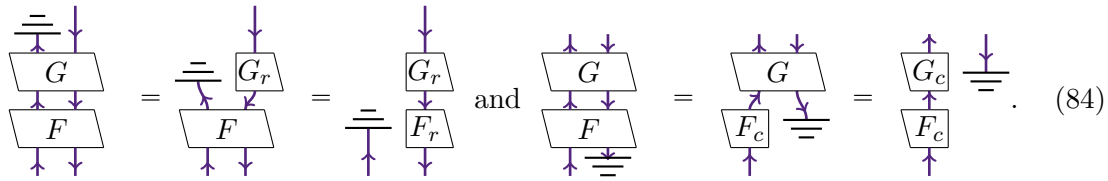

(84)

We have therefore identified the constraint that our toy model of particle physics – with particles represented by finite-dimensional representations of a group G and their antiparticle counterparts by conjugate representations – must satisfy to be time-symmetric while making sensible operational predictions. This leads us to the following definition:

Definition 3.3. We define the process theory $\mathbf{QPart}_G$ as the restriction of $\mathbf{QRep}_G$ to the subtheory of processes satisfying non-signalling from particles to antiparticles and vice versa (i.e., Eq. (80)).

Note that whilst Def. 3.3 imposes that there is non-signalling between causal and retrocausal systems in $\mathbf{QPart}_G$, it does not mean that the interactions between these types of systems are necessarily trivial. In particular, it is well known in quantum and classical theory that there can be processes that do not permit signalling but that are not of product form, e.g., PR-boxes [61].

The dagger for $\mathbf{QRep}_G$ is also a dagger for $\mathbf{QPart}_G$. Hence:

Proposition 3.4. $\mathbf{QPart}_G$ is a time symmetric process theory.

Returning now to general process theories, we take the non-signalling conditions (from causal to retrocausal and vice versa) as the definition of a well-behaved theory with causal and retrocausal systems:

Definition 3.5 (Dual-causal)**.** A process theory with dual systems is said to be dual-causal if it is non-signalling from causal to retrocausal and vice versa, namely, if

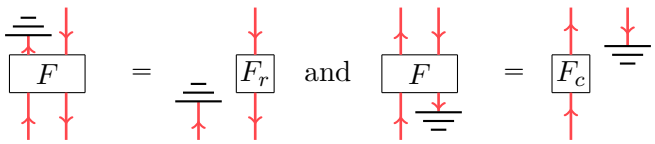

(85)

for all processes F.

Given any causal process theory, there is a straightforward way to construct such a time-symmetric theory by forbidding any interactions between the causal and retrocausal systems. In particular, we demand that:

$\forall F\ \exists F_c$ and F_r such that

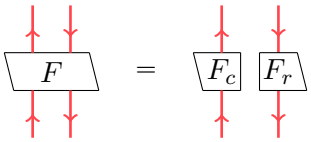

(86)

where F_c belongs to the original causal process theory **Proc** and F_r is from the associated retrocausal theory $\mathbf{Proc}_R$. We can then define a dagger by simply using the time reversal map R as

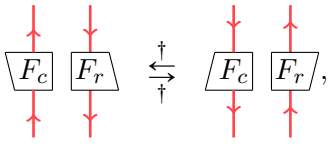

(87)

where now $F_c \in \mathbf{Proc}_R$ is a retrocausal process and $F_r \in \mathbf{Proc}$ is a causal process.

Remark 3.6. Categorically, what we are defining is simply $\mathbf{Proc} \times \mathbf{Proc}_R$ and noting that we can use the R functor to define a dagger as $\mathbf{Proc} \times \mathbf{Proc}_R \cong \mathbf{Proc}_R \times \mathbf{Proc}$.

This, however, does not seem like a particularly useful or insightful theory. For it to be of interest, we need non-trivial interactions between the causal and retrocausal systems. Hence, an important direction for further research is to ask whether there is a generic construction that takes a causal process theory to a dual-causal process theory such that there are non-trivial interactions between causal and retrocausal systems.

3.3 Discussion

In this section, we have seen two possible ways to try to construct time-symmetric theories, with a focus on quantum theory as our key example. Our third approach, which can moreover lead to time-neutral theories, will be discussed in the next section. All of these examples, however, show the utility of the process-theoretic approach to this problem. It provides an intuitive definition of time symmetry (in terms of the existence of a dagger) and hence a clear target to aim for when developing time-symmetric theories. Moreover, it provides a flexible framework in which such theories can be developed and compared.

The most important directions for future work stemming from these approaches are:

- To explore the connections between our first approach, i.e., $\mathbf{QPhys}|_{unital}$, and Refs. [19] and [18]. In particular, these connections should be formalised process-theoretically so that they are straightforward to identify.

- To explore whether there are physical consequences to our second approach, i.e., $\mathbf{QPart}_G$, in particular, whether it can be generalised to include physically relevant infinite-dimensional representations, such as those of the Poincaré group.

Another interesting question is this: if $\mathbf{QPart}_G$ were a faithful representation of nature, how could we distinguish it from **QPhys**? Is there some experimental test that could shed light on which is the correct theory?

If we were agents composed only of causal systems, the non-signalling conditions would always prevent us from directly probing the retrocausal systems. That is, some mechanism akin to spontaneous symmetry breaking can be seen to occur – we have a symmetric theory, but any instantiation of the theory will necessarily appear asymmetric!

Direct observation may therefore be impossible; however, there remains the possibility that the existence of these retrocausal counterparts could be probed more indirectly. An analogy can be drawn with spontaneous collapse theories – the spontaneous collapse cannot be directly observed, but its presence could manifest as an additional decoherence term which has no alternative explanation [62]. Similarly, if retrocausal systems were interacting with us, then their presence could be detected by otherwise inexplicable decoherence in experiments.

A more direct answer could come from the study of higher-order processes. The framework of higher-order processes can be viewed as another attempt to modify the notion of causality and causal structure, in particular, with the motivation of studying (quantumly) indefinite causal structure. The main tools in this study are higher-order processes, particularly those known as *process matrices* [31]. These are defined as maps, W, from pairs of channels to channels:

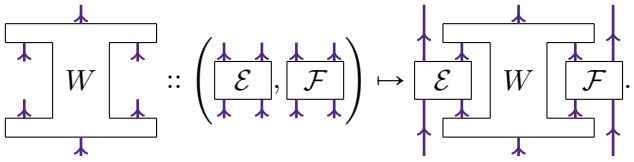

(88)

We know that quantum circuits can only realise a subset of logically possible process matrices, that is, those that can be written as [31]:

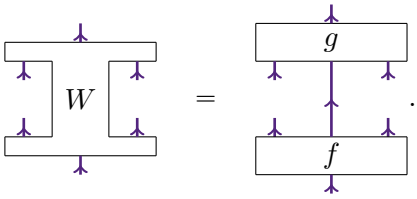

(89)

Now suppose that the theory also contains retrocausal systems. This would provide access to circuit realisable process matrices of the form:

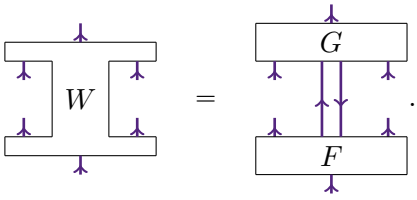

(90)

This leads us to our first open problem:

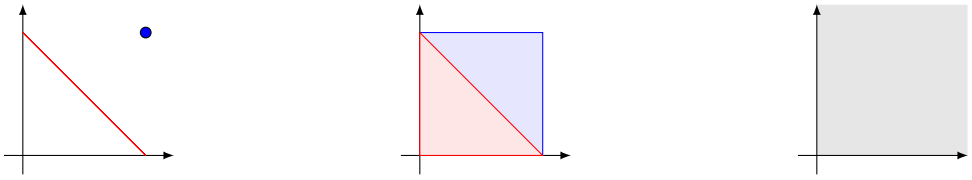

(a) The trace-preserving states (red line segment) and effects (blue point) for a bit. There is a clear asymmetry here as we can see that there is a unique effect, corresponding to marginalisation, whilst there are many states, corresponding to probability distributions over a two-element set.

(b) The trace-nonincreasing states (red triangle) and effects (blue square) for a bit. We can again see that there is an asymmetry, this time with more effects than states.

(c) Without any trace condition, the states and effects both form the same convex cone (grey octant), so there is no obstacle to time symmetry here. However, we can now obtain 'probabilities' greater than 1 by composing states and effects and so this theory does not make sensible predictions.

Figure 1: The states and effects for a classical bit.

Open Problem 1. Do there exist process matrices that are circuit realisable in $\mathbf{QPart}_G$ but not in **QPhys**?

4 Two equivalent approaches to time neutrality

The above approaches may lead to interesting time symmetric theories. However, they do not address time neutrality. The only time neutral theory that we have encountered so far is the process theory **QCalc**. Nevertheless, we noted that this is not a good physical theory as it does not always make valid probabilistic predictions. This section, therefore, asks whether it is possible to adapt **QCalc** such that it remains time neutral and at the same time makes sensible probabilistic predictions.

One candidate approach would be to restrict to trace-nonincreasing maps, as this has more processes than **QPhys** (which assumed trace preservation) but fewer than **QCalc** (which does not impose a trace condition at all). Although this may seem promising, it quickly fails, as for any (non-trivial) system the set of subnormalised states is not isomorphic to the set of effects in the theory. See Fig. 1 for a simple example in the case of a classical bit.

This has led to various other approaches in the literature that have, directly or indirectly, proposed time-neutral formulations of quantum theory. For example, see Refs. [10, 13–17, 29, 30]. Whilst there are many philosophical differences between them, all of these are – up to choices of convention – captured by the formalism we present below.

These approaches are typically presented as a modification of the measurement postulate of quantum theory. They begin by extending the set of measurements to allow for measurements, M, in which the POVM elements do not necessarily sum to the identity:

$$M = \{\{M_a\}_{a\in\mathbb{A}} | M_a \geqslant 0\}. \tag{91}$$

They then ensure that the probabilistic predictions are sensible via a suitable modification of the Born rule. That is, they are defined so that, if the "probability distribution" over measurement outcomes predicted by the standard quantum formalism is not normalised, i.e., if

$$N_M(\rho) := \sum_{a\in\mathbb{A}} \mathsf{tr}(M_a\rho) \neq 1, \tag{92}$$

then this distribution is renormalised by dividing by $N_M(\rho)$. This gives a new rule for computing probabilities:

$$\text{Prob}(a|\rho) := \frac{1}{N_M(\rho)}\mathsf{tr}(M_a\rho) \tag{93}$$

which differs from the standard Born rule whenever $N_M(\rho) \neq 1$. This new probability rule ensures that the theory makes valid probabilistic predictions. (At least, aside from the case of $N_M(\rho) = 0$ which is treated as a special case which we return to later.)

Simply modifying the probability rule is not, however, particularly satisfying from a process-theoretic perspective, since it is not manifestly compositional. What we show in this section is that we can recover this modified Born rule in a manifestly compositional way. We do so via the construction of two (equivalent) new process theories which naturally capture this modified Born rule in the appropriate situations.

The first process theory that we construct to achieve this can be defined by starting from **QCalc** and modifying composition so that it reduces to the modified Born rule when a state is composed with a measurement. The second process theory that we construct achieves this in a simpler, more elegant, and more adaptable way by defining an appropriate quotient of **QCalc**. We will then show how this process-theoretic viewpoint naturally leads to a new way to deal with probability-zero events, leading to a deterministic time neutral theory. Finally, we investigate the relationship between this theory and higher-order quantum processes.

4.1 Modified composition rule

Within **QCalc** the modified Born rule can be presented as:

Standard rule 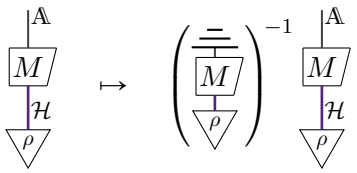 Modified rule (94)

for some CP map M from a quantum to a classical system and some quantum state ρ. Provided that this normalisation factor is non-zero, this defines a valid probability distribution over $\mathbb{A}$.

As mentioned above, we want to define a process theory with the same processes and systems as **QCalc** but in which composition is redefined. In particular, it needs to be such that when we compose a state ρ with a measurement M the renormalisation "automatically" occurs. More formally, we want to define a new composition rule $\bullet$ such that

$$M \bullet \rho := \frac{1}{\mathsf{tr}(M \circ \rho)} M \circ \rho. \tag{95}$$

Here, $\circ$ is the composition rule in **QCalc**. We will therefore denote the new process theory by $\mathbf{QCalc}^\bullet$.

To obtain a consistent process theory, however, we cannot simply redefine composition for these special cases. Instead, we must redefine composition in general, and obtain the

above as a particular instance of the new rule. Explicitly, we define $\bullet$ via:

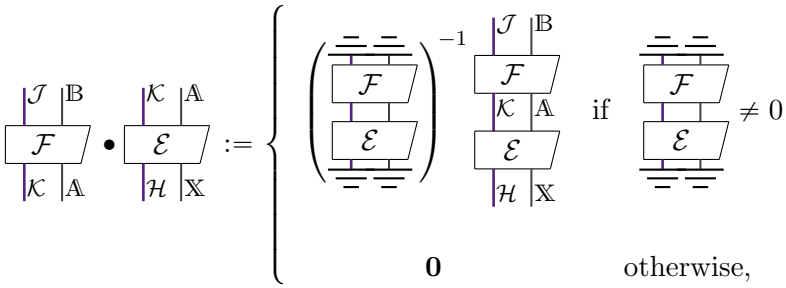

(96)

where now we are explicitly dealing with the probability-zero case. This is the same rule introduced in Refs. [10, 15] and used in Ref. [63] in the context of the P-CTC model for closed timelike curves. As an instance of this rule, we have:

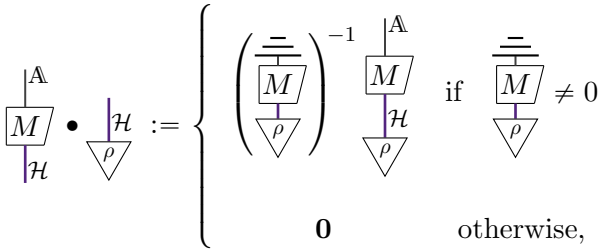

(97)

which reproduces the probabilities that we expect from the modified Born rule used in these time-neutral theories.

For this to define a valid process theory, various conditions must be satisfied. In particular, one can show by direct computation that $\bullet$ is associative and that it interacts suitably with parallel composition. The more interesting case, however, comes from considering the identity processes. In any process theory it must be the case that $\mathbb{1}_B \circ f = f = f \circ \mathbb{1}_A$ for every process $f : A \to B$. In our case, however, if we try to impose this condition with $\bullet$ we find that:

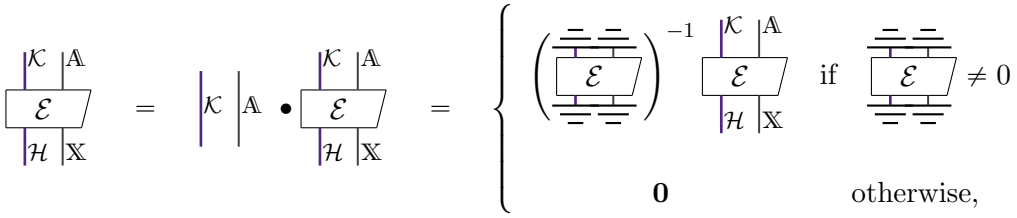

(98)

which only holds in the special cases that

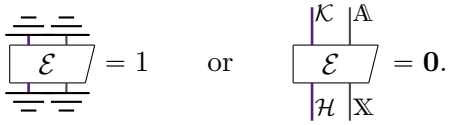

(99)

Therefore, to define the process theory **QCalc**$^\bullet$, we must both modify the composition rule *and* restrict the set of allowed processes to the above special cases. For example, for a classical bit, we obtain the states and effects as shown in Fig. 2.

Definition 4.1 (**QCalc**$^\bullet$)**.** This is the process theory which has the same objects as **QCalc**, and, as processes, the subset of processes in **QCalc** satisfying Eq. (99). Sequential composition is defined as in Eq. (96) whilst parallel composition is as in **QCalc**.

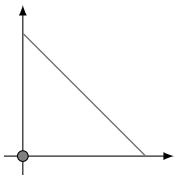

Figure 2: The states and effects for a classical bit in $\mathbf{QCalc}^\bullet$ are the same and therefore do not present an obstruction to time symmetry in this theory. They include the origin and the space of probability distributions. Note that this means that the state and effect spaces are non-convex.

Having modified **QCalc** in such a way as to obtain $\mathbf{QCalc}^\bullet$, it is no longer immediately clear that time neutrality or even time symmetry of **QCalc** have been preserved. It turns out that this is the case, but we will return to this in the following section once we have the more elegant characterisation of this process theory.

4.2 Modified processes

The construction of $\mathbf{QCalc}^\bullet$ within the previous section via a modified composition rule (and restriction of the allowed processes) can be more elegantly captured by a particular quotienting of **QCalc**. Specifically, we define a new theory in which the processes correspond to particular equivalence classes of processes in **QCalc**.

Specifically, we define two processes to be equivalent if they are equal up to a non-zero scalar:

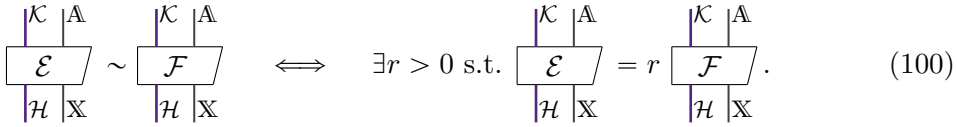

(100)

We denote the equivalence class of a process $\mathcal{E}$ by $\tilde{\mathcal{E}}$. It is then straightforward to note that we can construct the quotient process theory $\mathbf{QCalc}/\!\sim$. This has the same systems as **QCalc** but has processes that correspond to equivalence classes of processes in **QCalc** under the above equivalence relation. Composition within $\mathbf{QCalc}/\!\sim$ can then be defined, for instance, as

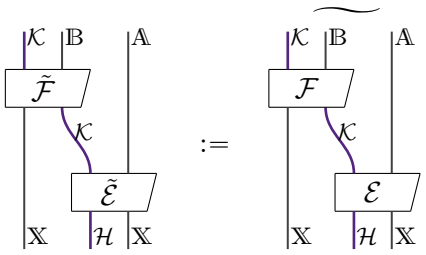

(101)

where $\mathcal{E}$ is an arbitrary element of $\tilde{\mathcal{E}}$ and $\mathcal{F}$ is an arbitrary element of $\tilde{\mathcal{F}}$. It is straightforward to verify that this is well defined as this definition is independent of the choice of these elements.

Note that the scalars of $\mathbf{QCalc}/\!\sim$ are severely restricted: the scalars in **QCalc** are $\mathbb{R}^+$ whilst in $\mathbf{QCalc}/\!\sim$ they are equivalent to $\mathbb{Z}_2$ since we have two equivalence classes $\{\tilde{0}, \tilde{1}\}$, where $\tilde{0} = \{0\}$ and $\tilde{1} = (0, \infty)$. This may at first glance seem problematic as we still want our time symmetric theory to make probabilistic predictions, which within **QCalc** are encoded in the scalars. The resolution to this is (as we saw with **QPhys** which is a causal theory and hence has only a single scalar) that the probabilistic predictions of a

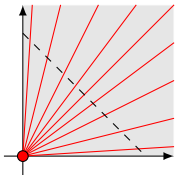

Figure 3: A schematic depiction of the states in **QCalc**/$\sim$ for a classical bit. The red dot at the origin represents the zero state whilst the red rays going out (which do not include the origin) are examples of equivalence classes. The union of these equivalence classes fills the grey shaded cone. The dashed line indicates normalised probability distributions. Note that each of the rays intersects the dashed line at a single unique point, hence can be unambiguously identified with this probability distribution.

theory are encoded into classical states

(102)

rather than into individual scalars. We can therefore ask what predictions can be made by the theory **QCalc**/$\sim$, i.e., what are the processes of the form:

?

(103)

It is not hard to see that these will be in one-to-one correspondence with probability distributions with one extra classical state left over, namely, the zero-state:

.

(104)

See Fig. 3 for the example where $\mathbb{X}$ is a two-element set. This theory, therefore, makes predictions that can be interpreted probabilistically most of the time: We simply view the equivalence class containing a probability distribution p as describing the same prediction as the probability distribution itself.

The exception to this interpretation is the zero state for which there is no obvious interpretation as a probability distribution. If we want the process theory to represent physical processes, then this serves as a challenge as it is difficult to give a physical interpretation of a zero-process, which, by definition, is something that does not occur. In other words, the existence of the zero state means that the process theory is not deterministic. We return to this shortly. Before that, however, we will show that this theory is time-symmetric and time-neutral:

Proposition 4.2. The Hermitian adjoint $\dagger_H$ is a dagger for **QCalc**/$\sim$.

Proof. It is simple to see that $\dagger_H$ preserves equivalence classes as $\dagger_H(r\mathcal{F}) = r\dagger_H(\mathcal{F})$ for all processes $\mathcal{F}$ and scalars r. □

Proposition 4.3. **QCalc**/$\sim$ is time neutral.

Proof. The equivalence class containing the cup and the equivalence class containing the cap will define cups & caps for **QCalc**/$\sim$. These satisfy Eq. (33), as follows immediately from the definition of composition of equivalence classes in Eq. (101). □

4.3 Equivalence of these approaches

We now compare the process theory **QCalc**/∼ (that we have defined via quotienting) to the process theory **QCalc**$^\bullet$ (that we defined in the previous subsection via a modification of the composition rule). In particular, we show that these are equivalent process theories.

To begin, note that any process in **QCalc** is in the same equivalence class as a process satisfying one of the conditions in Eq. (99), and, moreover, that each equivalence class will contain a unique such element. The crux of this is the following property of processes in **QCalc**,

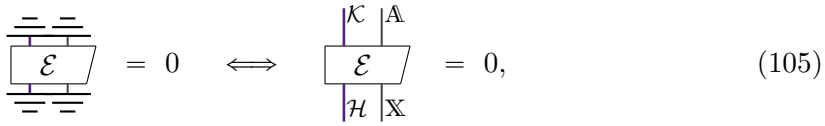

(105)

which follows from the fact that **QCalc** is locally tomographic, and that the trace and maximally mixed states are internal to the cone of effects and states respectively:

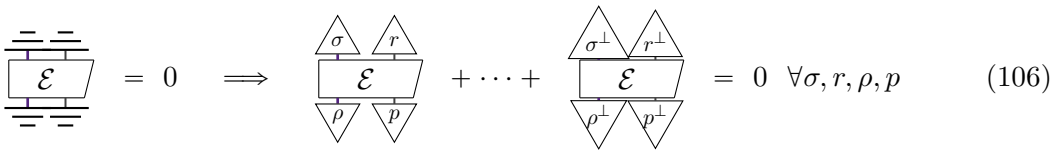

(106)

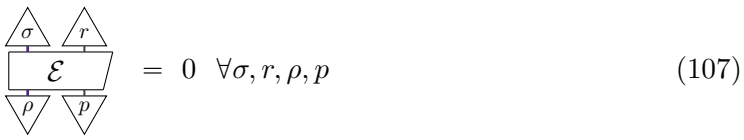

(107)

(108)

where in the first step we use that the maximally mixed state and the discarding effect are internal, in the second step we use that if a sum of non-negative terms is zero then each term is zero, and in the final step we use that **QCalc** is locally tomographic.

The processes in **QCalc**$^\bullet$ can therefore be viewed as a particular conventional choice of representative elements for the equivalence classes. The modified composition rule can then be derived by: i) replacing each process with its equivalence class, ii) composing the equivalence classes, and iii) choosing the representative element for the equivalence class of the composite.

What is therefore clear from this is that:

Theorem 4.4. These two time neutral theories are in fact equivalent to one another,

$$\textbf{QCalc}/\sim \;\cong\; \textbf{QCalc}^\bullet. \tag{109}$$

Specifically, **QCalc**$^\bullet$ is simply a way to describe **QCalc**/∼ using representative elements of the equivalence classes, rather than just working with the equivalence classes themselves.

The inelegant nature of **QCalc**$^\bullet$ can be viewed as a consequence of the somewhat arbitrary way in which the representative elements are chosen. Picking some other convention for how to pick a representative element would lead to a distinct composition rule (and hence a different probability rule) but would ultimately be describing the same process theory.

This equivalence, together with Props. 4.2 and 4.3, immediately tells us that **QCalc**$^\bullet$ is also time symmetric and time neutral as we claimed earlier.

Note that, in practice, if we want to compute anything within $\mathbf{QCalc}/\sim$ it will be much more convenient to work with representative elements rather than the equivalence classes themselves. In this case, one may be tempted to work within $\mathbf{QCalc}^\bullet$; however, there is no real benefit to choosing these particular representative elements. What is usually simplest is i) to pick *arbitrary* representatives from each equivalence class, ii) compose them within **QCalc**, and iii) quotient at the end to obtain the equivalence class again.

4.4 Restoring determinism

As mentioned earlier, in $\mathbf{QCalc}/\sim$ (or equivalently $\mathbf{QCalc}^\bullet$) we have a problem with determinism – we have zero-processes that describe things that cannot occur. In other approaches this has been dealt with in a fairly arbitrary way by simply stating that when $N_M(\rho) = 0$ then $\text{Prob}(a|\rho) = 0$ for all $a \in \mathbb{A}$. However, how are we to operationally understand a measurement in which all of the possible outcomes occur with probability zero?

This manifests in our approach by the fact that we have a pair of scalars in $\mathbf{QCalc}/\sim$ rather than the single scalar that we would expect in a deterministic theory (Def. 2.2). Furthermore, our constructions have a somewhat arbitrary nature to them: i) the modified composition rule has to have two cases depending on whether or not a zero appears, and ii) when we quotient, we do so with respect to non-zero scalars rather than arbitrary scalars.

Here we present an alternative solution that will result in a deterministic theory in which the zero-cases naturally do not arise. We do so by moving away from **QCalc** as the starting point for our constructions.

To begin, we observe that in any real-world experiment, we never manage to completely suppress all sources of noise. That is, in the lab we never actually prepare a pure state or perform a projective measurement (at least not on the system of interest). More generally, we can view the experimentally realisable quantum processes as those of the form:

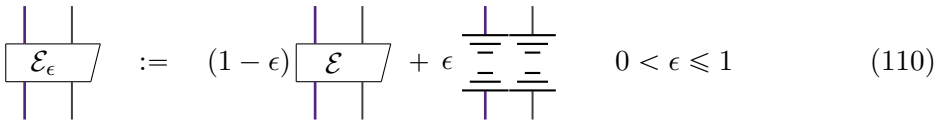

(110)

that is, processes that have some non-zero epsilon of noise. Indeed, we can define a restriction of **QCalc** denoted $\mathbf{QCalc}|_{noise}$ which has only processes of this form. It is straightforward to verify that this indeed defines a valid process theory as these are closed under composition.

$\mathbf{QCalc}|_{noise}$ and **QCalc**, at least from an operational point of view, are equivalent to one another. No real-world experiment could distinguish between some process $\mathcal{E}$ and the noisy version $\mathcal{E}_\epsilon$ provided that ϵ is small enough – we can approximate $\mathcal{E}$ arbitrarily well by using processes $\mathcal{E}_\epsilon$ and taking the limit of $\epsilon \to 0$. In some sense then, choosing between **QCalc** (as is standard in the literature) and $\mathbf{QCalc}|_{noise}$ is purely a matter of convenience. However, when we move over to the respective quotiented theories, $\mathbf{QCalc}/\sim$ and $\mathbf{QCalc}|_{noise}/\sim$, we obtain strikingly different theories. This is highlighted by considering the scalars of the theory. Note that the scalars in **QCalc** are $[0, \infty)$ whilst in $\mathbf{QCalc}|_{noise}$ they are $(0, \infty)$. As noted before, this means that in $\mathbf{QCalc}/\sim$ we have two scalars $\{\tilde{0}, \tilde{1}\}$, whereas in $\mathbf{QCalc}|_{noise}/\sim$ we have only the scalar $\tilde{1}$ (as 0 is not a scalar in $\mathbf{QCalc}|_{noise}$). This means that $\mathbf{QCalc}|_{noise}/\sim$ is a deterministic process theory, and so every process has a valid operational interpretation. To see this more clearly, note that the states for a

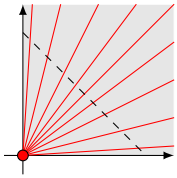

(a) The space of states for $\mathbf{QCalc}/\sim$

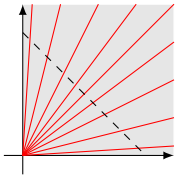

(b) The space of states for $\mathbf{QCalc}|_{noise}/\sim$

Figure 4: Note that the only difference is whether or not the zero state appears. Every state (i.e., ray) on the right-hand side can be uniquely identified with the probability distribution at the point where it intersects the dashed line.

classical system are now only those of the form:

(111)

as the zero state,

, (112)

is no longer part of the theory. This means that *every* classical state can be thought of as a probability distribution. For the case of a classical bit this is illustrated in Fig. 4.

This theory remains time symmetric as it obtains a dagger from the Hermitian adjoint of **QCalc**:

Proposition 4.5. $\mathbf{QCalc}|_{noise}/\sim$ is time symmetric.

As per our current definition, however, it is not time neutral as it does not have cups & caps, nor does it even have identities or swaps! It is simple to see this as the cups & caps, identities, and swaps in **QCalc** are not processes of the form of Eq. (110).

Remark 4.6. As discussed above, in $\mathbf{QCalc}|_{noise}$ we do not have identities, swaps, and other processes which we could construct from them. From a categorical point of view, this would mean that we do not have identity morphisms and symmetries. This highlights a subtle difference between process theories (as we define them here) and SMCs which previously have been taken to be equivalent. We conjecture, however, that these identities and swaps can be freely reintroduced without changing the theory's semantics. If true, this would restore the equivalence between process theories and SMCs.

In this case, we can freely add all of these *wiring processes* back into the theory without changing anything that we have discussed so far.

Definition 4.7 (**QNeut**)**.** We define **QNeut** as the process theory generated by

(113)

and quotiented by arbitrary scalars.

Note that we do not need to restrict to quotienting by scalars $r > 0$ because, by construction, all scalars will satisfy this condition as we see in the following.

Theorem 4.8. **QNeut** is a deterministic and time neutral process theory.

Proof. Time neutrality follows immediately from the existence of the cups & caps that we just added. We can then note that the set of noisy processes (i.e., those of the form in Eq. (110)) is closed under composition with wiring processes, e.g.:

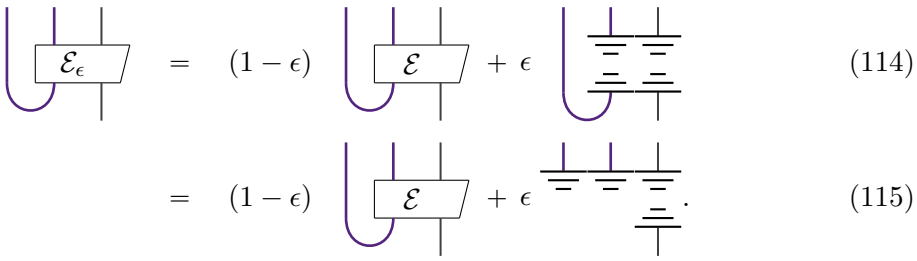

(114)

(115)

If we have some generic diagram, we can therefore always absorb these wiring processes into noisy processes unless the wiring processes are disconnected, as in:

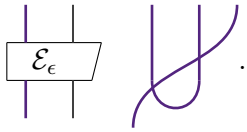

(116)

A general scalar (i.e., closed diagram) can therefore be written as some scalar r_ϵ in $\mathbf{QCalc}|_{noise}$ composed in parallel with some closed wiring, for example:

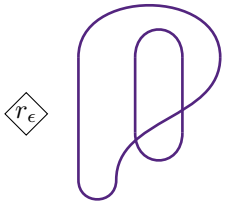

(117)

One can then note that such closed wiring is simply the parallel composition of closed loops and, because each such closed loop equals the dimension of its system, the result is a strictly positive number w. Hence, $r_\epsilon w > 0$ and, when we quotient, we obtain the scalar $\tilde{1}$. We have therefore shown that there is a unique scalar in the theory. □

4.5 Connections to higher-order processes

Recent work has connected the study of process matrices [31] (briefly introduced in section 3.3) to the time-symmetric formulation of quantum theory [29] and the time-neutral formulation of quantum theory [15]. Here we re-derive part of this in a particularly simple way, using the cups and caps in our time-neutral formalism (Def. 2.10), as instantiated by **QNeut** (Def. 4.7).

Proposition 4.9. Any time-neutral quantum theory, in the sense of Def. 2.10, that contains all CPTP maps also contains all process matrices [15, 29].

Proof. A process matrix is a higher-order linear map, W, that acts on a pair of CPTP maps, A and B, to give a new CPTP map:

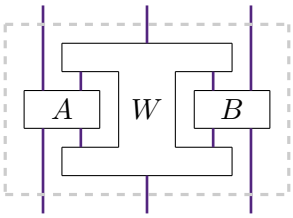

(118)

The question we want to answer is whether such a map can be implemented in our time-neutral formalism. To see that this is indeed possible we will use the cups & caps that we introduced earlier. Consider the special case in which A and B are both swaps. The definition of the process matrices implies that this must be a CPTP map which we will denote $\mathcal{W}$:

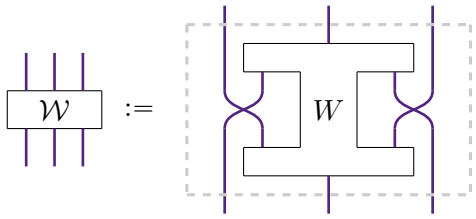

(119)

Since all CPTP maps are in the theory by assumption, $\mathcal{W}$ is also in the theory. It is now simple to show, using cups & caps, that this means that W itself is in our formalism:

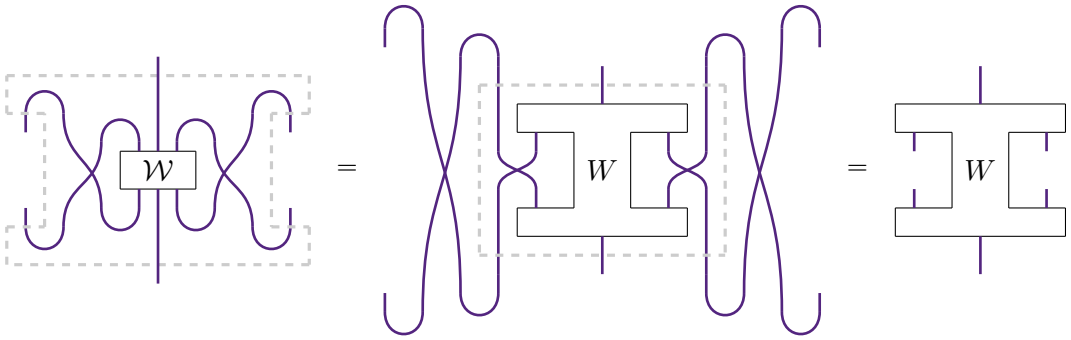

(120)

The LHS is a valid diagram in our theory because it is built from $\mathcal{W}$, which is a CPTP map, and from cups and caps, all of which are valid processes. The first equality follows from the definition of $\mathcal{W}$, and the second from the definition of cups and caps. Hence, the RHS is also valid in our formalism. Strictly speaking, these kinds of "diagrams with holes" are not part of the process-theory framework, so we should not say that W is literally "in our formalism". Nevertheless, the action of W on arbitrary CPTP maps is represented within the formalism, since we can use the above idea to write:

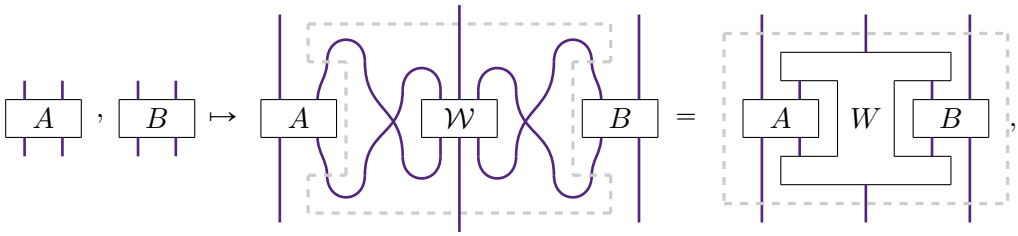

(121)

for arbitrary CPTP maps A and B. □

Such reasoning also applies to more general higher-order processes [33, 64]. They can all be viewed as particular classes of CP maps between quantum systems; what varies is how they are 'normalised' and how they can be composed, particularly while preserving the normalisation constraints. As our formalism does not impose any constraints on normalisation and does not impose any constraints on composition, it seems clear that it should be able to simulate any such higher-order process. However, we will leave the details and explore the consequences of this in future work.

Open Problem 2. Does **QNeut** subsume the frameworks for higher-order processes of [33, 64]? Moreover, can **QNeut** be viewed as a higher-order process theory in the sense of [65]?

4.6 Beyond quantum theory

So far in this section, we have focused on the example of quantum theory, in particular, with the aim of constructing our deterministic time-neutral quantum theory **QNeut**. It is clear, however, that a great deal of this is much more broadly applicable.

In particular, the key construction of quotienting by non-zero scalars makes sense for any process theory in which the scalars are $\mathbb{R}^+$. This includes arbitrary generalised probabilistic theories (GPTs) [20, 21], operational probabilistic theories (OPTs) [7, 22, 23], categorical probabilistic theories [24], and certain kinds of causal-inferential theories [66]. Moreover, the construction can also apply when the scalars belong to an arbitrary commutative semiring with no zero divisors; it therefore applies to many of the more exotic theories found in, for example, [28]. The only subtlety to watch out for here is that to obtain anything interesting in any of these cases we need to go beyond the causal processes (i.e., the analogue of **QPhys**), and instead quotient the supertheory that does not impose any causality constraint (i.e., the analogue of **QCalc**).

Even after this issue is addressed, it is not obvious whether these alternative theories yield a time-neutral theory. It is therefore interesting to investigate what properties a physical theory, **G**, must satisfy so that $\mathbf{G}/\!\sim$ is a time neutral theory.

Exploring the properties of time-neutral physical theories is therefore an important direction for future research. In particular, one can ask whether we can find a set of principles that single out **QNeut** from the space of all possible time neutral theories. Or, in other words, can we find a reconstruction of time symmetric quantum theory from well-motivated principles? The reconstruction of [39] may well be a useful starting point for this, as it introduces a time symmetric version of the purification postulate which was itself used to great effect in the reconstruction of [22].

4.7 Discussion

By taking a process-theoretic approach to the problem of time neutrality we have subsumed seemingly different approaches to the problem, and shown that they amount to different conventional choices within the same underlying physical theory. Moreover, we have shown that this can be elegantly captured via a particular type of quotienting of the theory. However, the approach also yields several further benefits:

- Arguably, this is a much simpler presentation of the theory, at least, for those who are already familiar with categorical quantum mechanics, and means that tools such as the ZX calculus [67, 68] can now be applied to the study of time neutral theories.
- The process-theoretic nature of the construction lends itself to generalisation to alternative physical theories such as GPTs. Although the construction can be applied to any GPT, it need not yield a time-neutral theory. Identifying the GPT principles that imply time neutrality under this construction is therefore an interesting direction for future research.
- The process-theoretic nature also means that it should be more directly comparable to the frameworks for higher-order processes such as Refs. [33, 64, 65]. Section 4.5 provides a simple proof-of-principle application in this direction.
- This approach may well also more readily extend beyond the finite-dimensional setting which has so far dominated the literature on time symmetric theories, in particular, by considering the categorical formalism of [69, 70].

- It can also be considered in the context of the functorial approach to field theories presented in [71]. This may provide a formal way to study how time neutral theories behave in a space-time context, and, in particular, whether they lead to contradictions with the non-signalling principle.

- Finally, one can ask whether it is possible to find a reconstruction of quantum theory that reproduces the time neutral theory, **QNeut**, rather than the standard causal theory, **QPhys**. Recent process-theoretic reconstructions such as Refs. [39] and [47] provide natural starting points for this question.

The biggest issue with the current formalism, however, is that whilst the theory is time neutral, the presentation of the theory still involves an artificial division of the systems associated with a given process into input and output systems. In [44] we show how this can be overcome by moving from a categorical formulation of process theories to an operadic formulation based on the work of Refs. [43, 72].

Acknowledgements

JHS would like to thank Tony Short for an interesting discussion that helped to shape the early stages of the project. JHS and MES would also like to thank Lucien Hardy for insightful discussions about his work on time symmetry in operational theories. BC thanks Carlo Rovelli and Augustin Vanrietvelde for discussions on time symmetry. The authors also thank Arthur Parzygnat for helpful feedback on the work.

JHS and MES were supported by the Foundation for Polish Science (IRAP project, ICTQT, contract no. MAB/2018/5, co-financed by EU within Smart Growth Operational Programme). BC and SG were supported by the John Templeton Foundation through grant 61466, The Quantum Information Structure of Spacetime (qiss.fr). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

All of the diagrams within this manuscript were prepared using Aleks Kissinger's TikZit.

Author contribution statement

All authors contributed to the first draft of the manuscript. JHS and MES handled journal submission and subsequent revisions. AI tools were only used in the final stage to check for spelling, grammar, and typographic errors – the scientific content of this paper is all human generated.

References

[1] Eugene Paul Wigner. *Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren*. Springer, 1931. DOI: 10.1007/978-3-663-02555-9.

[2] Gerhart Lüders. On the equivalence of invariance under time reversal and under particle-antiparticle conjugation for relativistic field theories. *Dan. Mat. Fys. Medd.*, 28:1–17, 1954.

[3] John Stewart Bell. Time reversal in field theory. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 231(1187):479–495, 1955. DOI: 10.1142/9789812795854˙0011.

[4] Yakir Aharonov, Peter G Bergmann, and Joel L Lebowitz. Time symmetry in the quantum process of measurement. *Physical Review*, 134(6B):B1410, 1964. DOI: 10.1103/physrev.134.b1410.

[5] Satosi Watanabe. Symmetry of physical laws. part iii. prediction and retrodiction. *Reviews of Modern Physics*, 27(2):179, 1955. DOI: 10.1103/revmodphys.27.179.

[6] Andrew Thomas Holster. The criterion for time symmetry of probabilistic theories and the reversibility of quantum mechanics. *New Journal of Physics*, 5(1):130, 2003. DOI: 10.1088/1367-2630/5/1/130.

[7] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Probabilistic theories with purification. *Physical Review A*, 81(6):062348, 2010. DOI: 10.1103/physreva.81.062348.

[8] B. Coecke. Terminality implies non-signalling. arXiv:1405.3681, 2014.

[9] Bob Coecke, Stefano Gogioso, and John H Selby. The time-reverse of any causal theory is eternal noise. *arXiv preprint arXiv:1711.05511*, 2017. DOI: 10.48550/arXiv.1711.05511.

[10] Ognyan Oreshkov and Nicolas J Cerf. Operational formulation of time reversal in quantum theory. *Nature Physics*, 11(10):853–858, 2015. DOI: 10.1038/nphys3414.

[11] L. Hardy. Towards quantum gravity: a framework for probabilistic theories with non-fixed causal structure. *J. Phys. A*, 40(12):3081, 2007. DOI: 10.1088/1751-8113/40/12/s12.

[12] Matthew S Leifer and Robert W Spekkens. Towards a formulation of quantum theory as a causally neutral theory of bayesian inference. *Physical Review A*, 88(5):052130, 2013. DOI: 10.1103/physreva.88.052130.

[13] Robert Oeckl. General boundary quantum field theory: Foundations and probability interpretation. *Advances in Theoretical and Mathematical Physics*, 12(2):319–352, 2008. DOI: 10.4310/atmp.2008.v12.n2.a3.

[14] Robert Oeckl. A local and operational framework for the foundations of physics. *arXiv preprint arXiv:1610.09052*, 2016. DOI: 10.4310/atmp.2019.v23.n2.a4.

[15] Ognyan Oreshkov and Nicolas J Cerf. Operational quantum theory without predefined time. *New Journal of Physics*, 18(7):073037, 2016. DOI: 10.1088/1367-2630/18/7/073037.

[16] Yakir Aharonov, Sandu Popescu, Jeff Tollaksen, and Lev Vaidman. Multiple-time states and multiple-time measurements in quantum mechanics. *Physical Review A*, 79(5):052110, 2009. DOI: 10.1103/physreva.79.052110.

[17] Yakir Aharonov and Lev Vaidman. The two-state vector formalism: an updated review. In *Time in quantum mechanics*, pages 399–447. Springer, 2008. DOI: 10.1007/978-3-540-73473-4˙13.

[18] Andrea Di Biagio, Pietro Donà, and Carlo Rovelli. The arrow of time in operational formulations of quantum theory. *Quantum*, 5:520, 2021. DOI: 10.22331/q-2021-08-09-520.

[19] Lucien Hardy. Time symmetry in operational theories. *arXiv preprint arXiv:2104.00071*, 2021. DOI: 10.48550/arXiv.2104.00071.

[20] L. Hardy. Quantum theory from five reasonable axioms. *arXiv:quant-ph/0101012*, 2001. DOI: 10.48550/arXiv.quant-ph/0101012.

[21] J. Barrett. Information processing in generalized probabilistic theories. *Physical Review A*, 75:032304, 2007. DOI: 10.1103/physreva.75.032304.

[22] G. Chiribella, G. M. D'Ariano, and P. Perinotti. Informational derivation of quantum theory. *Physical Review A*, 84(1):012311, 2011. DOI: 10.1103/physreva.84.012311.

[23] Giacomo Mauro D'Ariano, Giulio Chiribella, and Paolo Perinotti. *Quantum theory*

from first principles: an informational approach. Cambridge University Press, 2017. DOI: 10.1017/9781107338340.002.

[24] Stefano Gogioso and Carlo Maria Scandolo. Categorical probabilistic theories. *EPTCS*, 266:367–385, 2018. DOI: 10.4204/EPTCS.266.23.

[25] S. Abramsky and B. Coecke. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 415–425, 2004. DOI: 10.1109/lics.2004.1319636. arXiv:quant-ph/0402130.

[26] B. Coecke, B. Edwards, and R. W. Spekkens. Phase groups and the origin of non-locality for qubits. *Electronic Notes in Theoretical Computer Science*, 270(2):15–36, 2011. DOI: 10.1016/j.entcs.2011.01.021. arXiv:1003.5005.

[27] B. Coecke, R. Duncan, A. Kissinger, and Q. Wang. Generalised compositional theories and diagrammatic reasoning. In G. Chiribella and R. W. Spekkens, editors, *Quantum Theory: Informational Foundations and Foils*, Fundamental Theories of Physics. Springer, 2016. DOI: 10.1007/978-94-017-7303-4˙10. arXiv:1203.4988.

[28] Stefano Gogioso. Fantastic quantum theories and where to find them. *arXiv preprint arXiv:1703.10576*, 2017. DOI: 10.48550/arXiv.1703.10576.

[29] Ralph Silva, Yelena Guryanova, Anthony J Short, Paul Skrzypczyk, Nicolas Brunner, and Sandu Popescu. Connecting processes with indefinite causal order and multi-time quantum states. *New Journal of Physics*, 19(10):103022, 2017. DOI: 10.1088/1367-2630/aa84fe.

[30] Ralph Silva, Yelena Guryanova, Nicolas Brunner, Noah Linden, Anthony J Short, and Sandu Popescu. Pre-and postselected quantum states: Density matrices, tomography, and kraus operators. *Physical Review A*, 89(1):012121, 2014. DOI: 10.1103/physreva.89.012121.

[31] Ognyan Oreshkov, Fabio Costa, and Časlav Brukner. Quantum correlations with no causal order. *Nature communications*, 3(1):1–8, 2012. DOI: 10.1038/ncomms2076.

[32] Giulio Chiribella, G Mauro D'Ariano, and Paolo Perinotti. Quantum circuit architecture. *Physical review letters*, 101(6):060401, 2008. DOI: 10.1103/physrevlett.101.060401.

[33] Alessandro Bisio and Paolo Perinotti. Theoretical framework for higher-order quantum theory. *Proceedings of the Royal Society A*, 475(2225):20180706, 2019. DOI: 10.1098/rspa.2018.0706.

[34] Sander Uijlen and Aleks Kissinger. A categorical semantics for causal structure. *Logical Methods in Computer Science*, 15, 2019. DOI: 10.23638/lmcs-15(3:15)2019.

[35] Matt Wilson and Giulio Chiribella. A mathematical framework for transformations of physical processes. *arXiv preprint arXiv:2204.04319*, 2022. DOI: 10.48550/arXiv.2204.04319.

[36] Bob Coecke. A universe of processes and some of its guises. *Deep beauty: understanding the quantum world through mathematical innovation*, pages 129–186, 2011. DOI: 10.1017/cbo9780511976971.004.

[37] B. Coecke and A. Kissinger. Categorical quantum mechanics I: causal quantum processes. In E. Landry, editor, *Categories for the Working Philosopher*. Oxford University Press, 2016. DOI: 10.1093/oso/9780198748991.003.0012. arXiv:1510.05468.

[38] B. Coecke and A. Kissinger. *Picturing Quantum Processes. A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, 2016. DOI: 10.1017/9781316219317.002.

[39] John H Selby, Carlo Maria Scandolo, and Bob Coecke. Reconstructing quantum theory from diagrammatic postulates. *Quantum*, 5:445, 2021. DOI: 10.22331/q-2021-04-28-445.

[40] John H Selby. *A process theoretic triptych*. PhD thesis, PhD thesis, Imperial College London, 2017.

[41] S. Abramsky and B. Coecke. Categorical quantum mechanics. In K. Engesser, D. Gabbay, and D. Lehmann, editors, *Handbook of quantum logic and quantum structures: quantum logic*, pages 261–324. Elsevier, 2009. ISBN 0444528695. DOI: 10.1016/b978-0-444-52869-8.50010-4.

[42] S. Mac Lane. *Categories for the working mathematician*. Springer-verlag, 1998. DOI: 10.1007/978-1-4612-9839-7.

[43] Evan Patterson, David I Spivak, and Dmitry Vagner. Wiring diagrams as normal forms for computing in symmetric monoidal categories. *arXiv preprint arXiv:2101.12046*, 2021. DOI: 10.4204/EPTCS.333.4.

[44] John H Selby, Maria E Stasinou, Matt Wilson, and Bob Coecke. Generalised process theories. *arXiv preprint arXiv:2502.10368*, 2025. DOI: 10.48550/arXiv.2502.10368.

[45] Bob Coecke, Chris Heunen, and Aleks Kissinger. Categories of quantum and classical channels. *Quantum Information Processing*, 15(12):5179–5209, 2016. DOI: 10.1007/s11128-014-0837-4.

[46] John Selby and Bob Coecke. Leaks: quantum, classical, intermediate and more. *Entropy*, 19(4):174, 2017. DOI: 10.3390/e19040174.

[47] Sean Tull. A categorical reconstruction of quantum theory. *arXiv preprint arXiv:1804.02265*, 2018. DOI: 10.48550/arXiv.1804.02265.

[48] B. Coecke and R. Lal. Causal categories: relativistically interacting processes. *Foundations of Physics*, 43:458–501, 2013. DOI: 10.1007/s10701-012-9646-8. arXiv:1107.6019.

[49] Aleks Kissinger, Matty Hoban, and Bob Coecke. Equivalence of relativistic causal structure and process terminality. *arXiv preprint arXiv:1708.04118*, 2017. DOI: 10.48550/arXiv.1708.04118.

[50] S. Abramsky and B. Coecke. Abstract physical traces. *Theory and Applications of Categories*, 14(6):111–124, 2005. DOI: 10.70930/tac/vgpjookq. arXiv:0910.3144.

[51] P. Selinger. Dagger compact closed categories and completely positive maps. *Electronic Notes in Theoretical Computer Science*, 170:139–163, 2007. DOI: 10.1016/j.entcs.2006.12.018.

[52] B. Coecke. Kindergarten quantum mechanics. In A. Khrennikov, editor, *Quantum Theory: Reconsiderations of the Foundations III*, pages 81–98. AIP Press, 2005. DOI: 10.1063/1.2158713. arXiv:quant-ph/0510032.

[53] John H Selby and Bob Coecke. A diagrammatic derivation of the hermitian adjoint. *Foundations of Physics*, 47(9):1191–1207, 2017. DOI: 10.1007/s10701-017-0102-7.

[54] Stefano Gogioso. *Categorical quantum dynamics*. PhD thesis, University of Oxford, 2017.

[55] Stefano Gogioso. A Diagrammatic Approach to Quantum Dynamics. In Markus Roggenbach and Ana Sokolova, editors, *8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019)*, volume 139 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:23, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-120-7. DOI: 10.4230/LIPIcs.CALCO.2019.19. URL `http://drops.dagstuhl.de/opus/volltexte/2019/11447`.

[56] Bob Coecke, Tobias Fritz, and Robert W Spekkens. A mathematical theory of resources. *Information and Computation*, 250:59–86, 2016. DOI: 10.1016/j.ic.2016.02.008.

[57] John C Baez and Aaron Lauda. A prehistory of n-categorical physics. *Deep Beauty: Understanding the Quantum World Through Mathematical Innovation*, pages 13–128, 2011. DOI: 10.1017/cbo9780511976971.003.

[58] H. Scutaru. Some remarks on covariant completely positive linear maps on C*-algebras. *Reports on Mathematical Physics*, 16(1):79–87, 1979. ISSN 0034-4877. DOI: 10.1016/0034-4877(79)90040-5. URL https://www.sciencedirect.com/science/article/pii/0034487779900405.

[59] Dominic Verdon. A covariant stinespring theorem. *Journal of Mathematical Physics*, 63(9):091705, 2022. DOI: 10.1063/5.0071215.

[60] Markus P Müller and Cozmin Ududec. Structure of reversible computation determines the self-duality of quantum theory. *Physical review letters*, 108(13):130401, 2012. DOI: 10.1103/physrevlett.108.130401.

[61] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994. DOI: 10.1007/bf02058098.

[62] Daniel Goldwater, Mauro Paternostro, and PF Barker. Testing wave-function-collapse models using parametric heating of a trapped nanosphere. *Physical Review A*, 94(1): 010104, 2016. DOI: 10.1103/physreva.94.010104.

[63] Nicola Pinzani, Stefano Gogioso, and Bob Coecke. Categorical semantics for time travel. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–20. IEEE, 2019. DOI: 10.1109/LICS.2019.8785664.

[64] Aleks Kissinger and Sander Uijlen. A categorical semantics for causal structure. In *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, pages 1–12. IEEE, 2017. DOI: 10.1109/lics.2017.8005095.

[65] Matt Wilson and Giulio Chiribella. Causality in higher order process theories. *arXiv preprint arXiv:2107.14581*, 2021. DOI: 10.4204/EPTCS.343.12.

[66] David Schmid, John H Selby, and Robert W Spekkens. Unscrambling the omelette of causation and inference: The framework of causal-inferential theories. *arXiv preprint arXiv:2009.03297*, 2020. DOI: 10.48550/arXiv.2009.03297.

[67] B. Coecke and R. Duncan. Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics*, 13:043016, 2011. DOI: 10.1088/1367-2630/13/4/043016. arXiv:quant-ph/09064725.

[68] Bob Coecke, Dominic Horsman, Aleks Kissinger, and Quanlong Wang. Kindergarden quantum mechanics graduates ...or how i learned to stop gluing lego together and love the zx-calculus. *Theoretical Computer Science*, 897:1–22, 2022. ISSN 0304-3975. DOI: 10.1016/j.tcs.2021.07.024. URL https://www.sciencedirect.com/science/article/pii/S0304397521004308.

[69] Stefano Gogioso and Fabrizio Romano Genovese. Infinite-dimensional categorical quantum mechanics. *EPTCS*, 236:51–69, 2017. DOI: 10.4204/EPTCS.236.4.

[70] Razin A. Shaikh and Stefano Gogioso. Categorical semantics for feynman diagrams. *arXiv preprint arXiv:2205.00466*, 2022. DOI: 10.48550/arXiv.2205.00466.

[71] Stefano Gogioso, Maria E Stasinou, and Bob Coecke. Functorial evolution of quantum fields. *Frontiers of Physics*, 9, 2020. DOI: 10.3389/fphy.2021.534265.

[72] David I Spivak, Patrick Schultz, and Dylan Rupel. String diagrams for traced and compact categories are oriented 1-cobordisms. *Journal of Pure and Applied Algebra*, 221(8):2064–2110, 2017. DOI: 10.1016/j.jpaa.2016.10.009.

[73] Bob Coecke, Ross Duncan, Aleks Kissinger, and Quanlong Wang. Strong complementarity and non-locality in categorical quantum mechanics. In *2012 27th Annual IEEE Symposium on Logic in Computer Science*, pages 245–254. IEEE, 2012. DOI: 10.1109/LICS.2012.35.

[74] Stefano Gogioso and Aleks Kissinger. Fully graphical treatment of the quantum algorithm for the hidden subgroup problem. *arXiv preprint arXiv:1701.08669*, 2017. DOI: 10.48550/arXiv.1701.08669.

[75] Stefano Gogioso and William Zeng. Generalised mermin-type non-locality arguments. *Logical Methods in Computer Science*, 15, 2019. DOI: 10.23638/LMCS-15(2:3)2019.

[76] Jamie Vicary. Topological structure of quantum algorithms. In *2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 93–102. IEEE, 2013. DOI: 10.1109/LICS.2013.14.

[77] Stefano Gogioso and William Zeng. Fourier transforms from strongly complementary observables. *arXiv preprint arXiv:1501.04995*, 2015. DOI: 10.48550/arXiv.1501.04995.

A Diagrams for groups and their representations

Below is a recap of diagrammatic notation for group representations. The correspondence between strongly complementary structures and finite-dimensional group algebras was introduced in [73] for the Abelian case and in [74, 75] for the general case. Diagrammatic notation for group representations previously appeared in [76, 77]—in the context of quantum algorithms and the quantum Fourier transform—and was extensively discussed in [54, 55]—in the context of quantum dynamics.

A.1 Diagrams for groups

We denote a group G by a new kind of wire,

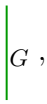

(122)

such that the elements of the group, $g \in G$, are states of this wire,

(123)

In particular, we denote the identity element of the group by:

(124)

We then introduce the defining group operations, namely group multiplication and inversion:

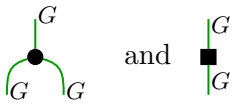

(125)

respectively. These can be defined by their action on the group elements via:

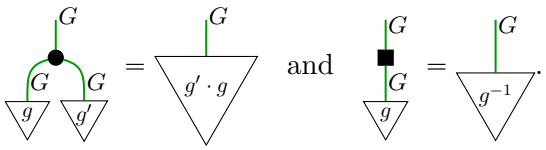

(126)

Various important properties of groups, including associativity of multiplication, can be captured elegantly by these diagrams:

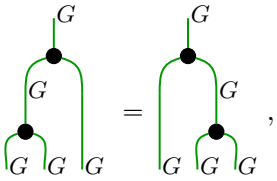

(127)

that the identity element is the unit,

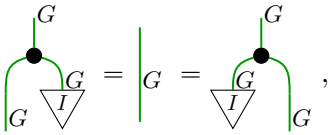

(128)

that the group inverse is involutive,

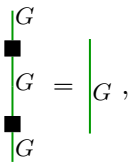

(129)

and that multiplication is “antisymmetric”,

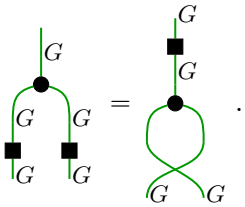

(130)

To capture more of the group structure diagrammatically, we introduce a copy map and a discard map for these new systems, denoted:

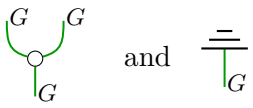

(131)

respectively. These, too, can be defined by their action on group elements via:

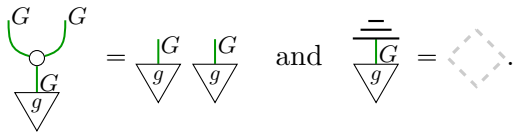

(132)

Important properties of these processes, including associativity of copying, can also be captured diagrammatically:

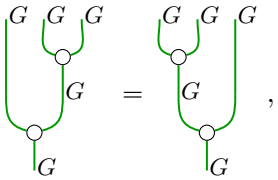

(133)

that the discard is the counit for copy,

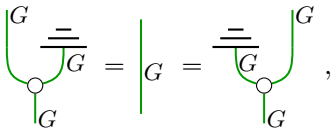

(134)

and that copying is symmetric,

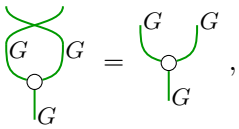

(135)

Finally, we can use the interaction of these new processes with the group multiplication and inverse in order to capture more structure of the group. In particular,

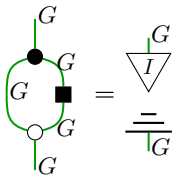

(136)

ensures that the group inverse behaves as expected. Moreover, one can show that copying and multiplication form a bialgebra,

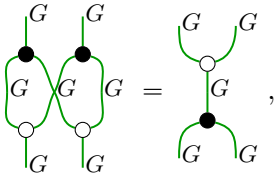

(137)

that group multiplication is causal,

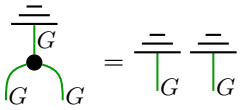

(138)

and that the inverse is copied by the copy operation,

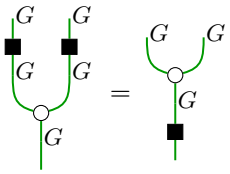

(139)

A.2 Group representations

In this section, we define group representations on systems in **QCalc** (or **QPhys**) as an interaction between the new systems G and systems from **QCalc**. Formally we can think of the systems G as being particular (potentially infinite-dimensional) classical systems in which the point distributions correspond to the group elements.

A *causal* group representation of G on some system Q in **QCalc** is a process, π, of the form:

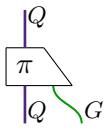

(140)

such that the following equations are satisfied:

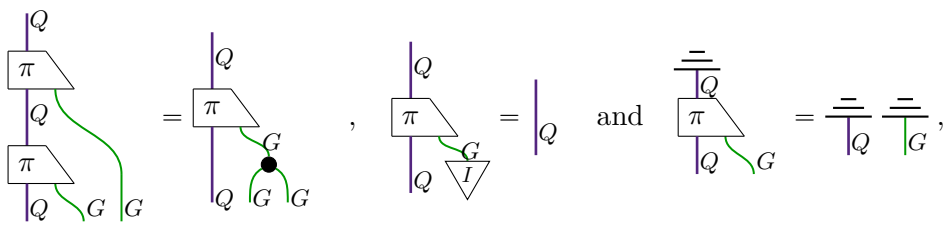

(141)

The first two equations guarantee that this is a valid representation and the third guarantees that this is a *causal* representation. Note that, if we are working with **QPhys** rather than **QCalc**, then this last condition is automatically satisfied.

If Q is strictly quantum, that is, a system $\mathcal{H}$, then these causal representations are necessarily unitary; that is, they satisfy

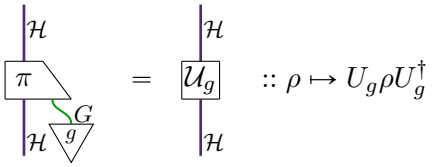

(142)

for all $g \in G$. To see this, simply note that the axioms for group representations imply that the processes on the left-hand side are necessarily reversible and trace-preserving (i.e., causal), and reversible CPTP maps are necessarily unitary supermaps.

For finite-dimensional classical systems $\mathbb{X}$, the only representation that we will typically use is the trivial representation, namely:

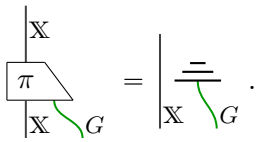

(143)

For many groups, this will in fact be the only possible representation for finite-dimensional classical systems.

We can compose representations of single systems to define representations of composite systems via:

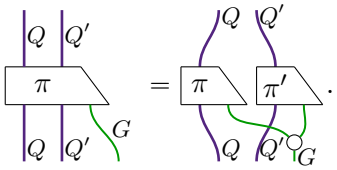

(144)

Thus, if we compose a quantum and a classical representation we end up with a representation:

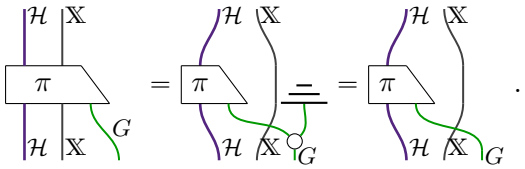

(145)

That is, the quantum part of the composite system may transform non-trivially under the action of the group, but the classical part is left invariant. Finally, note that the representation on a trivial system is necessarily trivial, that is:

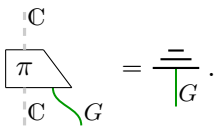

(146)

A.3 Intertwiners

Next we introduce *intertwiners*, i.e., processes in **QCalc** (or **QPhys**) that are symmetric with respect to the group action. In particular, they are characterised by the covariance condition:

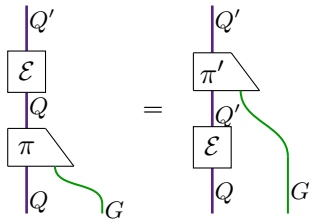

(147)

Using the fact that representations on the trivial system and classical systems are trivial representations, one can easily show that states and measurements are intertwiners if and only if they are invariant under the group action, namely:

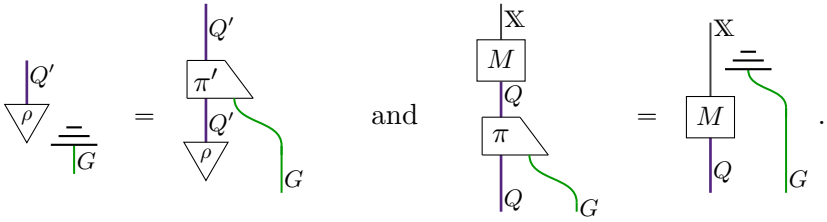

(148)

A.4 Dual systems and conjugate representations

Finally, we introduce dual systems and their representations. Systems in **QCalc** have duals, which we will now denote explicitly by arrows. That is, a generic process is denoted by

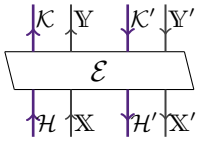

(149)

and given by a CP map

$$\mathcal{E} : \mathcal{B}[\mathcal{H}] \otimes \left(\bigoplus_{x \in \mathbb{X}} \mathcal{B}[\mathbb{C}] \right) \otimes \mathcal{B}[\mathcal{H}']^* \otimes \left(\bigoplus_{x' \in \mathbb{X}'} \mathcal{B}[\mathbb{C}] \right)^* \to \mathcal{B}[\mathcal{K}] \otimes \left(\bigoplus_{y \in \mathbb{Y}} \mathcal{B}[\mathbb{C}] \right) \otimes \mathcal{B}[\mathcal{K}']^* \otimes \left(\bigoplus_{y' \in \mathbb{Y}'} \mathcal{B}[\mathbb{C}] \right)^* . \tag{150}$$

That is, when the arrow points up, we use the primal vector space; when it points down, we use the dual vector space.

It is straightforward to verify that this is time symmetric and time neutral in the sense of section 2.7.

Given a representation on a system $\mathcal{Q}^{\uparrow}$, we define the conjugate representation π^* on the system $\mathcal{Q}^{\downarrow}$ as follows:

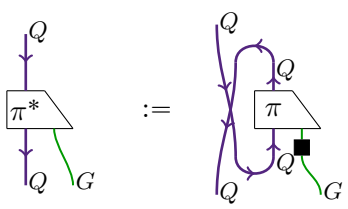

(151)

Note that if $Q^{\uparrow}$ is classical, $\mathbb{X}^{\uparrow}$, then just as the representations on $\mathbb{X}^{\uparrow}$ were necessarily trivial, so is the conjugate representation on $\mathbb{X}^{\downarrow}$.

From this definition it is easy to see that cups & caps are intertwiners for composite representations on the inputs. That is,

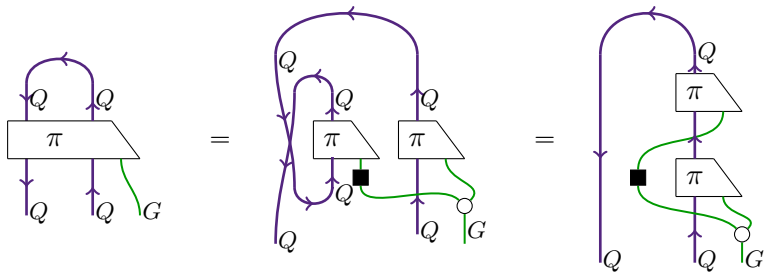

(152)

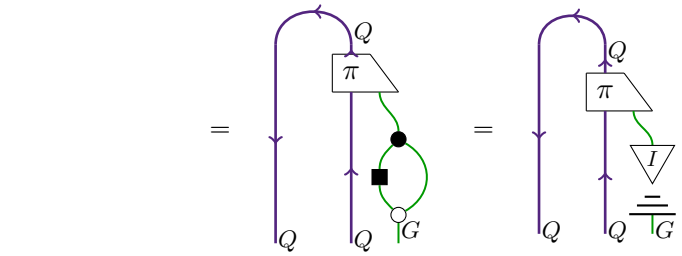

(153)

(154)

In the first equality, we use the assumption that the representation on the input is a composite of representations together with the definition of the conjugate representation. In the second equality, we use the definition of cups and caps; in the third, the definition of a representation; in the fourth, a key result about groups; and in the final equality, the definition of a representation again.