Certifying Quantum Signatures in Thermodynamics and Metrology via Contextuality of Quantum Linear Response

Lostaglio, Matteo

DOI
10.1103/PhysRevLett.125.230603

Publication date
2020

Document Version
Final published version

Published in
Physical Review Letters

Certifying Quantum Signatures in Thermodynamics and Metrology via Contextuality of Quantum Linear Response

Matteo Lostaglio[1,2,*]
[1]ICFO-Institut de Ciencies Fotoniques, The Barcelona Institute of Science and Technology, Castelldefels (Barcelona) 08860, Spain
[2]QuTech, Delft University of Technology, P.O. Box 5046, 2600 GA Delft, Netherlands

(Received 26 May 2020; accepted 3 November 2020; published 2 December 2020)

I identify a fundamental difference between classical and quantum dynamics in the linear response regime by showing that the latter is, in general, contextual. This allows me to provide an example of a quantum engine whose favorable power output scaling *unavoidably* requires nonclassical effects in the form of contextuality. Furthermore, I describe contextual advantages for local metrology. Given the ubiquity of linear response theory, I anticipate that these tools will allow one to certify the nonclassicality of a wide array of quantum phenomena.

DOI: 10.1103/PhysRevLett.125.230603

Linear response theory describes the reaction of a quantum system to a small perturbation. The theory finds countless applications in many fields of quantum physics, including molecular, atomic, and nuclear physics, quantum optics, and statistical mechanics. In this Letter, I present a general test to certify whether the linear response of a quantum system necessarily requires contextuality.

Using these results, I identify quantum signatures in the power of a heat engine. In the context of quantum thermodynamics, the issue of identifying truly quantum signatures has been a long-standing problem in the field. Several theoretical claims have been made that quantum coherence can offer improvements over certain incoherent thermodynamic engines and refrigerators (e.g., Refs. [1–10], and references therein), followed by recent experimental effort [11]. However, similar signatures can be observed in classical engines as well [12,13]. Hence, such claims should be backed by a no-go theorem that (i) defines a precise notion of nonclassicality and (ii) shows that this notion leads to statistical predictions incompatible with the corresponding quantum statistics.

Here, I adopt a stringent notion of nonclassicality, namely, generalized contextuality [14]. I prove that the power output of a two-stroke quantum engine in the weak coupling regime cannot be achieved in any noncontextual model reproducing the operational features of the quantum experiment [15]. As a second application, I turn to local metrology and consider the archetypal example of phase estimation using a qubit system. I show that, given the phenomenology of the phase estimation experiment, a nonzero Fisher information is incompatible with all classical (noncontextual) models. This complements a recent result showing that certain features of postselected metrology are nonclassical [17], but in my case I do not have to consider any postselection.

The tools developed here, applicable as they are to any quantum system in the linear response regime, can find applications in the identification of genuine quantum signatures in a wide range of different platforms.

Noncontextual ontological models.—Great care needs to be taken when claiming that the performance of a device *requires* nonclassicality. For example, Ref. [12] shows that a short-time cooling enhancement, which in quantum theory is attributed to the presence of quantum coherence, also occurs in classical models where a set of oscillators undergoes Hamiltonian evolution. Here, I want to identify signatures in the dynamics of a quantum system which *unavoidably* signal that quantum effects are at play. Formally, I will identify phenomena which cannot occur within any noncontextual ontological model (OM). Let me describe in detail this broad class of models [18].

We may start from the operational description of preparations, transformations, and measurements, understood as sets of laboratory instructions according to which these operations are performed. To each, we associate the corresponding physical description in the OM, as summarized in Table I below:

(1) To every preparation procedure P, one assigns a probability distribution $\mu_P(\lambda)$ over some (measurable) set of physical states λ. For example, if P involves leaving the system alone for a long time and $\lambda = (x_1, \ldots, x_N, p_1, \ldots, p_N)$ are phase space points, $\mu_P(\lambda)$ may be a thermal distribution.

(2) A transformation procedure T is described by an update rule giving the probability that any final state λ' is reached, given that the initial state was λ. I denote this transition probability by $\mathcal{T}_T(\lambda'|\lambda)$. For example, $\mathcal{T}_T(\lambda'|\lambda)$ may be generated by a rate equation among a discrete set of λ's, as in the classical model described in Ref. [13]. Or, in Hamiltonian dynamics, $\lambda = (x, p)$ and, after time t,

TABLE I. An operational theory and its description within an ontological model.

Operational description	OM description
P preparation procedure	$\mu_P(\lambda)$ probability distribution
T transformation procedure	$\mathcal{T}_T(\lambda'\vert\lambda)$ transition probabilities
M (with outcomes k) measurement procedure	$\xi_M(k\vert\lambda')$ response function
$p[k\vert T(P), M]$ operational statistics	$\int d\lambda d\lambda' \mu_P(\lambda)\mathcal{T}_T(\lambda'\vert\lambda)\xi_M(k\vert\lambda')$

$\mathcal{T}_T[x', p'|x(0), p(0)] = \delta[x' - x(t)]\delta[p' - p(t)]$, where $x(t)$, $p(t)$ is the solution of Hamilton's equations with initial conditions $x(0)$, $p(0)$.

(3) A measurement procedure M with outcomes k is associated to a response function $\xi_M(k|\lambda')$, giving the probability that an outcome k is returned by M if the physical state is λ'. For example, in classical mechanics, if M is a measurement of the energy E of a single particle of mass m and momentum p in a potential $V(x)$, $\xi_M(E|x, p) = \delta\{E - [(p^2/2m) + V(x)]\}$. In a general OM, $\xi_M(k|\lambda')$ may be nondeterministic.

Let $p[k|T(P), M]$ be the statistics collected in an experiment where P is prepared, a transformation T is applied, and, finally, a measurement M with outcomes k is performed. The OM predicts

$$p[k|T(P), M] = \int d\lambda d\lambda' \mu_P(\lambda)\mathcal{T}_T(\lambda'|\lambda)\xi_M(k|\lambda'), \quad (1)$$

as naturally follows from the propagation of probabilities. Hamiltonian mechanics is just a member of a class of OM.

Since arbitrary operational statistics, quantum or otherwise, can be reproduced by an appropriate OM, here we consider the subclass of OMs that are noncontextual [19]. An OM is noncontextual, in the generalized sense introduced by Spekkens [14,20], when it has the property of *assigning identical physical descriptions to operationally indistinguishable procedures*. Specifically, two preparations P and P' are operationally indistinguishable (denoted $P \simeq_{\mathrm{op}} P'$) when $p(k|P, M) \equiv p(k|P', M)$ for every measurement procedure M. This means no experiment is able to distinguish between P and P'. A noncontextual model requires

$$P \simeq_{\mathrm{op}} P' \Rightarrow \mu_P(\lambda) = \mu_{P'}(\lambda) \quad \forall \ \lambda. \quad (2)$$

The same has to hold for operationally equivalent measurements and transformations: If we define $M \simeq_{\mathrm{op}} M'$ as $p(k|P, M) \equiv p(k|P, M') \, \forall \, P$ and $T \simeq_{\mathrm{op}} T'$ as $p[k|T(P), M] \equiv p[k|T'(P), M] \, \forall \, P, M$, a noncontextual OM is one for which

$$M \simeq_{\mathrm{op}} M' \Rightarrow \xi_M(k|\lambda) = \xi_{M'}(k|\lambda) \quad \forall \ k, \lambda, \quad (3)$$

$$T \simeq_{\mathrm{op}} T' \Rightarrow \mathcal{T}_T(\lambda'|\lambda) = \mathcal{T}_{T'}(\lambda'|\lambda) \quad \forall \ \lambda, \lambda'. \quad (4)$$

Noncontextuality, in the general form presented here, can be understood as an extension of the original Kochen-Speckers notion ([18], Appendix C). One can easily see that classical Hamiltonian dynamics is a class of noncontextual OM (see Supplemental Material, Sec. A [21]). Noncontextual models include as special cases the classical models previously considered in the literature: e.g., discrete models with jump probabilities generated by rate equations [13]; Hamiltonian dynamics obtained via classical limit [12]; quantum mechanics in a fixed basis obtained via dephasing in the energy basis [5]. Other examples include Spekken's toy model [22] or Hamiltonian mechanics with an uncertainty principle (the latter is equivalent to Gaussian quantum mechanics [23]). These examples show that noncontextual OMs allow one to reproduce features normally attributed to quantum measurement disturbance, superposition, and entanglement.

In this Letter, I will hence adopt the same stringent notion of *quantum signature* used to analyze several quantum information primitives [17,24–27]: a set of *operational features* that unavoidably require contextuality.

Quantum linear response.—Consider a quantum state $|\psi(t)\rangle$ in a finite-dimensional Hilbert space evolving according to the Schrödinger equation under a time-dependent perturbation $V(t)$:

$$i\hbar \frac{d}{dt}|\psi(t)\rangle = [H_0 + gV(t)]|\psi(t)\rangle. \quad (5)$$

I develop my considerations here for pure states, but the extension to mixed states by linearity is straightforward. We are interested in the change of expectation value of an observable $O = \sum_k o_k |o_k\rangle\langle o_k|$ due to the perturbation (for technical convenience, without loss of generality, I shift O so that $o_k \geq 0$). It is convenient to work in the interaction picture (label "I"), $O_I(t) = e^{iH_0 t/\hbar} O e^{-iH_0 t/\hbar}$, $|\psi_I(t)\rangle = e^{iH_0 t/\hbar}|\psi(t)\rangle \coloneqq U_I(t)|\psi(0)\rangle$, and study

$$\langle \Delta O_I \rangle_t^Q \coloneqq \langle \psi_I(t)|O_I(t)|\psi_I(t)\rangle - \langle \psi_I(0)|O_I(t)|\psi_I(0)\rangle, \quad (6)$$

where "Q" stands for "quantum." Operationally, this corresponds to

$$\langle \Delta O_I \rangle_t \coloneqq \sum_k o_k p[k|T_t(P), M_t] - \sum_k o_k p(k|P, M_t), \quad (7)$$

where P, T_t, and M_t are the preparation, transformation, and measurement procedures described in quantum mechanics by initial state $|\psi(0)\rangle$, unitary dynamics $U_I(t)$, and measurement of the observable $O_I(t)$, respectively. From Dyson's series,

$$U_I(t) = \mathbb{1} - \frac{ig}{\hbar}\int_0^t dt' V_I(t') + \mathcal{O}(g^2), \tag{8}$$

where $V_I(t) = e^{iH_0 t/\hbar} V(t) e^{-iH_0 t/\hbar}$. Quantum linear response gives

$$\langle \Delta O\rangle_t^Q = \frac{ig}{\hbar}\int_0^t dt' \langle\psi(0)|[V_I(t'), O_I(t)]|\psi(0)\rangle + \mathcal{O}(g^2). \tag{9}$$

The most important aspect of this formula is that one can have a response of $\mathcal{O}(g)$ if there are no pairwise commutations among $|\psi(0)\rangle\langle\psi(0)|$, $O_I(t)$, and $\int dt' V_I(t')$.

Another crucial fact is encoded in the following channel equality [28]. Suppose that for g small enough

$$\frac{1}{2}\mathcal{U}_t + \frac{1}{2}\mathcal{U}_t^\dagger = (1-\tilde{p}_d)\mathcal{I} + \tilde{p}_d \mathcal{C}_t, \tag{10}$$

where $\mathcal{U}_t(\cdot) \coloneqq U_I(t)(\cdot)U_I^\dagger(t)$, $\mathcal{I}$ is the identity channel, $\mathcal{C}_t$ is some other channel, and $\tilde{p}_d = \mathcal{O}(g^2)$ as $g\to 0$. I will later give tools to verify if a quantum experiment under consideration admits this decomposition in linear response. For now, it suffices to say that in the case of a single qubit this decomposition holds for every nontrivial perturbation.

Equation (10) underlies the fact that the transformation T_t (represented by U_t in quantum mechanics) can be reversed, to first order in g, by convex combination with another transformation T_t^* [represented by $U_I^\dagger(t)$ in quantum mechanics]. In particular, tossing a fair coin and performing either T_t or T_t^* is operationally indistinguishable from doing nothing with probability $1-\mathcal{O}(g^2)$. These facts can be summarized as

$$\frac{1}{2}T_t + \frac{1}{2}T_t^* \simeq_{\text{op}} (1-p_d)T_{\text{id}} + p_d T_t', \tag{11}$$

where T_{id} denotes the "do-nothing" operation and T_t' denotes some other transformation. As we will see, this approximate "reversibility by mixing" or "stochastic reversibility" tells us that the perturbation T_t cannot be "too far" from the do-nothing operation in any noncontextual model. I stress that Eq. (11) will be required, not Eq. (10). Crucially, Eq. (11) *does not assume the dynamics* T_t *is reversible*. This means my results are applicable beyond exactly unitary dynamics. For example, if $T_t = (1-s)\mathcal{U}_t + s\mathcal{D}$, with $\mathcal{D}$ depolarizing noise [$\mathcal{D}(\rho) = \mathbb{1}/d$ for all ρ], $s\in[0,1]$, and $\mathcal{U}_t$ satisfying Eq. (10), then Eq. (11) holds with $p_d = \tilde{p}_d + s(1-\tilde{p}_d)$.

I now prove the weakness encoded operationally in Eq. (11), together with the observation of a $\mathcal{O}(g)$ response of a quantum system, can occur *only* in the presence of contextuality.

Main theorem.—From Eqs. (1) and (7), an OM predicts that $\langle\Delta O_I\rangle_t$ equals

$$\sum_k o_k\left[\int d\lambda d\lambda' \mu_P(\lambda)\mathcal{T}_{T_t}(\lambda'|\lambda)\xi_{M_t}(k|\lambda') - \int d\lambda \mu_P(\lambda)\xi_{M_t}(k|\lambda)\right]. \tag{12}$$

In other words, when the initial state $|\psi(0)\rangle$ is prepared, a λ is sampled with probability $\mu_P(\lambda)$; when the unitary $U_I(t)$ is performed, the state is updated to λ' with probability $\mathcal{T}_{T_t}(\lambda'|\lambda)$; and, finally, a measurement of the observable $O_I(t)$ returns outcome o_k with probability $\xi_{M_t}(k|\lambda')$. Then

Theorem 1: Noncontextual bound on linear response.—Suppose the operational equivalence in Eq. (11) is observed. Then in any noncontextual OM

$$|\langle\Delta O\rangle_t^{\text{NC}}| \le 4 p_d o_{\max}, \tag{13}$$

where $o_{\max}$ is the largest eigenvalue of O.

For the proof, see Supplemental Material, Sec. B [21]. The only remaining idealization in Theorem 1 is that Eq. (11) holds exactly, which will not be the case for generic noise. Luckily, this can be circumvented by deploying the array of techniques developed in Ref. [29]. In summary, the experimentally realized channels may satisfy Eq. (11) only approximately (see Supplemental Material, Sec. E [21]).

Let us now discuss the claim of the theorem. Of course, in general, one can have a classical linear response of $\mathcal{O}(g)$. What the main theorem proves is that a $\mathcal{O}(g)$ response, together with the phenomenology described in Eq. (11), cannot be reproduced by classical models. This is because Eq. (11) forces noncontextual models to have a response at most of $\mathcal{O}(g^2)$. A central question is then whether Eq. (11) will be observed in a quantum experiment for g small enough. The next lemma gives a sufficient condition.

Lemma 2: Operational condition test.—Fix $t>0$ and suppose there exists $C>0$ such that the following matrix is positive definite:

$$\tilde{J}_{kj} = 1 - \frac{c_{kj}}{C}, \qquad c_{kj} = (\alpha_k - \alpha_j)^2, \tag{14}$$

where α_k are the eigenvalues of $\int_0^t V_I(t')dt'$. Then Eq. (10) holds for g small enough.

For the proof, see Supplemental Material, Sec. C [21]. Note that to construct $\tilde{J}$ we need only to use linear response operators. For example, in the case of a single qubit

$$\tilde{J} = \begin{bmatrix} 1 & 1-\frac{c_{01}}{C} \\ 1-\frac{c_{01}}{C} & 1 \end{bmatrix},$$

which has eigenvalues $x_1 = c_{01}/C$ and $x_2 = 2 - c_{01}/C$. Hence, for C large enough, one has $\tilde{J} > 0$ for any nondegenerate perturbation ($\alpha_0 \neq \alpha_1$). For a qutrit, there are nontrivial counterexamples to $\tilde{J} > 0$, so one needs to perform the test for the specific scheme under consideration.

The above gives a general method to identify quantum signatures (certified against *arbitrary* noncontextual models) in arbitrary quantum systems in the linear regime:

(1) Compute whether $\tilde{J} > 0$. If that is the case, by carrying out the experiment, one will be able to verify Eq. (11) (using the tools of Supplemental Material, Sec. E [21], to deal with noise and imperfections).

(2) Check whether the response in Eq. (7) is of $\mathcal{O}(g)$.

When the two conditions above are satisfied, Theorem 1 returns a proof of contextuality for g small enough. This algorithm provides a powerful tool to identify quantum signatures. Here, I apply these considerations to quantum thermodynamics and metrology.

A contextual advantage in a quantum engine.—What is the role played by nonclassicality on the performance of thermodynamic devices? Conversely, what is the "thermodynamics of nonclassical properties" required for the understanding of quantum devices in which thermal effects cannot be neglected? The contextuality framework offers the opportunity to rigorously investigate both questions [30,31].

Despite recent theoretical and experimental advances, and a large number of proposals for quantum mechanical heat engines, a central outstanding question remains: Are there thermodynamic machines whose performance *unavoidably* requires quantum effects? The standard comparison with a "stochastic engine" obtained by simple dephasing of the quantum protocol [3–5,8–10] is insufficient to tackle this question. The elementary example discussed in Supplemental Material, Sec. D [21], shows that, in and by itself, the dephasing criterion is not a good notion of nonclassicality. I alternatively suggest to take contextuality as one's notion of nonclassicality and provide an upper bound on the power output of any noncontextual engine. This shows that quantum engines display a power output advantage over *every* noncontextual counterpart [32].

A heat engine is a machine that works between two baths at different temperatures and whose aim is to extract work from the heat flow between the two baths. It is useful to study the functioning of an engine as a sequence of "strokes," in which only some of the elements are involved. While the reasoning is applicable to more general models [5], I focus here on the two-stroke engine:

(1) The first stroke couples subsets of energy levels of the system to a hot and a cold bath to generate a nonequilibrium steady state $\rho(0)$.

(2) The second stroke is a unitary driving to implement work extraction.

We will assume that $\rho(0)$ is a two-level system, as in Ref. [11]. Consider the work extraction process over a unitary cycle lasting an amount of time τ:

$$H(t) = H_0 + gV(t), \qquad V(0) = V(\tau) = 0. \quad (15)$$

If $U(t)$ is the unitary process generated by $H(t)$ from time 0 to t, the work W extracted over the cycle is

$$\begin{aligned} W^{Q} &= \mathrm{Tr}[\rho(0)H_0] - \mathrm{Tr}[U(\tau)\rho(0)U^{\dagger}(\tau)H_0] \\ &= \mathrm{Tr}[\rho(0)H_0] - \mathrm{Tr}[U_I(\tau)\rho(0)U_I^{\dagger}(\tau)H_0]. \end{aligned} \quad (16)$$

Equation (9) returns

$$W^{Q} = \frac{2g\tau}{\hbar}\mathrm{Im}\mathrm{Tr}[\rho(0)XH_0] + \mathcal{O}(g^2), \quad (17)$$

where we set $X \coloneqq (1/\tau)\int_0^{\tau} V_I(t)dt$ (for an interesting relation to the so-called anomalous weak values, see Supplemental Material, Sec. F [21]). Division by τ gives the power of the unitary stroke, which can be $\mathcal{O}(g)$ in the coupling strength. Furthermore, as already noted, the operational equivalence of Eq. (11) is satisfied generically by the unitary driving, since $\rho(0)$ is a qubit system. Hence, setting $E^{\max} = \max_i E_i$, Theorem 1 applies. In every noncontextual model, Eq. (13) holds:

$$W \leq 4E^{\max}p_d \coloneqq W_{\mathrm{NC}}. \quad (18)$$

Hence, $W \leq \mathcal{O}(g^2)$ as $g \to 0$ in any noncontextual model, and the same holds for power. Since we can have $W^{Q} > W^{\mathrm{NC}}$ for g small enough, a quantum advantage in the power output of the work stroke emerges in the weak coupling limit. In fact, the bound relies only on setting a finite upper bound on the maximum energy $E^{\max}$ the noncontextual model can access and not on how it represents H_0, $V(t)$, etc. A gap will emerge at sufficiently small g (or at sufficiently short pulses for fixed g).

The quantum advantage is exhibited in the difference between the $\mathcal{O}(g)$ scaling possible in quantum mechanics as compared with the $\mathcal{O}(g^2)$ bound of any noncontextual model. This proves that the gap analyzed on the basis of the dephasing criterion in Refs. [5,10,11,33] signals a true separation between classical and quantum thermodynamics. Specifically, in the presence of the phenomenology featured in the quantum experiment, the gap unavoidably requires nonclassicality in the form of contextuality.

A contextual advantage in local metrology.—Local metrology is a paradigm to study the ultimate limits of parameter estimation. We look here at the archetypal case of phase estimation, where the relevant parameter is the phase η in the dynamics $U_{\eta} = e^{-iH\eta}$ for some observable H. An initial qubit state $|\psi(0)\rangle$ is prepared, undergoes the dynamics U_{η}, and is measured according to some arbitrary positive operator-valued measure M_x ($M_x \geq 0$, $\sum_x M_x = \mathbb{1}$). After N trials, there exists a measurement such that the error (variance) $\mathrm{Var}(\eta)$ in the estimated phase scales as $\mathcal{O}[1/(4N\Delta H^2)]$, where $\Delta H^2 \coloneqq \langle\psi(0)|H^2|\psi(0)\rangle - [\langle\psi(0)|H|\psi(0)\rangle]^2$. This is finite only if the state is a (nontrivial) superposition of eigenstates of H; otherwise, $\mathrm{Var}(\eta) = +\infty$. Hence, dephasing trivially prevents sensing in this scheme. But what about other noncontextual models, which as already discussed can be much more complex than quantum mechanics plus dephasing? Here I show $\mathrm{Var}(\eta) = +\infty$ in every noncontextual model reproducing the operational phenomenology of quantum sensing.

Let $p(x|\eta)$ be the probability of getting outcome x from a measurement M when the true value of the parameter is η. So $p(x|\eta) = p[x|T_\eta(P), M]$ if P, T_η, and M are the operational descriptions of preparation, transformation, and measurement procedures, represented in quantum theory by $|\psi(0)\rangle$, U_η, and $\{M_x\}$, respectively. Recall that an estimator $\hat{\eta}(x_1, x_2, \ldots)$ maps the measurement outcomes $(x_1, x_2, \ldots)$ to a guess η for the unknown parameter. For independent observations, the variance of any unbiased estimator is lower bounded by $1/(NF_\eta^{P,M})$, with $F_\eta^{P,M}$ the Fisher information

$$F_\eta^{P,M} = \sum_x p(x|\eta)\left[\frac{\partial}{\partial\eta}\ln p(x|\eta)\right]^2. \tag{19}$$

The best strategy involves optimizing over all allowed preparations P and measurements M, where for simplicity we will assume x runs over a bounded, while possibly extremely large, set of indexes. In any OM, from Eq. (1),

$$p(x|\eta) = \int d\lambda d\lambda' \mu_P(\lambda)\mathcal{T}_{T_\eta}(\lambda'|\lambda)\xi_M(x|\lambda'), \tag{20}$$

where $\mu_P(\lambda)$, $\mathcal{T}_{T_\eta}(\lambda'|\lambda)$, and $\xi_M(x|\lambda')$ are the OM descriptions of P, T_η, and M, respectively. Using the relation $p(x|\eta+\delta) = p[x|T_\delta(P_\eta), M]$, where $P_\eta = T_\eta(P)$, and the fact that Eq. (11) is satisfied with $p_d = \mathcal{O}(\delta^2)$, we can prove

$$F_\eta^{P,M} = 0 \tag{21}$$

for any P and any measurement M with a finite number of outcomes and $p(x|\eta) \neq 0$ (see Supplemental Material, Sec. G [21]). Hence, $\mathrm{Var}(\eta) = +\infty$, as anticipated. This again is a consequence of the weakness of linear response in noncontextual models.

Outlook.—I proved that the linear response of quantum systems driven by small external perturbations has a scaling that unavoidably requires nonclassicality. While the quantum response can scale linearly in the strength of the perturbation parameter g, noncontextual models reproducing the operational phenomenology in Eq. (11) respond only quadratically. [Curiously, one can note that, since classical models display a quadratic response in the presence of the operational equivalence in Eq. (11), a phenomenon such as the quantum Zeno effect is naturally expected on the basis of the assumption of noncontextuality—see Supplemental Material, Sec. H [21]].

The $\mathcal{O}(g)$ vs $\mathcal{O}(g^2)$ gap is a certifiable quantum signature highlighting a central dynamical differences between noncontextual models and quantum mechanics. I gave readily applicable tools to analyze arbitrary linear response experiments. As an application, I showed that the improved performance in the power output of a quantum engine *necessarily requires* nonclassical effects in the form of contextuality.

Building up on this work, it will be desirable to use the tools introduced here to reanalyze in detail the experimental heat engine signature of Ref. [11], as well as to develop flexible certification tools applicable to larger-scale systems and based on more compelling operational constraints than "stochastic reversibility." While I presented an example of a performance boost unavoidably connected to contextuality, this does not settle the question of the superiority of quantum engines as practical devices. The latter requires discussing issues of scalability, control, and efficiency in the implementation of basic operations. One can also envision this work as a first result in the "thermodynamics of contextuality," meaning how a central property signaling the departure from classical reality interacts with actual thermodynamic processes at the operational level.

I am indebted to Matt Pusey, Ravi Kunjwal, Kavan Modi (who asked about the Zeno effect), Antonio Acin, Mohammad Mehboudi, Felix Huber, Luis García-Pintos, Gabriel Senno, Amikam Levy, Joseph Bowles, David Jennings (who proposed the "quantum toaster criterion"), Nicole Yunger Halpern, David Arvidsson Shukur, Mauro Paternostro, and Alessio Belenchia for helpful discussions. I acknowledge financial support from the EU Marie Sklodowska-Curie individual fellowships (H2020-MSCA-IF-2017, GA794842), Spanish MINECO (Severo Ochoa SEV-2015-0522, Project No. QIBEQI FIS2016-80773-P), Fundacio Cellex, and Generalitat de Catalunya (CERCA program and SGR 875), ERC Grant EQEC No. 682726.

[*]lostaglio@protonmail.com

[1] M. O. Scully, K. R. Chapin, K. E. Dorfman, M. B. Kim, and A. Svidzinsky, Proc. Natl. Acad. Sci. U.S.A. **108**, 15097 (2011).

[2] N. Killoran, S. F. Huelga, and M. B. Plenio, J. Chem. Phys. **143**, 155102 (2015).

[3] J. B. Brask and N. Brunner, Phys. Rev. E **92**, 062101 (2015).

[4] M. T. Mitchison, M. P. Woods, J. Prior, and M. Huber, New J. Phys. **17**, 115013 (2015).

[5] R. Uzdin, A. Levy, and R. Kosloff, Phys. Rev. X **5**, 031044 (2015).

[6] A. Ü. C. Hardal, N. Aslan, C. M. Wilson, and Ö. E. Müstecaplıoğlu, Phys. Rev. E **96**, 062120 (2017).

[7] K. E. Dorfman, D. Xu, and J. Cao, Phys. Rev. E **97**, 042120 (2018).

[8] A. Levy and D. Gelbwaser-Klimovsky, in *Thermodynamics in the Quantum Regime* (Springer, New York, 2018), pp. 87–126, https://doi.org/10.1007/978-3-319-99046-0_4.

[9] P. A. Camati, J. F. G. Santos, and R. M. Serra, Phys. Rev. A **99**, 062103 (2019).

[10] R. Dann and R. Kosloff, New J. Phys. **22**, 013055 (2020).

[11] J. Klatzow, J. N. Becker, P. M. Ledingham, C. Weinzetl, K. T. Kaczmarek, D. J. Saunders, J. Nunn, I. A. Walmsley, R. Uzdin, and E. Poem, Phys. Rev. Lett. **122**, 110601 (2019).

[12] S. Nimmrichter, J. Dai, A. Roulet, and V. Scarani, Quantum **1**, 37 (2017).

[13] J. O. González, J. P. Palao, D. Alonso, and L. A. Correa, Phys. Rev. E **99**, 062102 (2019).

[14] R. W. Spekkens, Phys. Rev. A **71**, 052108 (2005).

[15] *Operational* here refers to statements that can be, in principle, verified from the conditional probabilities of observing certain outcomes for given input measurements, without relying on a detailed model of the inner workings of the devices at hand. The terminology refers to the "black-box" approach successfully applied in quantum foundations [14], device-independent quantum cryptography, and self-testing [16].

[16] I. Šupić and J. Bowles, Quantum **4**, 337 (2020).

[17] D. R. Arvidsson-Shukur, N. Y. Halpern, H. V. Lepage, A. A. Lasek, C. H. Barnes, and S. Lloyd, Nat. Commun. **11**, 3775 (2020).

[18] M. Leifer, Quanta **3**, 67 (2014).

[19] Another property that is typically required is locality, but since here we are not dealing with spacelike separated systems such an assumption is inconsequential.

[20] R. W. Spekkens, arXiv:1909.04628.

[21] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.125.230603 for further discussions and technical proofs.

[22] R. W. Spekkens, Phys. Rev. A **75**, 032110 (2007).

[23] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, Phys. Rev. A **86**, 012103 (2012).

[24] D. Schmid and R. W. Spekkens, Phys. Rev. X **8**, 011015 (2018).

[25] D. Saha, P. Horodecki, and M. Pawłowski, New J. Phys. **21**, 093057 (2019).

[26] M. Lostaglio and G. Senno, Quantum **4**, 258 (2020).

[27] A. Tavakoli and R. Uola, Phys. Rev. Research **2**, 013011 (2020).

[28] A channel is a completely positive and trace-preserving map, describing the most general quantum evolution of a system.

[29] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, and R. W. Spekkens, Nat. Commun. **7**, 11780 (2016).

[30] K. V. Hovhannisyan and A. Imparato, New J. Phys. **21**, 052001 (2019).

[31] A. Levy and M. Lostaglio, PRX Quantum **1**, 010309 (2020).

[32] The dephasing criterion claims an advantage when the quantum protocol outperforms its dephased version. Note that the dephased protocol is a special case of a noncontextual model, so my proposed criterion is more demanding than the dephasing criterion.

[33] L. P. García-Pintos, A. Hamma, and A. Del Campo, Phys. Rev. Lett. **125**, 040601 (2020).