Kolmogorov extension theorem for (quantum) causal modelling and general probabilistic theories

Simon Milz^{1,2}, Fattah Sakuldee^{3,4}, Felix A. Pollock², and Kavan Modi²

¹Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, 1090 Vienna, Austria

²School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia

³International Centre for Theory of Quantum Technologies, University of Gdańsk, Wita Stwosza 63, 80-308 Gdańsk, Poland

⁴MU-NECTEC Collaborative Research Unit on Quantum Information, Department of Physics, Faculty of Science, Mahidol University, Bangkok 10400, Thailand.

October 7, 2019

In classical physics, the Kolmogorov extension theorem lays the foundation for the theory of stochastic processes. It has been known for a long time that, in its original form, this theorem does not hold in quantum mechanics. More generally, it does not hold in any theory of stochastic processes – classical, quantum or beyond – that does not just describe passive observations, but allows for active interventions. Such processes form the basis of the study of causal modelling across the sciences, including in the quantum domain. To date, these frameworks have lacked a conceptual underpinning similar to that provided by Kolmogorov’s theorem for classical stochastic processes. We prove a generalized extension theorem that applies to *all* theories of stochastic processes, putting them on equally firm mathematical ground as their classical counterpart. Additionally, we show that quantum causal modelling and quantum stochastic processes are equivalent. This provides the correct framework for the description of experiments involving continuous control, which play a crucial role in the development of quantum technologies. Furthermore, we show that the original extension theorem follows from the generalized one in the correct limit, and elucidate how a comprehensive understanding of general stochastic pro-

cesses allows one to unambiguously define the distinction between those that are classical and those that are quantum.

1 Introduction

Stochastic processes are ubiquitous in nature. Their theory is used, among other applications, to model the stock market, predict the weather, describe transport processes in cells and understand the random motion of particles suspended in a fluid [1, 2]. Intuitively, when we speak of stochastic processes, we often mean joint probability distributions of random variables at a finite set of times: the probability for a stock to have prices P_1, P_2 and P_3 on three subsequent days, or the probability to find a particle undergoing Brownian motion in regions R_1 and R_2 when measuring its position at times t_1 and t_2 .

This finite description of stochastic processes is motivated by both experimental and mathematical considerations. On the experimental side, temporal resolution is generally limited and digital instruments always record a finite amount of data. Hence, the only accessible information we are left with is encoded in probability distributions with a finite number of arguments. On the mathematical side, it is much less cumbersome to model stochastic processes on discrete times – for example, by defining transition probabilities $\mathbb{P}(Y|X)$ between random variables at a fixed set of different times – than modelling probability densities on the space of all possible ‘trajectories’ of random variables.

Simon Milz: simon.milz@oeaw.ac.at

These motivations notwithstanding, the fundamental laws of physics are continuous in nature and one always implicitly assumes that there is an *underlying* process that leads to the experimentally accessible finite distributions. Put differently, one assumes that there exists an infinite joint probability distribution that has *all* the finite ones as marginals. For classical stochastic processes, these two points of view, the finite and the infinite one, are reconciled by the Kolmogorov extension theorem (**KET**), which lays bare the minimal requirements for the existence of an underlying process, given a family of measurement statistics for finite sets of times [3–6]. It bridges the gap between experimental reality and mathematically rigorous theoretical underpinnings and, as such, enables the definition of stochastic processes as the limit of finite – and hence observable – objects. Additionally, the KET enables the modelling of continuous processes based on finite probability distributions. As a consequence, in the classical setting, stochastic processes over a continuous set of times, and families of finite probability distributions are two sides of the same coin.

The validity of the KET hinges crucially on the fact that the statistics of observations at a time t do not depend on the kind of measurements that were made at any time $t' < t$. Put differently, just like the Leggett-Garg inequalities for temporal correlations [7–9], the KET is based on the assumption of noninvasive measurements and realism *per se*. While the latter property means that any measurement merely reveals a well-defined pre-existing value, the former implies that said measurements can be carried out without disturbing the state of the measured system. For example, in a classical stochastic process, measuring the position of a particle undergoing Brownian motion, reveals pre-existing information, but does not actively change the state of the particle. Both of these conditions together ensure the existence of compatible measurement statistics for different sets of times, which form the basis for the derivation of the KET.

On the other hand, the assumptions of non-invasiveness or realism *per se* are not fulfilled in many experimental scenarios, leading to a breakdown of the KET, at the cost of a clear connection between an underlying process

and its finite time manifestations. This is the case whenever an experimenter chooses to actively interfere with a process to uncover its causal structure or to investigate the reaction to different inputs. For example, instead of just observing the progress of a disease, a pharmacologist tries to find out how the course of a disease changes with the administration of certain drugs. More generally, agent based modelling investigates how systems behave when they can not only be monitored, but actively influenced [10]. Experimental situations where interventions are actively used to uncover causal relations fall within the field of *causal modelling* [11].

Interventions appear naturally in quantum mechanics, where measurements necessarily perturb a system's state; in fact, a complete description of quantum processes without interventions is not possible [12]. As in the classical case, interventions can also be used to actively probe the causal structure of a quantum process, and the description of quantum processes with interventions has been recently used to develop the field of *quantum causal modelling* [13–15]. Importantly, as in the case of classical processes with interventions, the invasiveness of measurements means that the KET does not hold for quantum processes [16]. This is analogous to the violation of Leggett-Garg inequalities [7, 9] in quantum mechanics.

The fundamental lack of an extension theorem in quantum mechanics (or any other theory with interventions) would be problematic for several reasons: Firstly, it would suggest a lack of consistency between descriptions of a process for different sets of times; for example, the description of a process for three times t_1, t_2 , and t_3 would *not* already include the description of the process for the two times t_1 and t_3 only. In other words, we would need seven different independent descriptors for each of the seven subsets of times to describe all possible events! This lack of consistency would render the study of (quantum) causal models in multi-step experiments impossible; if local interventions lead to a completely different process, it is not meaningful to try to deduce causal relations by means of active manipulations of the system at hand.

Furthermore, the present situation (without

an extension theorem) implies an incompatibility between existing frameworks to describe processes with interventions (both classical and quantum) and the classical theory of stochastic processes, even though they should converge to the latter in the correct limit. This then suggests that the mere act of interacting with a system over time introduces a fundamental divide between the continuity of physical laws and the finite statistics that can be accessed in reality, thus begging the question: What do we generally mean by a (quantum) ‘stochastic process’, and how can we reconcile causal modelling frameworks with the idea of an underlying process?

In this Paper, we answer these questions by generalizing the KET to the framework of (quantum) causal modelling, thus closing the apparent divide between the finite and the continuous point of view. To this end, in Sec. 2 we reiterate the relation between classical stochastic processes and classical causal modelling and show the breakdown of the KET when we allow for active interventions in Sec. 3. We analyze the quantum case in Sec. 4 and find that stochastic processes can only be defined properly by taking interventions into account. Consequently, the framework of quantum stochastic processes is equivalent to quantum causal modelling. In Sec. 4.2, we prove our main result, that the KET can be generalized to quantum stochastic processes, and this *generalized extension theorem (GET)* reduces to the classical one in the correct limit. The breakdown of the KET is a breakdown of formalism only, not a fundamental property of quantum processes. Our generalized extension theorem provides an overarching theorem that puts all processes with interventions and, in particular, quantum processes on an equally sound footing as their classical counterpart.

We discuss the equivalence of quantum stochastic processes and quantum causal modelling in Sec. 4.4. As a direct application, in Sec. 4.5, we use the GET to provide a distinction between general, *i.e.*, non-Markovian, classical and quantum processes, as has been recently introduced in [17] for the restricted case of processes without memory. While we phrase our results predominantly in the language of causal modelling, they apply to a wide range of current

theories of quantum processes and beyond. The relation of our results to other frameworks, in particular to the work of Accardi, Frigerio and Lewis [18], is discussed in Sec. 5. We conclude the Paper in Sec. 6.

2 Classical stochastic processes and Causal Modelling

2.1 Classical stochastic processes

A classical stochastic process can be described by joint probability distributions $\mathbb{P}_{\Lambda_k}(i_k, \dots, i_1) := \mathbb{P}_{\Lambda_k}(\mathbf{i}_{\Lambda_k})$ of random variables that take values $\{i_\alpha\}$ at time t_α [5], where Λ_k is a collection of times with cardinality $|\Lambda_k| = k$. For instance, for a k step process, the set of times could be $\Lambda_k = \{t_k, \dots, t_1\}$. We will employ the convention that subscripts signify the time as well as the particular value of the respective random variable. For example, i_α signifies a value of the random variable at time t_α .

The distribution $\mathbb{P}_{\Lambda_k}(\mathbf{i}_k)$ could express the probability for a particular length- k sequence of heads and tails when flipping a coin, or the probability to find a particle undergoing Brownian motion at positions $\mathbf{i}_k$ when measuring it at times Λ_k . Importantly, this description of a stochastic process is sufficient to describe the behaviour on any subset of the times considered; for instance, the distribution over all but the j th time is found by marginalizing the larger distribution: $\mathbb{P}_{\Lambda_k \setminus \{t_j\}}(i_k, \dots, i_{j+1}, i_{j-1}, \dots, i_1) = \sum_{i_j} \mathbb{P}_{\Lambda_k}(\mathbf{i}_k)$. This property implicitly assumes that there is only one instrument that is used to interrogate the system of interest, and this interrogation does not influence its state. Neither of these assumption are fulfilled in more general processes, such as the ones employed in causal modelling.

2.2 Classical causal modelling

Observing the statistics for measurement outcomes reveals correlations between events, but no information about causal relations. For instance, correlations of two events A and B could stem from A influencing B , B influencing A , or both A and B being influenced by an earlier event C [11, 13, 15] (see Fig. 1). Reiterating an example from Ref. [15], events A and B could be the occurrence of sunburns and the sales of

ice cream, respectively. While these two variables are highly correlated, this correlation alone would not fix a causal relation between them. Inferring the causal structure of a process is the aim of *causal modelling*. Here, active interventions are used to uncover how different events can influence each other. In the example above, one could suspend the sale of ice cream to see how it affects the occurrence of sunburns, and would find out that ice cream sales have no direct effect on sunburns (and vice versa, as the correlations of ice cream sales and sunburns stem from a common cause, sunny weather, and not from any direct causal relation).

Mathematically, causal modelling for k events $A_k, \dots, A_1$ necessitates the collection of all joint probability distributions $\mathbb{P}_{\Gamma_k}(i_k, \dots, i_1 | j_k, \dots, j_1) := \mathbb{P}_{\Gamma_k}(\mathbf{i}_{\Gamma_k} | \mathcal{J}_{\Gamma_k})$ to measure the outcomes $i_k, \dots, i_1$ given that the *interventions* $j_k, \dots, j_1$ were performed. Here, Γ_k is a set of labels for events; *a priori*, there is no particular order imposed on the elements of Γ_k , and we use a different symbol for the set of event labels to distinguish it from the set of times Λ_k used above. For example, Γ_k could contain labels for different laboratories where experiments are performed. $\mathcal{J}_{\Gamma_k}$ are the *instruments* that were used at each of the events; these can be seen as rules for how to intervene upon seeing a particular outcome (we will formalize the notion of an instrument in Sec. 4). For example, when investigating Brownian motion, an instrument could be a deterministic replacement rule: upon finding the particle at i_a replace it by a particle at i'_a . It could also be probabilistic: upon finding the particle at i_a , with probability $p_{i'_a}$ replace it by a particle at i'_a .

One possible instrument is the trivial idle instrument $\mathcal{J}_a = \text{id}_a$, the instrument that only measures the particle but doesn't change it. For classical stochastic processes, the corresponding joint distribution over outcomes can be thought of as the instrument-independent underlying distribution of the random variables describing the process:

$$\mathbb{P}_{\Gamma_k}(\mathbf{i}_{\Gamma_k} | \mathbf{id}_{\Gamma_k}) = \mathbb{P}_{\Gamma_k}(\mathbf{i}_{\Gamma_k}), \quad (1)$$

where $\mathbf{id}_{\Gamma_k}$ denotes the idle instrument at each of the events in Γ_k . If we chose Γ_k to be a set of times Λ_k , the right-hand side of Eq. (1) has the

form of a k -step stochastic process. This directly leads to the following (well-known) observation:

Observation 1. *Classical causal modelling contains classical stochastic processes as a special case.*

As mentioned, this statement follows by choosing the set of events Γ_k to be a set Λ_k of ordered times and the instruments to be the idle instrument. We emphasize that causal modelling does not impose a temporal ordering *per se*, but deduces the ordering of events from the obtained correlation functions (finding this order, or, more precisely, the underlying *directed acyclic graph* (DAG) that defines the causal relations of the events, is the original aim of causal modelling [11, 13]). As neither the proof of the KET, nor the proof of the GET makes use of the notion of *a priori* temporal ordering (see Sec. 4.2 for a discussion), in what follows, we will drop the distinction between sets of labels Γ_k and sets of times Λ_k . We now show that the introduction of interventions, that is crucial for the deduction of causal relations, leads to a breakdown of fundamental properties that are satisfied by classical stochastic processes.

3 The Kolmogorov extension theorem and interventions

3.1 The KET

The KET is concerned with the question of which properties a family of finite joint probability distributions have to satisfy in order for an underlying process to exist. As such, it defines the classical notion of a stochastic process. In what follows, we will distinguish between stochastic processes on a finite number of times – which are characterized by joint probability distributions with finitely many arguments – and the *underlying* stochastic process that leads to *all* of these finite distributions.

As already mentioned, a classical stochastic process is described by a family of joint probability distributions $\mathbb{P}_{\Lambda_k}(\mathbf{i}_{\Lambda_k})$ for different finite sets of times Λ_k . An underlying process on a set Λ (finite, countably or uncountably infinite) is a joint probability distribution $\mathbb{P}_{\Lambda}$, that has all finite ones as marginals. In detail, we have $\mathbb{P}_{\Lambda_k}(\mathbf{i}_{\Lambda_k}) = \sum_{\Lambda \setminus \Lambda_k} \mathbb{P}_{\Lambda}(\mathbf{i}_{\Lambda}) := \mathbb{P}_{\Lambda}^{\Lambda_k}(\mathbf{i}_{\Lambda_k})$

for all $\Lambda_k \subseteq \Lambda$, where $\mathbf{i}_{\Lambda_k}$ is the subset of $\mathbf{i}_\Lambda$ corresponding to the times Λ_k , $\sum_{\Lambda \setminus \Lambda_k}$ denotes a sum over realizations of the random variables at all times that are part of $\Lambda \setminus \Lambda_k$ (*i.e.*, all the times that lie in Λ but not in Λ_k), and $\mathbb{P}_\Lambda^{\Lambda_k}$ denotes the restriction of $\mathbb{P}_\Lambda$ to the times Λ_k . In the case where the set Λ is infinite, the marginalization procedure can correspond to an integral over the times in $\Lambda \setminus \Lambda_k$ (though, to avoid introducing too much notation, we will still use $\sum_{\Lambda \setminus \Lambda_k}$ to represent it). For example, if the process we are interested in is the Brownian motion of a particle, $\mathbb{P}_\Lambda$ would be the probability density of all possible trajectories that the particle could take in the time interval Λ , and all finite distributions could in principle be obtained from $\mathbb{P}_\Lambda$.

If the finite joint probability distributions stem from an underlying process, it is easy to see that probability distributions for any two finite subsets of times $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$ fulfill a *consistency condition* (or *compatibility condition*) amongst each other, *i.e.*, $\mathbb{P}_{\Lambda_k}$ is a marginal of $\mathbb{P}_{\Lambda_\ell}$. Expressed in the notation introduced above, we have $\mathbb{P}_{\Lambda_k} = \mathbb{P}_{\Lambda_\ell}^{\Lambda_k}$ for all $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$. Intuitively, this means that $\mathbb{P}_{\Lambda_\ell}$, the descriptor of the stochastic process on the times Λ_ℓ , contains all information about subprocesses on fewer times.

While an underlying process leads to a family of compatible finite probability distributions, the KET shows that the converse is also true. Any family of consistent probability distributions implies the existence of an underlying process. Specifically, the Kolmogorov extension theorem [3–6] defines the minimal properties finite probability distributions have to satisfy in order for an underlying process to exist:

Theorem. [Kolmogorov] *Let Λ be a set of times. For each finite $\Lambda_k \subseteq \Lambda$, let $\mathbb{P}_{\Lambda_k}$ be a (sufficiently regular) k -step joint probability distribution. There exists an underlying stochastic process $\mathbb{P}_\Lambda$ that satisfies $\mathbb{P}_{\Lambda_k} = \mathbb{P}_\Lambda^{\Lambda_k}$ for all finite $\Lambda_k \subseteq \Lambda$ iff $\mathbb{P}_{\Lambda_k} = \mathbb{P}_{\Lambda_\ell}^{\Lambda_k}$ for all $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$.*

In other words, if a family of joint probability distributions on finite sets of times satisfies a consistency condition there is an underlying stochastic process on Λ that has the distributions $\{\mathbb{P}_{\Lambda_k}\}_{\Lambda_k \subseteq \Lambda}$ as marginals. More precisely, the Kolmogorov extension theorem guarantees the existence of a probability *measure* on an infinite

product of measurable spaces if the respective measures on said spaces are compatible with each other in the above sense, and each of the measurable spaces is *inner regular* [6]. That is, the measure of any set can be approximated by that of compact subsets. As we will consider our value spaces (*i.e.*, the spaces of possible outcomes) to be $\mathbb{R}$ or $\mathbb{N}$ throughout this paper, the requirement of inner regularity of the considered probability distributions will always be automatically satisfied.¹

As stated above, the KET defines the notion of a classical stochastic process and reconciles the existence of an underlying process with its manifestations for finite times. It also enables the modelling of stochastic processes: Any mechanism that leads to finite joint probability distributions that satisfy a consistency condition is ensured to have an underlying process. For example, the proof of the existence of Brownian motion relies on the KET as a fundamental ingredient [19–22].

We emphasize that, in the (physically relevant) case where Λ is an infinite set, the probability distribution $\mathbb{P}_\Lambda$ is generally not experimentally accessible. For example, in the case of Brownian motion, the set Λ could contain all times in the interval $[0, t]$ and each realization $\mathbf{i}_\Lambda$ would represent a possible continuous trajectory of a particle over this time interval. While we assume the existence of these underlying trajectories (and hence the existence of $\mathbb{P}_\Lambda$) in experiments concerning Brownian motion, we often only access their finite time manifestations, *i.e.*, $\mathbb{P}_{\Lambda_k}$ for some Λ_k . The KET bridges the gap between the finite experimental reality and the underlying infinite stochastic process.

Loosely speaking, the KET holds for classical stochastic processes, because there is no difference between ‘doing nothing’ and conducting a measurement but ‘not looking at the outcomes’ (*i.e.*, summing over the outcomes at a time). Put differently, the validity of the KET is based on the fundamental assumption that the interrogation of a system does not, on average, influence its state. Consequently, marginalization is the correct way to obtain the descriptor for fewer times and any classical

¹In particular, both $\mathbb{R}$ and $\mathbb{N}$ equipped with their standard topology are *Borel spaces*. A probability measure on such a space is inner regular [6].

Figure 1: *(Quantum) Causal network.* Performing different interventions allows for the causal relations between events to be probed. For example, in the figure the event B_1 directly influences the events C_3 and A_2 , while A_3 influences only B_4 . Depending on the degrees of freedom that can be accessed by the experimenter, these causal relations can or cannot be detected. For example, the influence of A_3 on B_4 could not be discovered if only the degrees of freedom in the gray areas were experimentally accessible. Independent of the accessible degrees of freedom, the generalized extension theorem (GET) that we derive below holds for *any* process. On the other hand, the statistics of events do in general not satisfy the requirements of the KET. For example, the events D_3, D_4, B_5 could be successive (e.g., at times t_3, t_4 and t_5) spin measurements in z -, x - and z -direction, respectively. Summing over the results of the spin measurement in x -direction at t_4 would not yield the correct probability distribution for two measurements in z -direction at t_3 and t_5 only (see also Sec. 4.1).

stochastic process leads to compatible finite joint probability distributions; this compatibility is independent of whether the system was interrogated or not, and the converse also holds. This fails to be true in causal modelling scenarios.

3.2 The KET and causal modelling

The compatibility of joint probability distributions for different sets of times hinges on the fact that observations in classical physics do not alter the state of the system that is being observed. In contrast to passive interrogations, that merely reveal information, active interventions, like they are used in the case of causal modelling, on average change the state of the interrogated system. Thus the future statistics after an intervention crucially depends on how the system was manipulated and the prerequisite of compatible joint probability distributions is generally not fulfilled anymore.

Figure 2: **(a)** A possible three step process without intervention is the drawing with replacement of colored balls from an urn. In our example, independent of the actions of the experimenter, a red ball drops into the urn at t_2 (this could, e.g., represent the interaction with an uncontrollable environment.) The experimenter can deduce the joint probability distribution $\mathbb{P}_{\{t_3, t_2, t_1\}}(c_3, c_2, c_1)$, to draw different sequences of colors. $\mathbb{P}_{\{t_3, t_2, t_1\}}$ contains all distributions for fewer times, for example $\mathbb{P}_{\{t_3, t_1\}}$, $\mathbb{P}_{\{t_3, t_2\}}$, and $\mathbb{P}_{\{t_1\}}$. **(b)** Instead of putting the same ball back in the urn, the experimenter could exchange it with a different color (for example, upon drawing yellow, they could replace it with green at t_1 , replace blue with white at t_2 and replace red with blue at t_3). The respective replacement rules are encapsulated in the instruments $\mathcal{J}_3$, $\mathcal{J}_2$, and $\mathcal{J}_1$. Now, from the probability distribution $\mathbb{P}_{\{t_3, t_2, t_1\}}(c_3, c_2, c_1 | \mathcal{J}_3, \mathcal{J}_2, \mathcal{J}_1)$, it is generally not possible to deduce probability distributions for fewer steps, like, e.g., $\mathbb{P}_{\{t_3, t_1\}}(c_3, c_1 | \mathcal{J}_3, \mathcal{J}_1)$, or $\mathbb{P}_{\{t_3, t_2\}}(c_3, c_2 | \mathcal{J}_3, \mathcal{J}_2)$. This lack of consistency can not be remedied by simple relabeling of the times due to the red ball that drops into the urn at t_2 . Note that if all instruments are the idle instrument, the case with interventions coincides with the case without interventions.

Consider, for example, the case of a pharmacologist that tries to understand the effect of different drugs they developed on a disease. In our simplified example, let the disease have two different symptoms S_a and S_b , and denote the absence of symptoms S_c . Whenever the pharmacologist observes S_a , they administer drug D_a , whenever they observe

symptom S_b they administer drug D_b , and whenever they observe S_c they do nothing; this choice of actions defines an instrument $\mathcal{J}$. Running their trial with sufficiently many patients, the pharmacologist can deduce probability distributions for the occurrence of symptoms over time, given the drugs that were administered. For example, if the drugs were administered on three consecutive days, they would have obtained a probability distribution of the form $\mathbb{P}_{\Lambda_3}(s_3, s_2, s_1 | \mathcal{J}_3 = \mathcal{J}, \mathcal{J}_2 = \mathcal{J}, \mathcal{J}_1 = \mathcal{J})$, where $s_\alpha \in \{S_a, S_b, S_c\}$, and the instruments (*i.e.*, the drug administration rule) are the same each day. However, this data would not allow them to find out what would have happened, had they *not* administered drugs on day two, *i.e.*, $\sum_{s_2} \mathbb{P}_{\Lambda_3}(s_3, s_2, s_1 | \mathcal{J}_3, \mathcal{J}_2, \mathcal{J}_1) \neq \mathbb{P}_{\Lambda_3 \setminus \{t_2\}}(s_3, s_1 | \mathcal{J}_3, \mathcal{J}_1)$; intermediate interventions change the state of the interrogated system, and hence the future statistics that are being observed; for another, more numerically tangible example, see Fig. 2. Consequently, probability distributions do generally not satisfy compatibility conditions when interventions are allowed.

Compatibility can fail to hold whenever the system of interest is actively interrogated. In particular, it fails to hold in quantum mechanics, where even projective measurements in general change the state of a system on average, and interventions are not just an experimental choice but unavoidable.

4 (Quantum) Stochastic processes with interventions

4.1 The KET in QM

As hinted at throughout this work, descriptions of quantum mechanical processes must necessarily account for the fundamental invasiveness of measurements, which renders the KET invalid for the same reason that some choices of intervention do in the case of classical causal modelling. To see how even projective measurements in quantum mechanics lead to families of probability distributions that do not satisfy the KET, consider the following concatenated Stern-Gerlach experiment: Let the initial state of a spin- $\frac{1}{2}$ particle be $|+\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$, where $|\uparrow\rangle$ and $|\downarrow\rangle$ are the spin-up and

spin-down state in the z -direction, respectively. Now, we measure the state sequentially in the z -, x - and z -directions at times t_1, t_2 and t_3 . These measurements have the possible outcomes $\{\uparrow, \downarrow\}$ and $\{\rightarrow, \leftarrow\}$ for the measurement in z - and x -direction, respectively. It is easy to see that the probability for any possible sequence of outcomes is equal to $1/8$. For example, we have

$$\begin{aligned} \mathbb{P}_{\Lambda_3}(\uparrow, \rightarrow, \uparrow | \mathcal{J}_z, \mathcal{J}_x, \mathcal{J}_z) \\ = \mathbb{P}_{\Lambda_3}(\uparrow, \leftarrow, \uparrow | \mathcal{J}_z, \mathcal{J}_x, \mathcal{J}_z) = \frac{1}{8}, \end{aligned} \quad (2)$$

where $\mathcal{J}_z$ and $\mathcal{J}_x$ represent the instruments used to measure in the z - and x - direction respectively, and $\Lambda_3 = \{t_3, t_2, t_1\}$. Summing over the outcomes at time t_2 , we obtain the marginal probability $\mathbb{P}_{\Lambda_3}^{\{t_1, t_3\}}(\uparrow, \uparrow | \mathcal{J}_z, \mathcal{J}_z) = 1/4$. However, by considering the case where the measurement is not made at t_2 , it is easy to see that $\mathbb{P}_{\{t_3, t_1\}}(\uparrow, \uparrow | \mathcal{J}_z, \mathcal{J}_z) = 1/2$. The intermediate measurement changes the state of the system, and the corresponding probability distributions for different sets of times are not compatible anymore [16, 23].

It is important to highlight the close relation of this breakdown of consistency and the violation of Leggett-Garg inequalities in quantum mechanics [7, 9]. The assumption of consistency between descriptors for different sets of times that is crucial for the derivation of the KET subsumes the assumptions of realism *per se* and noninvasive measurability that are the basic principles leading to the derivation of Leggett-Garg inequalities: While realism *per se* implies that joint probability distributions for a set of times can be expressed as marginals of a joint probability distribution for more times, non-invasiveness means that all finite distributions are marginals of the *same* distribution. For example, the two-step joint probability distributions $\mathbb{P}_{\{t_2, t_1\}}$, $\mathbb{P}_{\{t_3, t_2\}}$, and $\mathbb{P}_{\{t_3, t_1\}}$, that are considered in the Leggett-Garg scenario are all marginals of a three-step distribution $\mathbb{P}_{\{t_3, t_2, t_1\}}$. As soon as non-invasiveness and/or realism *per se* do not hold, the KET can fail and Leggett-Garg inequalities can be violated.

Nevertheless, there should be some compatibility between descriptors for different sets of times; the breakdown of the KET should be a problem of the formalism rather than a physical fact. We

now show that a change of perspective enables one to prove an extension theorem in quantum mechanics and any theory with interventions.

4.2 Instruments and Combs

Processes involving interventions, including quantum processes and those in classical causal modelling, do not lead to compatible joint probability distributions for different sets of times in general. This problem can be remedied by assuming the standpoint of quantum causal modelling, and choosing a description of such stochastic processes that takes interventions and their corresponding change of the system into account. With this description, it is possible to recover a compatibility property that is satisfied by any process with interventions, and a generalized extension theorem can be derived.

As in the classical causal modelling case, in quantum mechanics, an experimenter can choose an instrument $\mathcal{J}_\alpha$ at each time t_α , and every outcome i_α corresponds to a particular transformation of the system that is interrogated. Denoting the Hilbert space of the system at t_α by $\mathcal{H}_\alpha$, mathematically, an observation of outcome i_α given the instrument $\mathcal{J}_\alpha$ corresponds to a (trace non-increasing) completely positive (CP) map $\mathcal{M}_{i_\alpha} : \mathcal{B}_1(\mathcal{H}_\alpha^i) \rightarrow \mathcal{B}_1(\mathcal{H}_\alpha^o)$ that describes the change of the state of the system [24, 25]. Here, $\mathcal{B}_1(\mathcal{H}_\alpha)$ denotes the space of trace class operators on $\mathcal{H}_\alpha$ – which in the finite dimensional case coincides with the set of bounded operators $\mathcal{B}(\mathcal{H}_\alpha)$ on $\mathcal{H}_\alpha$ – and we account for the possibility that $\mathcal{M}_{i_\alpha}$ creates or discards degrees of freedom by distinguishing between its input (i) and output (o) spaces. The set of possible CP maps an instrument comprises add up to a completely positive trace preserving (CPTP) map $\mathcal{M}_\alpha = \sum_{i_\alpha} \mathcal{M}_{i_\alpha}$, which describes the overall average transformation applied by the instrument. While the CP maps the instrument comprises can only be implemented probabilistically, the corresponding overall CPTP map can be performed deterministically, *i.e.*, with unit probability. In contrast to classical physics without interventions, where the introduction of maps (or *events* more generally [23]) is superfluous, it is fundamentally unavoidable in quantum mechanics, as well as in more general probabilistic theories [23, 26]; every measurement alters the state of the system of interest, and a

full description of a temporal process necessitates knowledge of how the system is changed at each time.

As in the example from Sec. 4.1, an experimenter could choose to measure a system in different bases. The projective measurement in the basis $\{|i\rangle\}$ at t_α of a state ρ that yields outcome i_α would be described by a CP map of the form $\mathcal{M}_{i_\alpha}[\rho] = \langle i_\alpha | \rho | i_\alpha \rangle |i_\alpha\rangle\langle i_\alpha|$, where $|i_\alpha\rangle \in \{|i\rangle\}$. More generally, the measuring instrument need not preserve the measured state of the system, but could replace it entirely; upon measuring an outcome i_α (corresponding to a projection on a state $|i_\alpha\rangle$), a different instrument could leave the system in the state ρ_{i_α} , with a resulting CP map $\mathcal{M}_{i_\alpha}[\rho] = \langle i_\alpha | \rho | i_\alpha \rangle \rho_{i_\alpha}$. In the most general case, the experimenter – at a time t_α – could perform *any* (trace non-increasing) CP map, including deterministic operations such as unitary transformations. We will employ the convention that, for a given instrument $\mathcal{J}_\alpha$, the CP map corresponding to the outcome i_α is denoted by $\mathcal{M}_{i_\alpha}$, and we will denote the complex vector space that is spanned by all of these CP maps by $\mathcal{L}_\alpha$.

In this language, each realization of an experiment corresponds to a (possibly temporally correlated, see below) sequence of CP maps that transform the system at a series of times. The set of possible CP maps that could be applied is dictated by the choice of instruments used to interrogate the system in question. A quantum process is fully characterized once *all* of the probabilities $\mathbb{P}_{\Lambda_k}(i_k, \dots, i_1 | \mathcal{J}_k, \dots, \mathcal{J}_1)$ for each such sequence with all possible instruments are known. Having all of these probability distributions at hand allows one to deduce the causal structure of a process, *i.e.*, it is the basis of quantum causal modelling [13].

Written more succinctly, a k -step quantum process is fully characterized by an object $\mathcal{T}_{\Lambda_k}$ that maps sequences of CP maps to probabilities. Specifically, this means that $\mathcal{T}_{\Lambda_k}$ is a CP map $\mathcal{T}_{\Lambda_k} : \mathcal{L}_{\Lambda_k} \rightarrow \mathbb{C}$, where $\mathcal{L}_{\Lambda_k} = \bigotimes_{\alpha \in \Lambda_k} \mathcal{L}_\alpha$, such that $\mathcal{T}_{\Lambda_k}[\mathcal{M}_{i_k}, \dots, \mathcal{M}_{i_1}]$ yields the probability $\mathbb{P}_{\Lambda_k}(i_k, \dots, i_1 | \mathcal{J}_k, \dots, \mathcal{J}_1)$ to obtain the outcomes $i_k, \dots, i_1$ given the choices of instruments $\mathcal{J}_k, \dots, \mathcal{J}_1$ (see Fig. 3). In this sense, $\mathcal{T}_{\Lambda_k}$ represents a Born rule for temporal processes [27, 28]. The mapping $\mathcal{T}_{\Lambda_k}$ is a completely positive multilinear functional

that can be reconstructed in a finite number of experiments [12, 29–31].

For example, a k -step quantum process could be of the form

$$\begin{aligned} \mathcal{T}_{\Lambda_k}[\mathcal{M}_{i_k}, \dots, \mathcal{M}_{i_1}] \\ = \text{tr}\{\mathcal{M}_{i_k} \circ \mathcal{E}_{k-1} \circ \mathcal{M}_{i_{k-1}} \circ \mathcal{E}_{k-2} \cdots \\ \cdots \circ \mathcal{M}_{i_2} \circ \mathcal{E}_1 \circ \mathcal{M}_{i_1}[\rho]\}, \end{aligned} \quad (3)$$

where $\{\mathcal{E}_i\}$ are CPTP maps and ρ is a fixed state of the system of interest. More generally, $\mathcal{T}_{\Lambda_k}$ can describe a process with memory – *i.e.*, a non-Markovian process [30] – in which case it would be of the form

$$\begin{aligned} \mathcal{T}_{\Lambda_k}[\mathcal{M}_{i_k}, \dots, \mathcal{M}_{i_1}] \\ = \text{tr}\{(\mathcal{M}_{i_k} \otimes \mathcal{I}_e) \circ \mathcal{U}_{k-1} \circ (\mathcal{M}_{i_{k-1}} \otimes \mathcal{I}_e) \circ \cdots \\ \cdots \circ \mathcal{U}_1 \circ (\mathcal{M}_{i_1} \otimes \mathcal{I}_e)[\rho_{se}]\}, \end{aligned} \quad (4)$$

where ρ_{se} is a (possibly correlated) state on the system of interest and an additional environment, and the maps $\{\mathcal{U}_i\}$ can be chosen to be unitary evolutions on the system-environment space. As information from the past can be propagated through the additional environment, processes that are described via Eq. (4) generally display non-trivial memory effects [30].

If Λ_k corresponds to an ordered set of times, then *every* k -step process on Λ_k can be written in the form of Eq. (4) [30, 32]. On the other hand, if the process at hand were to not abide by a clear causal order, then it would not possess a representation of this form, but could still display non-trivial correlations between different events [33, 34]. The respective causal ordering that $\mathcal{T}_{\Lambda_k}$ is compatible with imposes further requirements on its properties; besides being CP, $\mathcal{T}_{\Lambda_k}$ has to be properly normalized. Naturally, $\mathcal{T}_{\Lambda_k}$ has to yield unit probability when acting on an operation that can be implemented deterministically. Consequently, we have

$$\mathcal{T}_{\Lambda_k}[\mathcal{M}_k, \dots, \mathcal{M}_1] = 1 \quad (5)$$

for all sequences $\mathcal{M}_k, \dots, \mathcal{M}_1$ of CPTP maps. Additionally, depending on the underlying causal structure, the respective deterministic operations an experimenter performs could be correlated – both in a classical and a quantum way. For example, if Λ_k is an ordered set of times, the choice of instrument used to interrogate the system at a time $t_{\alpha'} \in \Lambda_k$ could be conditioned

on the outcomes of all measurements at times $\Lambda_k \ni t_\alpha < t_{\alpha'}$. Put more generally, in quantum mechanics, deterministic operations are *all* operations that can be realized by preparing ancillary states, performing unitary operations, and discarding degrees of freedom in a temporally ordered fashion that agrees with the ordering given by Λ_k .

On the other hand, in situations where Λ_k does not correspond to a temporally ordered set, or, more generally, a partially ordered set corresponding to a definite causal order [32], but rather a set of labels for different laboratories with an unclear causal ordering, then such non-trivially correlated operations are not considered to be experimentally implementable, and the corresponding linear functionals $\mathcal{T}_{\Lambda_k}$ only have to satisfy Eq. (5) (and be CP), with no additional restrictions [34].

Finally, in operational probabilistic theories (OPTs), the set of deterministically implementable operations is the set of operations obtained from concatenating deterministic *preparations*, deterministic *transformations*, and deterministic *effects* in a well-defined order that agrees with the causal structure of the respective OPT [35].

In anticipation of later proofs, it is worth briefly discussing the set of operations that proper physical linear functionals on k times can be meaningfully applied to. As mentioned above, in quantum mechanics, a deterministically implementable operation is one that can be decomposed as a causally ordered concatenation of state preparations, unitary operations and a final discarding of degrees of freedom. Analogously, a general (temporally correlated) quantum measurement that an experimenter performs can always be decomposed as a causally ordered concatenation of state preparations, unitary operations and a final projective measurement on some degrees of freedom [32].

More concretely, setting $\mathcal{H}_{\Lambda_k} = \bigotimes_{\alpha \in \Lambda_k} \mathcal{H}_\alpha$, such scenarios would be described by a collection of temporally non-local (trace non-increasing) CP maps $\mathcal{M}_\gamma^{\Lambda_k} : \mathcal{B}_1(\mathcal{H}_{\Lambda_k}^i) \rightarrow \mathcal{B}_1(\mathcal{H}_{\Lambda_k}^o)$, where each of the maps $\mathcal{M}_\gamma^{\Lambda_k}$ corresponds to a possible measurement outcome. Overall, the respective probabilities have to add to unity, which implies that $\mathcal{M}_{\Lambda_k} = \sum_\gamma \mathcal{M}_\gamma^{\Lambda_k}$ corresponds to a deterministic operation. Such a collection

$\{\mathcal{M}_\gamma^{\Lambda_k}\}$ of maps is known as a *generalized instrument* or *tester* in the literature [27, 32, 36] (see [32] for a thorough discussion of their properties).

Naturally, a physically reasonable mapping $\mathcal{T}_{\Lambda_k}$ for such a causal scenario has to satisfy $0 \leq \mathcal{T}_{\Lambda_k}[\mathcal{M}_\gamma^{\Lambda_k}] \leq 1$ for all possible tester elements on Λ_k and $\mathcal{T}_{\Lambda_k}[\mathcal{M}_{\Lambda_k}] = 1$ for any temporally correlated operation that can be implemented deterministically. Likewise, in scenarios where no *a priori* causal order is assumed, $\mathcal{T}_{\Lambda_k}$ only has to satisfy Eq. (5) as well as $0 \leq \mathcal{T}_{\Lambda_k}[\mathcal{M}_{i_1}, \dots, \mathcal{M}_{i_k}] \leq 1$ for any collection $\{\mathcal{M}_{i_k}\}$ of CP maps. Finally, in OPTs, ‘tester elements’ are given by all operations that can be realized by concatenating deterministic preparations, deterministic transformations, and *probabilistic* effects in an order that agrees with Λ_k . As before, a physically reasonable mapping $\mathcal{T}_{\Lambda_k}$ would have to yield unit probability on deterministically implementable operations, and a probability $p \in [0, 1]$ for any operation that cannot be implemented deterministically.

In each case, the set of operations that a physically reasonable functional $\mathcal{T}_{\Lambda_k}$ can be applied to forms a convex set: if $\{\mathcal{M}_\gamma^{\Lambda_k}\}$ is a set of (probabilistically and deterministically) implementable operations, then

$$\mathcal{T}_{\Lambda_k} \left[\sum_\gamma \mu_\gamma \mathcal{M}_\gamma^{\Lambda_k} \right] \in [0, 1]. \quad (6)$$

where $\mu_\gamma \geq 0$ and $\sum_\gamma \mu_\gamma \leq 1$. We will denote this set of operations on times Λ_k that $\mathcal{T}_{\Lambda_k}$ can meaningfully act on $\mathcal{K}_{\Lambda_k}$.

Since the respective normalization of the families of linear functionals $\mathcal{T}_{\Lambda_k}$ that we consider, as well as the particular underlying theory is not of importance for the proof and/or the applicability of the main theorem of the paper, we will assume an agnostic standpoint and adhere to the following convention:

Throughout the remainder of this Paper, a multilinear functional $\mathcal{T}_{\Lambda_k} : \mathcal{L}_{\Lambda_k} \rightarrow \mathbb{C}$ that is positive on $\mathcal{K}_{\Lambda_k}$, and which yields unit probability on *all* deterministically implementable operations that the underlying causal structure (and underlying theory) allows, will be referred to as a *k-step comb*, following Refs. [32, 37, 38]. For the particular case of quantum mechanics, we would additionally demand that the comb is CP.

Note that, in the convention above, $\mathcal{K}_{\Lambda_k}$ and

Figure 3: *Graphical representation of a four step quantum comb.* A four step comb can be represented as an object with four slots (each slot corresponds to a time $t_\alpha \in \Lambda_4$); it encodes all multi-time correlations between observables at those times. The probability to observe the outcomes $i_4, \dots, i_1$ given the instruments $\mathcal{J}_4, \dots, \mathcal{J}_1$ is obtained by inserting the corresponding CP maps into the comb, *i.e.*, by letting $\mathcal{T}_{\Lambda_4}$ act on $\mathcal{M}_{i_4}, \dots, \mathcal{M}_{i_1}$.

its associated set of k -step combs depend on the respective scenario one considers. Additionally, as $\mathcal{T}_{\Lambda_k}$ is positive on all elements of $\mathcal{K}_{\Lambda_k}$, we automatically have $\mathcal{T}_{\Lambda_k}[\mathcal{M}_\gamma^{\Lambda_k}] \leq 1$, since for any $\mathcal{M}_\gamma^{\Lambda_k} \in \mathcal{K}_{\Lambda_k}$, there exists a $\mathcal{M}_{\gamma'}^{\Lambda_k} \in \mathcal{K}_{\Lambda_k}$ such that $\mathcal{M}_\gamma^{\Lambda_k} + \mathcal{M}_{\gamma'}^{\Lambda_k} \in \mathcal{K}_{\Lambda_k}$ is a deterministically implementable operation. Importantly, outside the set $\mathcal{K}_{\Lambda_k}$, $\mathcal{T}_{\Lambda_k}$ can yield ‘probabilities’ that exceed 1, even when acting on maps that may be deterministically implementable in other scenarios. This is due to the fact that the action of a comb $\mathcal{T}_{\Lambda_k}$ on a deterministic map that is not causally ordered in a compatible way can lead to causal loops, and thus non-sensical results; for example, letting a comb that is ordered t_1 before t_2 act on an operation that is ordered t_2 before t_1 will lead to causal loops and is, as such, not meaningful.

In what follows, we will predominantly phrase our statements with respect to CP maps (*i.e.*, for the case of quantum mechanics), with the understanding that generalization to other theories is always straightforwardly possible.

A comb $\mathcal{T}_{\Lambda_k}$ contains all the multi-time correlations necessary to fully characterize a k -step quantum process. While the CP maps $\mathcal{M}_{i_\alpha}$ change the state of the system, they do not change the k -step process given by $\mathcal{T}_{\Lambda_k}$. Loosely speaking, the comb contains all parts of the dynamics that are not manipulated by the experimenter. This is analogous to the way in which the preparation of an initial state and the measurement of the final state in

quantum process tomography do not influence the underlying dynamics (*i.e.*, the CPTP map connecting input and output state).

Just as in the classical case, the knowledge of all relevant joint probability distributions (*i.e.*, the knowledge of $\mathcal{T}_{\Lambda_k}$) allows one to deduce causal relations between the k events in Λ_k . We emphasize that classical causal modelling is included in this quantum causal modelling framework as a special case. Whenever a system is measured and prepared in a fixed basis (using a classical instrument), and the process $\mathcal{T}_{\Lambda_k}$ also preserves this basis, the result is a set of joint distributions consistent with a classical causal model. From Observation 1, it also follows that classical stochastic processes without interventions can be described by the same framework.

4.3 Generalized extension theorem (GET)

With the complete description on finite sets of times at hand, we can determine the compatibility condition between related combs. A family of combs that stems from an underlying (open) dynamics fulfills a natural consistency condition [30]; for any two sets of times $\Lambda_k \subseteq \Lambda_\ell$, the comb $\mathcal{T}_{\Lambda_k}$ can be obtained from $\mathcal{T}_{\Lambda_\ell}$ by letting it act on identity operations $\mathcal{I}_{t_\alpha}$ (with $\mathcal{I}_{t_\alpha}[\rho] = \rho$ for any state of the system ρ at time t_α) at times $t_\alpha \in \Lambda_\ell \setminus \Lambda_k$, *i.e.*,

$$\mathcal{T}_{\Lambda_k}[\cdot] = \mathcal{T}_{\Lambda_\ell} \left[\bigotimes_{t_\alpha \in \Lambda_\ell \setminus \Lambda_k} \mathcal{I}_{t_\alpha}, \cdot \right] := \mathcal{T}_{\Lambda_\ell}^{\Lambda_k}[\cdot], \quad (7)$$

where we have employed the shorthand notation $\bigotimes_{t_\alpha \in \Lambda_\ell \setminus \Lambda_k} \mathcal{I}_{t_\alpha}$ to signify that the identity operation was ‘implemented’ at each time $t_\alpha \in \Lambda_\ell \setminus \Lambda_k$. The graphical representation of Eq. (7) is depicted in Fig. 4.

It is important to note the difference between Eq. (7) and the consistency condition for classical stochastic processes, stemming from the stronger notion of ‘doing nothing’ in the quantum case. If there is an underlying process, any comb can be obtained from one that applies to a larger set of times by letting it act on the identity map, which leaves any state unchanged, at the excessive times. This is by no means the same as computing the marginals of families of probability distributions that have been obtained for a fixed set of measurement instruments, which

Figure 4: *Consistency condition for combs.* If there is an underlying process, any comb $\mathcal{T}_{\Lambda_k}$ can be obtained from $\mathcal{T}_{\Lambda_\ell}$ by letting $\mathcal{T}_{\Lambda_\ell}$ act on the identity map at the excessive times. Here, for the sets of times $\Lambda_{13} = \{t_{13}, \dots, t_1\}$, $\Lambda_8 = \{t_{13}, t_{12}, t_{11}, t_9, t_7, t_6, t_3, t_1\}$ and $\Lambda_5 = \{t_{13}, t_{12}, t_6, t_3, t_1\}$ we depict the containment of the comb $\mathcal{T}_{\Lambda_8}$ in $\mathcal{T}_{\Lambda_{13}}$ and the containment of $\mathcal{T}_{\Lambda_5}$ in both $\mathcal{T}_{\Lambda_{13}}$ and $\mathcal{T}_{\Lambda_8}$.

will only preserve states which are diagonal in a fixed basis. We recover descriptors for different sets of times that are compatible with each other only when we switch to a causal modelling description of the process. From this, we obtain our main result, the generalized extension theorem (GET):

Theorem (Generalized extension theorem). *Let Λ be a set of times. For each finite $\Lambda_k \subseteq \Lambda$ let $\mathcal{T}_{\Lambda_k}$ be a k -step comb. There exists a general stochastic process $\mathcal{T}_\Lambda$, *i.e.*, a multilinear CP functional defined on all times in Λ , that satisfies $\mathcal{T}_{\Lambda_k} = \mathcal{T}_\Lambda^{\Lambda_k}$, as defined in Eq. (7), for all finite $\Lambda_k \subseteq \Lambda$ iff $\mathcal{T}_{\Lambda_k} = \mathcal{T}_{\Lambda_\ell}^{\Lambda_k}$ for all $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$.*

Note that we make no assumption about the cardinality of the number of outcomes at each time t_α . The GET holds independently of whether there are finitely, countably or uncountably many possible outcomes.

Importantly, the GET is qualitatively distinct from quantum marginal problems [39]. The question it answers is *not* whether, for a given collection of combs $\mathcal{T}_{\Lambda_k}, \mathcal{T}_{\Lambda_\ell}, \dots$, there exists a comb $\mathcal{T}_{\Lambda_k \cup \Lambda_\ell \cup \dots}$ that has all of them as marginals. Rather, it starts from the assumption that there exists a family of combs $\{\mathcal{T}_{\Lambda_k}\}$ defined on *all* finite subsets of Λ and shows that this family can be extended to a comb $\mathcal{T}_\Lambda$ if the family $\{\mathcal{T}_{\Lambda_k}\}$ satisfies the consistency conditions laid out in the above theorem. For the case that Λ is a finite set, the theorem is thus trivial, as the desired comb is simply the comb of the family $\{\mathcal{T}_{\Lambda_k}\}$

that is defined on the largest set of times (that is, on Λ itself). The importance of the GET lies in the case where Λ is an infinite set. There, it shows that a family of finite combs satisfying proper consistency conditions suffices to define an underlying comb on all times in Λ .

The proof of the GET proceeds analogously to that of the original Kolmogorov extension theorem, presented, *e.g.*, in Ref. [6]. It can be broken up into two main parts: (i) The consistency property is used to define a unique comb $\mathcal{T}_\Lambda^\sharp$ on a ‘sufficiently large’ container space $\mathcal{L}_\Lambda^\sharp$. (ii) It is shown that $\mathcal{T}_\Lambda^\sharp$ is linear and bounded on a subset $\mathcal{K}_\Lambda^\sharp$ of said container space and can thus be extended to a linear functional $\overline{\mathcal{T}}_\Lambda^\sharp$ fulfilling the properties of the desired comb $\mathcal{T}_\Lambda$ on the closure of $\mathcal{K}_\Lambda^\sharp$. As in the classical case [5], the underlying stochastic process characterized by $\mathcal{T}_\Lambda$ is – unlike $\overline{\mathcal{T}}_\Lambda^\sharp$ – not necessarily unique. Since the action of all possible $\mathcal{T}_\Lambda$ coincides with the unique $\overline{\mathcal{T}}_\Lambda^\sharp$ on the correct set of operations, and hence yields the correct finite combs $\mathcal{T}_{\Lambda_k}$, this non-uniqueness cannot be detected experimentally and does not constitute a practical problem.

Proof. To begin with, a short comment on the spaces that the respective objects we will deal with are defined on is in order. Firstly, for ease of notation, we assume all CP maps that the combs act on to have the same input and output space, *i.e.*, they do not create or discard degrees of freedom.² Consequently, at each time t_α , we have $\mathcal{M}_{i_\alpha} : \mathcal{B}_1(\mathcal{H}_\alpha) \rightarrow \mathcal{B}_1(\mathcal{H}_\alpha)$. As before, for any finite Λ_k , we will employ the naming conventions $\mathcal{H}_{\Lambda_k}$, $\mathcal{B}_1(\mathcal{H}_{\Lambda_k})$, and $\mathcal{L}_{\Lambda_k}$ for the Hilbert space $\mathcal{H}_{\Lambda_k} = \bigotimes_{\alpha \in \Lambda_k} \mathcal{H}_\alpha$, the space of trace class operators thereon, and the vector space spanned by the CP maps on $\mathcal{B}_1(\mathcal{H}_{\Lambda_k})$, respectively. Notably, in the finite dimensional case, $\mathcal{B}_1(\mathcal{H}_{\Lambda_k})$ coincides with $\mathcal{B}(\mathcal{H}_{\Lambda_k})$, the space of bounded operators on $\mathcal{H}_{\Lambda_k}$, while in general we have $\mathcal{B}_1(\mathcal{H}_{\Lambda_k}) \subset \mathcal{B}(\mathcal{H}_{\Lambda_k})$. As the generalization to infinite dimensional Hilbert spaces $\mathcal{H}_\alpha$ does not bring additional technical difficulties, we will not assume $\dim(\mathcal{H}_\alpha)$ to be finite in what follows.

Since we will deal with operations at *each* time $t_\alpha \in \Lambda$, the natural Hilbert space to consider

is the – possibly uncountably infinite – tensor product $\mathcal{H}_\Lambda := \bigotimes_{\alpha \in \Lambda} \mathcal{H}_\alpha$. Such infinite products of Hilbert spaces have been defined in Def. 3.5.1 of [40]. For their construction, let $\{f_\alpha | f_\alpha \in \mathcal{H}_\alpha\}_{\alpha \in \Lambda}$ be a C -sequence, *i.e.*, a sequence of elements of the Hilbert spaces $\mathcal{H}_\alpha$ such that $\prod_{\alpha \in \Lambda} \|f_\alpha\|$ converges.³ On any such C -sequence, one can define particular linear functionals

$$\Phi^r(f_\alpha; \alpha \in \Lambda) := \prod_{\alpha \in \Lambda} (f_\alpha^r, f_\alpha), \quad (8)$$

where $\{f_\alpha^r | f_\alpha^r \in \mathcal{H}_\alpha\}_{\alpha \in \Lambda}$ is a C -sequence, and $(\cdot, \cdot)$ is the respective inner product of $\mathcal{H}_\alpha$. In a slight abuse of notation, this functional can be written $\Phi^r = \bigotimes_{\alpha \in \Lambda} f_\alpha^r$. With this, we can define the space $\mathcal{F}^0$ as the finite linear span of all such elements Φ^r , *i.e.*,

$$\mathcal{F}^0 = \{\Theta | \Theta = \sum_{r=1}^N \Phi^r\}. \quad (9)$$

This space is equipped with a well-defined scalar product [40]: Taking two elements $\Phi_1 = \sum_{r=1}^{N_1} \bigotimes_{\alpha \in \Lambda} f_\alpha^r$ and $\Phi_2 = \sum_{s=1}^{N_2} \bigotimes_{\alpha \in \Lambda} g_\alpha^s$ of $\mathcal{F}^0$, we can set

$$(\Phi_1, \Phi_2) := \sum_{r=1}^{N_1} \sum_{s=1}^{N_2} \prod_{\alpha \in \Lambda} (f_\alpha^r, g_\alpha^s). \quad (10)$$

With this, we obtain the space $\mathcal{H}_\Lambda$ as the set of all limits of Cauchy sequences in $\mathcal{F}^0$. Specifically, all elements $\varphi \in \mathcal{H}_\Lambda$ are such that there exists a Cauchy sequence $\{\Theta_s; \Theta_s \in \mathcal{F}^0\}$ that converges to φ in the sense that

$$\varphi(f_\alpha; \alpha \in \Lambda) = \lim_{s \rightarrow \infty} \Theta_s(f_\alpha; \alpha \in \Lambda), \quad (11)$$

for all C -sequences $\{f_\alpha; \alpha \in \Lambda\}$, where convergence is understood with respect to the metric induced by the scalar product in Eq. (10). To put it more intuitively, the space $\mathcal{F}^0$ is the space spanned by all infinite tensor products $\bigotimes_{\alpha \in \Lambda} f_\alpha$ which have a finite norm, and $\mathcal{H}_\Lambda$ is its completion. On this Hilbert space, we can define the space of trace class operators $\mathcal{B}_1(\mathcal{H}_\Lambda)$, and we shall denote the complex vector space spanned by CP maps $\mathcal{M} : \mathcal{B}_1(\mathcal{H}_\Lambda) \rightarrow \mathcal{B}_1(\mathcal{H}_\Lambda)$ by $\mathcal{L}_\Lambda$.

Intuitively, the general stochastic process $\mathcal{T}_\Lambda$, whose existence we will prove below, should be a

²A generalization to maps with distinct input and output spaces is straightforward.

³See [40] for a detailed discussion of the convergence of uncountably infinite products and sums.

functional of the form $\mathcal{T}_\Lambda : \mathcal{L}_\Lambda \rightarrow \mathbb{C}$. However, as we will see, the space $\mathcal{L}_\Lambda$ is slightly ‘too big’ for two distinct reasons, and $\mathcal{T}_\Lambda$ will only be uniquely defined on a smaller, naturally arising space. With these preliminary definitions of the involved spaces out of the way, we can now prove the first part of the theorem:

(i) Existence of a unique comb $\mathcal{T}_\Lambda^\sharp$ on the set Λ .

Let Λ be a (possibly uncountable) set, $\{\Lambda_k\}_{\Lambda_k \subseteq \Lambda}$ the set of all finite subsets of Λ , and let $\{\mathcal{T}_{\Lambda_k}\}_{\Lambda_k \subseteq \Lambda}$ be the corresponding family of combs. Consequently, we have $\mathcal{T}_{\Lambda_k} : \mathcal{L}_{\Lambda_k} \rightarrow \mathbb{C}$. Now, let the family of combs satisfy the consistency condition of Eq. (7) for all finite $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$.

With this, we can ‘lift’ the family of combs to a comb $\mathcal{T}_\Lambda^\sharp$ that acts on the full space $\mathcal{L}_\Lambda$ and contains all of the finite combs as ‘marginals’. To this end, we first define the inverse projection $\pi_{\Lambda_k}^{-1} : \mathcal{L}_{\Lambda_k} \rightarrow \mathcal{L}_\Lambda$ which trivially extends every map $\xi_{\Lambda_k} \in \mathcal{L}_{\Lambda_k}$ that is only defined on a finite number of times to an operator

$$\pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}] = \xi_{\Lambda_k} \otimes \bigotimes_{t_\alpha \in \Lambda \setminus \Lambda_k} \mathcal{I}_{t_\alpha} \quad (12)$$

that is defined on all times.⁴ In other words, $\pi_{\Lambda_k}^{-1}$ maps any ξ_{Λ_k} to a corresponding operator that lies in $\mathcal{L}_\Lambda$ and only acts non-trivially on $\mathcal{L}_{\Lambda_k}$. The operator $\pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}]$ exists and is unique for any finite $\Lambda_k \in \Lambda$ and all $\xi_{\Lambda_k} \in \mathcal{L}_{\Lambda_k}$ [40].

In the same way, we can define a partial inverse projection $\pi_{\Lambda_k \leftarrow \Lambda_\ell}^{-1} : \mathcal{L}_{\Lambda_k} \rightarrow \mathcal{L}_{\Lambda_\ell}$ for any two finite sets $\Lambda_k \subseteq \Lambda_\ell \subseteq \Lambda$, *i.e.*,

$$\pi_{\Lambda_k \leftarrow \Lambda_\ell}^{-1}[\xi_{\Lambda_k}] = \xi_{\Lambda_k} \otimes \bigotimes_{t_\alpha \in \Lambda_\ell \setminus \Lambda_k} \mathcal{I}_{t_\alpha}. \quad (13)$$

Employing these partial inverse projections, the consistency property of the family $\{\mathcal{T}_{\Lambda_k}\}$ reads

$$\mathcal{T}_{\Lambda_k}[\xi_{\Lambda_k}] = \mathcal{T}_{\Lambda_\ell}[\pi_{\Lambda_k \leftarrow \Lambda_\ell}^{-1}[\xi_{\Lambda_k}]]. \quad (14)$$

All of the lifted operators $\pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}]$ are elements of $\mathcal{L}_\Lambda$. Let $\mathcal{L}_\Lambda^\sharp \subset \mathcal{L}_\Lambda$ denote the set of all lifted operators, *i.e.*, for all Λ_k finite, we have

$$\mathcal{L}_\Lambda^\sharp = \{\xi \in \mathcal{L}_\Lambda \mid \xi = \xi_{\Lambda_k} \otimes \bigotimes_{t_\alpha \in \Lambda \setminus \Lambda_k} \mathcal{I}_{t_\alpha}\}. \quad (15)$$

⁴We will denote elements of $\mathcal{L}_{\Lambda_k}$ by ξ_{Λ_k} instead of $\mathcal{M}_{\Lambda_k}$ to emphasize that they are not necessarily CP maps, but rather complex linear combinations of CP maps.

On this space, we can define a comb $\mathcal{T}_\Lambda^\sharp$ via

$$\mathcal{T}_\Lambda^\sharp[\xi] := \mathcal{T}_{\Lambda_k}[\xi_{\Lambda_k}], \text{ where } \xi = \pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}]. \quad (16)$$

It remains to show that $\mathcal{T}_\Lambda^\sharp$ is well-defined in the sense that it maps every $\xi \in \mathcal{L}_\Lambda^\sharp$ to a unique value. Specifically, if there are two different operators $\xi_{\Lambda_k} \in \mathcal{L}_{\Lambda_k}$ and $\xi_{\Lambda_\ell} \in \mathcal{L}_{\Lambda_\ell}$ that are lifted to the same $\xi \in \mathcal{L}_\Lambda$, such that

$$\xi = \pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}] = \pi_{\Lambda_\ell}^{-1}[\xi_{\Lambda_\ell}], \quad (17)$$

then $\mathcal{T}_\Lambda^\sharp[\xi] = \mathcal{T}_{\Lambda_k}[\xi_{\Lambda_k}]$ but also $\mathcal{T}_\Lambda^\sharp[\xi] = \mathcal{T}_{\Lambda_\ell}[\xi_{\Lambda_\ell}]$ and $\mathcal{T}_\Lambda^\sharp$ might not be well-defined. However, uniqueness of $\mathcal{T}_\Lambda^\sharp[\xi]$ is ensured by the consistency property; from Eq. (17), it is straightforward to see that

$$\begin{aligned} \pi_{\Lambda_k \leftarrow \Lambda_k \cup \Lambda_\ell}^{-1}[\xi_{\Lambda_k}] &= \pi_{\Lambda_\ell \leftarrow \Lambda_k \cup \Lambda_\ell}^{-1}[\xi_{\Lambda_\ell}] \\ &=: \xi_{\Lambda_k \cup \Lambda_\ell}. \end{aligned} \quad (18)$$

Employing the consistency condition (14) then yields

$$\begin{aligned} \mathcal{T}_{\Lambda_k}[\xi_{\Lambda_k}] &= \mathcal{T}_{\Lambda_k \cup \Lambda_\ell}[\pi_{\Lambda_k \leftarrow \Lambda_k \cup \Lambda_\ell}^{-1}[\xi_{\Lambda_k}]] \\ &= \mathcal{T}_{\Lambda_k \cup \Lambda_\ell}[\xi_{\Lambda_k \cup \Lambda_\ell}] \\ &= \mathcal{T}_{\Lambda_k \cup \Lambda_\ell}[\pi_{\Lambda_\ell \leftarrow \Lambda_k \cup \Lambda_\ell}^{-1}[\xi_{\Lambda_\ell}]] = \mathcal{T}_{\Lambda_\ell}[\xi_{\Lambda_\ell}]. \end{aligned} \quad (19)$$

Consequently, $\mathcal{T}_\Lambda^\sharp[\xi]$ is independent of the representation of ξ . Additionally, by construction we have $\mathcal{T}_\Lambda^{\sharp|\Lambda_k} = \mathcal{T}_{\Lambda_k}$, *i.e.*, $\mathcal{T}_\Lambda^\sharp$ contains all finite combs of the family $\{\mathcal{T}_{\Lambda_k}\}$ as ‘marginals’. However, so far $\mathcal{T}_\Lambda^\sharp$ is only defined on the set $\mathcal{L}_\Lambda^\sharp \subset \mathcal{L}_\Lambda$. In the second part of the proof, we show that $\mathcal{T}_\Lambda^\sharp$ can be uniquely extended to a linear functional on a bigger space $\overline{\mathcal{L}}_\Lambda^\sharp \supseteq \mathcal{L}_\Lambda^\sharp$.

(ii) Extension of $\mathcal{T}_\Lambda^\sharp$.

In order to extend $\mathcal{T}_\Lambda^\sharp$ to a linear functional $\overline{\mathcal{T}}_\Lambda^\sharp$ that acts on elements of (a subset of) the closure $\overline{\mathcal{L}}_\Lambda^\sharp$ of $\mathcal{L}_\Lambda^\sharp$, we will make use of the fact that any linear bounded mapping from a subset $\mathcal{K}_\Lambda^\sharp$ of a normed vector space X to a normed complete vector space Y can be uniquely extended to a linear transformation from the completion $\overline{\mathcal{K}}_\Lambda^\sharp$ of $\mathcal{K}_\Lambda^\sharp$ to Y (see, *e.g.*, Thm. 2.7-11 of [41]).

So far, we have considered $\mathcal{T}_\Lambda^\sharp$ as a mapping $\mathcal{T}_\Lambda^\sharp : \mathcal{L}_\Lambda^\sharp \rightarrow \mathbb{C}$. However, $\mathcal{T}_\Lambda^\sharp$ is not necessarily bounded on $\mathcal{L}_\Lambda^\sharp$; as we discussed in Sec. 4.2, the action of a k -step comb $\mathcal{T}_{\Lambda_k}$ is only meaningfully

defined on the set $\mathcal{K}_{\Lambda_k}$ of maps that agree with the causal ordering of $\mathcal{T}_{\Lambda_k}$. In general, there will be many maps in $\mathcal{L}_{\Lambda_k}$ that have, for example, an opposite causal ordering than the one $\mathcal{T}_{\Lambda_k}$ abides by. Thus, the action of $\mathcal{T}_{\Lambda_k}$ on such a map would create causal loops and lead to ‘probabilities’ that exceed 1. The norm of $\mathcal{T}_{\Lambda_k}$ on $\mathcal{L}_{\Lambda_k}$ would thus depend on the number of possible causal loops, which, in turn, depends on the number of times in Λ_k . This, finally, implies that, in principle, for every positive number r , we could find a CPTP map $\mathcal{M}^r \in \mathcal{L}_{\Lambda}^{\sharp}$ with unit norm, such that $\mathcal{T}_{\Lambda}^{\sharp}[\mathcal{M}^r] > r$, rendering $\mathcal{T}_{\Lambda}^{\sharp}$ unbounded on $\mathcal{L}_{\Lambda}^{\sharp}$. While this unboundedness is not a problem for the first part of the proof, it keeps us from uniquely extending $\mathcal{T}_{\Lambda}^{\sharp}$ to a linear functional on the closure of $\mathcal{L}_{\Lambda}^{\sharp}$.

However, this is not a conceptual problem, as a CPTP map $\mathcal{M}^r$ that yields a probability higher than 1 cannot be implemented within the causal order the combs $\mathcal{T}_{\Lambda_k}$ abide by. Consequently, without losing generality, we can restrict ourselves to the respective subsets of operations that the combs $\mathcal{T}_{\Lambda_k}$ are meaningfully defined on. Specifically, we set

$$\mathcal{K}_{\Lambda}^{\sharp} := \{\xi \in \mathcal{L}_{\Lambda}^{\sharp} \mid \xi = \xi_{\Lambda_k} \otimes \bigotimes_{t_{\alpha} \in \Lambda \setminus \Lambda_k} \mathcal{I}_{t_{\alpha}}, \quad \xi_{\Lambda_k} \in \mathcal{K}_{\Lambda_k}\}, \quad (20)$$

i.e., $\mathcal{K}_{\Lambda}^{\sharp}$ contains all trivial extensions of operators that the finite combs $\mathcal{T}_{\Lambda_k}$ are meaningfully defined on. As $\mathcal{K}_{\Lambda}^{\sharp} \subset \mathcal{L}_{\Lambda}^{\sharp}$, $\mathcal{T}_{\Lambda}^{\sharp}$ is uniquely defined on $\mathcal{K}_{\Lambda}^{\sharp}$. Now, considering the mapping

$$\mathcal{T}_{\Lambda}^{\sharp} : \mathcal{K}_{\Lambda}^{\sharp} \rightarrow \mathbb{C}, \quad (21)$$

we can show that there exists a unique extension $\overline{\mathcal{T}}_{\Lambda}^{\sharp} : \overline{\mathcal{K}}_{\Lambda}^{\sharp} \rightarrow \mathbb{C}$ that has the desired properties. To this end, we have to show that $\mathcal{L}_{\Lambda}^{\sharp}$ is a normed vector space, and that $\mathcal{T}_{\Lambda}^{\sharp}$ is linear, and bounded on $\mathcal{K}_{\Lambda}^{\sharp}$ (the space $\mathbb{C}$ is well-known to be a complete vector space). We start with the former:

Let $\beta, \gamma \in \mathbb{C}$ and $\xi = \pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}], \zeta = \pi_{\Lambda_{\ell}}^{-1}[\zeta_{\Lambda_{\ell}}]$. It follows directly from the definition (12) of the inverse projection that

$$\beta \xi + \gamma \zeta = \pi_{\Lambda_k \cup \Lambda_{\ell}}^{-1}[\beta \xi_{\Lambda_k} + \gamma \zeta_{\Lambda_{\ell}}] \in \mathcal{L}_{\Lambda}^{\sharp} \quad (22)$$

Now, we define $\Gamma_{\Lambda_k \cup \Lambda_{\ell}} \in \mathcal{L}_{\Lambda_k \cup \Lambda_{\ell}}$ as

$$\Gamma_{\Lambda_k \cup \Lambda_{\ell}}$$

$$= \pi_{\Lambda_k \leftarrow \Lambda_k \cup \Lambda_{\ell}}^{-1}[\beta \xi_{\Lambda_k}] + \pi_{\Lambda_{\ell} \leftarrow \Lambda_k \cup \Lambda_{\ell}}^{-1}[\gamma \zeta_{\Lambda_{\ell}}]. \quad (23)$$

One immediately sees that

$$\beta \xi + \gamma \zeta = \pi_{\Lambda_k \cup \Lambda_{\ell}}^{-1}[\Gamma_{\Lambda_k \cup \Lambda_{\ell}}] \in \mathcal{L}_{\Lambda}^{\sharp}, \quad (24)$$

which implies that $\mathcal{L}_{\Lambda}^{\sharp}$ is a complex vector space. Additionally, $\mathcal{L}_{\Lambda}^{\sharp}$ becomes a *normed* vector space by setting $\|\xi\| = \|\pi_{\Lambda_k}^{-1}[\xi_{\Lambda_k}]\| := \|\xi_{\Lambda_k}\|_{\text{op}}$, where $\|\cdot\|_{\text{op}}$ is the norm on $\mathcal{L}_{\Lambda_k}$ induced by the trace norm on $\mathcal{B}_1(\mathcal{H}_{\Lambda_k})$, *i.e.*,

$$\|\xi_{\Lambda_k}\|_{\text{op}} = \sup_{|A|_{\text{tr}}=1} \{|\xi[A]|_{\text{tr}} \mid A \in \mathcal{B}_1(\mathcal{H}_{\Lambda_k})\}. \quad (25)$$

To prove linearity, and boundedness on $\mathcal{K}_{\Lambda}^{\sharp}$ of $\mathcal{T}_{\Lambda}^{\sharp}$, we make use of the linearity and boundedness of the finite combs $\mathcal{T}_{\Lambda_k}$: For all Λ_k , we have $\mathcal{T}_{\Lambda_k}[\mathcal{M}_{\Lambda_k}] \leq 1$ for all $\mathcal{M}_{\Lambda_k} \in \mathcal{K}_{\Lambda_k}$. As this bound is uniform, *i.e.*, independent of the set of times Λ_k , we immediately see that $\mathcal{T}_{\Lambda}^{\sharp}[\mathcal{M}_{\Lambda}] \leq 1$ for all $\mathcal{M}_{\Lambda} \in \mathcal{K}_{\Lambda}^{\sharp}$, rendering $\mathcal{T}_{\Lambda}^{\sharp}$ bounded on $\mathcal{K}_{\Lambda}^{\sharp}$.

The linearity of $\mathcal{T}_{\Lambda}^{\sharp}$ follows in a similar vein; due to the linearity of $\mathcal{T}_{\Lambda_k}$ and the linearity of the inverse projection operators, for all $\beta, \gamma \in \mathbb{C}$ and all $\eta, \xi \in \mathcal{L}_{\Lambda}^{\sharp}$ we have

$$\mathcal{T}_{\Lambda}^{\sharp}[\beta \xi + \gamma \eta] = \beta \mathcal{T}_{\Lambda}^{\sharp}[\xi] + \gamma \mathcal{T}_{\Lambda}^{\sharp}[\eta]. \quad (26)$$

Consequently, there exists a unique comb $\overline{\mathcal{T}}_{\Lambda}^{\sharp}$ defined on the completion $\overline{\mathcal{K}}_{\Lambda}^{\sharp}$ of $\mathcal{K}_{\Lambda}^{\sharp}$ that has, by construction, the family $\{\mathcal{T}_{\Lambda_k}\}_{\Lambda_k \subseteq \Lambda}$ as ‘marginals’. As $\mathcal{T}_{\Lambda}^{\sharp}$ is positive and bounded by 1 on $\mathcal{K}_{\Lambda}^{\sharp}$, by continuity so is $\overline{\mathcal{T}}_{\Lambda}^{\sharp}$ on $\overline{\mathcal{K}}_{\Lambda}^{\sharp}$. This concludes the proof. $\square$

The space $\overline{\mathcal{K}}_{\Lambda}^{\sharp}$ is a proper subset of $\overline{\mathcal{L}}_{\Lambda}^{\sharp}$ (this latter space is sometimes called *quasilocal algebra* in the literature [42, 43]). There are two important points to note about these spaces. On the one hand, it is clear that in most relevant cases, $\overline{\mathcal{K}}_{\Lambda}^{\sharp}$ cannot coincide with $\overline{\mathcal{L}}_{\Lambda}^{\sharp}$. As mentioned above, $\mathcal{L}_{\Lambda}^{\sharp}$ generically contains CPTP maps that do not agree with the causal order of the finite combs $\mathcal{T}_{\Lambda_k}$, and extending $\mathcal{T}_{\Lambda}^{\sharp}$ to a linear functional on $\overline{\mathcal{L}}_{\Lambda}^{\sharp}$ would neither be possible (as outlined above) nor meaningful. In this sense, $\overline{\mathcal{K}}_{\Lambda}^{\sharp}$ is the ‘biggest’ space that we can extend the finite functionals to, and it can be understood as the set of *all* operations on times in Λ that abide with the causal order of the finite combs Λ_k . As such, it is not just the ‘biggest’ space we

can extend the action of $\mathcal{T}_\Lambda^\sharp$ to, but also the most meaningful one.

On the other hand, it is important to note that the space $\bar{\mathcal{L}}_\Lambda^\sharp$ does *not* coincide with $\mathcal{L}_\Lambda$ (they coincide iff Λ is finite [40]). Consequently, there might be different combs $\mathcal{T}_\Lambda$ defined on $\mathcal{L}_\Lambda$ with coinciding action on all elements of $\bar{\mathcal{K}}_\Lambda^\sharp$. This, however, is not problematic; first, $\bar{\mathcal{L}}_\Lambda^\sharp$ “is in a way more important than” $\mathcal{L}_\Lambda$ because its elements arise from the ones of $\mathcal{L}(\mathcal{H}_\alpha)$ “by extension and algebraical and topological processes” [40]. Second, just as for the KET [5], the different possible combs on $\mathcal{L}_\Lambda$ all lead to the same measurement statistics on any experimentally accessible set of times, so this non-uniqueness is not accessible/detectable in practice.

We emphasize that, even though we have phrased the above in the language of quantum mechanics, there is nothing particularly quantum mechanical about the GET. The proof of the theorem uses the compatibility, linearity and boundedness of the combs $\mathcal{T}_{\Lambda_k}$, as well as the assumption that the spaces they act on span a vector space. Consequently, it holds for *any* probabilistic theory (with interventions) satisfying these minimal assumptions.

Furthermore, the input and output spaces of the CP maps the comb acts on do not have to be of the same dimension. In this case, the identity map used for the consistency condition has to be slightly generalized: A CPTP map $\mathcal{M}_\alpha : \mathcal{B}_1(\mathcal{H}_\alpha^i) \rightarrow \mathcal{B}_1(\mathcal{H}_\alpha^o)$ is implemented via a corresponding unitary U_α , a *fixed* ancillary state $\eta_\alpha \in \mathcal{B}_1(\mathcal{H}_{A_\alpha})$, and a partial trace tr_{B_α} that is such that the resulting state $\mathcal{M}_\alpha[\rho] = \text{tr}_{B_\alpha} [U_\alpha(\rho \otimes \eta_\alpha)U_\alpha^\dagger]$ lies in $\mathcal{B}_1(\mathcal{H}_\alpha^o)$. With this, we can define a *generalized identity map*

$$\mathcal{I}_{t_\alpha}^{(i \rightarrow o)}[\rho] = \text{tr}_{B_\alpha}(\rho \otimes \eta_\alpha), \quad (27)$$

and the GET still holds. The only difference being that the inverse projections used in its derivation, and given in Eqs. (12) and (13), have to be changed accordingly to account for the altered identity map. Consequently, our theorem accounts for the case where particles are created/annihilated in the process, as well as the case where different degrees of freedom are manipulated at each time t_α , or where the number of measurement outcomes and active interventions differ.

Even more generally, the particular form of the ‘do nothing’ operation, *i.e.*, the action on the system of interest in the absence of active experimental intervention is of no importance for the derivation of the GET. In case that it does not coincide with the identity map $\mathcal{I}$ (or the more general identity map discussed above) but is represented by some map $\mathcal{M}$ (for example, one could imagine a theory where nature constantly measures the system of interest) the logic of the proof of the GET would still follow through. Again, the only change in its derivation would be an adjustment of the inverse projection maps of Eqs. (12) and (13), with the rest of the proof staying the same.

In the derivation of the GET, we make the implicit assumption that the employed CP maps only depend on the measurement outcome they correspond to, but not on the particular instrument that was used to carry out the respective measurement. This property has been dubbed ‘instrument non-contextuality’ [28, 44] or ‘operational instrument equivalence’ [23]. In principle, our derivation could be straightforwardly adapted to any theory, where this assumption is no longer satisfied, but in which probabilities are still a linear function of the maps *and* their respective contexts (*i.e.*, instruments). Instead of the identity map, one would then use a pair $(\mathcal{I}, \mathcal{J}_\mathcal{I})$ of identity map and identity context for marginalization, and the GET would still hold.

It is important to clearly distinguish between the classical Kolmogorov extension theorem and the GET. The KET hinges on the fact that, in classical physics, a measurement does not change the average state of a system. This fails to hold in quantum mechanics, or any theory with interventions. More concretely, in the language of quantum maps, the sum over the outcomes i_α of a measurement in a basis $\{|i\rangle\}$ at time t_α corresponds to the CPTP map $\mathcal{M}_\alpha = \sum_{i_\alpha} \mathcal{M}_{i_\alpha}$, where $\mathcal{M}_{i_\alpha}[\rho] = \langle i_\alpha | \rho | i_\alpha \rangle | i_\alpha \rangle \langle i_\alpha |$. In a classical stochastic process, the state ρ is diagonal in the basis $\{|i\rangle\}$, and we have $\mathcal{M}_\alpha[\rho] = \rho$; the average over measurement outcomes has the same effect as the classical ‘do nothing’ operation. As soon as ρ is not diagonal in the measurement basis, we have $\mathcal{M}_\alpha[\rho] \neq \rho$; on average, a measurement in quantum mechanics changes the state of the system and the future measurement statistics will

depend on the measurement that was performed. Consequently, joint probability distributions in classical physics (without interventions) exhibit a consistency condition, while quantum processes (and theories with interventions) generally do not.

As in the classical case, the proof of the GET does not assume an *a priori* temporal ordering. The sets Λ_k could be sets of times, but also labels of different laboratories. We have the following remark:

Remark. *The proof of the GET does not assume any ordering of the sets Λ_k , and only uses the generalized consistency property of Eq. (7) as its main ingredient.*

As alluded to above, this implies that the GET also applies to causally indefinite processes [33, 34, 45], as the descriptors for different sets of laboratories would still satisfy a compatibility condition. However, these processes do not admit a Stinespring dilation that is compatible with a fixed causal order [33, 34] and the interpretation of an underlying ‘process’ becomes much less clear in the absence of a definitive causal ordering. We will briefly remark on this further in our conclusions, but leave a full exploration of this interpretation as an open question for future work. Next, we will see that the distinction between stochastic processes and causal modelling does not exist in the general case.

4.4 Quantum stochastic processes and quantum causal modelling

Using an instrument at some intermediate time t_α alters the state of a quantum system (even when averaging over all outcomes) and influences the statistics of later measurements in a non-negligible way. Nevertheless, the *full* descriptor of an ℓ -step process, *i.e.*, $\mathcal{T}_{\Lambda_\ell}$, contains all descriptors $\mathcal{T}_{\Lambda_k}$ for fewer times $\Lambda_k \subseteq \Lambda_\ell$, and a family of compatible combs implies the existence of an underlying stochastic process $\mathcal{T}_\Lambda$.

Like in the classical case, the GET provides the mathematical underpinnings for the theory of stochastic processes in quantum mechanics, or any other theory with interventions, and fixes the minimal necessary requirements for the existence of an underlying process. As we have seen, in quantum mechanics, it is unavoidable to

employ a description that takes interventions into account, when attempting to obtain a consistent description of a quantum process; if one wants to properly define quantum stochastic processes, one is forced to use the framework of causal modelling where active interventions are used to infer the causal relations between different events. This motivates the following observation:

Observation 2. *The theory of quantum causal modelling and the theory of quantum stochastic processes are equivalent.*

In contrast to Observation 1, the set of quantum causal models does not just contain the set of quantum stochastic processes but coincides with it; in classical physics, we obtain a consistent description of stochastic processes without taking interventions into account, and we can *choose* to intervene whenever we want to probe the causal structure of a process. In quantum mechanics, a consistent description of stochastic processes can *only* be recovered if interventions are included in the description from the start. Interventions are not a choice but a necessity in quantum mechanics, which leads to the equivalence of quantum causal modelling and quantum stochastic processes.

This implies that the breakdown of the KET in quantum mechanics is *fundamental*, while it can in principle be removed by changing perspective in a classical process with interventions. In the latter case, a *super-observer*, that observes both the experimenter manipulating the system of interest as well as the stochastic process itself, would obtain families of joint probability distributions that display a compatibility property. Put differently, for classical processes, by incorporating the experimenter and their choice of instrument into the stochastic process, the KET always applies on a higher level. In quantum mechanics, this is generally not true. No matter the level at which a super-observer observes a process, the respective joint probability distributions do not satisfy a compatibility property, and the KET fails to hold. This fundamental breakdown of the KET in quantum mechanics is mirrored by no-go theorems that show that non-contextual theories cannot reproduce the predictions of quantum mechanics; for many of these theorems, the notion of ontic latent variables [46, 47] or ontic processes [23] are introduced, and the

basic assumption is made that the distributions over observable outcomes can be obtained by marginalization of a larger joint distribution over the values of the ontic variable. Subsequently, it is shown that, together with other assumptions, this prerequisite fails to reproduce predictions made by quantum mechanics. The GET dictates how to correctly compute marginals in quantum mechanics, such that all resulting probability distributions ‘fit together’ and are the marginals of one common comb $\mathcal{T}_\Lambda$. It is therefore conceivable that a derivation starting from the assumption of compatibility in the sense of the GET would lead to theories that can indeed reproduce quantum mechanics.

We reiterate that classical stochastic processes are a very special subset of general stochastic processes, namely the ones where the system of interest is never rotated out of its fixed (pointer) basis, and the experimenter can only perform projective measurements in this basis.⁵ We now show that the KET can be derived in a straight forward way as a corollary of the GET.

4.5 GET $\Rightarrow$ KET

Our generalised extension theorem applies to a strictly larger class of theories than the standard KET and includes the latter as a corollary. Specifically, in the language introduced above, a classical process is one where the experimenter can only perform measurements in a fixed basis, and the resulting joint probability distributions satisfy Kolmogorov consistency conditions. With this – under the aforementioned assumption that all considered value spaces are $\mathbb{R}$, $\mathbb{N}$, or, more generally, Borel spaces – we have the following proposition:

Proposition 1. *The GET implies the KET.*

Proof. In order to prove this statement, we will show that any family of classical joint probability distributions that satisfies the consistency property of the KET can be mapped onto a family of quantum combs that satisfies the consistency condition of the GET – albeit with a slightly different identity map. The GET then guarantees that there exists an underlying

⁵The set of quantum processes that can be described by only classical means is in fact slightly bigger [48, 49]. We will comment on this subtlety below.

classical comb $\mathcal{T}_\Lambda^{\text{cl}}$, and thus also an underlying classical process $\mathbb{P}_\Lambda$.

Let $\{\mathbb{P}_{\Lambda_k}\}_{\Lambda_k \subset \Lambda}$ be a family of joint probability distributions on all finite subsets of Λ that satisfies the consistency conditions of the KET, i.e., $\mathbb{P}_{\Lambda_k} = \mathbb{P}_{\Lambda_\ell}^{\Lambda_k}$ for all $\Lambda_k \subset \Lambda_\ell \subset \Lambda$. We denote the set of perfectly distinguishable possible outcomes at time t_α by $\{i_\alpha\}$. With this, we can define a Hilbert space $\mathcal{H}_\alpha$ spanned by an orthogonal set of states $\{|i_\alpha\rangle\}$, and projective CP operators $\mathcal{P}_{i_\alpha}$ that correspond to a measurement with outcome i_α . The action of these operators on a state $\rho \in \mathcal{B}_1(\mathcal{H}_\alpha)$ is given by $\mathcal{P}_{i_\alpha}[\rho] = \langle i_\alpha | \rho | i_\alpha \rangle |i_\alpha\rangle \langle i_\alpha|$. The complex vector space spanned by these projective operators will be denoted by Ω_α , and, correspondingly, we set $\Omega_{\Lambda_k} = \bigotimes_{\alpha \in \Lambda_k} \Omega_\alpha$. On said vector space, we can define a classical comb $\mathcal{T}_{\Lambda_k}^{\text{cl}}$, with its action on every $\mathcal{P}_{\mathbf{i}_{\Lambda_k}} = \bigotimes_{\alpha \in \Lambda_k} \mathcal{P}_{i_\alpha}$ given by

$$\mathcal{T}_{\Lambda_k}^{\text{cl}}[\mathcal{P}_{\mathbf{i}_{\Lambda_k}}] = \mathbb{P}_{\Lambda_k}(\mathbf{i}_{\Lambda_k}). \quad (28)$$

To stay closer in spirit to the proof of the GET, we could extend $\mathcal{T}_{\Lambda_k}^{\text{cl}}$ to a CP linear functional on the whole space $\mathcal{L}_{\Lambda_k} \supset \Omega_{\Lambda_k}$, but as this step is not necessary for the proof of the KET, we will not carry it out here. As the family of probability distributions $\{\mathbb{P}_{\Lambda_k}\}$ satisfies a consistency condition, we have for $\Lambda_k \subset \Lambda_\ell$

$$\sum_{\Lambda_\ell \setminus \Lambda_k} \mathcal{T}_{\Lambda_\ell}^{\text{cl}}[\mathcal{P}_{\mathbf{i}_{\Lambda_\ell}}] = \mathbb{P}_{\Lambda_k}(\mathbf{i}_{\Lambda_k}) = \mathcal{T}_{\Lambda_k}^{\text{cl}}[\mathcal{P}_{\mathbf{i}_{\Lambda_k}}], \quad (29)$$

where $\mathbf{i}_{\Lambda_k}$ is the restriction of $\mathbf{i}_{\Lambda_\ell}$ to Λ_k . Setting $\Delta_{t_\alpha} := \sum_{i_\alpha} \mathcal{P}_{i_\alpha}$, we see that the family of combs $\{\mathcal{T}_{\Lambda_k}\}_{\Lambda_k \subset \Lambda}$ satisfies a consistency condition with respect to the operators Δ_{t_α} (in contrast to the GET, where the corresponding operator was $\mathcal{I}_{t_\alpha}$). As discussed above, the proof of the GET can be straightforwardly generalised to any choice of the ‘do-nothing’ operation. Analogous to the proof of the GET, setting $\Omega_\Lambda^\# = \{\xi | \xi = \xi_{\Lambda_k} \otimes_{t_\alpha \in \Lambda \setminus \Lambda_k} \Delta_{t_\alpha}, \xi_{\Lambda_k} \in \Omega_{\Lambda_k}\}$ and

$$\omega_\Lambda^\# = \{\xi \in \Omega_\Lambda^\# | \xi = \bigotimes_{\alpha \in \Lambda_k} \mathcal{P}_{i_\alpha} \otimes \bigotimes_{\Lambda \setminus \Lambda_k} \Delta_{t_\alpha}\}, \quad (30)$$

we see that there exists a unique comb $\overline{\mathcal{T}}_\Lambda^{\text{cl},\#}$, defined on the closure $\overline{\omega}_\Lambda^\#$ of $\omega_\Lambda^\#$, that has the family $\{\mathcal{T}_{\Lambda_k}^{\text{cl}}\}$ as marginals (with respect to the operators Δ_{t_α}).

It remains to show how to obtain a probability distribution $\mathbb{P}_\Lambda$ from $\overline{\mathcal{T}}_\Lambda^{\text{cl},\#}$ that contains all finite

distributions as marginals. To this end, we note that the classical comb $\bar{\mathcal{T}}_{\Lambda}^{\text{cl.}\#}$ is well-defined and yields probabilities on all $\xi \in \bar{\omega}^{\#} \subset \bar{\Omega}_{\Lambda}^{\#}$. Every $\xi \in \bar{\omega}_{\Lambda}^{\#} \setminus \omega_{\Lambda}^{\#}$ that lies in the extension of $\omega_{\Lambda}^{\#}$ has a unique restriction $\xi|_{\Lambda_k} = \mathcal{P}_{i_{\Lambda_k}}$, to any *finite* set Λ_k , and as such uniquely fixes a set of outcomes $\mathbf{i}_{\Lambda_k}$. Taking the union of all these sets of outcomes, we see that every $\xi \in \bar{\omega}^{\#}$ defines a set $\mathbf{i}_{\Lambda}$ (we will thus add an additional subscript and denote them by $\xi_{\mathbf{i}_{\Lambda}}$) of corresponding outcomes at all times in Λ . Setting $\mathbb{P}_{\Lambda}(\mathbf{i}_{\Lambda}) = \bar{\mathcal{T}}_{\Lambda}^{\text{cl.}\#}[\xi_{\mathbf{i}_{\Lambda}}]$, we obtain a probability distribution $\mathbb{P}_{\Lambda}$ that – by construction – yields the correct probability distributions $\mathbb{P}_{\Lambda_k}$ when restricted to finite sets $\Lambda_k \subset \Lambda$. $\square$

While the original version of the KET does not hold for quantum processes, it is important to note that the breakdown of the compatibility property of joint probability distributions is not a signature of quantum mechanics *per se*; as we have already seen, any framework that allows for interventions will exhibit this feature. The GET provides a proper theoretical underpinning for the corresponding experimental situations. On the other hand, the breakdown of the compatibility property can happen in quantum mechanics even if only projective measurements in a fixed basis $\{|i_{\alpha}\rangle\}$ are performed [16, 17].

As already mentioned, the absence of compatibility is tantamount to the absence of either realism *per se*, or non-invasiveness (or both). Consequently, it can be used as a definition of non-classicality, as proposed in Ref. [17]. There, the authors employ the breakdown of the consistency condition on the level of probability distributions, when measuring in a fixed basis, as a means to define the non-classicality of Markovian processes. Using the framework of quantum combs for the description of quantum stochastic processes the ideas of [17] can be extended to general processes with memory, *i.e.*, non-Markovian processes [48, 49].

Following Ref. [17], we consider an ℓ -step process to be classical if its joint probability distributions with respect to measurements in a fixed basis $\{|i_{\alpha}\rangle\}$ satisfy a consistency condition. Put differently, an ℓ -step process $\mathcal{T}_{\Lambda_k}$ is classical (with respect to the basis $\{|i\rangle\}$) iff for all $\Lambda_k \subseteq \Lambda_{\ell}$

and all possible sequences of outcomes $i_k, \dots, i_1$

$$\mathcal{T}_{\Lambda_k}[\mathcal{P}_{i_k}, \dots, \mathcal{P}_{i_1}] = \sum_{\Lambda_{\ell} \setminus \Lambda_k} \mathcal{T}_{\Lambda_{\ell}}[\mathcal{P}_{i_{\ell}}, \dots, \mathcal{P}_{i_1}], \quad (31)$$

where $\mathcal{P}_{i_{\alpha}}$ corresponds to obtaining outcome i_{α} from a projective measurement in a fixed basis at time t_{α} , *i.e.*, $\mathcal{P}_{i_{\alpha}}[\rho] = \langle i_{\alpha} | \rho | i_{\alpha} \rangle | i_{\alpha} \rangle \langle i_{\alpha} |$.

The general structure of classical combs that satisfy Eq. (31) can then be analyzed using the Choi isomorphism between quantum processes and positive matrices [50, 51]. As combs can describe general processes with memory, Eq. (31) represents a consistent definition of classical processes with memory and allows a direct extension of the results obtained in Ref. [17] to the non-Markovian case [48, 49].

5 Relation to previous works

As already mentioned, the proof of the GET does not rely on any particularities that are exclusive to quantum mechanics or our formulation thereof. The GET constitutes a sound basis for the description of *any* conceivable (classical, quantum or beyond) theory of stochastic processes with interventions – independent of the employed framework.

While we referred throughout to the framework of *quantum combs* [32, 37, 38], originally derived as the most general representation of quantum circuit architectures, our results apply equally well to any other framework for describing quantum processes as linear functionals. Examples of the mathematical objects and frameworks (often the same thing under a different name) given a firm theoretical foundation by the GET include: *process tensors* [12, 29, 30] and *causal automata/non-anticipatory channels* [43, 52], which describe the most general open quantum processes with memory; *causal boxes* [53] that enter into quantum networks with modular elements; *operator tensors* [54, 55] and *superdensity matrices* [56], employed to investigate quantum information in general relativistic space-time; and, finally, *process matrices*, used for quantum causal modelling [13–15, 34]. In classical physics, as well as the standard causal modelling framework discussed in Sec. 2, our result applies to the ϵ -transducers used within the framework of computational mechanics [57, 58] to describe processes with active interventions.

Our theorem proves the existence of a container space for all of the aforementioned frameworks and allows for their complete and consistent representation in the continuous time limit, thus providing an overarching theorem for probabilistic theories with interventions. This is of particular importance for the field of open quantum mechanics where the lack of an extension theorem has been a roadblock to obtaining a framework that coincides with classical descriptions in the correct limit [16]. Here, switching perspective allows one to describe both classical as well as quantum open systems in a unified framework. This fact has recently been used to obtain an unambiguous definition of non-Markovianity in quantum mechanics that coincides with the classical one in the correct limit [59].

The GET goes beyond previous attempts to generalize the KET for quantum mechanics. An extension theorem for positive operator valued measures was derived in Ref. [60] and was used in Ref. [61] to show the existence of an ‘infinite composition’ of an instrument. This extension theorem is, however, limited to particular cases of positive operator valued measures, and not general enough to provide an underpinning for the description of stochastic processes with interventions.

More generally, a version of the KET for quantum processes was derived in Ref. [18]. In this work, the authors showed that any quantum stochastic “process can be reconstructed up to equivalence from a projective family of correlation kernels”. By decomposing the control operations $\mathcal{M}_{i_\alpha}$ into their component Kraus operators, it can explicitly be shown that these correlation kernels correspond to combs, and consequently, for quantum processes, the GET is equivalent to Thm. 1.3 in Ref. [18]. However, the mathematical structure of the latter does not tie in easily with recently developed frameworks for the description of quantum (or classical) causal modelling, nor does it lend itself in a straightforward way to the discussion of their key properties. Additionally, our proof – in contrast to the one presented in [18] – highlights the role that causal order plays for the domain of the resulting stochastic process $\mathcal{T}_\Lambda$. Specifically, while *independent* CP maps at different times t_i are considered in [18], our construction makes

explicit the set of *correlated* operations on Λ that $\mathcal{T}_\Lambda$ can be meaningfully applied to.

The structural features of combs render the investigation of fundamental features of a process, like their non-Markovianity [29, 59], their causal structure [13, 34, 53], and their classicality tractable. Furthermore, our formulation has the advantage that combs are defined in a clear-cut *operational* way, and allow for a generalized Stinespring dilation [30, 32], which makes their interpretation in terms of open quantum system dynamics straightforward. Finally, even though the GET is stated for combs that map sequences of CP maps to probabilities, its proof also applies – with slight modifications – to general quantum combs (*i.e.*, maps that map combs onto combs [32, 38]).

6 Conclusions

While the KET is the fundamental building block for the theory of classical stochastic processes, it does not hold in quantum mechanics, or any other theory that allows for active interventions. This breakdown goes hand in hand with the violation of Leggett-Garg inequalities: the violation of such an inequality *always* implies that compatibility conditions are not satisfied, and hence the KET does not hold.

In this work, we have proven a generalized extension theorem that applies to any process with interventions, including quantum ones. We have therefore shown that the roadblocks encountered when describing quantum processes in terms of joint probability distributions can be remedied by changing perspective; while the evolution of a density matrix over time does not contain enough statistical information for consistency properties to hold [16], considering a quantum stochastic process as a linear functional acting on sequences of CP maps allows one to formulate a fully fledged theory. Taking interventions into account is the only way to obtain a consistent definition and rigorous mathematical foundation for quantum stochastic processes. Put differently, without taking interventions into account, there is no way to consistently define quantum stochastic processes. In this sense, two seemingly different frameworks – the framework of causal modelling, and the theory of quantum stochastic processes – are

actually two sides of the same coin.

In the limit of continuous time, the sequence of CP maps becomes a continuous driving/control of the system of interest. Thus, the GET provides the theoretical foundation for these experimental scenarios, important for development of quantum technologies. Likewise, just as in the case of classical stochastic processes, the GET provides a toolbox for the modelling of quantum stochastic processes; any mechanism that leads to consistent families of combs automatically defines an underlying process.

It is important to emphasize the generality of our main result. Due to the linearity of mixing, *any* meaningful description of a stochastic process – quantum or not – must be expressible in terms of a linear function on the space of locally accessible operations [12]. The proof of the GET is versatile enough to account for any framework that aims to describe temporally ordered processes, and hence provides a sound mathematical underpinning for all of them.

The GET contains the original KET as the special case where the family of processes is diagonal in the reference basis, and the only allowed CP maps are projective measurements in the same basis. On the one hand, this implies that our extension of classical processes to the quantum realm is the correct one. On the other hand, this clear-cut definition of classical combs lends itself ideally to the investigation of the interplay of coherence and classicality, as proposed in Ref. [17], in the experimental observation of real-world processes with memory.

Finally, our discussion made transparent where causality and causal order enter into the proof of the GET, and what sets of operations the resulting stochastic process can meaningfully be applied to. While we have mostly discussed temporally ordered processes, in principle, even causally disordered processes could be described by families of functionals that satisfy a consistency requirement (Λ would then be thought of as a set of labels for different laboratories). However, there is no deterministic Stinespring dilation for causally disordered processes [33]. There are, on the other hand, dilations that include post-selection [37, 62], and we conjecture that an underlying

causally disordered stochastic process would be equivalent to post-selection on a class of trajectories resulting from continuous weak measurement.

Acknowledgments

We thank Andrea Collecchio, Tim Geroni and Top Notoh for valuable discussions, and Roderich Tumulka for clarifying remarks on the KET for POVMs. We are especially grateful to the referees of this paper for providing thorough and constructive advice on how to improve the proof of the GET. KM and FAP thank an anonymous referee (of another paper) for pointing them to the lack of a KET in quantum mechanics. SM acknowledges funding from the the Monash International Postgraduate Research Scholarship (MIPRS), the J L William Scholarship, the Austrian Science Fund (FWF): ZK3 (Zukunftkolleg), the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 801110, and the Austrian Federal Ministry of Education, Science and Research (BMBWF). FS is grateful to thank Sri-Trang Thong Scholarship, Faculty of Science and “Postgraduate Student Exchange Program 2016”, International Relations Division, Mahidol University, for financial support during his visit to Monash University, Australia. KM is supported through ARC FT160100073.

References

- [1] Z. Schuss, *Theory and Applications of Stochastic Processes: An Analytical Approach* (Springer, New York, 2009).
- [2] M. Liao, *Applied Stochastic Processes* (Chapman and Hall/CRC, Boca Raton, 2013).
- [3] A. N. Kolmogorov, *Grundbegriffe der Wahrscheinlichkeitsrechnung* (Springer, Berlin, 1933) [*Foundations of the Theory of Probability* (Chelsea, New York, 1956)].
- [4] W. Feller, *An Introduction to Probability Theory and Its Applications* (Wiley, New York, 1971).
- [5] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2007).

- [6] T. Tao, *An Introduction to Measure Theory* (American Mathematical Society, 2011).
- [7] A. J. Leggett and A. Garg, *Phys. Rev. Lett.* **54**, 857 (1985).
- [8] A. J. Leggett, *Rep. Prog. Phys.* **71**, 022001 (2008).
- [9] C. Emary, N. Lambert, and F. Nori, *Rep. Prog. Phys.* **77**, 016001 (2014).
- [10] N. Gilbert, *Agent-Based Models* (SAGE Publications Inc., Los Angeles, 2007).
- [11] J. Pearl, *Causality: Models, Reasoning and Inference* (Cambridge University Press, Cambridge, U.K.; New York, 2009).
- [12] S. Milz, F. A. Pollock, and K. Modi, *Open Sys. Info. Dyn.*, 1740016 (2017).
- [13] F. Costa and S. Shrapnel, *New J. Phys.* **18**, 063032 (2016).
- [14] O. Oreshkov and C. Giarmatzi, *New J. Phys.* **18**, 093020 (2016).
- [15] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, *Phys. Rev. X* **7**, 031021 (2017).
- [16] H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini, *Rev. Mod. Phys.* **88**, 021002 (2016).
- [17] A. Smirne, D. Egloff, M. G. Díaz, M. B. Plenio, and S. F. Huelga, *Quantum Sci. Technol.* **4**, 01LT01 (2018).
- [18] L. Accardi, A. Frigerio, and J. T. Lewis, *Publ. Rest. Inst. Math. Sci.* **18**, 97 (1982).
- [19] N. Wiener, A. Siegel, B. Rankin, and W. T. Martin, eds., *Differential Space, Quantum Systems, and Prediction* (The MIT Press, Cambridge (MA), 1966).
- [20] M. P. Lévy, *Am. J. Math.* **62**, 487 (1940).
- [21] Z. Ciesielski, *Lectures on Brownian motion, heat conduction and potential theory* (Aarhus Universitet, Matematisk Institutt, 1966).
- [22] R. Bhattacharya and E. C. Waymire, *A Basic Course in Probability Theory* (Springer, New York, NY, 2017).
- [23] S. Shrapnel and F. Costa, *Quantum* **2**, 63 (2018).
- [24] G. Lindblad, *Comm. Math. Phys.* **65**, 281 (1979).
- [25] E. B. Davis, *Quantum Theory of Open Systems* (Academic Press Inc, London; New York, 1976).
- [26] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Phys. Rev. A* **81**, 062348 (2010).
- [27] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Phys. Rev. Lett.* **101**, 180501 (2008).
- [28] S. Shrapnel, F. Costa, and G. Milburn, *New J. Phys.* **20**, 053010 (2018).
- [29] K. Modi, *Sci. Rep.* **2**, 581 (2012).
- [30] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, *Phys. Rev. A* **97**, 012127 (2018).
- [31] S. Milz, F. A. Pollock, and K. Modi, *Phys. Rev. A* **98**, 012108 (2018).
- [32] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Phys. Rev. A* **80**, 022339 (2009).
- [33] G. Chiribella, G. M. D'Ariano, P. Perinotti, and B. Valiron, *Phys. Rev. A* **88**, 022318 (2013).
- [34] O. Oreshkov, F. Costa, and Č. Brukner, *Nat. Commun.* **3**, 1092 (2012).
- [35] G. M. D'Ariano, G. Chiribella, and P. Perinotti, *Quantum Theory from First Principles: An Informational Approach*, 1st ed. (Cambridge University Press, Cambridge, United Kingdom ; New York, NY, 2017).
- [36] G. Chiribella and D. Ebler, *New J. Phys.* **18**, 093053 (2016).
- [37] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Europhys. Lett.* **83**, 30004 (2008).
- [38] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Phys. Rev. Lett.* **101**, 060401 (2008).
- [39] T. Tyc and J. Vlach, *Eur. Phys. J. D* **69**, 209 (2015).
- [40] J. v. Neumann, *Comp. Math.* **6**, 1 (1939).
- [41] E. Kreyszig, *Introductory Functional Analysis with Applications*, 1st ed. (Wiley, New York, 1989).
- [42] R. Haag and D. Kastler, *J. Math. Phys.* **5**, 848 (1964).
- [43] D. Kretschmann and R. F. Werner, *Phys. Rev. A* **72**, 062323 (2005).
- [44] C. M. Caves, C. A. Fuchs, K. K. Manne, and J. M. Renes, *Found. Phys.* **34**, 193 (2004).
- [45] M. Araújo, C. Branciard, F. Costa, A. Feix, C. Giarmatzi, and Č. Brukner, *New J. Phys.* **17**, 102001 (2015).

- [46] R. W. Spekkens, *Phys. Rev. A* **71**, 052108 (2005).
- [47] E. G. Cavalcanti, *Phys. Rev. X* **8**, 021018 (2018).
- [48] P. Strasberg and M. G. Díaz, *Phys. Rev. A* **100**, 022120 (2019).
- [49] S. Milz, D. Egloff, P. Taranto, T. Theurer, M. B. Plenio, A. Smirne, and S. F. Huelga, [arXiv:1907.05807](#) (2019).
- [50] M.-D. Choi, *Linear Algebra Appl.* **10**, 285 (1975).
- [51] A. Jamiołkowski, *Rep. Math. Phys.* **3**, 275 (1972).
- [52] F. Caruso, V. Giovannetti, C. Lupo, and S. Mancini, *Rev. Mod. Phys.* **86**, 1203 (2014).
- [53] C. Portmann, C. Matt, U. Maurer, R. Renner, and B. Tackmann, *IEEE Trans. Inf. Theory* **63**, 3277 (2017).
- [54] L. Hardy, [arXiv:1608.06940](#) (2016).
- [55] L. Hardy, *Phil. Trans. R. Soc. A* **370**, 3385 (2012).
- [56] J. Cotler, C.-M. Jian, X.-L. Qi, and F. Wilczek, *J. High Energy Phys.* **2018**, 93 (2018).
- [57] N. Barnett and J. P. Crutchfield, *J. Stat. Phys.* **161**, 404 (2015).
- [58] J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *npj Quantum Inf.* **3**, 6 (2017).
- [59] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, *Phys. Rev. Lett.* **120**, 040405 (2018).
- [60] R. Tumulka, *Lett. Math. Phys.* **84**, 41 (2008).
- [61] E. Haapasalo, T. Heinosaari, and Y. Kuramochi, *J. Phys. A* **49**, 33LT01 (2016).
- [62] S. Milz, F. A. Pollock, T. P. Le, G. Chiribella, and K. Modi, *New J. Phys.* **20**, 033033 (2018).