The Pragmatic QFT Measurement Problem
and the need for a Heisenberg-like Cut in QFT

Daniel Grimmer[1, 2, 3, *]

[1] *Reuben College, University of Oxford, Oxford, OX2 6HW United Kingdom*
[2] *Faculty of Philosophy, University of Oxford, Oxford, OX2 6GG United Kingdom*
[3] *Barrio RQI, Waterloo, Ontario N2L 3G1, Canada*

Despite quantum theory's remarkable success at predicting the (statistical) results of experiments, many philosophers worry that it nonetheless lacks some crucial connection between theory and experiment. Such worries constitute the Quantum Measurement Problem. We can identify two kinds of worries: 1) pragmatic: it's unclear how to model our experiments to extract theoretical predictions, and 2) realist: there is no realist narrative for the experiment underlying these theoretical predictions. I will argue that while both deserve attention, the pragmatic worries have worse consequences if left unanswered. Moreover, upon reflection, a satisfactory prediction of many of quantum theory's key experimental successes unavoidably involves modeling quantum fields at some point. Without a satisfactory pragmatic theory-to-experiment link for QFT, we are at risk of losing any right to claim evidential support for large parts of quantum theory. This is the *Pragmatic QFT Measurement Problem*.

But, what makes formalizing measurements in QFT so hard? As I will discuss, attempts to naively transplant projective measurements into QFT violate core tenets of relativity. Thus we need a new (or at least refined) measurement theory for QFT. However, as I will argue, aiming too directly at a new measurement theory is an incautious way to proceed and is apt to lead us astray. Instead, we ought to first review how our non-relativistic quantum measurement theory is rooted in notions of measurement chains and Heisenberg cuts. Next, we ought to generalize these notions and transplant them into QFT. My analysis suggests the need for a pragmatic *QFT-cut* analogous to the Heisenberg cut. Finally, we can then see what measurement theory (if any) we are led to for QFT. As I will discuss, the Unruh-DeWitt detector model is uniquely positioned to give us a wide-scoping physically-meaningful measurement theory for QFT.

Keywords: quantum measurement problem; quantum field theory; observables; Heisenberg cut; measurement chain;

I. INTRODUCTION: A QUANTUM MEASUREMENT PROBLEM

It is uncontestable that quantum theory has been remarkably successful at predicting the (statistical) results of a wide range of experiments. However, despite its many predictive successes, many philosophers and physicists are nonetheless worried that quantum theory lacks some crucial connection between theory and experiment. Various dissatisfactions with various theory-to-experiment disconnects each deserve the title[1] "A Quantum Measurement Problem": How are we to understand/model the measurement of quantum systems?

In order to differentiate the various kinds of dissatisfactions from each other, it's perhaps best to start from a telling of quantum theory which (hopefully nearly) every physicist and philosopher is dissatisfied with. I have in mind the parts of non-relativistic quantum theory which students are urged to focus on after they are told to "Shut up, and calculate!". Let us call this *the sopho-*

more's *quantum theory*. Students are here taught to model quantum experiments as follows: Given some initial conditions, $|\text{in}\rangle$, we apply the given unitary evolution, $\hat{U}$, and then the given projective measurement, $\langle \text{out}|$. This computation unambiguously gives a statistical prediction of the experiment's outcome via Born's rule, $p(\text{out}|\hat{U}, \text{in}) = |\langle \text{out}| \hat{U} |\text{in}\rangle|^2$. (A sophisticated sophomore may also learn about non-projective measurements, selective measurements and non-selective measurements as well as post-measurement state updates via Lüders rule.) Many physicists and philosophers are dissatisfied with the sophomore's quantum theory and claim that it lacks the right kind of connection between theory and experiment, and rightly so.

One can distinguish between two types of worries surrounding quantum measurement: pragmatic worries and realist worries. Pragmatic worries are aimed at clarifying how exactly these statistical predictions are to be pulled out of the theory. Specifically, how should the measurement process be *modeled*? By contrast, realist worries are aimed at establishing a realist narrative for the measurement process. How could the world be such that our experiments have these statistics? What explains the different definite outcomes I see in each seemingly identical run of my experiment? How should the measurement process be *understood*?

The sophomore's quantum theory fails on both fronts.

---

[1] This situation is rather like the mis-advertised efforts to "find *the* cure for cancer". There are many types of cancer each of which will need their own cure. Moreover, even considering a single type of cancer, it may be cured in several ways.

Its realist failures are well known, but its pragmatic failures deserve some further comment. The sophomore's quantum theory does not, in fact, give us unambiguous statistical predictions for an experiment's outcomes. While it is true that statistical predictions can be unambiguously associated with initial states, unitaries, and projectors, $p(\text{out}|\hat{U}, \text{in}) = |\langle\text{out}|\hat{U}|\text{in}\rangle|^2$, these themselves have not yet been suitably connected with our real-life experimental setups. Specifically, the sophomore has no answer to the question[2]: "In modeling this piece of lab equipment, exactly which projectors am I supposed to use and when?" There may be a ready-answer to this question pre-written on the sophomore's problem sheets, but show them a foreign piece of lab equipment and watch them falter.

Often the sophomore may intuitively guess which projector to use and when. They will not always be wrong, indeed they will often guess right. It is highly intuitive that in a two-slit experiment when the electrons hit the detection screen that just then a position measurement happens. And indeed for most pragmatic purposes, this is effectively what happens. Indeed, it may often be the case that the sophomore's guesses consistently give accurate-enough predictions. But ultimately they are nothing more than just that: guesses. As I will argue, in order to do better we need to explicitly model the measurement process in some dynamical detail [1–3].

Of these two types of failings, which has worse consequences if neglected? While both deserve attention, I argue the pragmatic worries are by far more important since if unanswered their consequences are far more severe. While it is at least debatable whether or we need to find a realist narrative underlying quantum theory [4], it is not debatable that quantum theory (indeed, any physical theory) must provide us with a robust account[3] of how predictions can and should be made from it (something better than intuitive guessing, even if our intuitive guesses are often right-enough). It is exactly this pragmatic link to real-life experimental practice which grants our physical theories evidential support, their connection with reality, and arguably their physicality [1] as well.

———

2 The sophomore's question will be answered in Sec. IV C. One answer involves decoherence theory and the Born rule. It should be noted, however, that contrary to popular belief, there may be other methods of extracting predictions from quantum theory in a principled way.

3 Above, I complained that the sophomore's quantum theory doesn't give us unambiguous predictions. However, on reflection, demanding complete unambiguity is asking too much. As I will discuss later, even for non-relativistic quantum theory it is necessary to make approximations when extracting concrete experimental predictions from our theory. Often we will have freedom (and hence ambiguity) regarding what kind of approximations we should make and where to make them. There is nothing wrong with this so long as the consequences of these choices are well understood and controllable. Indeed, when attempting to explain our experimental data, the goal is not to give definite unambiguous predictions but rather to give predictions which are within the experiment's error bars.

Without a clear understanding of this pragmatic connection between theory and experiment, we are at risk of losing any right to claim evidential support (or worse physicality) for quantum theory. As interesting as the realist's worries are, I rather say: let us first work on bringing home the spoils of experimental success, then we can worry about what it all means later, once we have better footing.

While above I have had non-relativistic quantum theory in mind, the distinction between pragmatic and realist worries about measurement can be made for any physical theory. The relative severity of the pragmatic worries over the realist worries (if unanswered) holds in every case. But what, roughly, would count as a satisfying solution to these pragmatic worries? What is at risk is our theory's evidential support, that is, its empirically supported connection with reality. As such any minimal solution to these worries must give a *measurement framework*: a principled explanation about how experimental predictions are to be extracted from the theory at least for the theory's major experimental successes and potentially on a case-by-case basis. Let us call this the core pragmatic measurement problem. Solving this would restore empirical support to our theory's key experiments.

One might, however, set their sights higher than this. The core pragmatic measurement problem could be solved with a case-by-case measurement framework with minimal scope leaving us in the dark about how measurements work generally in our theory. One might further want a unified *measurement theory*: a principled explanation about how experimental predictions are to be extracted from the theory in a holistic and wide-scoping way. Let us call this the extended pragmatic measurement problem. While not strictly necessary for our theory to be physically meaningful and evidentially supported, solving the extended problem has some significant benefits. Solving the extended problem allows us to talk about the "observables" of the theory in a physically meaningful way. (Note, however, that one cannot have a physically meaningful characterization of a theory's "observables" without examining the dynamics of its measurement processes in at least some detail.)

Let's return our attention to these pragmatic worries in non-relativistic quantum theory. As discussed above, the sophomore's quantum theory fails on both the pragmatic and realist fronts, but how do more sophisticated tellings of quantum theory fair? I will argue (in Sec. IV) that for non-relativistic quantum theory the core pragmatic measurement problem has been solved in terms of measurement chains and pragmatic Heisenberg cuts. These notions give us a case-by-case measurement framework within which we can work to extract predictions experiments within the remit[4] of non-relativistic quantum theory.

———

4 As I will discuss in Sec. II the scope of quantum theory's key experimental successes which are analyzable entirely within non-

Moreover, for non-relativistic quantum theory the extended pragmatic measurement problem has also been solved. Using decoherence theory and the Born rule, one can give a unified wide-scoping account of the measurement process. Specifically, decoherence theory underwrites our usual projective measurement theory.

Consequently, with both the core and extended pragmatic measurement problem solved, within non-relativistic quantum theory almost all attention has shifted onto the realist's worries. And rightly so. This so much so that for non-relativistic quantum theory the quantum measurement problem has become synonymous with such realist worries.

However, as I will argue (in Sec. III) the same is not true for quantum field theory[5] (QFT) with severe consequences; Without a satisfactory theory-to-experiment link we are at risk of losing any right to claim any evidential support (or worse physicality) for quantum field theory. Given the severity of this situation, the focus of this paper is exclusively on such pragmatic worries in the context of quantum field theory. That is, my focus is on *The Pragmatic QFT Measurement Problem* (both core or extended). In particular, this paper will carefully consider which parts of our non-relativistic solution can and cannot be transferred over into QFT.

Outline of Paper

My above suggestion that we first work on "bringing home the spoils of experimental success" raises the following questions: for quantum theory generally, where were these metaphorical spoils won and what currently blocks their way home? In Sec. II, I will discuss where the evidential spoils of quantum theory lie so-to-speak. I will argue (paralleling Wallace [5, 6]) for the following perhaps surprising claim; Explaining all but maybe a small corner of quantum theory's experimental successes unavoidably involves modeling quantum fields at some point. (Importantly, this is for fairly basic conceptual reasons; I am not here demanding hyper-accuracy.) Thus, if we cannot establish an adequate pragmatic link between QFT and experimental practice then not only quantum field theory but nearly the whole of quantum theory is at risk of losing its evidential support. While

the argument in Sec. II is not essential for the main philosophical points made in the rest of the paper, it substantially raises the stakes.

Given that the route home for most of our experimental spoils runs through quantum field theory, our next question becomes: Why is it (at least currently) so hard to properly model measurements involving quantum fields? Sec. III will answer this question. As a growing portion of the physics community is becoming aware [7–23], one cannot naively transplant our projective measurement theory from non-relativistic quantum theory into QFT. Why is this so?

It is often said [7–23] that the these difficulties can be traced to certain non-trivial mathematical differences between QFT and non-relativistic quantum theory (e.g., new causal and algebraic structures at play). These differences disallow us from directly transplanting our non-relativistic quantum measurement theory into QFT. Proceeding naively leads us to make mathematical blunders which either violate the central 'commandments' of relativity (covariance, causality, and locality)[14, 18–23] or otherwise disrespect the QFT's local algebraic structure. For these reasons, it has been argued we need a new (or at least refined) measurement theory for quantum field theory, see for instance [7, 8, 11].

While all of this is true, I argue that laying the blame on these technical issues belies a deeper methodological issue. Specifically, I will argue in Sec. III that, even when one avoids all such mathematical blunders, other critical symptoms of a deeper methodological issue still remain. In fact, this root methodological issue is the same issue that plagues the sophomore's quantum theory: failure to give concrete dynamical models of our measurement processes, opting instead for simply guessing the right projector.

In this paper, I propose and carry out an alternate way forward: First, we ought to understand that our projective measurement theory for non-relativistic quantum theory has its roots in discussions of measurement chains and Heisenberg cuts. In Sec. IV, I will discuss how these notions automatically give us a case-by-case measurement framework for giving good models of quantum experiments. This already solves the core pragmatic measurement problem in the non-relativistic context. Further, decoherence theory and the Born rule then give us a solution to the extended pragmatic measurement problem, namely they underwrite our usual projective measurement theory.

The central message of this paper is that we should not transplant the end of this story (i.e. our projective measurement theory) into QFT. Rather we should start from the beginning. We ought to generalize these notions of chains and cuts (see Sec. V and Appendix A) and then transplant them into QFT. This will automatically give us a case-by-case measurement framework for QFT, solving its core pragmatic measurement problem. My

---

relativistic quantum theory and classical physics is much smaller than one might expect [5, 6]. Indeed, a robust account of our predictions for many of these experiments unavoidably involves modeling quantum fields at some point along the measurement chain. As such, solving the core pragmatic measurement problem for quantum theory generally ultimately requires an understanding of how to model measurements involving quantum fields.

[5] Throughout this paper, unless specified otherwise, QFT refers to relativistic quantum field theory (as opposed to non-relativistic quantum field theory). Similarly, by non-relativistic quantum theory I mean, unless specified otherwise, first quantized quantum theories (e.g., not non-relativistic quantum field theory).

analysis here reveals the need for a pragmatic *QFT-cut*[6] when modeling QFT experiments which is analogous to the need for a pragmatic Heisenberg cut when modeling quantum experiments generally.

Finally, having understood measurement chains and cuts for QFT, we can then see what might be done about its extended pragmatic measurement problem. We can then see what sort of unified wide-scoping measurement theory (if any) is currently available to us. This is our route to a physically meaningful identification of the observables of QFT. To this end, Sec. VI will provide an overview of the current state of the art in the physics literature. As I will discuss, our current best tool for getting a wide-scoping measurement theory for QFT is the Unruh-DeWitt detector model [15, 23–35], see [7].

II. WHERE THE SPOILS WERE WON

In the introduction, I claimed that while the pragmatic portion of the quantum measurement problem is satisfactorily solved for non-relativistic quantum theory, it's nowhere near to being solved for quantum field theory. I will argue for these claims in Sec. IV and Sec. III respectively. For now, however, let's take these two claims on faith, so that we can estimate the magnitude of their consequences.

If it is indeed the case that we have a pragmatically satisfactory theory-to-experiment link in one case and not the other, it becomes relevant to ask the following. Are the foundational experimental successes of quantum theory more so grounded:

a) exclusively within the domain of non-relativistic quantum phenomena (and so easily-recoverable) or,

b) at least partly within the domain quantum field theoretic phenomena (and so non-recoverable)?

While I expect many philosophers and physicists believe the answer to be a), I rather think the answer is b). Supposing the answer is a) one might respond to the difficulties we currently face in modeling measurements involving QFTs as follows. "Perhaps we can't properly model measurements of quantum particles moving at relativistic speeds and the esoterica of CERN and the LHC. This is unfortunate. However, we can still adequately model the measurements underlying the major foundational experimental successes of quantum theory: the ultraviolet catastrophe, the spectrum of hydrogen, the double slit experiment, etc. As such, the core of quantum theory still has its evidential support. This is non-ideal but it is not a foundational crisis."

―――――

[6] This concept was first introduced in [15] under the name *relativistic cut*. However, after publication we realized this name is apt to cause confusion. For reasons I will discuss in Sec. V I believe the name *QFT-cut* to be more appropriate.

If, rather, the answer is b) then we do have a potential foundational crisis on our hands. Concretely, suppose things are as I claim and explaining all but maybe a small corner of quantum theory's experimental successes unavoidably involves modeling quantum fields. Then without a proper understanding of how to model measurements involving quantum fields, the very core of quantum theory would then be at risk of losing its evidential support.

In a recent talk [5], Wallace has similarly argued that explaining the majority of quantum theory's experimental successes requires quantum field theory:

> Quantum field theory is not just for the esoterica of CERN. Quantum field theory is what we need to understand the scattering of light in the sky.

Several of the following examples are directly inspired by his talk and subsequent paper [6]. As these examples will show, while many experimentally relevant calculations can be done in terms of non-relativistic quantum theory, the route from these calculations to our experimental observations almost always goes through quantum field theory at some point. In particular, even when QFT does not feature prominently in the middle of an experiment, it often will be crucially important for either the initialization or (more importantly) the measurement steps.

However, before getting to these examples some comments are needed on exactly what kind of explaining a theory must do of an successful experimental prediction in order to claim it as evidential support.

A. Explanation and Evidential Support

My claim in this section is that in order for quantum theory to re-secure many of its claims to evidential support, we must re-explain many of its canonical experimental successes now using quantum fields to model at least part of the experiment. But what kind of explaining is relevant for securing evidential support?

One should recall that our concern here is with pragmatic worries, not realist ones. As such, the relevant notion of explanation (i.e., evidential-support-securing explanation) does not involve giving a realist narrative for the experiment. Rather, here our theory explains an experiment when it can give a robust account of how it predicts the (statistical) outcomes seen in this particular experiment. This is the kind of explaining which is relevant for securing evidential support. Unless stated otherwise, this is the kind of explanation relevant throughout this paper.

Moreover, it should be noted that securing evidential support is not a binary issue: Incrementally better explanations of an experiment give us an incrementally more secure right to claim it as evidential support. By contrast, any faults, gaps, or hand-waving in our explanation

put us *at risk* of losing this right; just as a home-made car driving with loose screws is *at risk* of falling apart. I have been careful throughout this text to always speak in terms of risk and security. As in life, we can never remove all risks. Instead, as in life, we ought to reduce our risk to a high (but contextually reasonable) standard. But what should go into this standard?

One's first though is perhaps accuracy. However, as I will now discuss this cannot be all that is relevant. When I earlier claimed that QFT has a pragmatic theory-to-experiment disconnect which puts it at risk of losing evidential support, I anticipate there were some complaints along the following lines: "What do you mean QFT has a pragmatic disconnect? QFT has made many fantastically accurate and precise predictions. For example, the muon $g-2$ experiment at Fermilab has confirmed the Standard Model to a precision of 0.46 parts per million [36]. If something were rotten or ambiguous in the way we make predictions from QFT, it would be a miracle that so many of them are so well confirmed."

While, yes, it is true that QFT has been used to make many extremely accurate predictions, so too has the sophomore's quantum theory discussed in the introduction. Or rather it does so when we have correctly guessed which $\langle \text{out}|$, $\hat{U}$, and $|\text{in}\rangle$ correspond to our real-life lab equipment. For various reasons, we may be very good at guessing. (It is highly intuitive that in a double slit experiment for most pragmatic purposes we can take a position measurement to happen when the electrons hit the screen.) However, in order to properly explain our experiments we need to do better than just guessing the right projector. As I will suggest, we will generally need to model our lab equipment [1–3] in some dynamical detail.

Clearly, we must ask more of our explanations than just accurate predictions. In fact, for the purposes of this paper, nothing hinges on issues of accuracy or precision. For securing evidential support, increased accuracy is only helpful up to within the experiment's error bars. While it is true that a known quantum phenomena being experimentally reconfirmed at ultra-high precision would subsequently require a more precise explanation, this is not relevant here. The experiments in question here are not ultra-high precision, rather I consider here the roughest experiments which clearly demonstrate core quantum phenomena. When I say "we need QFT" in what follows it is never for reasons of accuracy or precision.

If not accuracy, what is at issue here? In what follows my focus will be on methodological issues (namely explanatory gaps and hand-waving) as well as feasibility issues. A good evidential-support-securing explanation can withstand a thorough conceptual audit: What justifies this approximation? How does the measurement apparatus work? How does the experiment's initialization work? What happens to the signal between it being sent and received? How does the sending work? How does the receiving work? One may complain that such questions can continue endlessly, to which I have two responses.

Firstly, given that we are only seeking to explain a given experiment to within its error bars, it's not clear such questions can continue endlessly. And secondly, even if they could, we ought not demand perfection. Our standards ought to be high but still contextually reasonable.

Moreover, a good evidential-support-securing explanation must be computationally tractable. Recall that what is at stake here is the empirical support our theory supposedly gains from certain real-life experiments. We need to be able to deliver real-life fleshed out explanations for our theory's predictions for these experiments in order to be able to claim them as evidential support. (Whether or not computational tractability is required for solving the extended pragmatic measurement problem (i.e., identifying the theory's observables) will be discussed in Sec. VI.)

One may ask: Is it necessary that each individual scientist be able to withstand such an stringent audit about their field's key experiments? Clearly this is too much to ask. The biologist may rightfully point you to an organic chemist when you start asking too many detailed questions about their measurement processes. Just the same, one subfield of physics may outsource explanations of its measurement processes to another subfield of physics. What is important is that we can withstand such an audit collectively. Similarly regarding providing a fully fleshed out computation of our theory's predictions; no individual scientist needs to be able to do this alone. The situation here is much like proof in mathematics, we do not require mathematicians to individually give all of the details of their proofs in terms of elementary logical operations. We do, however, demand that if we were to press the issue then they would collectively be able to give us such a long detailed proof.

Before getting into the examples there is one more important thing to note. In the following examples, what is at issue is not whether an explanation with QFT is better or worse than one without. Rather, the question is whether we need QFT in order to withstand a rigorous conceptual audit and to meet the explanatory standard set above for this particular experiment.

B. Example 1: Ultraviolet Catastrophe

One of the first major successes of quantum theory was resolving the ultraviolet catastrophe. The catastrophe was that classical electromagnetism predicted that thermal bodies ought to emit unbounded amounts of energy through radiation, especially in the low-wavelength (i.e., ultraviolet) regime. Of course, in reality this is not the case. Casual observation of thermal bodies tells us they radiate energy at only a finite rate.

While Planck himself was not motivated by the ultraviolet catastrophe [37], in 1900 he took the first steps towards its eventual resolution. Namely, he did so by adding to classical electromagnetism an *assumption* that electromagnetic radiation can only be emitted or absorbed in discrete energy packets with

$\Delta E_\text{atom} = \Delta E_\text{light} = \hbar\omega$ where ω is the frequency of the light in question. Equipped with this assumption, Planck was able to derive a thermal emission distribution in close close alignment with all known thermal spectroscopic data at both high and low frequencies. Repeating Planck's calculation is a central part of most non-relativistic quantum theory courses.

Is this textbook Planck-inspired resolution of the ultraviolet catastrophe a good one? Does this modified theory of electromagnetism get to claim solving the ultraviolet catastrophe as evidential support? Certainly at the time this explanation was the best available. If we were to then infer to the best explanation, we would have likely inferred to this one. However, in contrast with our modern explanations this Planck-inspired explanation has some serious explanatory gaps.

One comparative improvement is that the second part of Planck's assumption (that $\Delta E_\text{light} = \hbar\omega$) is today no longer viewed as an assumption. To be clear this is still an assumption in non-relativistic quantum theory, but it is a derivable result within quantum field theory. Thus, here quantum field theory provides the best explanation. This is not because it's more accurate or fundamental but because it can derive what the others assume.

However, as discussed above, the question at hand is not whether we can get a better explanation by using QFT than we could otherwise. Rather, the question is whether we can meet the explanatory standard set above *without modeling some part of the experiment within QFT*. Making use of a general QFT result in a non-relativistic context does not require us to model any quantum fields. However, let us next consider the other two parts of Planck's assumption.

Firstly, there is the assumption that the energy change in the atom matches the energy change in the field, $\Delta E_\text{atom} = \Delta E_\text{light}$. The obvious explanation for this (available to nearly any theory) is simply energy conservation ultimately stemming from time-translation invariance. However, a better explanation would provide us a model or mechanism for this energy exchange. But once again, it should be stressed that the question isn't which explanation is best, it's whether we can meet the explanatory standard set above without modeling something with QFT. I will leave the issue open as to whether mechanism-less conservation arguments meet our standards.

The final part of Planck's assumption is that the light-matter interaction can only exchange energy in discrete packets. Why? Answering this question satisfactorily requires us to give some model of the light-matter interaction.

At this point some brief discussion is warranted about the possibility of modeling light within non-relativistic quantum theory. It is impossible [38]. To elaborate, light is fundamentally relativistic. It has a mass of zero and therefore always travels at the speed of light. Moreover, having zero mass means that every energy scale is significantly above the field's mass-energy scale, $mc^2 = 0$.

Thus, there are only two possibilities for faithfully modeling the dynamics[7] of light: 1) if quantum effects are relevant, then within QFT, otherwise 2) classically within either special or general relativity. There is no conceptually coherent treatment of the dynamics of light within non-relativistic quantum theory.

As argued above, to satisfactorily flesh out this Planck-inspired resolution of the ultraviolet catastrophe, we need a model of the light-matter interaction. For this interaction, do we need to model the light as quantum (and hence as a quantum field)? Firstly, we know that light is ultimately quantum and is well-approximated as classical only in a certain regime. Roughly, we can model light classically when its quantum fluctuations are small compared to the field's overall amplitude. That is (even more roughly) whenever there are many photons. For macroscopic objects at relatively high temperatures we can expect many photons to be emitted thermally. In such cases we can model their emission and absorption classically.

Unfortunately, this doesn't work for the ultraviolet catastrophe. The ultraviolet catastrophe is not only about the total emission rate of thermal bodies. Rather, it is about bringing the predicted amount of UV-emissions (infinite) in line with the observed amount (near zero). For any body at any temperature, if we look deep enough into the UV we will eventually find individual photons being emitted sporadically. Explaining this is the key to resolving the ultraviolet catastrophe. These individual photons were emitted by individual atoms. To resolve the ultraviolet catastrophe we need to understand why those individual emission events are so rare. To properly understand these isolated one-photon-one-atom emission events requires[8] a quantum treatment of the dynamics of light and hence quantum field theory.

C. Example 2: Spectrum of Hydrogen

Another major success of quantum theory was explaining the spectrum of Hydrogen. One may have hopes that

———

[7] It is possible to address some of the steady-state properties of light without relativity. For instance, Snell's law of refraction. But how do such steady state situations arise? What interactions produce or absorb the light? Answering any such dynamical questions requires relativity.

[8] In the period between Planck's solution to the ultraviolet catastrophe and Einstein's proposal of quantum light to explain the photoelectric effect, much effort was made to confine the consequences of Planck's discrete assumption [37]. In particular, Planck sought to confine this discreteness to just the matter (and not the light) or alternatively to confine it to only the interaction between the matter and light. However, over the next decade, other physicists began to see that the consequences of his assumption could not be so confined. Of course, we now know that light itself is quantum. Planck-style quantum-matter classical-light description of the micro-physical light-matter interaction cannot withstand a modern conceptual audit.

this can be explained within non-relativistic quantum theory. Indeed, calculating the spectrum of Hydrogen is a standard part of any non-relativistic quantum theory course. But does this undergraduate calculation alone explain our experimental observations? Is it enough to say, "Look, I have calculated the energetic structure of Hydrogen using my non-relativistic theory. Its energy gaps line up one-to-one with the readings picked up by my spectrograph." Case closed?

Unfortunately, this explanation fails to meet the explanatory standards discussed in Sec. II A. Ultimately, it suffers from the same kind of error as the sophomore's quantum theory. It is not enough to simply match up a pattern in your theory with the patterns in the experimental data. That is, it is not enough to simply find some collection of theory objects which give you the right predictions and guess that this is what is going on in your experiment.

Let's begin auditing this explanation. How does the spectrograph work? Classical optics tells us that spectrographs separate light into its different frequencies. Thus, the pattern in our spectrograph is really a pattern among frequencies. Why are we comparing this with a pattern of energies in the atom? Well using $E_{\text{light}} = \hbar\omega$ we can turn frequencies into energies. This is a general QFT result, but that doesn't mean we have to model the light's dynamics using QFT. Fair enough. How did the light get to the spectrograph? Classical electromagnetism tells us about the free propagation of light; let's assume there is enough light to handle this part classically.

So far so good, but where did the light come from? The light was emitted by atoms in the way Planck assumed: in discrete packets with $\Delta E_{\text{atom}} = \Delta E_{\text{light}}$. Here we will begin to run into the same issues as in the previous example. We might accept a mechanism-less energy conservation argument for the equality, but the discreteness of the emission will require some dynamical model of the light-matter interaction. While we may ultimately have a classical amount of light emitted, this was surely itself built up by the emission of individual photons by individual atoms. Why do individual atoms emit individual photons only at certain frequencies? To properly understand this requires a quantum treatment of the dynamics of light and hence quantum field theory.

D. Example 3: Double Slit Experiment

As the previous two examples have shown, a proper treatment of the light-matter interaction in a quantum setting requires quantum field theory. Taking this to heart, let's try to get away from the light-matter interaction. Another major success of quantum theory was explaining the double slit experiment. For variety, let's consider two versions of this experiment done with either neutrons or electrons.

One may have hopes that this can be explained within non-relativistic quantum theory. Indeed, calculating the double slit interference pattern is a standard part of any non-relativistic quantum theory course. But does this undergraduate calculation alone explain our experimental observations? Is it enough to say: "Look, I have calculated the neutron/electron wavefunction just before they hit the screen using my non-relativistic theory. Its squared amplitude lines up one-to-one with the readings picked up by my detectors." Case closed?

For much the same reasons as in the previous example, this explanation fails to meet the explanatory standards discussed in Sec. II A. The error here is of the same kind as in the sophomore's quantum theory. Mere pattern matching between theory and experiment is not a good enough explanation. It is highly intuitive that a position measurement happens when the electrons/neutrons hit the final screen, and moreover for most practical purposes this is in fact what happens effectively. However, as discussed above, we need to do more than simply find some collection of theory objects which give us the right predictions and then guess that this is what is going on in our experiment.

Let's begin auditing this explanation, first considering the case with neutrons. For the concrete real-life experiment in question, how was the neutron beam created? In practice, neutron beams are created by fusing isotopes of hydrogen. We need quantum field theory to properly model fusion. In fairness, one might dismiss this issue[9] claiming we ought to have lower standards for explaining initializations than measurement processes. Fair enough, but what about the measurement of neutrons?

We can also ask: for the concrete real-life experiment in question, how were the neutrons detected? Neutrons do not interact electromagnetically and are too light to be detected gravitationally. It turns out that in practice, neutrons are detected via nuclear interactions, e.g. beta decay. We need quantum field theory to properly model nuclear interactions. Thus, it appears that to satisfactorily explain any neutron double slit experiment we need to model at least part of it using QFT.

Let's next audit this explanation for the case with electrons. For the concrete real-life experiment in question, how was the electron beam created? In practice, there are many ways of creating electron beams. Let us consider two of them: shining a high-intensity laser on metal and cathode ray tubes. Modeling the creation of an electron beam via laser is going to require some description

----

[9] Indeed, there is an intuitive sense in which initialization and measurement measurement processes are different. If our measurement procedure somehow unavoidably involves QFT phenomena, one might think "Fair enough, we have to model that with QFT". However, if the initialization of our experiment involves thermal sunlight, do we really need to trace this back to a QFT-model of fusion in the sun? Of course, we can always ask for the initial conditions of our initial conditions. In practice we have to stop somewhere, hopefully in a principled way. Such considerations need to be built into our high but contextually reasonable standards for explanation. Sec. IV discusses this point further.

of the quantum light-matter interaction. As discussed above, this requires quantum field theory. However, the other option is more hopeful. To produce an electron beam via a cathode ray tube: a metal coil is heated until electrons start to jump off of it. Then these free electrons are accelerated by a classical electric field through a small aperture. As far as I am aware, this can all be modeled without QFT without conceptual error or critical loss of precision.

So far so good. If the electron beam is produced by a cathode ray tube then we can use non-relativistic quantum models to carry us through the initialization step and through the double slits right up to the point when the electrons are detected. In practice there are many ways of detecting electrons. Let us consider two of them: via fluorescent screens and via a cascading avalanche in a semiconductor.

Modeling electron detection via fluorescent screens is going to require some description of the quantum light-matter interaction. As discussed above, this requires quantum field theory. However, the other option is more hopeful. Fast-moving electrons in a semiconductor are detectable as follows. The initial fast-moving electrons knock other electrons out of their valence. These now-free electrons are moving fast enough to go on to disturb further electrons. This ultimately causes a cascading avalanche of electrons. This small current could then be used to activate a transistor and change a bit of memory in a computer. As far as I am aware, this can all be modeled without QFT without conceptual error or critical loss of precision.

Thus, we have found a potential counter example to my claim. In a double slit experiment with a certain kind of initialization and a certain kind of detection mechanism, we can conceivably predict the result of this experiment in a satisfactory way without modeling any quantum fields. I am not sure whether or not anyone has done this particular experiment, but it's conceivable that they could. Does this doom the central claim of this section? I think not.

In two ways this is rather a case of "the exception that proves the rule". Firstly, the existence of an exception proves that the rule doesn't automatically cover everything by some trick of mathematics or logic. The real possibility of exceptions makes each time the rule does apply a substantive mark in its favor. Secondly, looking for an exception has led us to such an odd gerrymandered counterexample (i.e., an incredibly specific combination initialization, evolution, and detection) that it all but proves the rule for normal cases.

As these examples have shown, upon closer examination, explanation of a great many of quantum theory's experimental successes requires us to model them (at least partly) using quantum field theory. It is true that often a key experimentally relevant calculation can be done in terms of non-relativistic quantum theory. However, upon modeling the initialization and (more importantly) the measurement steps, we very often find the need for quantum field theory. In particular, modeling the quantum aspects of the light-matter interaction requires quantum field theory. Moreover, modeling any detection mechanism which involves nuclear interaction (fusion, fission, decay, etc.) will require quantum field theory. With all these measurement procedures removed from our toolbox, not much is left.

While it's conceivable some of the core quantum phenomena could potentially be demonstrated from this limited toolbox, it's not clear anyone has done this. We are hear faced with a choice: 1) one could try to extend this limited non-QFT experimental toolbox as much as possible and then try to recreate as many core quantum phenomena as one can, or 2) one could try to regain a license to use the full toolbox by solving the Pragmatic QFT Measurement Problem.

III. WHAT BLOCKS THEIR WAY HOME

Given the conclusion of the last section (that the route home for most of quantum theory's experimental spoils runs at least momentarily through quantum field theory) our next question becomes: Why is it so hard to properly model measurements involving quantum fields, at least currently?

The main pitfall which people fall into when modeling measurements in QFT is assuming (often implicitly) that the projective measurement theory we had for non-relativistic quantum theory can still be applied in QFT. It cannot in general. Luckily, this pitfall is not impossible to avoid, however doing so requires one to model their experiment's measurement process in some dynamical detail. This section will focus on this pitfall itself, with possible ways around it (for instance, in particle collider experiments, e.g., the LHC) being discussed further in Sec. VI. The question now is why can't we do projective measurements in QFT?

The stock answer to this question in the literature [7–23] is as follows. There are certain non-trivial mathematical differences between QFT and non-relativistic quantum theory which disallow us from directly transplanting our non-relativistic measurement theory into QFT. The first major difference is that we now have field theoretic degrees of freedom. Secondly, we now have a locally Lorentzian spacetime background whose causal structure we must respect. Thirdly, QFT has a very different local algebraic structure than non-relativistic quantum theory (now a Type III rather Type I von Neumann algebra, see [39, 40] although the technicalities are not relevant here). Attempting to transplant our old measurement theory into QFT directly leads us to make mathematical blunders which either violate the central 'commandments' of relativity (covariance, causality, and locality) [14, 18–23] or otherwise disrespect the QFT's local algebraic structure. For these reasons, it has been argued we need a new (or at least refined) measurement theory

for quantum field theory, see for instance [7, 8, 11].

I agree with all of the above and will overview some of these technical issues in this section. However, laying the blame on these technical issues belies two deeper methodological issues. Firstly, even when one avoids all such mathematical blunders, we end up having the same issues that plague the sophomore's quantum theory in the introduction: a failure to give concrete dynamical models of our measurement processes, opting instead for simply guessing the right projector. Intuitively guessing at projectors is wrong in a non-relativistic context and it is wrong in QFT.

Secondly in trying to understand how to model measurement in QFT, our starting point shouldn't be the projective measurement theory we found in the non-relativistic case. In Sec. IV, I will tell a non-relativistic story beginning with measurement chains and Heisenberg cuts and ending with decoherence theory and projective measurements. In QFT we ought not assume this story ends the same way. Rather, we ought to start from the beginning and see what changes.

To demonstrate the above-discussed technical issues and to show how navigating around them does not solve the root methodological issues, I would like to consider a dialogue between the following two characters: Alice and Bob. Alice is a brilliant mathematical-physicist with deep expertise in QFT. Bob, on the other hand, does not know much of anything about QFT. Bob is an experimental-physicist with an interesting new quantum experiment which he wants to explain. Bob has been convinced (perhaps by skimming Sec. II) that some part of his experiment needs to be explained using quantum field theory. Luckily, Bob has a good working relationship with Alice.

In what follows I will sketch a long conversation between Alice and Bob. We hope that with Alice's mathematical expertise and with Bob's familiarity with his experiment, they will together be able to come to a satisfying explanation of how Bob's experiment works. To tip things as much in their favor as possible, let's assume that Alice will be able to prove/disprove any theorem and do any calculation which Bob requests. The only vice we will assume that either of them has is Bob's complete unwillingness to give a concrete model of his measurement device, opting instead for simply guessing the right projector. If Alice and Bob cannot satisfactorily explain his experiment in these extremely favorable conditions, we can place the blame squarely on Bob's methodological vice.

Of course, we can see from the outset that they will not be able to satisfactorily explain Bob's experiment. Bob's vice is exactly what causes pragmatic worries for the sophomore's quantum theory and it will cause the same issues here too. The purpose of this dialogue is to review the technical issues Bob will face and show that even once he overcomes them his deeper methodological issue remains.

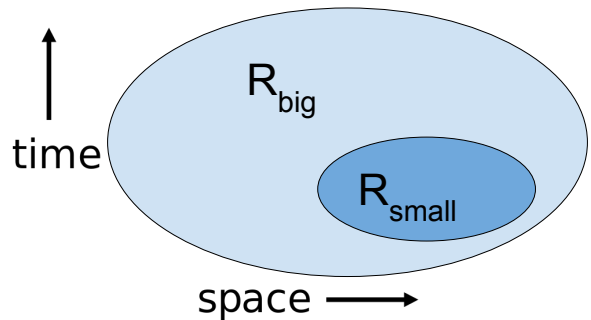For concreteness, specific details will be given about

Figure 1. A spacetime diagram of a late stage in Bob's experiment. Bob's lab equipment is localized in region R_{big} and appears to Bob's ignorant-of-QFT perspective to be counting the number of electrons in some spacetime region R_{small}.

Bob's experiment and Alice and Bob's conversation. However, I stress that these are not instrumental to their failure to satisfactorily explain Bob's experiment.

A. Alice and Bob's Dialogue

Bob begins by explaining the relevant part of his experiment to Alice. Towards the end of Bob's experiment there is a piece of lab equipment localized in spacetime region R_{big} which from Bob's ignorant-of-QFT perspective appears to count the number of electrons in some spacetime region R_{small}. See Fig. 1. These electrons are moving fast enough that Bob thinks he needs to model them with QFT (but adamantly not his measurement procedure itself).

Bob's first thought is that his lab equipment in R_{big} tells him something like: "There are exactly three particles in R_{small} at locations $x_1(t)$, $x_2(t)$, and $x_3(t)$" and that this exactly determines what is going on in R_{small}. Bob asks Alice if there is any projective operator in QFT which fits this description, hoping that if so, maybe his measurement device does that.

Alice says no, for many reasons that sort of thing is not possible in QFT. The first issue that comes to Alice's mind is that this violates the algebraic structure of QFT. In particular it is a theorem [18] for Type III von Neumann algebras (the type relevant for QFT) that all local projectors are of infinite rank. Since all human-doable measurements are local (e.g., my experiment takes place in this room for an hour) such measurement must be of infinite rank.

As a quick refresher on what the rank of a measurement is: Suppose I measure the angular momentum of an atom and find that the quantum number $\ell = 1$. To update the state post-measurement I should apply the rank-3 projector:

$$\hat{P}_{\ell=1} = |1,1\rangle \langle 1,1| + |1,0\rangle \langle 1,0| + |1,-1\rangle \langle 1,-1| \quad (1)$$

where $|\ell, m\rangle$ are the atom's angular momentum states.

This measurement is rank-3 because there are three perfectly distinguishable (i.e., orthogonal) possibilities consistent with the outcome $\ell = 1$, namely $m = 1, 0$, and -1. Contrast this with a measurement of whether a quantum harmonic oscillator is in an "even" or "odd" energy state. This is an infinite-rank measurement since, in either outcome, there are still an infinite number of distinguishable possibilities left after the measurement.

Alice's theorem thus tells us that in QFT all local projective measurements must leave not only multiple options open, but an infinite number. As such, no local projective measurement can "exactly determine what is going on in R_{small}" as Bob has requested. Bob takes this lesson to heart. However, per his methodological vice, Bob really doesn't want to provide a model of his measurement device. Rather Bob opts to look for a loophole.

Bob comes to Alice the next day and asks her if there is any projective operator in QFT which could tell us "There are three particles in R_{small} (wherever they may be)" thinking this would leave an infinite number of options open. Unfortunately, Alice says that this still doesn't work citing a theorem that there are no well-defined local number operators, $\hat{N}(R_{small})$, in QFT [18].

To give Bob some intuition she tells him that roughly this is because $\hat{N}(R_{small})$ would count $+1$ particles for all of the particle-antiparticle pair production events in R_{small} and would consequently diverge terribly. Bob takes this lesson to heart. However, per his methodological vice, Bob really doesn't want to provide a model of his measurement device. Rather, as before, Bob opts to look for yet another loophole.

Bob comes back a week later and says to Alice, "Maybe we can fix the divergence if we count particles and antiparticles with different signs. Then pair production events will count as $+0$ and therefore won't cause a divergence. Maybe when I said 'There are three particles in region R_{small}' if I was being more careful I should have said 'There are exactly three excess un-paired particles in R'". Alice says "Yes Bob that's very clever and it does fix the divergence issue... However, when you try to define a (signed) local number operator something goes wrong with the commutators which leads to faster-than-light signaling [18]."

Alice and Bob's conversation continues on like this for a long time. Let's go ahead and skip to the end.

B. Limit of Alice and Bob's Dialogue

What is the limit of Alice and Bob's dialogue? Ultimately, Bob is looking for a local operation defined on the spacetime region R_{big} which fits the description of the behavior he sees in his lab. Specifically he is looking for a local projective operator. As their conversation continues, Bob will learn a lot about QFT and he will make increasingly careful descriptions of what he is looking for. (Perhaps, ultimately he will be forced to stop talking in terms of localized particles, but this is besides
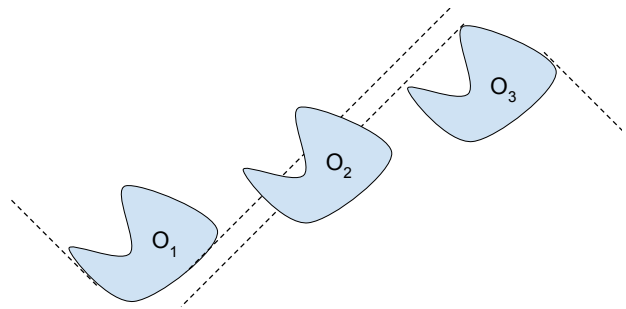

Figure 2. The impossible measurement scenario considered in Sorkin's [14]. Notice that regions O_1 and O_3 are space-like separated from one another. There is a mathematically well-defined local projection operation in the algebra associated with O_2 such that enacting this projector allows for signaling from O_1 to O_3. Thus not all mathematically well-defined local operations are physically allowed.

the point.) This conversation seems to have been very productive and clarifying for Bob.

Ultimately, what Alice is doing is telling Bob what local operations are and are not well-defined in QFT consistent with its causal and algebraic structure. A quick aside about the possibility of causality violations. Sorkin [14] was the first to notice that if one naively does projective operations in QFT then faster-than-light signaling is possible. See Fig. 2. Roughly, there are some mathematically well-defined local projection operations in region O_2 which will allow for signaling from region O_1 to region O_3. Thus projective operators in QFT must be physically as well as mathematically permissible.

Suppose Alice understands all of these potential causality issues. The best that she could possibly do in this conversation is to provide Bob with a full menu of which local operations are relativistically safe in QFT, i.e. consistent with its causal and algebraic structures. Let us suppose this happens. That is, Alice provides Bob with the relevant necessary and sufficient conditions. (We do in fact have these conditions [8, 9] at least for real scalar QFT in a globally hyperbolic spacetime.)

Notice, however, that Alice cannot yet in good faith claim that these relativistically-safe local operations characterize the "observables" of QFT. So far these operations have been identified based on the following principle: they don't allow for faster-than-light signalling. However, in order to properly identify a theory's observables, one must rather give some constructive/dynamical story about which measurement processes pick out these observables. One cannot have a physically meaningful characterization of a theory's "observables" without examining the dynamics of its measurement processes. More will be said about this in Sec. VI.

Overall, this conversation seems to have gone very well. Bob has learned a lot about QFT and Alice has proven a beautiful interesting theorem outlining the scope of what local operations are relativistically-safe in QFT. But is Bob any closer to giving us a satisfying account of what

quantum field theory predicts for his experiment? What can Bob do given Alice's theorem?

Bob wants a relativistically-safe local operation which lines up well with the observed/intuitive behavior of his lab equipment. Specifically he is looking for a local projective operator. He might search for this using Alice's theorem. There are two issues with Bob's methodology.

Firstly, we currently have no reason to believe that Bob's search will be successful. Indeed, we have no guarantee that any given physical measurement processes can be accurately modeled using a relativistically-safe projective operator. Of course, whenever a physical measurement process can be modeled as a projective operator, it must be one of the relativistically-safe one's identified by Alice's theorem. But what reason do we have to expect that all physical measurement processes in QFT can be modeled as projective? We may have good reasons to expect this, but these reasons must be grounded in detailed dynamical analysis of generic measurement processes (something Alice has not yet done).

Secondly, suppose that Bob's measurement process can in fact be modeled by some relativistically-safe projection operation. Moreover, suppose that Bob finds a relativistically-safe projective operator which lines up well with the observed/intuitive behavior of his lab equipment. Has Bob then explained how QFT predicts his experimental results? Of course not. As discussed in Sec. II, one needs to do more than intuitively guess the right projector. It is not enough to find a pattern among theory objects which matches up with the patterns in your experimental data. Some further explanation must be given to meet our explanatory standards.

To review: Yes, there are non-trivial mathematical issues which disallow us from straightforwardly transplanting our usual projective measurement theory into QFT. However, once these issues are carefully analyzed and we have identified the relativistically-safe local operations, we are still no better off. The issues are the same here as with the sophomore's quantum theory. We understand which local operations are conceptually allowed by our theory (i.e., the relativistically-safe ones) but past this we are just guessing how these allowed local operations relate to our concrete real-life experiments.

As I have argued above, ultimately we should not be trying to directly transplant our projective measurement theory into QFT in the first place. As I will discuss in the next section, our non-relativistic measurement theory has its roots in discussions of measurement chains and Heisenberg cuts. We ought to instead transplant these notions of chains and cuts into QFT. We can then see what measurement theory (if any) we are led to.

IV. MEASUREMENT CHAINS AND HEISENBERG CUTS

As the previous section has shown, despite the technical complications of quantum field theory, the core issue in the Pragmatic QFT Measurement Problem is at heart the same as it was in the non-relativistic case. As I have claimed throughout this paper, for non-relativistic quantum theory the pragmatic portion of the quantum measurement problem has been satisfactorily solved in terms of measurement chains and Heisenberg cuts[10]. This section will review this situation in preparation to transplant it into QFT.

One ought to here recall the distinction between the core and extended pragmatic measurement problems discussed in the introduction. In the core problem we demand a (potentially case-by-case) explanation of how our theory makes its predictions for certain key experimental successes. In the extended problem, we demand a unified account of how our theory makes predictions across a wide range of experiments. As I will now discuss, for non-relativistic quantum theory the solution to the core problem is given in terms of measurement chains and cuts. Past this, a solution to the extended problem is given in terms of decoherence and the Born rule: these underwrite our usual projective measurement theory for non-relativistic quantum theory. We ought not transplant the end of this story into QFT, but rather ought to start at the beginning and see what changes.

A. Heisenberg Cuts in Theory

How have measurement chains and Heisenberg cuts helped us overcome the core pragmatic measurement problem for non-relativistic quantum theory? First allow me to review these concepts.

We can often make sense of our experiments in terms of a measurement chain. Roughly, a measurement chain is the sequence of interactions which carries the measured information from the systems being measured to some record-keeping device. An abstract example: System A interacts with system B which then interacts with system C which then ... which then interacts with system R, for record-keeping device. More will be said about how one can choose the start and end of these chains momentarily.

In Appendix A, I generalize measurement chains to fit experiments which are not so sequential as this one. For our present purposes, however, restricting ourselves to the sequential A-then-B-then-C measurement chains will suffice.

———

[10] Although, as discussed in Sec. II the scope of quantum theory's key experimental successes which are analyzable entirely within non-relativistic quantum theory and classical physics is much smaller than one might expect [5, 6]. Indeed, a robust account of our predictions for many of these experiments unavoidably involves modeling quantum fields at some point along the measurement chain. As such, solving the core pragmatic measurement problem for quantum theory generally will ultimately require an understanding of how to model measurements involving quantum fields.
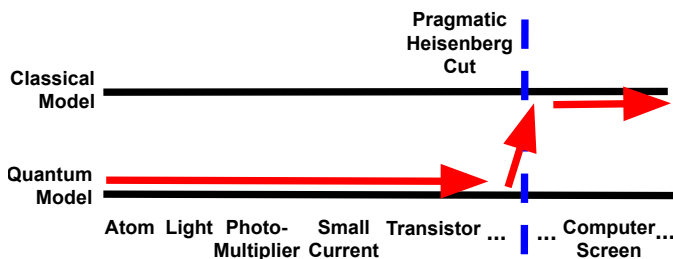
Figure 3. The measurement chain of a simple atomic experiment. The black lines show two possible types of models: quantum or classical. The red arrow shows which part of the experiment we are modeling with which theory. The dashed blue line shows where we are taking the pragmatic Heisenberg cut. That is, where we switch from modeling the experiment in a quantum way to a a classical way.

For a more concrete example, we may be interested in a certain amplitude associated with an atom in a certain superposition. Our experiment may proceed as follows: An atom in a superposition emits a photon which is detected by a photo-multiplier which triggers a small current which turns on a transistor which ... which displays a number on a screen which the experimenter writes in her notepad.

To be clear, in this paper, the measurement chain does not refer to the linear sequence of *physical systems/interactions* which carry the measured information, per se. Rather, here, the measurement chain is a formalization of these systems which the experimenter invents for the purposes of modeling her experiment. There may be multiple acceptable ways of parsing a given physical scenario into a formalized measurement chain.

Indeed, given an experiment, it's not always clear where we ought to place either end of the measurement chain. Regarding initialization, we can always ask for the initial conditions of our initial condition. Regarding the late stages of measurement, it's unclear where to stop: the computer screen, the experimenter, her notepad, etc. Such considerations need to be built into our high but contextually reasonable standards for evidential-support-securing explanations. The results of this paper do not depend sensitively on how this is done.

The above-discussed concrete example of a measurement chain is laid out horizontally in Fig. 3. It's important to note that while in this example, moving horizontally happens to move us into larger, more complex systems with more degrees of freedom, this is accidental. Horizontal movement in this diagram indicates only that we are moving from one system to another sequentially towards the end of the experiment. One can easily imagine experiments where advancing forward in the experiment temporarily moves the measured information into a smaller system. (Indeed, such a scenario is displayed in Fig. 4.) The two horizontal black lines in Fig. 3 represent two types of model that we could have for each part of our experiment (here, either classical and quantum).

The red arrows indicate how we are going to model each part of the experiment.

It's important to note that the path that the red arrow takes through this diagram is, in large part, a free choice of the experimenter. The location of the red arrow does not mean that this or that system *is* quantum/classical. All that this indicates is that, for the purposes of modeling this particular experiment, this particular experimenter has chosen to model this system as such.

However, importantly, it is not the case that any part of an experiment can be successfully modeled using any theory. In practice, there are always going to be some restrictions. Sometimes for the sake of accuracy it will be necessary to model a given system in a quantum way. Sometimes for computational or technological reasons it will not be feasible to model a given system in a quantum way (forcing us to model it classically). Sometimes it will be conceptually necessary to model a given system in a quantum way. (Perhaps, sometimes for conceptual reasons it will be necessary to model a given system in a classical way.) For these and other reasons, the possible routes which the red arrows may make through these diagrams are limited.

With these restrictions in mind, it may occur that modeling some part of our measurement chain in a quantum way is both conceptually mandatory and technically infeasible. In this case we simply cannot (yet) explain this experiment satisfactorily.

Perhaps it's best to work through an example. Consider the double slit experiment with electrons discussed in Sec. II D. The measurement chain for this experiment is shown in Fig. 4. Note that two red arrows are shown. The bottom red line opts for a quantum model whenever possible, whereas the top red line opts for a classical model wherever possible.

The experiment begins with many lab operations. As discussed above, there is some freedom in picking where exactly the measurement chain starts. However, whatever one chooses, the preliminary lab operations can be described classically. Indeed, it is infeasible to model these lab operations with quantum theory. Recall that our purpose here is to provide actual fully-modeled explanations of real-life experiments. Thus both of the red arrows in Fig. 4 *must* start on the top line.

These lab operations set up a current which travels through a filament in our cathode ray tube. This heats the filament which begins to thermally emit electrons. These electrons are then grabbed by an electric field and accelerated through a small aperture. As far as I know, all of these steps can be modeled classically without conceptual error or critical loss of accuracy. Hence the upper red arrow in Fig. 4 stays on the top row. As far as I know, all of these steps can be feasibly modeled quantumly. Hence the bottom red arrow in Fig. 4 jumps to the bottom row.

The next part of the experiment (the motion of these electrons through the double slit apparatus) must be modeled with quantum theory. There are both concep-
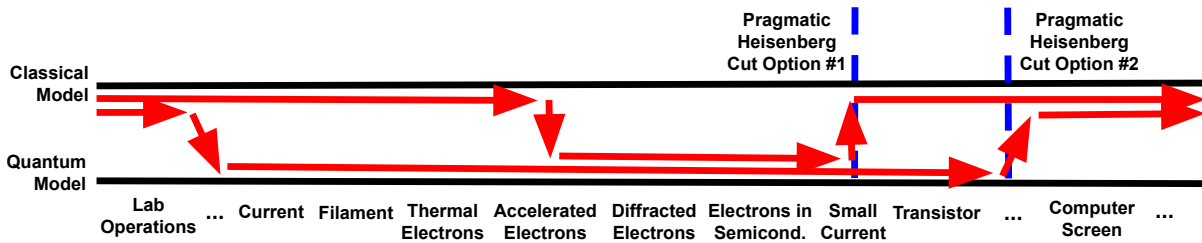
Figure 4. The measurement chain of one of the electron double slit experiments discussed in Sec. II D. The black lines show two possible types of models: quantum or classical. Each of the red arrows shows which parts of the experiment we are modeling with which theory. The bottom red line opts for a quantum model whenever possible, whereas the top red line opts for a classical model wherever possible. The dashed blue line shows where these two approaches to modeling this experiment place their respective pragmatic Heisenberg cuts.

tual issues and accuracy issues with modeling this part classically. Hence both of the red arrows *must* be on the bottom row here. Note that there is nothing per se quantum about electrons moving through an aperture. Whether we can ignore quantum effects present at this point in the experiment depends sensitively on what's coming later.

When the electrons reach the final screen they enter into a semiconductor. There they are detected by causing an cascading avalanche of electric discharge as discussed in Sec. II D. These electrons jumping over the semiconductor's band gap requires a quantum model. Hence both red arrows *must* be on the bottom row here.

However, once enough electrons are moving we can describe them collectively as a small (but classical) current. This current activates a transistor. Some (but not all) transistors make use of quantum effects, but let's assume this one doesn't. As far as I know, all of these steps can be modeled classically without conceptual error or critical loss of accuracy. Hence the upper red arrow in Fig. 4 moves to the top row. As far as I know, all of these steps can feasibly be modeled quantumly. Hence the bottom red arrow in Fig. 4 moves along the bottom row.

The sequence of events which follows the activation of this transistor can all be described classically. Indeed just as at the start of the experiment, it is infeasible to model the end of an experiment quantumly. As discussed above, there is some freedom in picking where exactly the measurement chain ends. However, whatever one chooses this part of the experiment can and must be described classically. Computer screens and humans and notepads are simply too large and complicated to model in a quantum way (at least for now and possibly forever).

As this example hopefully makes clear, whenever we have a measurement chain, part of which requires a quantum model, we will have to at some point after this switch from modeling the measurement chain quantumly to nonquantumly (i.e., classically). This claim is formalized and proved in Appendix A. In connection with the histori-

cal term[11], let us call wherever we happen to make this switch the *(pragmatic) Heisenberg cut.* As this is the only type of Heisenberg cut considered in this paper, I may sometimes drop the parenthetical.

This example should hopefully also make clear that there is nothing fundamental about the placement of the pragmatic Heisenberg cut. Indeed, one can believe the world to be quantum through-and-through and still make use of this cut for modeling purposes. Past the cut, we are no longer *modeling* the measurement apparatus using quantum theory; This is very different from the measurement apparatus no longer *being* quantum past the pragmatic Heisenberg cut.

At this point one may wonder: if the application of a pragmatic Heisenberg cut is a matter of non-fundamental pragmatic concern only, then do we really need it to make sense of the quantum measurement problem? One may ask: If we believe that the world is quantum through-and-through, then why would it be necessary to connect our quantum model of reality with a (known-to-be-incorrect) classical model of reality in order to model measurements within it? Can't we have a quantum-native understanding of quantum measurements?

Firstly, this conflates the realist and pragmatic worries about quantum measurements which I have taken care to distinguish: i.e., modeling versus understanding. If one wants to make such all-quantum-all-the-time demands on the realist side of the debate, you are more than welcome. However, as the above discussion has hopefully shown, this attitude is not tenable in the pragmatic side. The goal here is an actual tractable connect-the-dots model-to-model account of real-life quantum experiments. This is necessary in order to claim them as evidential support for quantum theory. My claim is that (at least currently

----

[11] I make no claim that the way that I am using the term "Heisenberg cut" here has any connection with Heisenberg's understanding of quantum theory. Rather, I employ the term here as it is used colloquially in modern discussions. No exegesis is attempted or implied.

and likely forever) a pragmatic Heisenberg cut is necessary for this.

It is perhaps possible (although I strongly doubt it) that we will one day be able to model the late parts of our experiments (including the experimenter) as quantum systems. However, even this possibility would not necessarily avoid the need for a Heisenberg cut. Suppose that somehow I can model an experiment up to and including the experimenter in a quantum way. It could still be the case that I can only parse the result of that experiment by means of taking some sort of classical approximation (i.e., taking a Heisenberg cut) on the experimenter right at the end [41]. We cannot have a quantum-native understanding of measurement without a quantum-native understanding of the observer. Thus, in the absence of both tremendous computing capabilities and a quantum-native understanding of observers, taking a pragmatic Heisenberg cut is necessary for any satisfactory explanation of any quantum experiment.

In fact, not only is it necessary to take a pragmatic Heisenberg at some point, we ought to do so explicitly and in a carefully formalized way. Indeed, a mishandling of the pragmatic Heisenberg cut is one of the main dangers in trying to explain quantum experiments. It is at the interface between our quantum and classical models that we need the most care both mathematically and conceptually. Handling this cut somewhere explicitly in the terms of either the dynamics or kinematics of our models is far superior to the sophomore's strategy of intuitively guessing the right projector. Indeed, as I will soon discuss, the success of our hand-waving about projective measurements is largely underwritten in terms of measurement chains and Heisenberg cuts in combination with decoherence theory and the Born rule. However, before discussing this, it is worthwhile to provide a taxonomy of all the ways one might take a Heisenberg cut.

B. Taxonomy of Heisenberg Cuts

As discussed above, a pragmatic Heisenberg cut occurs wherever along the measurement chain we switch from modeling our experiment quantumly to classically. This section lays out three interdependent dichotomies for classifying Heisenberg cuts: forward vs reverse, vertical vs diagonal, and with or without back-reaction. These distinctions will be formalized in Appendix A.

The forward/reverse distinction just simply indicates which direction we are crossing the cut: here this means whether we are moving from a quantum model to a classical model or vice versa. Fig. 3 does not have a reverse Heisenberg cut, whereas Fig. 4 does. When no qualifier is applied, reference to any cut should be assumed to mean a forward cut, rather than a reverse one.

The vertical/diagonal distinction is as follows. Given that a measurement chain is ultimately just a collection of interactions which are ordered in some way, there are only two ways to cross the cut: in between interactions,
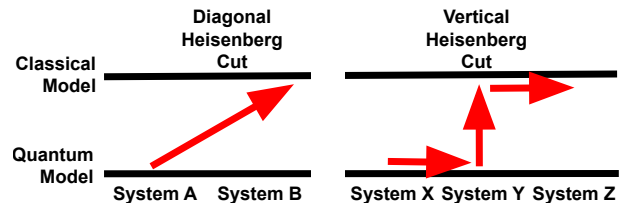

Figure 5. The two possible ways of taking a Heisenberg cut: diagonally (during an interaction) and vertically (in between interactions). On the left we have an example of a diagonal cut: system A is modeled quantumly and system B classically. Their interaction couples two systems modeled in different theories. On the right, we have an example of a vertical cut: The interaction between system X and Y is modeled within quantum theory, whereas the interaction between Y and Z is modeled classically. In between these interactions we apply some classical approximation scheme to Y while it is isolated.

or during an interaction. Let us call these vertical and diagonal respectively (for reasons which will become clear soon, see Fig. 5).

Regarding vertical cuts: consider an interaction between three systems: system X (which we model quantumly) and system Y (which we can model either quantumly or classically) and system Z (which we model classically). We model the interaction between X and Y quantumly and the interaction between Y and Z classically. In between these two interactions (after X and before Z) we take some classical approximation scheme on system Y in isolation. See the right side of Fig. 5 and notice that the red arrow moves vertically at system Y, hence the name "vertical cut".

Examples of a vertical Heisenberg cuts of varying quality from quantum to classical are:

1) taking some sufficiently decohered quantum state, and using the Born rule to map it onto a probability distribution,

2) taking a quantum state whose Wigner function [42] (i.e., the state's quasi-probability distribution in phase space) happens to be positive and reinterpreting it as a genuine probability distribution,

3) taking a minimum uncertainty quantum state and mapping onto the definite classical state with matching expectation values.

among many other possibilities [43]. Each of these may or may not be justified to differing degrees in different contexts.

The above are all examples of vertical forward cuts. We can also, of course, have vertical reverse cuts. That is, we may have principled ways of mapping classical states onto quantum states. For example, the reverse of each of the above discussed examples are sometimes justified.

Regarding diagonal Heisenberg cuts: Consider an interaction between system A (which we model quantumly)

and system B (which we model classically). See the left side of Fig. 5 and notice that the red arrow moves diagonally between systems A and B, hence the name "diagonal cut". The interaction between systems A and B is not entirely within quantum theory nor is it entirely within classical theory; It has a foot in both camps. That is, a diagonal cut involves a dynamical interaction between systems modeled in fundamentally different theories: here quantum and classical. This is in strong contrast with the vertical Heisenberg cuts discussed above which applied to *only one system* not two, and describe it *in isolation* not in interaction. There are reasons to favor vertical cuts over diagonal ones ceteris paribus, however as these are fairly general I will delay discussion of them to Sec. V. Before that, let's see some examples of diagonal Heisenberg cuts.

Albeit artificial, consider the following pair of coupled differential equations:

$$\partial_t |\psi_A(t)\rangle = \left(\frac{\hat{p}_A^2}{2\,m_A} + U_A(\hat{x}_A) + V(\hat{x}_A - y_B(t)) \right) |\psi_A(t)\rangle$$

$$m_B\,\partial_t^2 y_B = -\partial_{y_B} U_B(y_B) - \partial_{y_B} V(y_B - \langle \hat{x}_A(t) \rangle) \qquad (2)$$

for some potential functions U_A, U_B and V. Here we have a wavefunction, $|\psi_A(t)\rangle$, and a classical position, $y_B(t)$, each evolving under their free dynamics, U_A and U_B, plus an interaction term, V. The dynamics of $|\psi_A(t)\rangle$ depends on $y_B(t)$ while simultaneously the dynamics of $y_B(t)$ depends on $|\psi_A(t)\rangle$ through its expectation value, $\langle \hat{x}_A(t) \rangle = \langle \psi_A(t)| \hat{x}_A |\psi_A(t)\rangle$.

The above is an example of a diagonal forward cut. We can also, of course, have diagonal reverse cuts. For example, by reversing the indices $A \leftrightarrow B$.

In either a forward or reverse diagonal cut, a helpful approximation that one can make is to ignore back-reactions. In the above example this corresponds to removing the V term from system A's dynamical equation but not B's. We can only remove V from A's dynamics because in order for the measurement information to progress along the measurement chain, system B needs to be responsive to system A.

We can, of course, remove V from B's dynamics if we reverse the roles of A and B (i.e., if we take A to be after B in the measurement chain). Here we actually find something which is very commonly done in practice. Consider a quantum system which interacts with a classical electromagnetic field through a classical potential. The electromagnetic field may even evolve during this interaction under its own free dynamics, thereby changing the classical potential, and driving the quantum system. Note however, that in this example, there is no back-reaction on the classical electromagnetic field by the quantum system.

C. Heisenberg Cuts in Practice

Having reviewed the notion of measurement chains and Heisenberg cuts and classified them into various types, we are now ready to discuss how they work in practice. How do they help us overcome the core pragmatic measurement problem for non-relativistic quantum theory? How do they, as I claim, overcome the extended pragmatic measurement problem which combined with decoherence theory and the Born rule? How do they underwrite our usual projective measurement theory?

We have so far established that it is necessary to take a pragmatic Heisenberg cut at some point in any satisfactory explanation of any quantum experiment. Moreover, as we have seen, there will in general be many ways and points along the chain where we might take this pragmatic Heisenberg cut. Given the mathematical and conceptual dangers inherent in crossing the Heisenberg divide, it is important that we take our Heisenberg cut explicitly and in a carefully formalized way.

Analyzing a given quantum experiment in terms of measurement chains and various potential Heisenberg cuts gives us a road map to guide us in modeling its specific measurement processes. In particular, the road maps given to us by these notions have the dangerous areas ahead clearly marked out, as well as multiple possible routes to navigate them. Given our task, this is an indispensable tool.

By using these road maps we can go about giving satisfactory predictions for quantum experiments and gaining evidential support from them, at least on a case-by-case basis. This goes a long way towards solving the core pragmatic measurement problem. In essence, we now have a measurement framework for non-relativistic quantum theory. This is a much better position to be in than relying on the sophomore's strategy of merely intuitively guessing the right projector. We now have the tools to say, on a case-by-case basis, what it is that non-relativistic quantum theory predicts for any experiment within its remit (although this remit is much smaller than one might think, see Sec. II).

But what about the extended pragmatic measurement problem? These case-by-case road maps for modeling non-relativistic quantum experiments cannot help us here; they give us a case-by-case measurement framework, not a unified measurement theory. It would be nice if there was something we could say about quantum measurement processes generally (although, a priori we have no reason to expect this is possible).

We are tremendously lucky that such a thing is, in fact, possible for non-relativistic quantum theory. If our goal is to provide a wide-scoping analysis of non-relativistic quantum measurement processes, then we need to find one route across the Heisenberg divide which is available near-universal across all experiments. In terms of wide applicability, one way of crossing the Heisenberg divide stands out: namely, by using decoherence theory and the Born rule.

It should be stressed that within my analysis the only thing special about crossing the pragmatic Heisenberg divide in this decoherence way is its general applicability. As suggested above, we may be able to explain quantum

measurements on a case-by-case basis without it. However, it is decoherence theory's general consideration of measurement processes which allows it uniquely to underwrite a unified wide-scoping measurement theory. If one could find another equally general way of crossing the Heisenberg divide, one can equally well use that to underwrite a different complementary measurement theory. This is an interesting possibility which may shed new light on justifications for the Born rule, but this is outside of the scope of this paper.

With that said, let us restrict our attention to vertical Heisenberg cuts which are facilitated by decoherence theory and the Born rule. Where along the measurement chain are we justified in taking this specific kind of pragmatic Heisenberg cut? Luckily, to this we have a completely general answer: one can take such a pragmatic Heisenberg cut once enough decoherence has occurred that the possibility of spontaneous wide-scale recoherence (although not mathematically impossible) is practically inconceivable. That is, for modeling purposes, it does not matter where we put such a pragmatic Heisenberg cut so long as it is at a scale where quantum effects are (and will forever remain) irrelevant in practice.

The above "and will forever remain" caveat is a critically important one. It reinforces the warning that the pragmatic Heisenberg cut should not be thought of as being fundamental. Indeed, the validity of such a Heisenberg cut approximation will always depend on the context surrounding the measurement procedure under consideration. In particular, one cannot simply decide in the middle of modeling a measurement to take such a Heisenberg cut without knowing beforehand what the rest of the measurement procedure will be like.

No matter how small quantum coherence effects appear to be in the middle of an experiment, there is always a possibility that the coherence effects are brought back to their full force[12]. Moreover, even if within one experiment the quantum coherence effects never again become relevant, they may once again become relevant in other future measurements involving correlated systems. The consideration of measurements made by observers who themselves live inside of a giant quantum computer capable of wide-scale (observers included) recoherence, leads to interesting Wigner's friend-like puzzles.

With these caveats noted, in general once enough decoherence has occurred one can take such a pragmatic Heisenberg cut. The decoherence process naturally picks out the relevant basis against which we are going to take our classical approximation. That is, the projectors for our measurement are picked out by this dynamical process. The classical approximation is then facilitated by the Born rule (i.e., mapping diagonal density matrix elements to classical probabilities). We have thus arrived at an answer to the sophomore's question: In modeling this

———

[12] Such a carefully orchestrated large-scale recoherence is, in fact, exactly what quantum computers are designed to do.

piece of lab equipment, exactly which projectors am I supposed to use and exactly when? The relevant projectors are those approximately picked out by decoherence theory, and the relevant time is any time which is late enough.

The story about measuring quantum systems doesn't end here however. Taking this vertical Heisenberg cut has just moved the measured information into a now-classically-modeled system. How do we then model the measurement of this system? To answer this we can appeal to our classical measurement theory. Let us assume that we are satisfied with our classical measurement theory. In this case, the general considerations above effectively give us a wide-scoping measurement theory for non-relativistic quantum theory. That is, our non-relativistic quantum measurement theory is just whatever is induced when our classical measurement theory is applied after any late-enough decoherence-based Heisenberg cut. Working this out mathematically, one finds that generic measurement processes are associated with projector-valued measurements (PVMs) or positive-operator-valued measurements (POVMs).

This is the solution to the extended pragmatic measurement problem given by decoherence theory. In the next section I will discuss which parts of the above story can be carried over to QFT. We cannot a priori expect every feature of the above discussion to generalize to QFT. To preview, the ideas which will generalize are the measurement chain and the various kinds of Heisenberg cuts available to us. As above, these automatically give us a case-by-case measurement framework for QFT, i.e., a solution to its core pragmatic measurement problem. However, if we then want a solution to the extended pragmatic measurement problem for QFT then some more work must be done.

It is at this point that the QFT story begins to diverge. All of the later details of the above story (the decoherence process picking out a basis, the Born rule associating probabilities to elements of the density matrix, the association of measurements with PVMs and POVMs) do not transfer straightforwardly into QFT. If they did then we would be able to do generic projective measurements in QFT. As discussed in Sec. III we cannot.

In order to get a wide-scoping measurement theory for QFT we would need to find a near-universally applicable way of crossing between QFT-models and non-QFT models. We might then be able to induce a wide-scoping measurement theory for QFT.

V. GENERALIZED MEASUREMENT CHAINS AND CUTS

The above introduced conceptual apparatus of measurement chains and Heisenberg cuts is clearly going to be useful outside of its original context. Let's review some of the key features discussed above but stripping away any incidental details.

Firstly, we need an experiment which is analyzable in terms of a measurement chain. Until now we have assumed our measurement chains are sequential (i.e., System A interacts with system B which then interacts with system C which then ... which then interacts with system R, for record-keeping device). However, clearly not all experiments are able to be broken down in this way.

In Appendix A I generalize my discussion of measurement chains to include experiments which are not so sequential. Some of these generalizations are designed to better accommodate QFT. Roughly, an experiment is analyzable in terms of a measurement chain exactly when: 1) it can be understood in terms of a finite number of distinct systems/fields undergoing a finite number of distinct interactions, 2) at each interaction for every system there-involved it is being modeled in some theory, and 3) these interactions are related to each other by some notion of inputs and outputs. For our present purposes, however, restricting ourselves to the sequential A-then-B-then-C measurement chains will suffice.

Let's formalize sequential measurement chains. There may be many ways to parse such an experiment into a measurement chain especially at the ends (see Sec. IV). Whatever choices we end up making, they ought to stand up to a reasonable level of scrutiny. However, once parsed we will have the following sort of structure. A *(modeling of) a sequential measurement chain* is a 5-tuple $\langle S, <, T, H_{\text{early}}, , H_{\text{late}} \rangle$. S is a finite non-empty set containing labels for the systems (or fields) being modeled for our experiment. S comes along with some total ordering $<$ over S, i.e., $A < B < C < ... < R$. These systems interact pairwise between adjacent systems from start to end. T is a finite non-empty set containing labels for the theories which are being used to model parts of our experiment. $H_{\text{early}} : S \to T$ and $H_{\text{late}} : S \to T$ are functions which tell us within which theory $s \in S$ is modeled before and after its period of isolation, i.e., in its interaction with its predecessor or successor. A more general formalization of measurement chains is given in Appendix A.

In the above formalization, a vertical cut happens at $s \in S$ when $H_{\text{early}}(s) \neq H_{\text{late}}(s)$. That is, when s is modeled in different theories in its two interactions. Note that, between these we must somehow map its state from one theory to the other. Note that we could also have a vertical cut right at the end as well if $H_{\text{early}}(R) \neq H_{\text{late}}(R)$. In this case, our record keeping device is modeling in one way during its only interaction, then after this we take some approximation on it before its final analysis at the end of the experiment.

A diagonal cut occurs between $s, t \in S$ when s and t are adjacent with $s < t$ and $H_{\text{late}}(s) \neq H_{\text{early}}(t)$. That is, when in their interaction s and t are modeled in different theories. Diagrammatically, these two types of cuts are represented just as before, see Fig. 6. These definitions and diagrams are generalized in Appendix A.

Allow me to now introduce some generalized nomenclature. Suppose that somewhere along the measurement
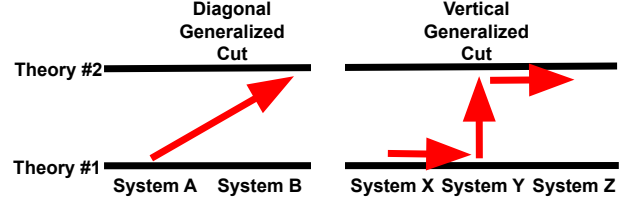

Figure 6. The two possible ways of taking a generalized cut: diagonally (during an interaction) and vertically (in between interactions). On the left we have an example of a diagonal cut: system A is modeled in Theory 1 and system B in Theory 2. Their interaction couples two systems modeled in different theories. On the right, we have an example of a vertical cut: The interaction between system X and Y is modeled within Theory 1, whereas the interaction between Y and Z is modeled within Theory 2. In between these interactions we apply some approximation scheme to Y while it is isolated.

chain we switch from modeling using theories in $V \subset T$ to another theory in $W \subset T$ with $V \cap W = \emptyset$. Let's call this a (V, W)-cut. In particular, a vertical (V, W)-cut occurs at $s \in S$ whenever $H_{\text{early}}(s) \in V$ and $H_{\text{late}}(s) \in W$. By contrast a diagonal (V, W) cut this occurs between two adjacent systems, s and t with $s < t$ when $H_{\text{late}}(s) \in V$ and $H_{\text{early}}(t) \in W$. In either case, when $W = T - V$ is the complement of V under T we can just call a $(V, T-V)$ cut simply a V cut. Moreover, when either V or W is a singleton we can just refer to it by its only element. Finally, we can call a (W, V) cut a reverse (V, W) cut.

Applying this terminology to the Heisenberg cut we find the following. Recall that a Heisenberg cut occurs whenever we jump from a quantum model to a non-quantum model (i.e., a classical model). Collecting all of our quantum theories into a set Q and all of our classical theories into a set C, the Heisenberg cut is a (Q, C) cut, that is a quantum-classical cut. Since presumably $C = T - Q$ we could alternately call the Heisenberg cut a Q cut, namely a quantum cut.

Under what conditions are generalized cuts (be they diagonal or vertical) necessary? As in the Heisenberg case, when choosing how to model our measurement chain we may find we are restricted in numerous ways. This could be for reasons of accuracy, conceptual coherence, technological feasibility, or many other reasons. We may find that, once these restrictions have been analyzed, that modeling some part of our measurement chain in a certain way is both conceptually mandatory and technically infeasible. In this case we simply cannot (yet) explain this experiment satisfactorily. Let us hope this doesn't happen.

More often, it may occur that one part of the experiment must be modeled some way and another later part must be modeled in some other way. In such cases, somewhere in between we need to make some kind of cut. Concretely, suppose that we are forced by such considerations to have $H(s) \in V \subset T$ and $H(t) \in W \subset T$ with

$V \cap W = \emptyset$ and $s \leq t$. It's easy to see that at some point between s and t we need to make a (V, W)-cut (be it vertical or diagonal). This claim is formalized and proved in Appendix A for non-sequential measurement chains.

Often then, we are forced to take some sort of cut, but which kind should we choose? What are the relative benefits of diagonal and vertical cuts in general? As I see it, vertical cuts have two principal advantages over diagonal cuts, and are therefore preferable ceteris paribus. Firstly, we often have a clearer prescription about how and when we are allowed to take vertical cuts than diagonal ones. It is generally expected in physics that our later theories reduce to our earlier theories under some approximation scheme in some relevant regimes. There is a wide literature on theoretic reduction [43, 44]. Such considerations give a prescription as to how and when we are allowed to take vertical cuts.

By contrast, however, it is generally not expected for our physical theories to tell us how they ought to dynamically couple to each other. As such some extra-theoretic work will often be needed to bridge the divide in a principled way. Significant care will need to be taken not to wildly violate the principles of either theory as they dynamically talk to each other.

Secondly, and relatedly, a diagonal cut involves, in a sense, mixed ontologies and mixed fundamental principles simultaneously present and interacting. By contrast, while a vertical cut makes reference to the ontologies/principles of both theories, it does so one at a time. Since our purpose here is only in pragmatics and modeling, this issue of mixed ontologies is not damning. However, we still demand that our models are conceptually coherent which is far from automatic for diagonal cuts.

As discussed in Sec. IV, in the Heisenberg case, measurement chains and cuts give us a case-by-case measurement framework for non-relativistic quantum theory thereby solving its core pragmatic measurement problem. Indeed, we there saw that making some sort of cut (i.e., diagonal or vertical) was necessary and moreover ought to be taken intentionally and explicitly.

The same is true for the core pragmatic measurement problem for any theory T_0. Let M be the set of all theories within which we know how to model measurements well and suppose $T_0 \notin M$. In any experiment involving T_0, at some point along the measurement chain we will need to make a cut into M. (This claim is formalized and proved in Appendix A.) Analyzing this experiment in terms of measurement chains and various potential generalized cuts gives us a road map to guide us in modeling its specific measurement process. In particular, as before, these road maps have the dangerous areas ahead clearly marked out, as well as multiple possible routes to navigate them. Hence for any theory we have a case-by-case measurement framework, thereby solving its core pragmatic measurement problem.

However, as in the Heisenberg case, the measurement framework alone does not solve the extended pragmatic measurement problem. For this, we would need to iden-

tify some near-universally available way of making the needed generalized cuts. Only then could we achieve a unified wide-scoping measurement theory. While a priori we have no reason to expect this is possible, such a measurement theory would give us a physically meaningful way to identify the theory's observables.

We are now ready to apply these generalized notions to the pragmatic QFT measurement problem.

A. The QFT-cut and its cousins

Given that we cannot feasibly model the late stages of any experiment using QFT, if we need to model some part of an experiment using QFT then at some later point we need to take some kind of cut away from QFT. Specifically, at some point in our modeling of this experiment we will need to switch from modeling our experiment using QFT, to not using QFT. According to the notation established above, this is a QFT-cut.

This idea has already been introduced in [15] although it was there called a *relativistic cut*. However, this name is apt to cause confusion because it's not relativity per se which we must cut away from. The following discussion will hopefully clarify this.

A few other related cuts deserve mentioning and naming at this point. If one feels that the particularly troublesome part of QFT is the fact that it describes things as a field, one might be interested in a field-cut. This is where we switch from modeling our measurement chain as a field (e.g., a relativistic quantum field, a non-relativistic quantum field, or a classical field) to not as a field (e.g., a qubit, a collection or classical point particles, or a nuclear spin degree of freedom).

Alternatively, one might feel that the troublesome part of QFT is the fact that things are relativistic, i.e., that our models are set in a locally Lorentzian spacetime. In this case one might be interested in a relativistic cut. If that name is already taken, we might instead call this a Lorentz-cut where we switch from modeling our measurement chain in a locally Lorentzian spacetime to something else, e.g., a locally Galilean spacetime. For instance, we might move from relativistic QFT to non-relativistic QFT.

Finally, one might feel that the troublesome part of QFT is the fact that its algebraic structure is that of a Type III rather than a Type I von Neumann algebra [39, 40]. In this case one might be interested in a Type III algebra-cut where we switch from modeling our measurement chain with a Type III algebraic structure to anything else, e.g. a Type I algebraic structure. For instance, we might move from relativistic QFT to non-relativistic quantum theory.

One might be interested in all of the above, or many other subtle variations thereof. However, if we are beginning from a QFT then all of the above are examples of a QFT-cut (or in the terminology of [15], a relativistic cut). As such, for the rest of this paper I will focus
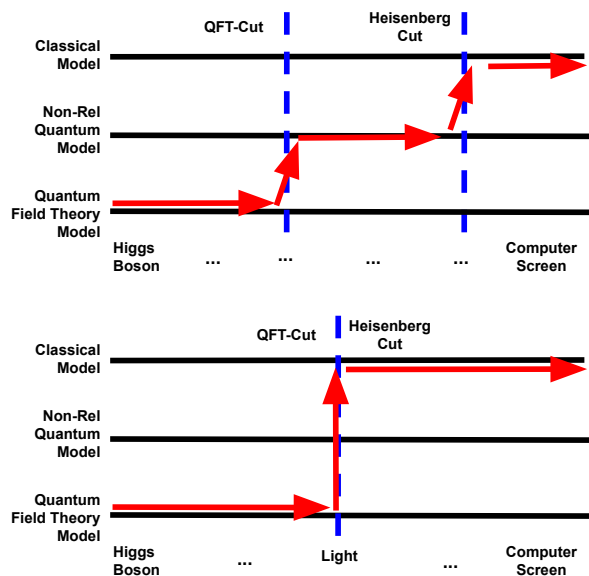
Figure 7. Two examples of QFT-cuts showing their possible relations to Heisenberg cuts. In the top half of the figure we first take a QFT-cut and then later take a Heisenberg cut. In the bottom half of the figure we take both cuts simultaneously.

exclusively on QFT-cuts, where we switch from modeling using a QFT to using anything else. Of particular interest, however, is switching from QFT to anything we know better how to model measurements of (i.e., classical physics, special relativity, general relativity, or even non-relativistic quantum theory).

As an example of a QFT-cut, consider the top half of Fig. 7. At some point in modeling the detection of a Higgs boson, we may switch from modeling things using QFT to modeling using non-relativistic quantum theory. This would be a QFT-cut (note this could be either a vertical or a diagonal cut). At some point later we may switch to modeling things classically. This would be a Heisenberg cut (again this could be either a vertical or diagonal cut).

It is possible to take both cuts simultaneously, consider the bottom half of Fig. 7. Suppose in the middle of modeling our Higgs boson experiment that we are modeling some state of light using QFT. We may be able to make some approximation on this state to then describe the light as a classical electromagnetic wave propagating in Minkowski space. This would be a QFT-cut since we are no longer using QFT to model the light. This would also be a pragmatic Heisenberg cut since we are no longer using anything quantum to model the light. This is not, however, a Lorentz cut. Note that because this theory switching happens in isolation (not during an interaction), both of these cuts are vertical.

It is an interesting question under what circumstances we are allowed to take certain kinds of cuts before or after other kinds of cuts. Could a Heisenberg cut ever be done before a QFT-cut? As these terms are defined in this paper, this is impossible: a Heisenberg cut would take us from a QFT into any classical theory, i.e., not a QFT. Hence, beginning from a QFT, all Heisenberg cuts are QFT-cuts. However, a Heisenberg cut may come either before or after a field cut for instance. It could also come before or after a Lorentz cut.

Let's review our situation: As I discussed in Sec. II, explaining all but maybe a small corner of quantum theories experimental successes unavoidably involves modeling quantum fields at some point. Moreover, as discussed in this section, whenever we are explaining an experiment which involves modeling quantum fields we need to at some point take a QFT-cut. Cuts, in general, are dangerous and tricky things both mathematically and conceptually. Thus, an explicit and carefully formalized handling of the QFT-cut is necessary for any satisfactory explanation of almost any quantum experiment. Given this, it seems extremely important that we have robust, flexible, widely applicable ways of making QFT-cuts.

To get a better handle on what types of QFT-cuts are available to us, and what their scopes are, both collectively and individually, I will next review the state of the art in the physics literature as it applies to QFT-cuts. As I will discuss, the current front-runner for getting us a wide-scoping measurement theory [7] for QFT is the Unruh-DeWitt detector model [15, 23–35].

VI. THE STATE OF THE ART

Hopefully, the above discussion raises a great many questions for you: Is the need for an explicit and formalized QFT-cut recognized in current scientific practice? Do physicists have good tools for making QFT-cuts? What are the current possibilities and limitations for various kinds of QFT-cuts? Diagonal or vertical cuts? With or without back-reaction? Crossing over into non-relativistic quantum theory or into classical physics? Are these tools collectively good enough to broadly cover all of quantum theory's QFT-model-needing experimental successes? Moreover, is any one of these tools of sufficient generality to allow us to induce a wide-scoping measurement theory for QFT from it? In order to answer these questions I will need to review the current state of the art in the physics literature.

Firstly, is this need recognized? Unfortunately not. For instance, in particle physics it is common to think in terms of scattering scenarios. Two input-states are assumed to come in from infinity and then interact in some unitary way. One can then compute various cross sections and the amount of amplitude which scatters out into some solid angle. It is common to end the story here and say that there is amplitude-squared probability that a detector covering that solid angle will detect something. This is not enough, it amounts to simply guessing that your particle detectors enact a certain intuitive projective operation. This guess may be effectively correct and give the right statistics for most pragmatic purposes, but it is
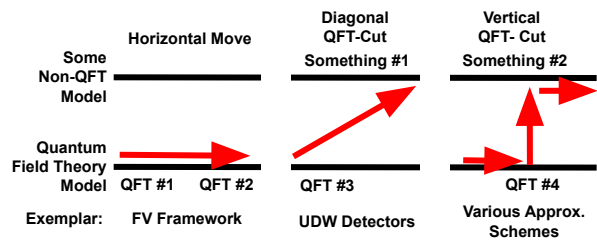
Figure 8. This figure shows the three ways that one can approach and cross a QFT-cut: horizontal moves, diagonal cuts and vertical cuts. Exemplars of these three types of moves are the Fewster Verch (FV) framework [10–12], the Unruh-DeWitt (UDW) detector model [7, 15, 23–35], and various approximation schemes [43, 45] respectively.

not a good explanation of the experiment. More must be said about the dynamics of the detection process. If we model this detection process in terms of a measurement chain, then at some point along this chain we will need to make a QFT-cut. Moreover, this QFT-cut ought to be taken intentionally and explicitly.

What about the experiments conducted at the LHC? Surely, the engineers who design and build particle detectors at the LHC are doing things right, right? I am not qualified to judge their methods, but I take it on faith that the engineers have detailed dynamical models of the detection process. Their accounts will meet the explanatory standards set in this paper if: 1) they do not implicitly assume that projective measurements are available in QFT, 2) they make an intentional and explicit QFT-cut somewhere in their modeling, 3) they do not make sophomore-like intuitive guesses about which projective measurements happen in a non-relativistic context.

As far as I am aware very few experimental groups make use of explicit formalized QFT-cuts in their modeling and only a few theoretical research groups [7, 9–13, 15–17] are engaged with these topics. What follows is an non-exhaustive literature review of some of the theoretical physics papers which touch on this topic.

Supposing that a theoretical or experimental physicist was interested in using an explicitly formalized QFT-cut in their modeling, what tools are currently available to them? What follows is a catalog of the various well developed ways of approaching and crossing a QFT-cut in the physics literature. This catalog will be organized into three sections: horizontal moves, diagonal QFT-cuts, and vertical QFT-cuts. See Fig. 8. Following this I will briefly discuss the scope of experiments that these tools cover collectively. Additionally, I will briefly discuss whether any individual tool has the wide-scoping applicability needed to induce a measurement theory for QFT.

A. Horizontal Moves

Before discussing how one might make a QFT-cut, allow me to first talk about how to approach one horizontally. The general shape of a horizontal move is shown on the left side of Fig. 8. Essentially, one QFT (QFT#1) couples to another QFT (QFT#2). Clearly, this is not a QFT-cut and so no collection of moves of this kind is the whole story.

However, of course, such moves may still be helpful in advancing us along the measurement chain until we are in a better position to make a QFT-cut. There is no issue with using a horizontal move as *part of* our measurement chain. The issue comes when one models an experiment involving QFT using only moves of this kind while neglecting to mention where exactly they take a QFT-cut. Or worse such an account might implicitly dismiss the need for a pragmatic QFT-cut altogether.

Before discussing what possibilities there are for describing interactions between QFTs, let's first talk about what sorts of isolated systems we know how to describe well with QFT. Neglecting momentarily their interaction, what types of systems QFT#1 and QFT#2 might go into the open slots in Fig. 8? (Or QFT#3 and QFT#4 for that matter?) Firstly, it should be said that we know well how to describe a wide variety of free systems using Lagrangian QFT: free scalar fields with or without mass, free electrons, free neutrinos, free photons, free Higgs particles, free gravitons (in the linearized gravity regime), etc. We can even model systems like a free proton or neutron as a free massive spinor field (assuming, of course, we ignore their quarky internal structure). Bose-Einstein condensates and some other condensed matter systems can also be treated within QFT. We can also include any small perturbative interaction term between any of these and calculate their joint evolution within perturbation theory.

What is much more difficult to do within QFT is to describe strongly interacting systems, including bound states such as atoms, or bound quark systems, or atomic nuclei. In principle, one ought to be able to consider the electromagnetic field interacting with the electron field and the proton field (pretend such a thing exists). We then ought to be able to find bound state solutions to this strongly interacting QFT which correspond to the various energy states of a first-quantized Hydrogen atom. However, these bound states of QED are remarkably difficult to either simulate or treat analytically. This is difficult for such systems in isolation, let alone interacting with an external field.

To summarize: even when we are just moving along the bottom line of Fig. 8 we have a rather limited mobility here currently. We have feasibility restrictions in terms of both what systems we can consider (i.e., QFT#1 and QFT#2) as well as how they might interact with each other.

Suppose that within computational feasibility, we have two QFTs in mind and an interaction between them.

What mathematical formalisms do we have for modeling this interaction? As is typical in physics, one can begin from either a Lagrangian or from a Hamiltonian formulation. However, in the case of QFT some opt to put the theory on even more secure mathematical footing by formulating it in algebraic terms, namely in Algebraic QFT [39]. For a recent philosophical debate about the differences in these approaches see [46–49]. A significant trade-off between these approaches are their differences in mathematical rigor and in practical utility.

As discussed in Sec. III, there are many technical issues which arise when one tries to apply our projective measurement theory to QFT, see Fig. 2. As fraught as this area is with mathematical stumbling blocks, some have looked to Algebraic QFT in hopes of a more secure way to approach modeling (at least parts of) measurement processes of quantum fields. In particular, the Fewster Verch (FV) framework [10–12] does this. Allow me to provide a brief overview.

Fewster Verch framework

Suppose that we can break at least a part of the measurement process down into a series of local interactions between QFTs. In particular, suppose that each of these interactions is localized in space and time, i.e. with one QFT acting as a local probe on another. The Fewster Verch (FV) framework [10–12, 16, 50] provides a model for such interactions within the mathematical rigor of Algebraic QFT. By doing so, one can be assured to be completely respectful of the central 'commandments' of relativity for at least part of the measurement chain. In particular, by describing this part of the measurement process entirely within Algebraic QFT, no causality violations of the kind shown in Fig. 2 are possible; Algebraic QFT has the fundamental principles of relativity built right into it.

To have something concrete in mind, let us consider a simple example (taken from [10, 16]) which just so happens to have an equivalent representation in Lagrangian QFT. Consider a scenario where one massive Klein Gordon field ("the probe field"), $\hat{\psi}(t,x)$, acts as a local probe another ("the system field"), $\hat{\phi}(t,x)$. The joint Lagrangian for our simple example is,

$$\mathcal{L} = \underbrace{\frac{1}{2}(\nabla_\mu \hat{\phi}(t,x))(\nabla^\mu \hat{\phi}(t,x)) - \frac{m_1^2}{2}\hat{\phi}^2(t,x)}_{\mathcal{L}_\phi} \quad (3)$$
$$+ \underbrace{\frac{1}{2}(\nabla_\mu \hat{\psi}(t,x))(\nabla^\mu \hat{\psi}(t,x)) - \frac{m_2^2}{2}\hat{\psi}^2(t,x)}_{\mathcal{L}_\psi}$$
$$- \underbrace{\lambda\,\rho(t,x)\,\hat{\psi}(t,x)\,\hat{\phi}(t,x)}_{\mathcal{L}_I}.$$

The first and second terms are the free Lagrangians of the system field, $\hat{\phi}(t,x)$, and the probe field, $\hat{\psi}(t,x)$, re-

spectively. The third term couples these two fields together. In the third term, λ determines the strength of the interaction and $\rho(t,x)$ is a spacetime function which determines the interaction profile. That is, $\rho(t,x)$ determines where in space and time the two fields interact. For the purposes of modeling localized interactions, we can take $\rho(t,x)$ to be compactly supported in some spacetime region K, (i.e., $\rho(t,x) = 0$ outside of K). See Fig. 9. Here N is some "processing region" in the future of K where the probe field undergoes further measurement processes.

It is important to note that the scope of interactions considered by the FV framework is much more general than this simple example. I have only specified the above interaction Lagrangian to have something concrete in mind for later comparison. In general in the FV framework, one quantum field acts as a local probe upon another quantum field. The nature of these two fields is left completely open so long as they can both be formulated within Algebraic QFT. The nature of their interaction is left open, except that it all happens within a bounded spacetime region, K, see Fig. 9. The spacetime background for such an interaction is even left open, i.e., it could be curved.

By using these sorts of local QFT-to-QFT interactions, one might be able to describe a part of a QFT-involved measurement chain. Of course, as discussed above, we cannot hope to provide a complete modeling of any real-life measurement process exclusively in terms of QFT-to-QFT interactions. This is for two reasons. Firstly, as discussed above, we currently have a rather limited technological and computational capacity for describing interacting bound states within Lagrangian QFT (let alone Algebraic QFT). Thus, the FV framework suffers here on grounds of feasibility, at least currently.

Secondly, once one QFT acts as a probe on another, we are still left with the problem of how to model the measurement of the second QFT. Indeed, the FV framework does not claim to solve the quantum measurement problem (pragmatic or realist) but rather their interest is "describing a link in the measurement chain, in a covariant spacetime context" [10]. In particular, they "take it for granted that the experimenter has some means of preparing, controlling and measuring the probe and sufficiently separating it from the QFT of interest" [10] or put more simply that "someone, somewhere, knows how to measure something" [51].

In total therefore the FV framework is potentially useful (within its presently limited computational feasibility) for modeling parts of the QFT-involved measurement chains for real-life experiments. In combination with the yet-to-be-discussed diagonal and vertical cuts, it may be helpful in solving the core pragmatic measurement problem for QFT.

But what about the extended pragmatic measurement problem? Our goal there is to give a unified wide-scoping account of measurements in QFT, i.e., to identify its observables. In this extended problem we care less about

computational feasibility. One might therefore expect the FV framework (and horizontal moves generally) to be more useful in the extended problem.

For instance, recent work claims to have used the FV framework to provide an "Asymptotic measurement schemes for every observable of a quantum field theory" [50] in order to "determine the set of system observables that can be measured by FV measurement schemes". Concretely, their objective is "to analyze how information about one physical structure (system) is transferred to another physical structure (probe) that is controlled by an external experimenter" [52]. In particular their interest is in the case where both the system and field are QFTs and the information is transferred via a local interaction. This is exactly the sort of thing that the FV framework is good at: working out how information moves between quantum fields which interact with each other in localized regions. This attempts to go beyond Alice's full menu discussed in Sec. III by giving a constructive/dynamical story about the measurement processes associated with each supposed observable.

The principal limitation in [50] however is that (as is always the case with the FV framework) it explicitly assumes that the experimenter has full control over the probe field. In particular, it is assumed that they know how to extract classical information from the probe, i.e. "someone, somewhere, knows how to measure something" [51]. While this is potentially a step in the right direction, these results ultimately end up assuming that we know what the observables in the probe field are. Contrary to this methodology, I argue that the only physically meaningful way to identify the observables within QFT is to connect them with observables outside of QFT by some measurement chain which includes a QFT-cut.

Allow me to briefly give the technical details of [50] by first introducing some terminology. Within Algebraic QFT, each bounded region of spacetime R is associated with an algebra, commonly called the "algebra of obesrevables". However, this remains to be justified as what is an observable is exactly what is at question here. This algebra includes the field operator $\hat{\phi}(t,x)$ integrated against all smooth functions compactly supported over R. Additionally, the algebra includes products and sums of these smeared field operators. An FV measurement scheme for a field $\hat{\phi}(t,x)$ ("the system field") specifies four things: a probe field $\hat{\psi}(t,x)$ labeled $\mathcal{P}$, an initial state for the probe field, $\rho_{\mathcal{P}}$, a unitary interaction, S, between $\hat{\phi}(t,x)$ and $\hat{\psi}(t,x)$ localized in some region K, (e.g., Eq. (3)) and finally an element of the probe algebra, B, associated with a processing region N in the future of K, see Fig. 9.

While the results of [50] are proven in terms of Algebraic QFT, it here suffices to give their translation into the usual language of Hilbert spaces. A FV measurement scheme gives a way of indirectly addressing elements of the system algebra associated with the region K. Namely
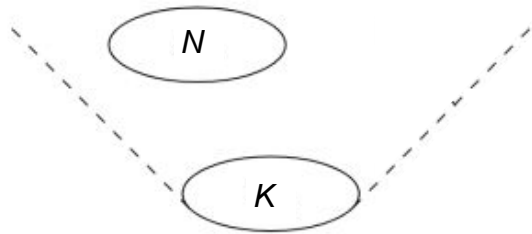

Figure 9. The bounded spacetime regions considered in [16]. The two quantum fields interact only within the coupling region K. In the future of this interaction is the "processing region" N where the probe field undergoes further measurement processes.

for every FV measurement scheme we have

$$\mathrm{Tr}_{\mathcal{SP}}(\rho_S \otimes \rho_P \, S^\dagger \mathbb{1} \otimes B \, S) = \mathrm{Tr}_S(\rho_S \, B_{\mathrm{ind}}) \qquad (4)$$

for some induced B_{ind} in the system algebra associated with K.

Ultimately, the question addressed by [50] is: Which elements of the system field's algebra at K can be indirectly measured via an FV measurement scheme (assuming we can measure any element of the probe algebra at N)? Their answer is roughly that for every element A of the system algebra at K there exists a sequence of FV measurement schemes which indirectly measure it arbitrarily well in the limit. Namely, for every A in the system algebra associated with K, we have [52]

$$\mathrm{Tr}_{\mathcal{SP}_\alpha}(\rho_S \otimes \rho_{\mathcal{P}_\alpha} \, S_\alpha^\dagger \mathbb{1} \otimes B_\alpha \, S_\alpha) \to \mathrm{Tr}_S(\rho_S \, A) \qquad (5)$$

for some sequence of FV measurement schemes indexed by an integer α. Thus, given full control over the probe system in the region N there is some sequence of probes, probe states and local interactions such that we can address any A in the system's algebra arbitrarily well.

In summary, the FV framework is a great tool for working out how information moves between quantum fields which interact with each other in localized regions. However, to solve either the core or extended pragmatic measurement problems for QFT it alone is not enough. We need to at some point take a step outside of QFT via a QFT-cut.

B. Vertical QFT-cuts

Let's next consider what vertical QFT-cuts are available to us currently. The general shape of a vertical QFT-cut is shown on the right side of Fig. 8. Essentially, we begin with something being modeled as a QFT (QFT#4) and then take some sort of approximation on this to arrive the very same thing now modeled as something other than a QFT (Something#2). Our freedoms in designing a vertical QFT-cut are: which theory we approximate into, what kind of QFT we begin from and

relatedly what kind of non-QFT system we land on, as well as the details of our approximation scheme.

We have discussed already in the previous subsection the sorts of systems which we have a grip on how to model in QFT: free systems plus perturbative interactions and some condensed matter systems but not strongly interacting systems or bound states. Our options for QFT#4 are fixed to be among these. As for which theory to cross into, the "nearest" theory to QFT would likely be non-relativistic QFT (i.e., QFT with $c \to \infty$, or rather QFT in a Galilean spacetime). For such a approximation to work our initial QFT#4 must be massive, i.e., not light nor gravity. Massive fields may limit onto particles in a non-relativistic limit, but massless ones will not [38, 43]. For massive free states, taking such a limit gives us non-relativistic quantum particles. This alone is not enough unless we understand well how to model the measurement of non-relativistic QFTs. I am not aware of any research in this direction, but it could be a fruitful way forward.

The next nearest theory we could cut into is non-relativistic quantum theory. The question then is what we should take as Something#2? An obvious experimentally relevant system would be a first-quantized Hydrogen atom. However, this would mean that QFT#4 needs to be some second quantized description of the Hydrogen atom. As I have already discussed, describing such bound states in QFT is difficult. More research in this direction is warranted.

Another intriguing option for what non-relativistic quantum system to put for Something#2 is an Unruh-DeWitt (UDW) detector [7, 15, 23–35]. These will be described in more detail in the next subsection, but roughly they are atom-like non-relativistic quantum systems which can be coupled to a quantum field in a way motivated by the light-matter interaction. In fact, recent work [45] has developed a "second quantized UDW detector", i.e., a QFT which reduces to a UDW detector in the non-relativistic limit. This an interesting avenue for future research, worthy of further development.

Finally, we can consider the possibility of approximating a QFT as a classical state of some sort. An example of this kind has already been discussed in Sec. V. One might have a state of light described as a QFT and then switch to describing this as a classical electromagnetic field in Minkowski space. Vertical cuts of this kind seem to be experimentally relevant and deserve to be developed further.

One might also be tempted to try to use decoherence theory and the Born rule to make a vertical cut into classical physics just as we did previously for non-relativistic quantum theory. If this is possible, it is not straightforwardly so. Indeed, if it were then we would be licensed to make generic projective measurements in QFT, which we are not, see Sec. III. At least one roadblock is the following: The Type III algebraic structure of QFT means that we cannot associate Hilbert spaces with bounded regions of spacetime. In particular this means we cannot associate a density matrix to bounded spacetime regions. As such, we cannot evaluate when via decoherence they become approximately diagonal. As such, we cannot identify their diagonal entries to apply the Born rule to.

In summary, vertical QFT-cuts are a promising possibility deserving of further research. If technical limitations surrounding second quantized atoms can be overcome, then these could have a substantial scope including many experimentally relevant systems.

C. Diagonal QFT-cuts

Finally, let's next consider what diagonal QFT-cuts are available to us currently. The general shape of a diagonal QFT-cut is shown in the center of Fig. 8. Essentially one QFT (QFT#3) couples directly to something which is not being modeled as a QFT (Something#1). As is the case for any diagonal cut, the direct coupling between system's described in fundamentally different theories poses conceptual challenges. Here, there is an inherent risk that coupling to a non-QFT system will end up breaking one of the central 'commandments' of relativity.

However, on the bright side, since quantum field theory itself provides us with no prescription for how such systems ought to interact, we have a great deal of freedom in how one might model such an interaction. In particular, our freedoms in designing a diagonal QFT-cut are: what kind of QFT we begin from, which theory we cut into, what kind of non-QFT system we cut into, and the nature of the interaction between the systems.

We have discussed in the previous subsection the sorts of systems which we have a grip on how to model in QFT: free systems plus perturbative interactions and some condensed matter systems but not strongly interacting systems or bound states. Our options for QFT#3 are fixed to be among these. As for which theory take a QFT-cut into, the "nearest" theory to QFT would likely be non-relativistic QFT (i.e., QFT with $c \to \infty$, or rather QFT in a Galilean spacetime). As far as I am aware, not much work has been put into the study of sensible dynamical couplings between QFT and non-relativistic QFT. One immediate concern is the possibility of uncontrolled faster-than-light signaling. We might have a relativistic field $\hat{\phi}_{\text{Rel}}(x)$ couples to $\hat{\phi}_{\text{Non-Rel}}(x)$ which can then send an instantaneous signal to $\hat{\phi}_{\text{Non-Rel}}(x+a)$ which is coupled to $\hat{\phi}_{\text{Rel}}(x + a)$. This seems problematic. More work may be needed in this direction.

The next "nearest" theory we could try to take a QFT-cut into is non-relativistic quantum theory. Significant work has been done in this direction, which will be discussed momentarily.

The last option which comes to mind is to couple our QFT to a classical system (e.g., something modeled in special relativity or general relativity). This doesn't seem terribly problematic, we could for instance have a classical Klein Gordon field interacting with the expectation

value of a quantum Klein Gordon field, perhaps without back reaction. More work would need to be done motivating why such an interaction is an accurate reflection of some part of a real-life experiment, but there don't seem to be insurmountable technical difficulties here.

Unruh-DeWitt Detector Models

Let's return to the possibility of diagonal cuts into non-relativistic quantum theory. We have a lot of freedom here in designing this interaction. Of course, there are also certain things we want from this diagonal theory-to-theory coupling if it's going to be a productive part of an experimental prediction. So what ought to guide us in designing this interaction? As a first guide, we may rely on a desire to preserve the central 'commandments' relativity (covariance, causality, and locality) as much as possible. Moreover, as a second guide we may rely on a desire to accurately model parts of real-life experiments.

For instance, we might take the Something#1 system in Fig. 8 to be something atom-like and we might take QFT#3 to be the electromagnetic field (or some scalar analog thereof). In this case, under such guidance, one is quickly led [53–55] to something very much like the Unruh-DeWitt (UDW) detector model first introduced in [24].

Alternatively, we might take QFT#3 to be a graviton field (in the linear gravity regime). In this case, one is quickly led to a certain variant of the Unruh-DeWitt detector model [56–58].

One can also take Something#1 to be a fermionic quantum system which interacts with QFT#3 being a neutrino field. In this case one is led to another variant of the Unruh-DeWitt detector model [59, 60].

The possibilities for which real-life interaction we might attempt to mimic here are very general. In each case, the resulting interaction model is within the family of Unruh-DeWitt-like models. Moreover, much of the above can be done in arbitrarily curved spacetime backgrounds as well [56].

Enough discussion of abstract possibilities, concretely what do these models look like? To have something concrete in mind, let us consider a simple example (taken from [15]). Consider a simple example in which a UDW detector $\hat{\mu}$ coupled to a massive Klein Gordon fields $\hat{\phi}(t,x)$ with joint Lagrangian,

$$\mathcal{L} = \underbrace{\frac{1}{2}(\nabla_\mu \hat{\phi}(t,x))(\nabla^\mu \hat{\phi}(t,x)) - \frac{m_1^2}{2}\hat{\phi}^2(t,x)}_{\mathcal{L}_\phi} \quad (6)$$
$$+ \mathcal{L}_{\text{UDW}} - \underbrace{\lambda\,\rho(t,x)\,\hat{\mu}(\tau)\,\hat{\phi}(t,x)}_{\mathcal{L}_I}.$$

The first term is the free Lagrangian of the field. The second term $\mathcal{L}_{\text{UDW}}$ is the free Lagrangian of the non-relativistic probe system, the UDW detector. We have total freedom to pick the internal dynamics of the non-relativistic probe system. For instance, it could be a qubit, or a quantum harmonic oscillator, or a first quantized Hydrogen atom.

In the third term, λ determines the strength with which the non-relativistic probe and field couple to each other. In the third term, $\rho(t,x)$ is a spacetime function determining the interaction profile. In this context, $\rho(t,x)$ is often called the probe's smearing function, and is taken to describe the size and shape of the probe through time (more will be said about this later). Just as in the FV framework, for the purposes of modeling local measurements, we can take $\rho(t,x)$ to be compactly supported in some spacetime region, K, (i.e., $\rho(t,x) = 0$ outside of K). In the third term, $\hat{\mu}$ is the degree of freedom of the non-relativistic probe which couples to the quantum field. For instance, if the probe is a harmonic oscillator, $\hat{\mu}(\tau)$ might be its number operator $\hat{n}$ or one of its quadrature operators $\hat{q}$ or $\hat{p}$.

The major difference between the FV and UDW approaches (Eq. (3) and Eq. (6)) is just that in the first case the probe system is a quantum field, $\hat{\psi}(t,x)$, and in the second case it is a non-relativistic quantum system, $\hat{\mu}$. For a more in depth comparison of the UDW detector model and the FV framework, see [15].

It's important to note that the UDW detector is designed as a "pointer" measurement device, i.e., one which translates a "needle" proportional to some targeted operator in the probed field. Rather, the UDW detector is designed to be atom-like. If one measures the UDW detector after its interaction with the field, one ought to interpret this roughly as one would if an atom had coupled to the field. For instance, if the UDW probe is initially in its ground state and then is later measured to be in an excited state, one might infer that it absorbed a photon from the quantum field.

More complexly, one can use UDW detectors to model an entanglement harvesting experiment [16, 29–35]. Roughly, in such an experiment two initially uncorrelated probe systems interact locally with a quantum field in such a way that they do not have time to signal to each other. Despite this, these two probes become entangled because there was already entanglement present between the two space-like separated regions they interacted with. The benefit of such an experiment is that the entanglement in the field has been transferred into more accessible systems, both physically and mathematically. We cannot associate a Hilbert space to bounded regions in QFT and as such cannot straightforwardly compute the entanglement between these regions. The final entanglement of these probes is a witness to the initial entanglement in the field.

A great many theoretical investigations of this sort have been carried out using UDW detectors [29–35]. Such studies can even be done in curved spacetimes: one can study the entanglement structure around a black hole near the event horizon for instance.

Thus, in addition to providing a good model for

many experimentally relevant systems, the UDW detector model covers a wide range of interesting hypothetical experiments, all while remaining computationally feasible.

In Sec.V I suggested that we ought to prefer vertical cuts over diagonal ones all else being equal. Have we here overcome the natural disadvantages of diagonal cuts? The first disadvantage was a lack of prescription from either theory about how they ought to couple to each other. However, as discussed above, we can use this freedom to design our theory-to-theory coupling as we wish. In particular, above we used this freedom to build an interaction which closely resembles certain experimentally relevant couplings. The second disadvantage was a clash of principles between the two coupled theories. How do we fare here? That is, how well do UDW-like detectors preserve the central 'commandments' of relativity: covariance, causality, and locality? Do they for instance lead to uncontrolled faster-than-light signaling in the QFT?

Before answering this, a distinction needs to be made between two modes of applications of the UDW model. It was mentioned above that we will often wish to localize the probe's smearing function $\rho(t, x)$ within some bounded spacetime region. For instance, when the UDW detector model is derived from the light matter interaction [53–55], the smearing function $\rho(t, x)$ turns out to be determined by the overlaps of certain atomic orbitals. That is, $\rho(t, x)$ is a near-literal description of the shape of the atom in space and time.

When the size of the UDW is much smaller than all other relevant scales and the detailed shape of the detector doesn't matter much, we can also approximate the detector as being point-like. Consider a point-like detector traveling through spacetime on some time-like trajectory, $z(t)$. We can localize the interaction to this trajectory by taking $\rho(t, x) = \chi(t)\delta(x - z(t))$. Here $\chi(t)$ controls where along the trajectory the probe couples to the field (it may turn on and off) and the δ function localizes the interaction to the detector trajectory, $z(t)$. Let us call these the point-like detectors.

With this established, let us return to the question of how well do UDW-like detectors preserve the central 'commandments' of relativity. In brief, they do so imperfectly, but with well understood and controllable issues [17, 61]. For smeared detectors (i.e., non-point-like detectors) there are some slight faster-than-light signaling issues. Essentially, the issue is that if some information is taken up by the left half of the detector it can "immediately" jump to the right half of the detector and then back into the field. Basically, signals can jump across the detector instantly. This breaks no-signaling and causes some issues with the relativity of simultaneity (this coupling does not treat all relativistically compatible time-orderings equally).

However, these issues are ultimately minor [17, 61]. The time ordering issues do not appear at the lowest orders of perturbation theory. If the light-matter interaction is weak enough, then the time-ordering issues are strongly suppressed. Moreover, the size of the no-signaling violations is set by the size of our detector. Recall $\rho(t, x)$ might have compact support. If the UDW detector has a width of 3 nm then signals can only arrive at most 10 atto-seconds early. Ultimately if we care about such time-scales (of the order of the light crossing time for the atom) then we shouldn't even be allowed to talk about first-quantized atoms in the first place. Indeed, all of these commandment-breaking issues go away when we use point-like detectors.

D. The Scope of these Tools

Having reviewed the state of the physic literature, for feasible ways of crossing the QFT-cut, what ultimately are the scope of these tools? In my assessment while each tool has its limitations and more development of each of them is needed, collectively these tools have a substantial scope. Thus, I believe that collectively these tools give us a good handle on a case-by-case measurement framework for QFT. That is, collectively they can give us a solution to the core pragmatic measurement problem for QFT.

But what about the extended pragmatic measurement problem for QFT? Can these tools help us identify the observables of QFT? As discussed in Sec. V, in order to get a wide-scoping unified measurement theory for QFT we would need some near-universally available way of making QFT-cuts. In particular, we would need for at least one of these tools to have a sufficiently wide range of applicability such that nearly all QFT-involving experiments can be modeled using it. Which of the above discussed tools has the widest scope?

As I have discussed above, at least currently the UDW detector model by far has the widest scope of applicability of any of the tools currently available. Thus, if one wants to develop a wide-scoping measurement theory for QFT and to identify its observables, this is currently the most promising way forward. Indeed, a recent paper claims to have used the UDW detector model to establish a detector-based measurement theory for quantum field theory [7].

VII. CONCLUSION

This paper began by distinguishing between the pragmatic and realist portions of the quantum measurement problem specifically for quantum field theory (QFT). Of these, I have argued that the pragmatic worries are by far more important since if unanswered their consequences are more severe. If we lose the pragmatic connection between theory and experiment, we are at risk of losing any right to claim evidential support (or worse physicality) for quantum field theory theory. This is the *Pragmatic QFT Measurement Problem*. To alleviate these pragmatic worries we must at a minimum demand a

case-by-case explanation of the QFT's predictions for key experimental successes (solving the core pragmatic measurement problem). One might further demand a wide-scoping explanation of QFT's predictions for nearly all possible experiments (solving the extended pragmatic measurement problem, and thereby identifying QFT's observables).

I have argued that for non-relativistic quantum theory, all of these pragmatic worries have been satisfactorily handled in terms of measurement chains and Heisenberg cuts. However, the same is not true for QFT, with severe consequences. Indeed, as I have argued in Sec. II, explaining all but maybe a small corner of quantum theory's experimental successes unavoidably involves modeling quantum fields [5, 6]. Thus the Pragmatic QFT Measurement Problem threatens to spread across almost all of quantum theory. In particular, if we cannot establish an adequate linkage between QFT and experiment then not only quantum field theory but the whole of quantum theory is at risk of losing its evidential support.

With the stakes so raised, I turned my attention towards what exactly the difficulty is with properly modeling the measurement of quantum fields. Blame is often placed on certain technical mathematical difficulties which prevent us from straightforwardly transplanting our projective measurement theory into QFT [7–23]. Naive attempts to apply our non-relativistic measurement theory lead one to commit mathematical blunders which either violate the central 'commandments' of relativity (covariance, causality, and locality) [14, 18–23] or otherwise disrespect the QFT's local algebraic structure. For these reasons, it has been argued we need a new (or at least refined) measurement theory for quantum field theory, see for instance [7, 8, 11]. However, focusing on these technical mathematical issues belies a deeper methodological issue: the need to provide a concrete model of our measurement processes. Acknowledging this fact shuts down an otherwise promising looking approach to solving the Pragmatic QFT Measurement Problem.

In particular, one might think we can more-or-less keep our non-relativistic measurement theory if we just cut out the bad bits (those which violate causality) and patch up any causality violations stemming from naive state update maps. Considerations of this sort might lead one to prove theorems [8, 9] which cleanly separate those projective measurements which violate causality from those that don't. This may usefully identify what measurement operations are relativistically-safe in QFT, but it falls short of giving us a principled way which of these measurement operations, if any, my real-life lab equipment is doing. Indeed, the authors of [8] recognize that "it would be useful to construct an explicit dictionary between update maps and specific probe models". If we are to offer a better explanation than simply guessing the right projector, we need to provide a concrete model of our measurement process.

Indeed, any such attempt to directly transplant our non-relativistic projective measurement theory into QFT will be methodologically unsatisfying. Rather than trying to transplant our existing projective measurement theory into QFT, I have in this paper proposed an alternate approach. Firstly, one ought to understand how our non-relativistic projective measurement theory has its roots in discussions of measurement chains and pragmatic Heisenberg cuts. Then, one ought to generalize these notions and transplant them into QFT. From this position we will be able to evaluate what measurement theory, if any, we are led to for QFT.

First in Sec. IV, I discussed how the core pragmatic measurement problem has been overcome for non-relativistic quantum theory, using measurement chains and Heisenberg cuts. I have discussed how in satisfactorily modeling any quantum experiment we need to explicitly and formally take a pragmatic Heisenberg cut at some point along the measurement chain. As I discussed, often in modeling a given experiment, we will have freedom regarding what kind of Heisenberg cut to take and where to take it. In fact, I have given a full taxonomy of the various kinds of pragmatic Heisenberg cuts available to us. Exactly which of these kinds of pragmatic Heisenberg cuts are appropriate and when must be determined on a case-by-case basis. Thus, for each experiment under consideration, we have a road map to guide us in modeling its specific measurement processes. These road maps have the dangerous areas ahead clearly marked out, as well as multiple possible routes to navigate them. Given our task, this is an indispensable tool.

Using these road maps we can explain our experiments and gain evidential support from them on a case-by-case basis. However, case-by-case considerations may give us a measurement framework, but they do not give us a wide-scoping measurement theory. That is, they do not address the extended pragmatic measurement problem. Luckily, for non-relativistic quantum theory one way of crossing the Heisenberg divide is of near-universal applicability and so gives us a wide-scoping measurement theory. Namely, by using decoherence theory and then the Born rule once a sufficient amount of decoherence has happened. Ultimately, this is what underwrites our non-relativistic quantum measurement theory. It is just whatever is induced when our classical measurement theory is applied after any late-enough decoherence-based Heisenberg cut.

The central message of this paper is that we should not transplant the end of this story (i.e. projective measurements) into QFT. Rather we should start from the beginning. We ought to generalize these notions of chains and cuts and then transplant them into QFT. This will automatically give us a case-by-case measurement framework for QFT, solving its core pragmatic worries. We can then see which (if any) measurement theory we are led to for QFT.

As I have shown in Sec. V, these notions of chains and cuts can be readily generalized to help orient ourselves in responding to the Pragmatic QFT Measurement Prob-

lem. My analysis has revealed that (in close analogy to the Heisenberg cut) whenever we are modeling experiments involving quantum fields, that at some point along the measurement chain we also need to make a pragmatic *QFT-cut* (i.e., a switch from a QFT model to a non-QFT model). As I have argued, in Sec. II such situations are far more common than one might at first think.

As with Heisenberg cuts, we will often have freedom regarding what kind of QFT-cut to take and where to take it. As above, these generalized notions of measurement chains and cuts provide us with a road map to guide us in modeling individual measurement processes. This is an indispensable tool in recovering quantum theory's experimental spoils. However, these considerations alone give us, at best, a case-by-case measurement framework for QFT, not a unified measurement theory. In order to get a wide-scoping measurement theory for QFT, we need some widely applicable principled way of making QFT-cut, with scope analogous to decoherence theory. This is necessary to identify the observables of QFT.

What then is the current state of the physics literature on these topics? Is the need for a QFT-cut recognized? Do physicists have good tools for making QFT-cuts? Are their tools collectively good enough to resecure evidential support for quantum theory? Is any individual tool wide-scoping enough to underwrite a measurement theory for QFT? In brief, my considered answers are: No, the need for a QFT-cut is not currently widely recognized. Nonetheless, the physicists who do care about this topic have several tools for approaching (the Few-ster Verch (FV) framework [10–12, 16]) and crossing the QFT-cut (UDW detector model [15, 23–35], and a handful of approximation schemes [43, 45].

In my assessment, collectively these tools have a substantial scope. Thus, collectively these tools give us a good handle on a case-by-case measurement framework for QFT solving its core pragmatic measurement problem. But is any one of these tools on its own sufficient to give us wide-scoping unified measurement theory for QFT? As I have discussed, the UDW detector model by far has the widest scope of applicability of any of the tools currently available. Thus, if one wants to develop a wide-scoping measurement theory for QFT or to identify its observables, this is the way to go. Indeed, a recent paper claims to have used the UDW detector model to establish a detector-based measurement theory for quantum field theory [7].

VIII.   ACKNOWLEDGEMENTS

[1] E. Curiel, "Schematizing the observer and the epistemic content of theories," (2020), arXiv:1903.02182 [physics.hist-ph].

[2] M. Giovanelli, ""but one must not legalize the mentioned sin". phenomenological vs. dynamical treatment of rods and clocks in einstein's thought," (2014).

[3] H. R. Brown, *Physical Relativity* (Oxford University Press, Oxford, 2005).

[4] W. Myrvold, in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta (Metaphysics Research Lab, Stanford University, 2018) Fall 2018 ed.

[5] D. Wallace, "The sky is blue, and other reasons physics needs the everett interpretation," (2021), [Oxford Philosophy of Physics Seminar, Michaelmas Term 2021, 4th Nov, Timestamp 37:34].

[6] D. Wallace, "The sky is blue, and other reasons quantum mechanics is not underdetermined by evidence," (2022).

[7] J. Polo-Gómez, L. J. Garay, and E. Martín-Martínez, "A detector-based measurement theory for quantum field theory," (2021), arXiv:2108.02793 [quant-ph].

[8] I. Jubb, Phys. Rev. D **105**, 025003 (2022).

[9] L. Borsten, I. Jubb, and G. Kells, Phys. Rev. D **104**, 025012 (2021).

[10] C. J. Fewster and R. Verch, Commun. Math. Phys. **378**, 851–889 (2020).

[11] C. J. Fewster, "A generally covariant measurement scheme for quantum field theory in curved spacetimes," (2019), arXiv:1904.06944 [gr-qc].

[12] H. Bostelmann, C. J. Fewster, and M. H. Ruep, Phys. Rev. D **103**, 025017 (2021).

[13] C. Anastopoulos and N. Savvidou, Entropy **24** (2022), 10.3390/e24010004.

[14] R. D. Sorkin, "Impossible measurements on quantum fields," (1993), arXiv:gr-qc/9302018 [gr-qc].

[15] D. Grimmer, B. d. S. L. Torres, and E. Martín-Martínez, Phys. Rev. D **104**, 085014 (2021).

[16] M. H. Ruep, Classical and Quantum Gravity **38**, 195029 (2021).

[17] J. de Ramón, M. Papageorgiou, and E. Martín-Martínez, Phys. Rev. D **103**, 085002 (2021).

[18] M. Redhead, Foundations of Physics **25**, 123 (1995).

[19] F. Dowker, "Useless qubits in "relativistic quantum information"," (2011), arXiv:1111.2308 [quant-ph].

[20] D. M. T. Benincasa, L. Borsten, M. Buck, and F. Dowker, Class. Quantum Gravity **31**, 075007 (2014).

[21] L. Borsten, I. Jubb, and G. Kells, "Impossible measurements revisited," (2019), arXiv:1912.06141 [quant-ph].

[22] A. Ortega, E. McKay, A. M. Alhambra, and E. Martín-Martínez, Phys. Rev. Lett. **122**, 240604 (2019).

[23] A. Teixidó-Bonfill, A. Ortega, and E. Martín-Martínez, Phys. Rev. A **102**, 052219 (2020).

[24] W. G. Unruh, Phys. Rev. D **14**, 870 (1976).

[25] S.-Y. Lin and B. L. Hu, Phys. Rev. D **76**, 064008 (2007).

[26] E. G. Brown, E. Martín-Martínez, N. C. Menicucci, and R. B. Mann, Phys. Rev. D **87**, 084062 (2013).

[27] M. Hotta, A. Kempf, E. Martín-Martínez, T. Tomitsuka, and K. Yamaguchi, Phys. Rev. D **101**, 085017 (2020).

[28] E. Tjoa and E. Martín-Martínez, Phys. Rev. D **101**, 125020 (2020).

[29] A. Valentini, Phys. Lett. A **153**, 321 (1991).

[30] B. Reznik, Found. Phys. **33**, 167 (2003).

[31] A. Pozas-Kerstjens and E. Martín-Martínez, Phys. Rev. D **92**, 064042 (2015).

[32] G. V. Steeg and N. C. Menicucci, Phys. Rev. D **79**, 044027 (2009).

[33] E. Martín-Martínez, A. R. H. Smith, and D. R. Terno, Phys. Rev. D **93**, 044001 (2016).

[34] E. Martín-Martínez and N. C. Menicucci, Class. Quantum Gravity **29**, 224003 (2012).

[35] L. J. Henderson, R. A. Hennigar, R. B. Mann, A. R. H. Smith, and J. Zhang, Class. Quantum Gravity **35**, 21LT02 (2018).

[36] B. Abi, T. Albahri, S. Al-Kilani, and et. al (Muon $g-2$ Collaboration), Phys. Rev. Lett. **126**, 141801 (2021).

[37] H. Kragh, "Max planck: The reluctant revolutionary," (2019).

[38] W. E. Lamb, Applied Physics B **60**, 77 (1995).

[39] F. Kronz and T. Lupher, in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta (Metaphysics Research Lab, Stanford University, 2021) Winter 2021 ed.

[40] E. Witten, Rev. Mod. Phys. **90**, 045003 (2018).

[41] D. Wallace, *The emergent multiverse [electronic resource] : quantum theory according to the Everett interpretation*, Oxford scholarship online (Oxford University Press, Oxford, 2012).

[42] C. K. Zachos, D. B. Fairlie, and T. L. Curtright, *Quantum Mechanics in Phase Space* (WORLD SCIENTIFIC, 2005) https://www.worldscientific.com/doi/pdf/10.1142/5287.

[43] J. Rosaler, *Inter-theory relations in physics: case studies from quantum mechanics and quantum field theory*, Ph.D. thesis, University of Oxford (2013).

[44] R. Batterman, in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta (Metaphysics Research Lab, Stanford University, 2020) Fall 2020 ed.

[45] F. Giacomini and A. Kempf, "Second-quantized unruhdewitt detectors and their quantum reference frame transformations," (2022), arXiv:2201.03120 [quant-ph].

[46] D. Wallace, Synthese **151**, 33 (2006).

[47] D. Wallace, "Taking particle physics seriously: a critique of the algebraic approach to quantum field theory," (2011).

[48] D. Fraser, Philosophy of Science **76**, 536 (2009).

[49] D. Fraser, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics **42**, 126 (2011), philosophy of Quantum Field Theory.

[50] C. J. Fewster, I. Jubb, and M. H. Ruep, "Asymptotic measurement schemes for every observable of a quantum field theory," (2022).

[51] C. J. Fewster, "Local measurement of quantum fields in curved spacetimes," (2021), [Relativistic Quantum information-Online 2020/21 - Waterloo Session 03: Wednesday February 10th, timestamp 09:45].

[52] M. Ruep, "Observing observables – causal measurement schemes for every observable of the linear real scalar field in curved spacetime," (2022), [Contribution to the Quantum Field Theory in Curved Spacetimes Workshop (23-27 May 2022)].

[53] E. Martín-Martínez and P. Rodriguez-Lopez, Phys. Rev. D **97**, 105026 (2018).

[54] R. Lopp and E. Martín-Martínez, Phys. Rev. A **103**, 013703 (2021).

[55] R. Lopp and E. Martín-Martínez, Phys. Rev. A **103**, 013703 (2021).

[56] E. Martín-Martínez, T. R. Perche, and B. de S. L. Torres, Phys. Rev. D **101**, 045017 (2020).

[57] R. Faure, T. R. Perche, and B. d. S. L. Torres, Phys. Rev. D **101**, 125018 (2020).

[58] J. P. M. Pitelli and T. R. Perche, Phys. Rev. D **104**, 065016 (2021).

[59] B. d. S. L. Torres, T. R. Perche, A. G. S. Landulfo, and G. E. A. Matsas, Phys. Rev. D **102**, 093003 (2020).

[60] T. R. Perche and E. Martín-Martínez, Phys. Rev. D **104**, 105021 (2021).

[61] E. Martín-Martínez, T. R. Perche, and B. d. S. L. Torres, Phys. Rev. D **103**, 025007 (2021).

[62] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*, Cambridge Series on Information and the Natural Sciences (Cambridge University Press, 2000).

Appendix A: Generalize, Formalize and Prove: Measurement Chains and Cuts

In Sec. IV and Sec. V I have made much use of the notions of sequential measurement chains and cuts. In this appendix these notions will be generalized and formalized. Ultimately, this appendix has four goals: First, to substantially widen the scope of experiments analyzable in terms of measurement chains. Second, to formalize this situation. Third, to show how this formalization can handle at least one non-trivial QFT experiment. And finally, to prove the main claims of the paper within this formalization.

1. Non-sequential Measurement Chains

In Sec. IV and V all measurement chains considered were sequential. In particular, the systems involved in the experiment were ordered in a definite non-repeating sequence of pairwise interactions: System A interacts with system B which then interacts with system C which then ... which then interacts with system R. This final system acts as the experiment's ultimate record-keeping device. Such measurement chains were visualized in the main text using figures along the lines of Fig. 10. In such figures, the sequence of systems is laid out horizontally and each system only ever interacts with the systems which are just before and just after it.

In this example, each system might (in either of its two interactions) be modeled in either Theory 1 or Theory 2. In Fig.10 the red arrow traces out which systems are being modeled in which theory when with Theory 1 on the top line and Theory 2 on the bottom line. Specifically,
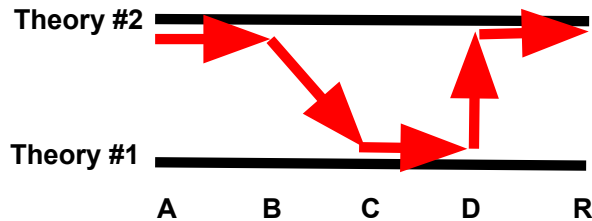
Figure 10. An example of how one might represent a sequential measurement chain of the kind considered in Sec. IV and Sec. V.
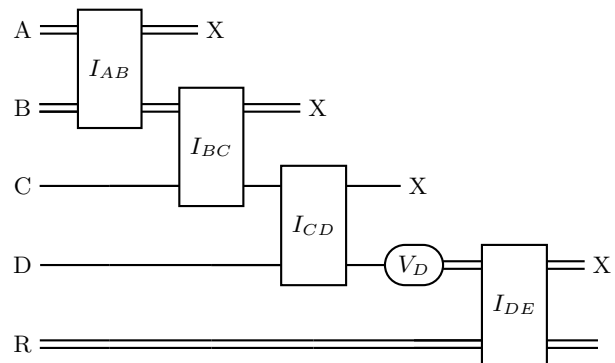

Figure 11. An improved representation of the sequential measurement chain shown in Fig.10. Here moving from one system to another is represented vertically and moving horizontally advances in time towards the end of the experiment. The boxes indicate which systems interact with each other. The type of line used (single or double) indicates which theory the system is being modeled in. The circles indicate where a system switches within which theory it is being modeled in between interactions.

in the measurement chain depicted by Fig. 10: System A is modeled in Theory 2 in its later interaction (it has no earlier interaction). System B is modeled in Theory 2 in both its earlier and later interactions. System C is modeled in Theory 1 in both its earlier and later interactions. System D is modeled in Theory 1 in its earlier interaction and Theory 2 in its later interaction. Finally, system R is modeled in Theory 2 in its earlier interaction (it has no later interaction).

Notice that system D switches from Theory 1 to Theory 2 between its interactions with C and with E. This is marked by the vertical arrow in Fig. 10. This is a vertical cut. Notice also that the BC interaction has system B modeled in Theory 2 and system C modeled in Theory 1. This is marked by the diagonal arrow in Fig. 10. This is a diagonal cut. All other interactions are between systems being modeled within the same theory. These are marked by the horizontal arrows in Fig. 10.

While these sorts of sequential measurement chains were good enough to spell out the central claims of this paper, they are of a rather limited scope for actual application. To address this, I will now generalize away from such sequential examples. The principal limitation of this way of representing measurement chains is that moving in the horizontal direction represents both moving from one system to another, and advancing in time. We can only represent measurement chains in this way if the experiment is arranged as sequentially as time is. In order to generalize upon such examples, we need to change how we are representing things.

In particular, we can revise the representation shown in Fig. 10 to the one shown in Fig. 11. The principle difference is that we now distinguish between moving from one system to another (moving vertically in Fig. 11), and moving forward in time (moving horizontally in Fig. 11). The incidental coincidence of these in the present example is made manifest in our new representation as things are arranged roughly diagonally. This overall diagonal arrangement is incidental and I will generalize away from it soon (see Fig. 12).

In this new representation, what do the lines, boxes, circles, and "X"s represent? The horizontal lines (whether they are single or double) represent the continu-

ity of the systems through time, progressing rightwards. The square box labeled "I_{AB}" which crosses lines A and B represents an interaction occurring between systems A and B. Similarly, for all other boxes. The "X"s at the end of each line indicate that these systems are no longer relevant for modeling the experiment.

The lines being either single or double represents the system being modeled in either Theory 1 or Theory 2 respectively. Notice that each of systems A, B, C, and E are modeled in the same way through the duration of the experiment. System D, however, is at first modeled with Theory 1 before later being modeled with Theory 2. This switch happens between system D's interactions with systems C and E and is marked in Fig. 11 by the circular box labeled "V_D".

In our old representation, vertical and diagonal cuts were represented by vertical and diagonal arrows (that's how they got their names). But how are they represented in our new representation? By definition, vertical cuts occur when, in between its interactions with other systems, we switch which theory we are using to model a system. Such situations are marked with circles in our new representation e.g., V_D in Fig. 11. By definition, diagonal cuts occur when we have an interaction between systems which are modeled in different theories. In our new representation, such situations are marked by boxes with multiple kinds of lines entering and exiting them, e.g., I_{BC} in Fig. 11.

To see what else this new representation is capable of, let us consider a more general example, e.g., Fig. 12. This is not just some contrived example, it comes from a considering the following quantum teleportation experiment [62] outlined below.
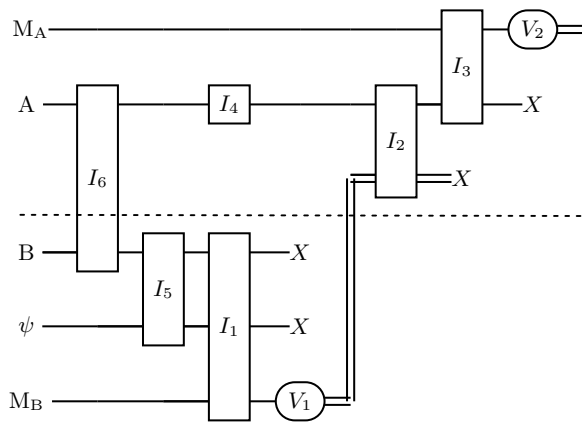
Figure 12. A representation of a non-sequential measurement chain, modeled after a quantum teleportation experiment [62]. The dashed line marks Alice's lab (top) off from Bob's lab (bottom). The various systems involved in this measurement chain are arranged vertically. Moving horizontally advances in time towards the end of the experiment. The boxes indicate which systems interact with each other. The type of line used (single or double) indicates which theory the system is being modeled in. The circles indicate where a system switches within which theory it is being modeled in between interactions.

There are two parties in this experiment, Alice and Bob. Alice's lab contains those systems above the dashed line, Bob's lab contains those below. Bob has a quantum state ψ which he wants to communicate to Alice, however he has no direct way to reliably send quantum information to her. He can, however, reliably send her classical information. Concretely, the journey from Bob's lab to Alice's lab is rough, such that all fragile quantum effects (superposition, entanglement, etc) are likely to be lost or distorted along the way. Given this situation, it would appear that there is no way for Bob to communicate his quantum state ψ to Alice.

However, surprisingly there is a way for Bob to "teleport" the quantum information to Alice. All that is required is for Alice and Bob to be able to reliably produce shared entangled states. Their procedure is as follows. Firstly, perhaps in the distant past, Bob and Alice's systems A and B undergo some interaction, I_6, which gets them maximally entangled. Bob then has B interact with ψ in a unitary way at I_5. At this point, A is maximally entangled with the joint system $B\psi$.

Bob then has B and ψ interact with some measurement device, M_B, in a unitary way at I_1. In the meantime Alice's A undergoes free evolution at I_4. The state of M_B following I_1 is such that it can be described classically without major error. Hence we are justified in taking the vertical cut at V_1. Throughout this example, a single wire means modeled in quantum theory, while a double wire means modeled classically.

States which can be described classically without ma-

jor error are just the kind of states that Bob can reliably send to Alice. He does so. Alice then interacts M_B with A in some unitary way at I_2. In particular, this interaction is a controlled unitary with a different unitary transformation happening to A conditionalized on the state of M_B. Incredibly, the result of this interaction is that the state of A after I_2 is identical to the initial state of ψ. In particular, whatever measurement system, M_A, Alice then interacts with A at I_3, the result will be just the same as if M_A had interacted with the initial state of ψ. Bob has thus successfully communicated the state of ψ to Alice without sending her any quantum information directly.

This example has three features which cannot be captured in our old representation. Firstly, this example includes a 1-party interaction, I_4, and a 3-party interaction, I_1, whereas before we could only handle 2-party interactions. Secondly, before all of our interactions were in a definite ordering whereas here I_4 might be simultaneous (or more generally space-like separated) with either I_5 or I_1. Thirdly, as I will now discuss, this example includes some non-Markovianity: systems may interact with things which they are already correlated with.

This experiment is non-Markovian in the following sense. The step-by-step evolution of system A cannot be understood in terms of its reduced state and the reduced state of the system's it directly interacts with. System A begins in a pure state but after I_6, A is maximally entangled with B such that its reduced state is maximally mixed. Its free evolution I_4 is unitary, such that after this its reduced state is still maximally mixed. The interaction I_2 is a controlled unitary, with a different unitary being applied depending on the state of M_B. Any such unitary would map the maximally mixed state to itself. As such after I_2 the reduced state of A should still be maximally mixed. However, it is not. The state of A after I_2 is pure, namely it is whatever pure state ψ initially had. The error in the above analysis is neglecting the fact that the states of A and M_B at the beginning of I_2 are already correlated with each other. This ability of previous states of a system to come back and non-trivially influence its later evolution, is non-Markovianity.

Thus, unlike figures such as Fig. 10, figures such as Fig. 11 and Fig. 12 can represent non-sequential measurement chains. However, it is important to distinguish three ways of being sequential:

- Strong: In a strongly sequential measurement chain we can place the systems involved in the experiment in a total ordering. For example, in Fig. 10 and Fig. 11 we have $A < B < C < D < R$. The only interactions between these systems are pairwise between adjacent systems.

- Global: In a globally sequential measurement chain, we can place the interactions involved in the experiment into a total ordering. For example in Fig. 12 if we neglect I_4 then we have $I_6 < I_5 < I_1 < I_2 < I_3$.

- Local: In a locally sequential measurement chain, for each system we can place *its* interactions in a total order. For example, Fig. 12 is not globally sequential because we cannot totally order all of the interactions (now including I_4). However, if we restrict our attention to any individual system, its interactions are totally ordered (from left to right). E.g., for system A we have $I_6 < I_4 < I_2$. This measurement chain is therefore locally sequential.

These three definitions are formalized in the following subsection.

So far every measurement chain which we have considered has been locally sequential. Indeed, this is often natural. In particular, when one of our systems are each localized around some spacetime trajectory, we can easily give a total ordering to its interaction: we just order them as they happen along the trajectory. However, when considering extended systems (such as QFTs) this will not always be possible. An extended system might be "simultaneously" be involved in two interactions at once. That is, two of its parts might be undergoing interactions in space-like separation. In QFT, we cannot easily dodge this issue by considering these parts themselves to be separate systems: due to its Type III algebra, we cannot partition the global Hilbert space into two tensor factors associated with bounded spacetime regions.

One might think that we therefore ought to generalize away from being locally sequential by allowing a system's interactions to be only partially ordered. However, as I will discuss in Appendix A 3, even this is not enough to cover many QFT experiments. We might for instance have one field undergo two interactions which are in some sense both before and after each other, see Fig. 14. Thus, I will actually assume nothing about the way that the interactions are ordered.

I have not yet found a good way to represent non-locally-sequential measurement chains pictorially. This, and the need to prove things about measurement chains, means we need to now formalize things.

2. Formalizing Measurement Chains and Cuts

The above subsection gestured at how we can move away from sequential measurement chains (strongly, globally, and locally) to no-sequential measurement chains. In this subsection I will formalize these generalized non-sequential measurement chains. The following subsection will then apply this formalization to a QFT experiment. After this I will prove the central claim of the paper within this formalization.

A *(modeling of) a measurement chain* is a 7-tuple $\langle S, I, T, @, H, \{\mapsto_s\}_{s \in S}, Z \rangle$ where:

1. S is a finite non-empty set containing labels for the systems (or fields) being modeled for our experiment.

2. I is a finite non-empty set containing labels for the interactions being modeled for our experiment.

3. T is a finite non-empty set containing labels for the theories which are being used to model parts of our experiment.

4. $@$ is a binary relation between S and I telling us which systems are involved in which interactions. That is, $s@i$ if and only if system s is involved in interaction i. Let $I_s := \{i \in I : s@i\}$ be the subset of interactions which s is involved in. Let $S_i := \{s \in S : s@i\}$ be the subset of systems which are involved in interaction i. We require that at least one system is involved in each interaction, $|S_i| > 0$ for all $i \in I$. We also require that every system is involved in some interaction, $|I_s| > 0$ for all $s \in S$.

5. $H : S@I \to T$ is a function into the set of theories, T, from the subset of $S \times I$ with $s@i$, i.e, $S@I := \{(s, i) \in S \times I : s@i\}$. This function tells us within which theory each system is being modeled at each interaction.

6. $\mapsto_s$ is, for each $s \in S$, a relation over I_s which coordinates the inputs and outputs of each interaction. In particular, $i \mapsto_s j$ if and only if some output of s at i is an input of s at j. More will be said about how to identify $\mapsto_s$ in practice later, especially regarding QFTs.

7. $Z : S@I \to T \cup \{X\}$ is a function from $S@I$ into T which can also take a special value X. After each of its interactions any system might have an output which is not relevant as an input for any further interactions, but is nonetheless relevant for the final result of the experiment. $Z(s, i) = X$ tells us that the output of s at i has no direct relevance for the end of the experiment (it may still be indirectly relevant however as an input to further interactions). Otherwise, $Z(s, i)$ tells us which theory s after i is to be modeled in for use at the end of the experiment.

Any 7-tuple satisfying the above constraints is a measurement chain.

Given a modeling of a measurement chain, we can identify its vertical and diagonal cuts within this formalization as follows:

V) A vertical cut occurs under two conditions. Firstly, a vertical cut occurs at $s \in S$ between i and $j \in I_s$ whenever $i \mapsto_s j$ and $H(s, i) \neq H(s, j)$. In this case, s switches theories between i and j. Secondly, a vertical cut occurs at $s \in S$ after $i \in I_s$ whenever $Z(s, i) \neq X$ and $H(s, i) \neq Z(s, i)$. In this case, s switches theories after one of its interactions and before the end of the experiment.

D) A diagonal cut occurs at $i \in I$ between s and $t \in S_i$ whenever $H(s,i) \neq H(t,i)$. In this case, interaction i is modeled as coupling together systems from different theories.

Connecting with the nomenclature of Sec. V, we say a vertical cut is a vertical (V,W)-cut just when $H(s,i) \in V \subset T$ and either $H(s,j) \in W \subset T$ or $Z(s,i) \in W \subset T$, whichever is relevant. Similarly, a diagonal cut is a diagonal (V,W)-cut just when $H(s,i) \in V \subset T$ and $H(t,i) \in W \subset T$.

The above description of a measurement chain presupposes no structure on $\mapsto_s$. That is, we have not yet made any restrictions as to how the measurement chain's inputs and outputs are allowed to be arranged. As I will soon discuss, this allows us to generalize away from the locally sequential measurement chains considered so far. However, before discussing this, let's see how this formalization works in a more familiar case. In particular, let's formalize the measurement chain shown in Fig. 12.

For this measurement chain the labels for our systems are $S = \{M_A, A, B, \psi, M_B\}$. The labels for our interactions are $I = \{I_1, I_2, I_3, I_4, I_5, I_6\}$. Finally, the labels for our theories are $T = \{T_1, T_2\}$ where T_1 is represented by a single line and T_2 is represented by a double line. The relation @ and the function $H : S@I \to T$ are given by the following table:

$H(s,i)$	I_1	I_2	I_3	I_4	I_5	I_6
M_A	$\emptyset$	$\emptyset$	T_1	$\emptyset$	$\emptyset$	$\emptyset$
A	$\emptyset$	T_1	T_1	T_1	$\emptyset$	T_1
B	T_1	$\emptyset$	$\emptyset$	$\emptyset$	T_1	T_1
ψ	T_1	$\emptyset$	$\emptyset$	$\emptyset$	T_1	$\emptyset$
M_B	T_1	T_2	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$

The presence of $\emptyset$ in row s and column i means that, $\neg s@i$, and that therefore H is not-valued here. In all other cases the value in row s and column i is the value that H takes at s and i. Thus, in this example, M_B at I_2 is modeled in T_2, but otherwise every system at every interaction is modeled within T_1.

From the above table we can read off this measurement chain's diagonal cuts. Its only one is at I_2 which couples A modeled in T_1 and M_B modeled in T_2.

For this measurement chain the arrow relations can be read off of the connectivity of the boxes. These relations are given by the following directed graphs:

S	Directed Graph of $\mapsto_s$ over I_s
M_A	I_3
A	$I_6 \mapsto_A I_4 \mapsto_A I_2 \mapsto_A I_3$
B	$I_6 \mapsto_B I_5 \mapsto_B I_1$
ψ	$I_5 \mapsto_\psi I_1$
M_B	$I_1 \mapsto_{M_B} I_2$

The fact that all of these directed graphs are linear (and consequently acyclic) are accidental facts about this example which we will generalize away from later. This only occurs in this example because the measurement chain is locally sequential.

The function $Z : S@I \to T \cup \{X\}$ is given by the following table:

| $Z(s,i)$	I_1	I_2	I_3	I_4	I_5	I_6
M_A	$\emptyset$	$\emptyset$	T_2	$\emptyset$	$\emptyset$	$\emptyset$
A	$\emptyset$	X	X	X	$\emptyset$	X
B	X	$\emptyset$	$\emptyset$	$\emptyset$	X	X
ψ	X	$\emptyset$	$\emptyset$	$\emptyset$	X	$\emptyset$
M_B	X	X	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$

As before, the presence of $\emptyset$ in row s and column i means that $\neg s@i$ and therefore that Z is not-valued here. In all other cases the value in row s and column i is the value that Z takes at s and i. As noted above, if $Z(s,i) = X$ tells us that (outside of being an input to further interactions) s is not relevant after i. In this example, only the state of M_4 after I_3 is directly relevant to the end of the experiment. Everything else in $S@I$ is only relevant through its further interactions.

From the above table and the above directed graphs we can read off this measurement chain's vertical cuts. It has two: one at M_B between interactions I_1 and I_2, and one at M_A between I_3 and the end of the experiment.

The fact that Z takes a value of X in all but one row is accidental to this example. This motivates the following definition. Let us define a measurement chain's *record-keepers* to be those systems $s \in S$, such that $Z(s,i) \neq X$ for at least one $i \in I_s$. Let us define a measurement chain's *record-keeping interactions* to be those interactions $i \in I$, such that $Z(s,i) \neq X$ for at least one $s \in S_i$. All examples considered up to this point have had a unique record keeper and a unique record-keeping interaction. This is in addition to all being locally-sequential.

But, what exactly changes when we move away from non-locally-sequential measurement chains? In order to address this question, I must first make a number of definitions leading to a characterization of strongly, globally and locally sequential measurement chains in this formalization.

For any measurement chain we can define a new relation, $\mapsto$, such that $i \mapsto j$ if and only if there exists an $s \in S$ with $i, j \in I_s$ and $i \mapsto_s j$. That is, $\mapsto$ tells whether any $\mapsto_s$ connects i and j. In the above example the directed graph for $\mapsto$ is:
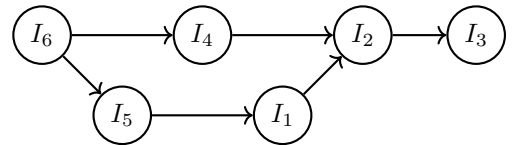
Note that, in this example, the directed graph of $\mapsto$ is not linear even though the graphs of each $\mapsto_s$ are. This reflects the fact that while this example is locally sequential, it is not globally sequential.

Before characterizing this distinction in this formalization another key definition is needed. Given any measurement chain we can define another new relation by taking

the reflexive transitive closure[13] of $\mapsto$ (notated $\mapsto^{\mathrm{RTC}}$). According to this new relation $i \mapsto^{\mathrm{RTC}} j$ just in case either $i = j$ or some sequence of systems and interactions connect i to j via their inputs and outputs.

In many cases, it will be natural to demand that each interaction i has $i \mapsto^{\mathrm{RTC}} j$ for some record keeping interaction j. Otherwise, it seems like i is irrelevant to the end state of the experiment. In our running example, I_3 is the only record keeping interaction and moreover every interaction i has $i \mapsto^{\mathrm{RTC}} I_3$.

We can now characterize strongly, globally, and locally sequential measurement chains in this formalization.

A measurement chain is globally sequential just in case $\mapsto^{\mathrm{RTC}}$ is a total ordering. Our running example is not globally sequential, but would be if we dropped interaction I_4.

A measurement chain is strongly sequential just in case it is globally sequential, all its interactions are pair-wise (with $|S_i| = 2$) and moreover $|S| = |I| + 1$. This allows us to define a total ordering over S.

By contrast, in this formalization, a measurement chain is locally sequential when for each $s \in S$ we have that $\mapsto_s^{\mathrm{RTC}}$ is a total ordering along with a compatibility constraint. Specifically, the set of relations $\{\mapsto_s\}$ is compatible if and only if $\mapsto^{\mathrm{RTC}}$ is at least a partial ordering. Note that our running example is locally sequential.

But what exactly is this compatibility constraint doing? In general, for any relation Q we have that Q^{RTC} is a partial order if and only if Q is acyclic: there is no $N > 1$ and no sequence $a_1, \ldots, a_N$ such that $a_n Q a_{n+1}$ and $a_N = a_1$. Thus, explicitly, the compatibility criteria here for $\{\mapsto_s\}_{s \in S}$ is just that $\mapsto$ is acyclic. The measurement chain having $\mapsto$ acyclic means that there are no circular dependencies.

How could this have failed for a locally sequential measurement chain? Consider Fig. 12, with each $\mapsto_s$ being linear but having different orientations, some right-to-left and some left-to-right. In such a case $\mapsto$ would have been cyclic. Compatibility ensures all of our systems agree about what "later" means.

Given this characterization, if we are to move away from locally sequential measurement chains, we need to consider $\mapsto_s$ which do not lead to $\mapsto_s^{\mathrm{RTC}}$ being a total ordering. As I will now discuss, it is useful to characterize measurement chains which are not locally-sequential into two types depending on whether $\mapsto_s^{\mathrm{RTC}}$ is partial ordering or not.

Let us then define a measurement chain to be *locally partially ordered* when for each $s \in S$ we have that $\mapsto_s^{\mathrm{RTC}}$ is a partial ordering and moreover the $\mapsto_s$ relations are

---

[13] Roughly, the reflexive transitive closure of a relation Q is this relation plus aQa for every a and aQb for every a and b where for some c we have $aQc \wedge cQb$. This last step is done repeatedly until convergence. Concretely, Q^{RTC} is the smallest extension of Q which is transitive and reflexive. Namely, it is the intersection of all of the extensions of Q which are reflexive and transitive.
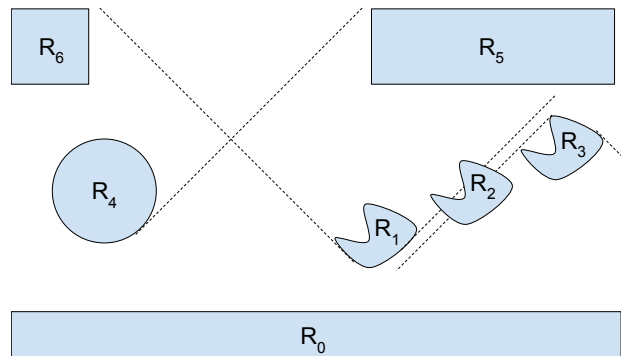

Figure 13. The space-time regions involved in the entanglement harvesting experiment formalized in Appendix A.

compatible such that $\mapsto^{\mathrm{RTC}}$ is also a partial ordering. Demanding that each $\mapsto_s^{\mathrm{RTC}}$ is a partial ordering is equivalent to demanding that each $\mapsto_s$ is acyclic. Thus an alternate name for such measurement chains might be to call them *locally acyclic* measurement chains. In the following subsection we will see a QFT-based example of a locally acyclic measurement chain which is not locally sequential.

By contrast, let us then define a measurement chain to be *locally cyclic* when for at least one $s \in S$ we have $\mapsto_s$ being cyclic. From this it follows that neither $\mapsto_s^{\mathrm{RTC}}$ nor $\mapsto^{\mathrm{RTC}}$ are partial orders. In the following subsection we will see a QFT-based example of a cyclic measurement chain.

3. Two QFT-based Examples

Let us now apply this formalization to a QFT experiment. In particular, let us consider an entanglement harvesting experiment [16, 29–35] of the kind described in Sec. VI C.

Consider the following experiment: A quantum field, ϕ, comes to a cold global thermal equilibrium by some process in the past of our experiment (in region R_0 in Fig.13). Following this entanglement is harvested from this field by local interactions with two initially uncorrelated probe systems, let's call them Probe 1 and Probe 2. Specifically, Probe 1 undergoes local interactions with ϕ in regions R_1, R_2, and R_3. Probe 2 interacts with ϕ in region R_4. Each of these probes later interact with some measurement device M_1 and M_2 in regions R_5 and R_6 respectively. The correlation between the outputs of these measurement devices confirms that the probes were entangled after their interaction with ϕ. Our experiment might then investigate how the amount of entangle harvested depends on various factors: the initial temperature of ϕ, the strengths of various interactions, etc.

Let us now design a measurement chain to help us guide us in modeling this experiment. Firstly, regarding the thermalization of the field, ϕ, we may reasonably leave this process out of our modeling and just assume

that ϕ starts in a global thermal state. Next we must consider how to model the probe systems. Roughly, our options are either to model them as QFTs (e.g., in the FV framework) or as non-relativistic quantum systems (e.g., as a UDW detector) localized around some trajectory. See Sec. VI for details.

Given the spacetime relations between R_1, R_2 and R_3 we cannot model Probe 1 as a non-relativistic system localized around some trajectory. No time-like trajectory can get us from R_1 to R_3. Thus, we are forced to model Probe 1 as a QFT, let's call it ψ. Regarding Probe 2, no such barrier exists. Let us suppose that we can model this probe well as a UDW detector.

The interactions in R_5 and R_6 are with measurement devices M_1 and M_2. It is doubtful we can model these systems or the measurement processes within QFT. Thus, between R_1-R_3 and R_5 we will need to switch to describing ψ within non-relativistic quantum theory. These measurements can then both be handled well within our already established non-relativistic quantum measurement theory.

Let's transfer the above considerations into a formalized measurement chain. We have five systems: ϕ the field the entanglement is being harvested from, ψ a probe field, UDW a Unruh-DeWitt detector, two measurement apparatuses, M_1 and M_2. Three theories are relevant to this measurement chain: T_1, quantum field theory, T_2, non-relativistic quantum theory, and T_3 being classical theory. These systems undergo six interactions as indicated by the following table:

| $H(s,i)$	I_1	I_2	I_3	I_4	I_5	I_6
ϕ	T_1	T_1	T_1	T_1	$\emptyset$	$\emptyset$
ψ	T_1	T_1	T_1	$\emptyset$	T_2	$\emptyset$
UDW	$\emptyset$	$\emptyset$	$\emptyset$	T_2	$\emptyset$	T_2
M_1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	T_2	$\emptyset$
M_2	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	T_2

Already we can see that our measurement chain involves a diagonal cut in I_4 between ϕ and our UDW detector. This is a diagonal QFT-cut because a QFT is coupled to a non-QFT.

Next we need to specify $\mapsto_s$. How do these spacetime relations in Fig. 13 determine the relations $\mapsto_s$? We can understand an interaction's "outputs" in the above formalization as meaning any future oriented time-like or light-like trajectories leaving the spacetime region associated with an interaction. We can similarly understand an interaction's "inputs". Thus, in this context[14] we can define $i \mapsto_s j$ if and only if $s@i$ and $s@j$ and there is future directed time-like or light-like curve beginning in

---

[14] Applied to other problems, one may want to add an extra condition here that these time-like or light-like trajectories are not allowed to travel through other interaction regions. If all paths from R_A to R_C lead through R_B, it makes more sense to understand I_A as having no output into I_C. Rather, we can then understand I_A to output into I_B which then outputs into I_C.

R_i and ending in R_j. Taking $\mapsto_s$ to be determined in this way, we have:

S	Directed Graph of $\mapsto_s$ over I_s
ϕ	$I_1 \mapsto_\phi I_2 \mapsto_\phi I_3, \quad I_4$
ψ	(See Below)
UDW	$I_4 \mapsto_{\text{UDW}} I_6$
M_1	I_5
M_2	I_6

The relation $\mapsto_\psi$ in this case is:
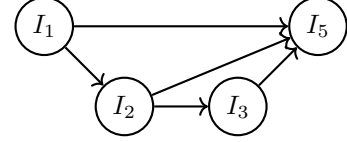
The first thing to notice here is that the directed graph for $\mapsto_\phi$ is composed of two disjoint subgraphs. It follows from this that $\mapsto_\phi^{\text{RTC}}$ is not a total order and therefore this measurement chain is not locally sequential. Next, notice that one interaction may have multiple outputs for a single system: $I_1 \mapsto_\psi I_2$ and $I_1 \mapsto_\psi I_5$. This has not occurred in any previous example. Secondly, notice that it is not the case that $I_1 \mapsto_\psi I_3$ even though $I_1 \mapsto_\psi I_2$ and $I_2 \mapsto_\psi I_3$. That is, $\mapsto_\psi$ is not transitive. The relation, $\mapsto_\psi^{\text{RTC}}$, however is transitive but importantly can't be interpreted straightforwardly in terms of spacetime regions.

The directed graph for $\mapsto$ is:
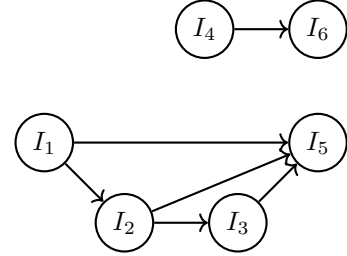
Notice that the directed graph for $\mapsto$ is composed of two disjoint subgraphs. This has not happened before in any previous examples. The fact that this happens here actually reflects a central consideration in this entanglement harvesting experiment. For the entanglement between ψ and UDW to be cleanly associated with space-like entanglement in the field ϕ their interactions and measurements need to be isolated from each other.

The last thing we need to do to specify our measurement chain is to define the Z function. Z here is given by:

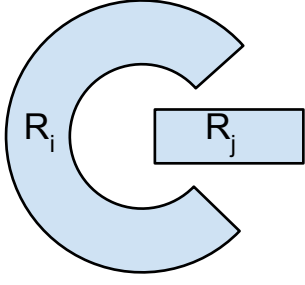| $Z(s,i)$	I_1	I_2	I_3	I_4	I_5	I_6
ϕ	X	X	X	X	$\emptyset$	$\emptyset$
ψ	X	X	X	$\emptyset$	X	$\emptyset$
UDW	$\emptyset$	$\emptyset$	$\emptyset$	X	$\emptyset$	X
M_1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	T_3	$\emptyset$
M_2	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	T_3

Figure 14. Spacetime regions which would lead to a cyclic measurement chain. Regions R_i and R_j are both before and after each other in a sense.

The only interactions which are directly relevant to the end of the experiment are the measurements at I_5 and I_6. As the above table indicates, the states of M_1 and M_2 following these interactions will be modeled in T_3, i.e., classically.

The vertical cuts in this measurement chain are as follows. Between I_1-I_3 and I_5 we have a vertical cut on ψ from T_1 to T_2. That is, after ψ is done probing ϕ we need to switch to modeling it as a non-relativistic quantum system. There are many options available as to how one might do this, see Sec. VI. In any case, this is a vertical QFT-cut. There are also two vertical Heisenberg cuts on M_1 and M_2 following the interactions I_5 and I_6 respectively.

We have thus formalized a non-locally-sequential measurement chain. This measurement chain was, however, acyclic. That is, $\mapsto$ is acyclic such that $\mapsto^{\mathrm{RTC}}$ is a partial order. As I suggested in the previous subsection, this will not always be the case. Consider for instance, the space-time regions shown in Fig. 14. Consider a field f which undergoes interactions i and j localized in these regions. In accordance with our above discussion we would have $i \mapsto_f j$ and $j \mapsto_f i$. This is because a part of R_i is in the future of R_j and likewise part of R_j is in the future of R_i. The formalization presented in the previous subsection is flexible enough to handle such examples.

4. Proving the Necessity of Cuts

Having so formalized modelings of measurement chains, let's now prove the claim in the main paper regarding under what conditions cuts are required. In particular, let us now prove the following claim: If one part of the experiment must be modeled some way and another later part must not be modeled in this way then somewhere in between these we need to make a certain kind of cut.

Formalizing this claim we understand the first part of the experiment as referring to a system s at one of its interactions i with $s@i$. The second part of the experiment mentioned above could is understood as a system

t *either at or after* one of its interactions j with $t@j$. In the first case we will be concerned with $H(t,j)$ and in the second case $Z(t,j)$. We understand the "must (must not) be modeled in some way" here as specifying some subset of our theories $V \subset T$ and fixing $H(s,i) \in V$ and $H(t,j) \notin V$ (or alternatively $Z(t,j) \notin V$). We understand the claim that the second part of the experiment is "later" as[15]. Finally, we can understand the "the certain kind of cut" here to mean either a vertical or a diagonal V cut.

Thus the formalized claim is as follows. Suppose we are given any modeling of a measurement chain and are given some $s_0, t_0 \in S$ and $i_0, j_0 \in I$ with $s_0@i_0$ and $t_0@j_0$ and $i_0 \mapsto^{\mathrm{RTC}} j_0$. Suppose further that $H(s_0, i_0) \in V$ and either $H(t_0, j_0) \notin V$ or if $Z(t_0, j_0) \neq X$ then $Z(t_0, j_0) \notin V$. In such cases, there must be some interaction k with $i_0 \mapsto^{\mathrm{RTC}} k \mapsto^{\mathrm{RTC}} j_0$ where either a diagonal or vertical V cut happens.

Proof: Suppose for contradiction that no such cut occurs. I will show that this implies that $H(t_0, j_0) \in V$ and if $Z(t_0, j_0) \neq X$ then $Z(t_0, j_0) \notin V$. This would be in explicit contradiction with our above assumptions.

First, note that no vertical V cuts implies that for all $s \in S$ and $i, j \in I_s$ with $i \mapsto_s j$ we have $H(s,i) \in V \implies H(s,j) \in V$. Similarly, if $Z(s,i) \neq X$ then $H(s,i) \in V \implies Z(s,i) \in V$. Similarly, no diagonal V cuts implies that for all $i \in I$ and $s,t \in S_i$ we have $H(s,i) \in V \implies H(t,i) \in V$. As I will now show, our desired result follows from a series of applications of these three implications.

To see this, begin by noting that $i_0 \mapsto^{\mathrm{RTC}} j_0$ means that there exists some sequence of interactions $k_1, \ldots, k_N \in I$ and some corresponding sequence of systems $x_0, \ldots, x_N \in S$ with,

$$i \mapsto_{x_0} k_1 \mapsto_{x_1} \cdots \mapsto_{x_{N-1}} k_N \mapsto_{x_N} j. \quad (A1)$$

This follows from the definition of $\mapsto^{\mathrm{RTC}}$ as the reflexive transitive closure of $\mapsto$ and the definition of $\mapsto$ in terms of $\mapsto_s$. For the above sequence to make sense it is required that $k_n \in I_{x_n} \cap I_{x_{n-1}}$ That is, the interaction k_n must involve both x_n and x_{n-1}.

We can decompose any such sequence into a series of "runs" where $x_n = \cdots = x_{n+m-1} = y$ for some $m > 0$ and some fixed $y \in S$. Such runs make repeated use of some $\mapsto_y$ as,

$$k_n \mapsto_y k_{n+1} \mapsto_y \cdots \mapsto_y k_{n+m}. \quad (A2)$$

Note that because no vertical V cuts are allowed we have $H(y, k_n) \in V \implies H(y, k_{n+1}) \in V$. Applying this logic repeatedly we ultimately have $H(y, k_n) \in V \implies H(y, k_{n+m}) \in V$. Thus, moving

---

[15] Note that given the discussion in the previous subsection, one must be careful when trying to interpret $\mapsto^{\mathrm{RTC}}$ in terms of spacetime relations.

along one of these $\mapsto_y$ runs cannot help us switch from theories in V to theories in $T - V$.

What about at the interface between such runs? Namely,

$$\cdots \mapsto_y k_{\ell-1} \mapsto_y k_\ell \mapsto_z k_{\ell+1} \mapsto_z \ldots. \qquad (A3)$$

with $y \neq z$ and $y, z \in S_{k_\ell}$. Must we have $H(z, k_{\ell+1}) \in V$ if $H(y, k_\ell) \in V$? Yes, because we are not allowed diagonal cuts we must have $H(y, k_n) \in V \implies H(z, k_n) \in V$. Since y and z are both involved in the interaction k_n they must both be modeled in V, anything otherwise would give rise to a diagonal V cut. However,

then by our above logic, without vertical cuts we have $H(z, k_\ell) \in V \implies H(z, k_{\ell+1}) \in V$.

Thus, without allowing for vertical cuts or diagonal V cuts, we ultimately have $H(s, i) \in V \implies H(t, j) \in V$ and therefore $H(t, j) \in V$. If $Z(t, j) \neq X$ then we have further that $Z(t, j) \in V$ since no vertical V cuts implies $H(t, j) \in V \implies Z(t, j) \in V$.

In order to avoid this conclusion we need to allow for vertical or diagonal V cuts. In particular, the above proof shows that we must either take a vertical V-cut right at the end, or otherwise at some point in sequence, Eq. (A1) connection i and j either type of cut must happen.