Counterfactual quantum sensing: What interaction-free measurement can and cannot buy

Christoph F. Wildfeuer*
FHNW University of Applied Sciences and Arts Northwestern Switzerland, School of Engineering and Environment, Klosterzelgstrasse 2, CH-5210 Windisch, Switzerland

(Received 29 July 2026; accepted 6 November 2026; published 20 November 2026)

Interaction-free measurement infers the presence of an absorbing object from a photon that, in the counterfactual sense, never interacted with it, and is widely described as a route to minimally invasive sensing. We ask what it actually buys, in estimation-theoretic terms. Writing the Elitzur-Vaidman interferometer as a channel-estimation problem, its Fisher information for the transmissivity T of an object in one arm is $[2T(1-T)]^{-1}$, exactly half that of direct transmission probing, and the two schemes deliver identical Fisher information per absorbed photon, $[T(1-T)^2]^{-1}$: for measuring how transparent something is, the interferometer buys nothing. The advantage lies in discrimination, and we show that what it requires is not that the object be opaque but that the competing hypothesis be the object's absence. Against empty space the Chernoff information per absorbed photon grows as $(8/\pi^2)[(1-\sqrt{T})/(1+\sqrt{T})]N\ln N$ in the number of Zeno cycles N, even for a weakly absorbing object, while between two partial transparencies it does not grow at all. Parasitic per-cycle loss ϵ caps the advantage: the conclusive interrogations per absorbed photon peak at $(0.2625/\epsilon)(1-\sqrt{T})/(1+\sqrt{T})$, attained at $N_{\mathrm{opt}} = 1.5936/\epsilon$ for an opaque object, which identifies the loss per cycle as the figure of merit governing how far interaction-free sensing can be pushed. Finally, the negative result is not special to the interferometer: for a single photon meeting a memoryless, nondispersive object any number of times through arbitrary fixed optics, the accessible quantum Fisher information never exceeds that of the same incident flux spent on independent single-pass probes, and is generically far below it. This recovers the bound of Massar, Mitchison, and Pironio for this class, by a short argument that also identifies when it is tight.

DOI: 10.1103/k4w9-8xwk

I. INTRODUCTION

Interaction-free measurement is the observation that an absorbing object placed in one arm of an interferometer alters the output statistics even in runs in which it absorbs nothing. Elitzur and Vaidman made the point vivid with a bomb that any single photon would detonate, and showed that its presence can nonetheless be certified with finite probability by a photon that was never absorbed [1]. Kwiat *et al.* realized the effect in the laboratory [2], then pushed its efficiency towards unity [3]. Their method was to interrogate the object through a chain of weakly rotating cycles: each pass presents only a small amplitude to the object, and the quantum Zeno effect suppresses the absorption. Whether such a measurement is properly called counterfactual has been debated at length [4,5]; nothing below depends on the answer.

The practical appeal is that a fragile sample might be interrogated without being damaged, and the idea has been carried into imaging [6–8], spectroscopy [9], and x-ray interrogation [10], and beyond optics into the microwave domain, where a coherent three-level protocol in a superconducting circuit detects a pulse interaction free [11] and has been proposed for noise detection [12]. The microwave protocol uses a driven three-level system rather than a single photon in fixed optics, so it lies outside the class of Sec. V; the dividing line of Sec. III between presence-or-absence and everything else should apply to it as well. It is also often called supersensitive. In quantum optical metrology that word has a settled technical meaning: a phase sensitivity $\delta\varphi$ that scales better than the shot-noise limit $1/\sqrt{\bar{n}}$ of $\bar{n}$ independent photons, down to the Heisenberg limit $1/\bar{n}$, and it is kept distinct from superresolution, the narrowing of fringes without a gain in sensitivity [13]. In the interaction-free literature it does two jobs at once. It can mean that the scheme extracts more information per probe, or that it deposits less energy in the sample. The two claims are independent, and only one of them survives. Careful versions say as much: Miller *et al.* obtain sub-shot-noise phase estimation in an induced-coherence interferometer, and state plainly that the gain comes from squeezing and not from the geometry of the interferometer [14].

Much of this ground was mapped early. Massar, Mitchison, and Pironio derived bounds relating the number of absorbed photons to the sensitivity of an absorption measurement. Their class of strategies is a wide one, admitting arbitrary photon number, ancillas, and interaction steps. They concluded that "using the quantum mechanical properties of light does not offer significant advantages over more traditional schemes such as counting absorbed photons or simple interferometric procedures that use one photon at a time" [15]. Mitchison

*Contact author: christoph.wildfeuer@fhnw.ch

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

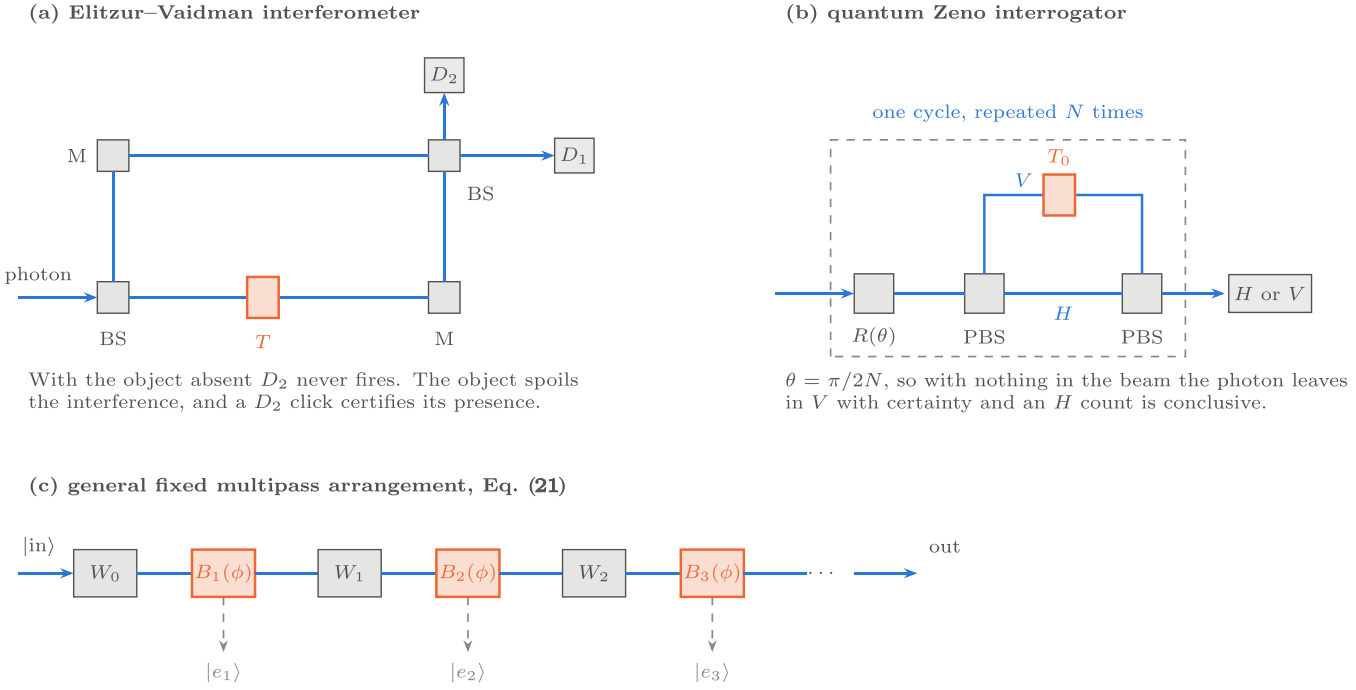


FIG. 1. The three arrangements considered. (a) is the interferometer of Sec. II, (b) the Zeno interrogator of Secs. III and IV, and (c) the general class covered by the bound of Sec. V, of which the first two are special cases.

and Massar showed separately that two partial transparencies cannot be discriminated absorption free, while complete transparency can be told from any partial one [16]; with Pironio they went on to bound the photon cost of the harder case [17]. Facchi *et al.* compared standard and quantum Zeno absorption tomography and found that, expressed per absorbed particle, the two Cramér-Rao bounds "reduce to the same bound" [18].

Closest to the present work, Thomas *et al.* studied semitransparent objects in interaction-free measurement by counting probe particles [19]. They reached both of the conclusions reached here, by other means: that for determining an unknown transparency such measurements "lose as many probe particles as classical measurements or more," and that "interaction-free measurements with zero loss are possible if one of the samples is perfectly transparent."

What this paper adds is quantitative. Section II treats the Elitzur-Vaidman interferometer as an estimation problem: it is a factor of two worse than direct probing, and exactly as good once we count the damage done. Section III locates the advantage that survives, in discrimination rather than estimation, and gives in closed form the rate at which it grows and the condition it requires. Section IV gives the limit that parasitic loss places on that advantage, also in closed form; this is the result an experiment can be designed against. Section V then shows that the negative result of Sec. II is no accident of one interferometer but holds for every fixed multipass arrangement of a single photon, which recovers the bound of Ref. [15] for this class by a short argument and identifies when it is tight.

II. INTERACTION-FREE MEASUREMENT AS A PARAMETER-ESTIMATION PROBLEM

Let the object have intensity transmissivity $T \in [0, 1]$ and amplitude transmission $t = \sqrt{T}$, with $T = 1$ corresponding to no object at all. A single photon enters one input port of the first beam splitter, $|\text{in}\rangle = |1\rangle_a|0\rangle_b$, and all probabilities below are per such photon. Writing each beam splitter as $\frac{1}{\sqrt{2}}\left(\begin{smallmatrix}1 & i\\ i & 1\end{smallmatrix}\right)$, the outcome probabilities of the Mach-Zehnder interferometer of Fig. 1(a) are

$$p_1(T) = \tfrac{1}{4}(1+\sqrt{T})^2, \tag{1}$$

$$p_2(T) = \tfrac{1}{4}(1-\sqrt{T})^2, \tag{2}$$

$$p_{\text{abs}}(T) = \tfrac{1}{2}(1-T), \tag{3}$$

where p_1 and p_2 refer to the bright and dark ports and p_{abs} is the probability that the object absorbs the photon. For $T = 0$ these reduce to $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$, recovering the Elitzur-Vaidman case, and for $T = 1$ to (1,0,0).

The relevant figure of merit is the Fisher information about T carried by one interrogation, $F(T) = \sum_i p_i^{-1}(\partial_T p_i)^2$. Substituting Eqs. (1)–(3), the $(1 \pm \sqrt{T})^2$ factors cancel and

$$F_{\text{MZI}}(T) = \frac{1}{4T} + \frac{1}{4T} + \frac{1}{2(1-T)} = \frac{1}{2T(1-T)}. \tag{4}$$

Compare the obvious classical alternative: send the photon straight at the object and record whether it survives. That outcome is Bernoulli with success probability T, so

$$F_{\mathrm{dir}}(T) = \frac{1}{T(1-T)}, \tag{5}$$

which also saturates the quantum Fisher information for a single-photon probe of a pure-loss channel, since T enters only through the populations of the {transmitted, lost } basis. Hence

$$F_{\mathrm{MZI}}(T) = \tfrac{1}{2} F_{\mathrm{dir}}(T) \tag{6}$$

for every T: as an estimator of transmissivity the interferometer sits a factor of two below direct probing. The reason is simple, since only half the amplitude ever visits the object. It is worth saying out loud anyway, because it is the opposite of what "supersensitivity" suggests.

The natural rebuttal is that the comparison is unfair: the point of IFM is that it absorbs fewer photons. That is true, and it is also exactly compensated. Throughout, n_{abs} denotes the mean number of photons the object absorbs in one interrogation, whatever the scheme. Here the object absorbs $n_{\mathrm{MZI}} = (1-T)/2$ photons per interrogation in the interferometer and $n_{\mathrm{dir}} = (1-T)$ in the direct scheme, so

$$\frac{F_{\mathrm{MZI}}}{n_{\mathrm{MZI}}} = \frac{F_{\mathrm{dir}}}{n_{\mathrm{dir}}} = \frac{1}{T(1-T)^2}. \tag{7}$$

The information gained per unit of damage done is identical, as Fig. 2 shows. The interferometer buys its lower mean number of absorbed photons, $n_{\mathrm{MZI}} = n_{\mathrm{dir}}/2$, at precisely the exchange rate one would expect, and nothing more.

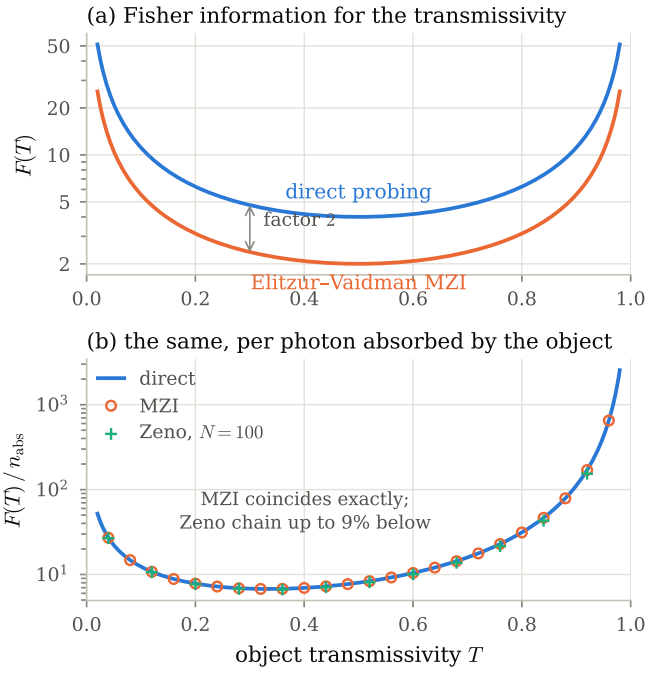


FIG. 2. The central negative result. (a) Fisher information about the transmissivity T of an object placed in one arm of the interferometer, Eq. (4), against direct transmission probing, Eq. (5). The interferometer sits a factor of two below the direct scheme at every T. (b) The same quantities normalized by the number of photons the object absorbs. The Elitzur-Vaidman interferometer (circles) falls exactly on the direct-probing curve $[T(1-T)^2]^{-1}$ (solid). The 100-cycle Zeno chain (crosses) lies between 1 and 9% below it. That shortfall is real and not a defect of the measurement: the classical Fisher information plotted here equals the quantum Fisher information of the reduced photon state, so the $\{H, V, \mathrm{lost}\}$ measurement is already optimal. Cycling discards information, as the bound of Sec. V allows it to.

III. ZENO CHAIN AND WHERE THE ADVANTAGE DOES LIVE

The high-efficiency scheme of Kwiat *et al.* [3] cycles the photon N times through a weakly rotating element, the object acting on one polarization component each pass [Fig. 1(b)]. With rotation angle $\theta = \pi/(2N)$, one cycle is

$$A = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{T} \end{pmatrix} \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}, \tag{8}$$

acting on the polarization amplitudes, and the state after the chain is $A^N(1,0)^{\mathsf{T}}$. We write P_H and P_V for the probabilities that the photon leaves unrotated or rotated, and P_{abs} for the probability that the object absorbs it. With the object absent ($T = 1$) the accumulated rotation is $N\theta = \pi/2$ and the photon emerges rotated with certainty. With an opaque object ($T = 0$) the rotated amplitude is removed at every pass and the photon survives unrotated with probability

$$P_{\mathrm{IFM}}(N) = \cos^{2N}\left(\frac{\pi}{2N}\right) = 1 - \frac{\pi^2}{4N} + \mathcal{O}(N^{-2}), \tag{9}$$

so the photons deposited in the object fall as $\pi^2/(4N)$ and the ratio of conclusive interrogations to absorbed photons grows without bound,

$$\mathcal{R}(N) \equiv \frac{P_{\mathrm{IFM}}}{P_{\mathrm{abs}}} = \frac{4N}{\pi^2} - \frac{1}{2} + \mathcal{O}(N^{-1}), \tag{10}$$

which we confirm numerically: $\mathcal{R} = 40.03$ at $N = 100$ and $\mathcal{R} = 404.78$ at $N = 1000$, against $4N/\pi^2 - \frac{1}{2} = 40.03$ and 404.78. Both limits are plotted in Fig. 3.

Equation (10) is often read as the supersensitivity the literature promises, and it cannot be. The information-per-damage exchange rate for estimating a transmissivity is fixed for the interferometer, Eq. (7), and bounded for every multipass arrangement, as Sec. V shows. What Eq. (10) describes is a different task, discriminating between two hypotheses about the object, and it is only there that the chain wins. Anisimov *et al.* recast Zeno interrogation in exactly these terms, as a statistical test between a clear path and a blocked one [20], and Zhou and Yung analyzed it as a channel-discrimination problem at finite cycle number, finding that entangled inputs give no advantage over single photons [21]. Hance, Matsushita and Hofmann quantify the information that the output statistics of a multipath interferometer carry about the presence or absence of an absorber, and trace the quantum part of the gain to the backaction of the absorber on the amplitudes in the paths it does not occupy [22]. That analysis says where the gain comes from; the present one prices it, per absorbed photon, and finds the same dividing line between presence or absence and everything else.

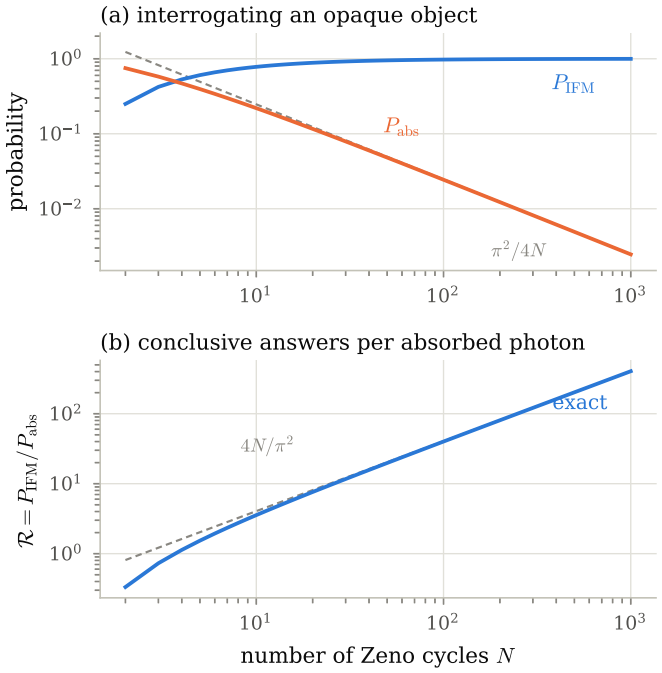


FIG. 3. Quantum Zeno interrogation of an opaque object. (a) The probability P_{IFM} of a conclusive interaction-free answer, Eq. (9), and the probability P_{abs} that the object absorbs the photon; the dashed line is the asymptote $\pi^2/4N$. (b) Their ratio $\mathcal{R}$, the number of conclusive interrogations per absorbed photon, approaching $4N/\pi^2$ (dashed) from below by the constant $\frac{1}{2}$ of Eq. (10). The growth is unbounded in a lossless chain.

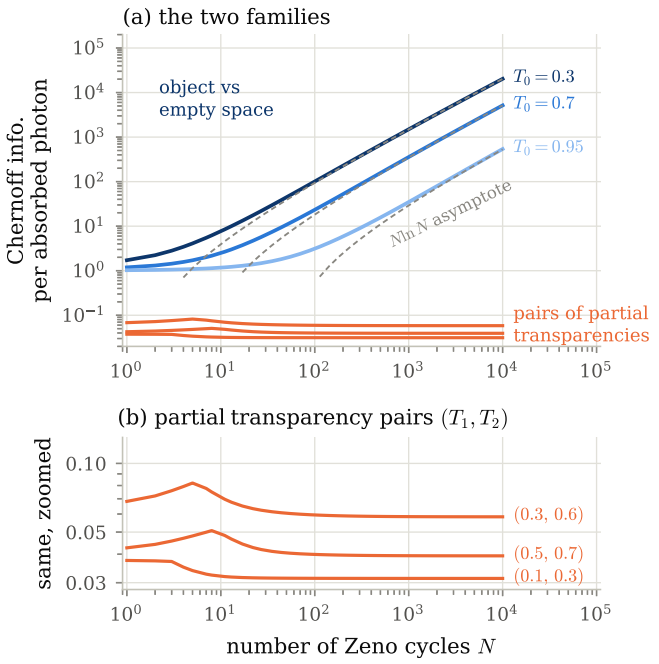


FIG. 4. What the discrimination advantage requires. Chernoff information for an N-cycle chain, normalized by the photons the object absorbs under whichever of the two hypotheses absorbs more. (a) Blue: an object of transmissivity T_0 interrogated against empty space; dashed lines are the closed-form asymptote of Eq. (14), which the exact curves join once N exceeds a few tens. The rate grows as $N \ln N$ for every T_0 shown, none of which is close to opaque. Orange: two partial transparencies interrogated against each other; the rate does not grow. The case $T_0 = 0$ is omitted from the blue family because the two outcome distributions then have disjoint support, a single interrogation decides, and the Chernoff information is infinite. (b) The orange family on its own scale, where the three pairs (T_1, T_2) can be separated and their flatness in N checked; note that the pair (0.3,0.6) outperforms (0.5,0.7), so the rate is set by the separation of the two candidates and not by how much either absorbs.

Equation (9) was derived at $T = 0$, so it is natural to credit the win to the object being opaque. That is wrong, and the error matters: it would restrict the advantage to far fewer samples than actually enjoy it. What the win really needs is that one of the two hypotheses be the object's absence.

The reason is best put in old-fashioned terms: the chain is a null experiment. With nothing in the beam the apparatus is tuned so that one detector never fires—the N rotations sum to exactly $\pi/2$ and the photon leaves in $|V\rangle$ with certainty, so $P(H \,|\, \text{absent}) = 0$ exactly, not merely small. A count in the H port is then not evidence but proof. What the Zeno effect contributes is that this null can be made deep while the sample is barely touched: the absorbed fraction falls as $\pi^2/4N$ while the null stays exact, so a conclusive click costs almost no damage. Note where that zero comes from. It is a property of the empty apparatus, not of the object, which is why the advantage survives at $T_0 = 0.95$ in Fig. 4. Any object at all spoils the perfect rotation and leaks amplitude into the forbidden port. It does not have to be opaque; it only has to be there.

Estimation has no null. Asking how transparent an object is means distinguishing T from $T + \mathrm{d}T$, and every outcome then carries nonzero probability under both hypotheses. Nothing is forbidden, so nothing is ever proof, and each run contributes a statistical increment rather than a certainty—which is to say one is back to counting photons, the currency that Eq. (7) and the bound of Sec. V govern. Two partial transparencies are in the same position: with T_1 and T_2 both below unity, both leak into both ports, and that is why the family in Fig. 4(b) is flat. In information-theoretic terms an outcome of probability zero under one hypothesis makes the divergence between the two distributions unbounded, so a single trial can carry unbounded information; once both distributions have full support every divergence is finite and the information per trial is capped.

Figure 4 makes the separation quantitative. We compare two hypotheses by their Chernoff information,

$$C = -\ln \min_{0 \leqslant s \leqslant 1} \sum_i p_i^s \, q_i^{1-s}, \tag{11}$$

where p_i and q_i are the outcome probabilities under the two hypotheses; the error probability of the optimal test on M independent interrogations decays as e^{-MC}. We normalize it as before by the photons the object absorbs—specifically by the absorption under whichever hypothesis absorbs more, since that is the damage the experiment must budget for. Against empty space, an object of transmissivity T_0 gives a rate that grows without bound for every $T_0 < 1$. This includes $T_0 = 0.95$, which absorbs five percent of what reaches it and is in no sense opaque. Against another partial transparency T_1

versus T_2, both below unity, the rate does not grow at all, even when one of the two candidates is perfectly opaque. What the advantage tracks is conclusiveness, not opacity.

The growth law can be given in closed form, and it is worth doing so because a fitted exponent is misleading here. With nothing in the beam the photon leaves in $|V\rangle$ with certainty, so V is the only outcome with support under both hypotheses and the sum in Eq. (11) reduces to the single term $P_V^s \times 1^{1-s}$, which is minimized at the boundary $s = 1$: the Chernoff exponent collapses to its boundary value, and $C = -\ln P_V$ exactly. The chain settles to a steady V amplitude $\sqrt{T}\,\theta/(1-\sqrt{T})$ before each encounter, giving

$$P_V \longrightarrow \frac{T\pi^2}{4N^2(1-\sqrt{T})^2}, \tag{12}$$

$$n_{\mathrm{abs}} \longrightarrow \frac{\pi^2}{4N}\,\frac{1+\sqrt{T}}{1-\sqrt{T}}, \tag{13}$$

so that $C = 2\ln N + \ln[4(1-\sqrt{T})^2/T\pi^2]$ and

$$\frac{C}{n_{\mathrm{abs}}} \longrightarrow \frac{8}{\pi^2}\,\frac{1-\sqrt{T}}{1+\sqrt{T}}\,N\ln N. \tag{14}$$

The rate grows as $N\ln N$, marginally faster than linearly. Both prefactors check out: at $N = 10^8$ we find $Nn_{\mathrm{abs}} = 8.447,\ 27.74,\ 192.42$ for $T = 0.3,\ 0.7,\ 0.95$ against the $8.444,\ 27.74,\ 192.43$ of Eq. (13), and the additive constant in C agrees to three decimals or better. We stress the closed forms rather than a fit because the subleading constant in C is large and negative, so a log-log slope measured over any accessible range of N reads about 1.1 rather than the asymptotic 1; the exponent is better read off Eqs. (12)–(13) than from the curves.

The null is exact only in the model, and it is worth saying what survives when it is not. Real polarizing optics have a finite extinction ratio; Kwiat *et al.* report crosstalk of order one percent [3]. Let a fraction δ of the amplitude leak into the wrong port, so that $P(H\,|\,\text{absent}) = \delta$ rather than zero. Two things change, and only one of them matters. An H count is no longer proof: it carries a likelihood ratio $P(H\,|\,\text{present})/\delta$, close to $1/\delta$ for a good chain—about 10^2 at one percent crosstalk and 10^3 at one part in 10^3. Those are odds of a hundred to one from a single click, which is not certainty but is not far from it. Second, the Chernoff information no longer grows without bound. For this model of the imperfection we find numerically that it saturates at

$$C_\infty = -\ln[2\sqrt{\delta(1-\delta)}] \xrightarrow{\delta\ll 1} \ln\frac{1}{2\sqrt{\delta}}, \tag{15}$$

which the exact chain approaches from above, agreeing to five significant figures at $N = 10^6$ for $\delta = 10^{-2}, 10^{-3}$, and 10^{-4} and for $T_0 = 0,\ 0.3$, and 0.7 alike.

The constant in Eq. (15) is specific to output-port extinction, in which a fraction δ is exchanged between the ports, and should not be read as universal. Imperfections that leave $P(V\,|\,\text{present}) \to 0$—a per-cycle retardance error, or a dark count in the H channel—saturate instead at $\ln(1/\delta)$, three times larger at $\delta = 10^{-2}$. Finite detection efficiency is a separate and often tighter ceiling, and one the null argument depends on more than it depends on crosstalk: $C = -\ln P_V$ requires the no-click channel as well as the H channel to be forbidden when the object is absent, which needs unit efficiency. At efficiency η the exponent saturates at $\ln[1/(1-\eta)]$, which is 4.61 at $\eta = 0.99$ and 2.30 at $\eta = 0.9$.

What all four models share is the important part. The rate per absorbed photon continues to grow linearly in N; only the logarithm freezes, $\ln N$ being replaced by a constant set by whichever imperfection dominates. The damage advantage, which is the useful part, survives imperfect optics and imperfect detectors alike. The logarithmic bonus does not, and at one percent crosstalk it is gone by N of order ten for $T_0 = 0.3$, though only near $N \simeq 300$ for $T_0 = 0.95$, where the ideal chain has its own transient of $(1-\sqrt{T_0})^{-1} \simeq 40$ cycles to get through first.

Neither half of this is new as a statement about particles. Mitchison and Massar proved the corresponding result for absorption-free discrimination—complete transparency can be told from any partial transparency, two partial transparencies cannot [16]—and Thomas *et al.* reached the same conclusion by counting probe particles through a semitransparent sample [19]. What Fig. 4 adds is the rate at which the two cases separate, in a quantity that bounds every measurement rather than a particular one. The flat family in that figure is the case whose photon cost Mitchison, Massar, and Pironio bounded [17]. This also explains why estimation is the harder problem. Estimating a continuous T means discriminating between neighboring partial transparencies, and that is precisely the case with no conclusive outcome, where the bound of Sec. V binds.

The discrimination task is naturally posed in Bayesian terms, and the single-cycle case is instructive. Let B be the event that the object is present, $\overline{B}$ its absence, and D_i a click in port i. Suppose an opaque object is present with prior probability $P(B) = \frac{1}{2}$. From Eqs. (1)–(3) at $T = 0$, $P(D_i|B) = \frac{1}{4}$ for each port, while $P(D_1|\overline{B}) = 1$ and $P(D_2|\overline{B}) = 0$. Hence $P(D_1) = \frac{5}{8}$ and $P(D_2) = \frac{1}{8}$, and

$$P(B|D_1) = \tfrac{1}{5}, \quad P(B|D_2) = 1. \tag{16}$$

A click in the dark port is conclusive. A click in the bright port is inconclusive, but it is not uninformative: it moves the posterior from $\frac{1}{2}$ down to $\frac{1}{5}$. Informal accounts often miss that second point, though the numbers of course depend on the prior assumed. In an N-cycle chain the conclusive outcome still carries posterior unity, while the absorbed fraction falls as $\pi^2/4N$. That is the content of Eq. (10).

The supersensitivity claim can now be stated precisely, and it depends entirely on the question asked. Ask how transparent an object is, and interaction-free measurement gains nothing: the information per absorbed photon is the same for every scheme in the family. Ask whether the object is there at all, and it gains a factor that grows as $N\ln N$, however weakly the object absorbs. Ask which of two semitransparent objects sits in the beam, and it gains nothing again.

The practical rule follows. Use interaction-free measurement to protect a fragile sample while checking whether something is present. Do not use it to measure how much is present, or to tell one semitransparent sample from another. Interaction-free polarimetry [23], which turns the Zeno chain towards measuring the polarization properties of a sample, sits on this boundary and is worth a remark. A nonabsorbing

birefringent sample imprints a unitary rotation, and a unitary parameter is like a phase: it accumulates coherently over passes, so the per-dose bound of Sec. V, which is about the loss parameter T, does not apply to it and multipassing can help, as it does for phase [24]. Only the absorptive part of a dichroic sample is covered by Eq. (26), and for that part the rule above stands.

IV. LOSS-LIMITED OPTIMUM

Equation (10) grows without bound only in a lossless chain. A lossy Zeno interrogator was studied by Anisimov *et al.*, who placed a beam-splitter loss in the interrogated arm and reported a nonmonotonic dependence of the transmission probability on it [20]. They did not extract an optimal cycle number, and did not compare conclusive detections with absorbed photons; both follow in closed form below. Let each cycle carry a survival probability $1-\epsilon$ from imperfect optics—mirrors, rotator, polarizing beam splitters—that act on both polarization components alike and are independent of the object; where the loss sits matters, and the unbalanced case is taken up at the end of this section. A photon lost this way at cycle j is also never absorbed by the object, so both channels are attenuated: the conclusive one over the whole chain, the absorbed one only up to the cycle where absorption happens. With $u=(1-\epsilon)\cos^2\theta$,

$$\mathcal{R}_\epsilon(N)=\frac{u^N(1-u)}{\sin^2\theta\,\left(1-u^N\right)}\;\simeq\;\frac{4N}{\pi^2}\,\frac{x}{e^x-1},\qquad x=N\epsilon. \tag{17}$$

Maximizing the large-N form gives the transcendental condition $2(e^x-1)=x\,e^x$, whose root is $x^\star=1.59362$, so that

$$N_{\rm opt}=\frac{x^\star}{\epsilon}=\frac{1.5936}{\epsilon},$$
$$\mathcal{R}_{\rm max}=\frac{4}{\pi^2}\,\frac{(x^\star)^2}{e^{x^\star}-1}\,\frac{1}{\epsilon}=\frac{0.26247}{\epsilon}. \tag{18}$$

Maximizing the exact expression in Eq. (17) over integer N reproduces this over four decades: for $\epsilon=10^{-2},10^{-3},10^{-4},10^{-5}$ we find $N_{\rm opt}=160$, 1594, 15937, 159362 and $\mathcal{R}_{\rm max}=25.73$, 261.95, 2624.1, 26246, against the predicted 159, 1594, 15936, 159362 and 26.25, 262.47, 2624.7, 26247. The extension to a partially transmitting object is asymptotic rather than exact, and the condition matters. Because the parasitic factor $\sqrt{1-\epsilon}\,\mathbb{1}$ is a scalar on the two-mode space it commutes with the rotations and with the object, so the null of Sec. III is untouched, $P(H\,|\,\text{absent})=0$ still, and

$$P_H^\epsilon(N)=(1-\epsilon)^N P_H^{(0)}(N),\qquad n_{\rm abs}^\epsilon(N)=\sum_{k=1}^{N}(1-\epsilon)^{k-1}d_k \tag{19}$$

exactly, with d_k the lossless per-cycle absorption. Replacing $\pi^2/4$ by the constant of Eq. (13) then gives

$$N_{\rm opt}\simeq\frac{x^\star}{\epsilon},\qquad \mathcal{R}_{\rm max}\simeq\frac{0.26247}{\epsilon}\,\frac{1-\sqrt{T_0}}{1+\sqrt{T_0}}, \tag{20}$$

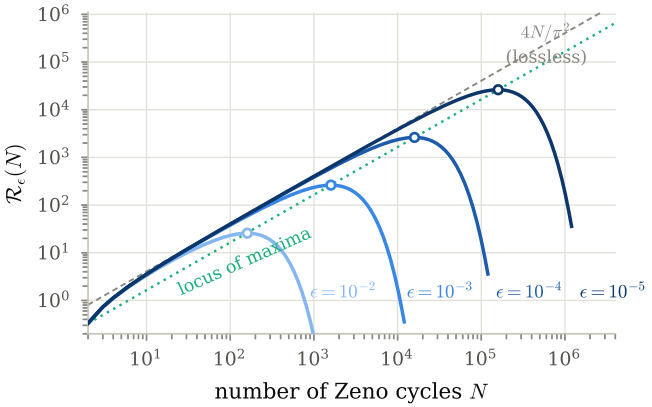


FIG. 5. The loss-limited optimum. Conclusive interrogations per absorbed photon, $\mathcal{R}_\epsilon(N)$ of Eq. (17), for four values of the parasitic per-cycle loss ϵ. Open circles mark the maxima, located on the integer grid; the dotted line is their locus, reached at $N_{\rm opt}=x^\star/\epsilon$, and the dashed line is the lossless asymptote $4N/\pi^2$. Cycling past $N_{\rm opt}$ is actively harmful: the chain leaks the photon before it can answer.

but only when d_k has reached its steady value and $P_H^{(0)}\to 1$, both of which require $N\gg(1-\sqrt{T_0})^{-1}$, that is $\epsilon\ll 1-\sqrt{T_0}$. Outside that regime the optimum moves. Exhaustive maximization of the exact chain over integer N gives so that

ϵ	$T_0=0$	0.3	0.7	0.95	0.99
10^{-2}	160	161	165	195	230
10^{-3}	1594	1595	1599	1636	1802

at $\epsilon=10^{-2}$ the peak has moved by 44% by $T_0=0.99$. The departure is governed by $\mu=\epsilon/(1-\sqrt{T_0})$: $N_{\rm opt}$ exceeds $x^\star/\epsilon$ by 22% at $\mu\simeq 0.4$ and by 44% at $\mu\simeq 2$. Equation (20) is therefore the small-μ limit, and the leading term $N_{\rm opt}=x^\star/\epsilon$ is exact only for an opaque object, where the earlier table is reproduced to the unit digit.

Figure 5 shows the four curves and the locus their maxima trace out for an opaque object. The practical message is plain: the loss per cycle sets the damage advantage of interaction-free sensing, and nothing else does. Improve ϵ by a factor of ten and you protect the sample ten times better. Note also that Eq. (20) gives $\mathcal{R}_{\rm max}<1$—fewer conclusive interrogations than absorbed photons, hence no advantage at all—whenever $(1-\sqrt{T_0})/(1+\sqrt{T_0})<\epsilon/0.26247$; at $\epsilon=10^{-2}$ and $T_0=0.95$ the exact chain gives $\mathcal{R}_{\rm max}=0.32$. Multipass electron microscopy runs into the same tradeoff, where the number of passes is limited by loss per pass rather than by the physics of the interrogation [25,26].

Where the loss sits

The null $P(H\,|\,\text{absent})=0$ is a property of the empty apparatus. It survives an imperfection only if that imperfection acts as a scalar on the two polarization modes; a loss that treats H and V differently breaks it. Three cases make this concrete.

Balanced loss, the case treated above, is such a scalar: the factor $\sqrt{1-\epsilon}\,\mathbb{1}$ commutes with everything, an H count

remains proof of presence, and the loss only makes such proofs rarer. That is the origin of the finite N_{opt}.

Loss in the object branch is different in kind. A per-cycle loss L there is indistinguishable from an object of transmissivity $1-L$, so the empty apparatus is no longer empty, and presence versus absence becomes the discrimination of $T_0(1-L)$ from $1-L$: the two-partial-transparencies problem of Fig. 4(b). The Zeno effect that protects the object now works against the measurement, freezing the photon in H even when only the loss is present. For $L=10^{-2}$ and no object, $P(H\,|\,\mathrm{absent})=0.020$, 0.12, and 0.43 at $N=100$, 300, and 1000, so beyond $N\simeq 1/L$ an H count carries almost no evidence. This ceiling is of the same order as $N_{\mathrm{opt}}=1.59/\epsilon$ but of the opposite kind: balanced loss makes answers scarce, unbalanced loss makes them meaningless.

Loss in the reference branch breaks the null as well, because the N rotations no longer accumulate exactly to $\pi/2$. For $L=10^{-2}$ in the H branch alone, $P(H\,|\,\mathrm{absent})$ is nonzero for every $N\geqslant 2$ and peaks at 0.026 near $N=300$, before the Zeno effect freezes the V component instead.

In both unbalanced cases the $N\ln N$ growth of Eq. (14) is lost and the bound of Mitchison and Massar [16] applies. An interaction-free sensor must therefore be specified against the extinction ratio of its polarizing optics and the loss asymmetry between the branches, not against their total loss.

V. HOW FAR THE NEGATIVE RESULT GENERALIZES

Equation (7) is a statement about two particular interferometers. We now show that no arrangement of a single photon does better, so that the exchange rate is a property of the problem and not of the device.

Setup. Put $T=\cos^2\phi$ and dilate the object to a beam splitter of angle ϕ coupling the object mode $|o\rangle$ to a fresh vacuum mode $|e_k\rangle$ at the kth of K encounters, so that the global photon-plus-environment state stays pure. An arbitrary scheme, sketched in Fig. 1(c), is then

$$U(\phi)=W_K B_K(\phi)\,W_{K-1}\cdots W_1 B_1(\phi)\,W_0, \tag{21}$$

with $B_k(\phi)=e^{-i\phi J_k}$ the kth encounter, $J_k=|o\rangle\langle e_k|+|e_k\rangle\langle o|$ in the single-excitation subspace, and the W_j arbitrary optics independent of T. Multipass geometries of this kind are standard in optical metrology; a Fabry-Perot resonator is the canonical example [27]. Writing $|\psi_k\rangle$ for the global state after the kth encounter and the optics that follow it, with $|\psi_0\rangle=W_0|\mathrm{in}\rangle$, let $v_k=\langle o|\psi_{k-1}\rangle$ be the amplitude incident on the object at pass k, so that the object absorbs

$$n_{\mathrm{abs}}=(1-T)\sum_{k=1}^{K}|v_k|^2 \tag{22}$$

photons on average. The argument needs three physical assumptions, which we state and use below.

Lemma. $\|\partial_\phi|\psi\rangle\|^2=\sum_k|v_k|^2$. We induct on K. From $|\psi_K\rangle=W_K B_K|\psi_{K-1}\rangle$, the product rule, and the unitarity of W_K and B_K (using $[J_K,B_K]=0$),

$$\|\partial_\phi\psi_K\|^2=\|J_K\psi_{K-1}\|^2+\|\partial_\phi\psi_{K-1}\|^2 -2\,\mathrm{Im}\,\langle J_K\psi_{K-1}|\partial_\phi\psi_{K-1}\rangle. \tag{23}$$

Because e_K is a mode the photon has not yet touched, $\langle e_K|\psi_{K-1}\rangle=0$, so $J_K|\psi_{K-1}\rangle=\langle e_K|\psi_{K-1}\rangle|o\rangle+v_K|e_K\rangle=v_K|e_K\rangle$ and the first term is $|v_K|^2$. That identity holds for all ϕ, so it may be differentiated: $\partial_\phi|\psi_{K-1}\rangle$ likewise has no e_K component and the cross term vanishes. This is a statement about the support of the state and is insensitive to the phases of the amplitudes. The second term is $\sum_{k<K}|v_k|^2$ by induction, and the base case is $\|\partial_\phi\psi_0\|^2=0$ because $|\psi_0\rangle=W_0|\mathrm{in}\rangle$ does not depend on ϕ. ∎

Purification identity. The generator $G=iU^\dagger\partial_\phi U$ has vanishing expectation. Peeling the tail unitaries, which cancel between bra and ket,

$$\langle\psi\,|\partial_\phi\psi\rangle=-i\sum_{k=1}^{K}\langle\psi_{k-1}|J_k|\psi_{k-1}\rangle=0, \tag{24}$$

each term vanishing by $\langle e_k|\psi_{k-1}\rangle=0$—the same freshness property that powers the Lemma. Hence $F_Q^{\mathrm{pur}}(\phi)=4\|\partial_\phi\psi\|^2$, and with $(\mathrm{d}T/\mathrm{d}\phi)^2=\sin^2 2\phi=4T(1-T)$ the Lemma gives

$$F_Q^{\mathrm{pur}}(T)=\frac{\bar n_{\mathrm{tot}}}{T(1-T)},\qquad \bar n_{\mathrm{tot}}=\sum_{k=1}^{K}|v_k|^2, \tag{25}$$

exactly, for every K, every choice of intervening optics and every $T\in(0,1)$. ∎

Theorem. The photon-plus-environment state is pure, and the state an experimenter actually holds is its reduction, obtained by discarding the environment. Write F_Q for the quantum Fisher information of that reduced photon state—the quantity any apparatus is limited by. Since the quantum Fisher information is monotone under the partial trace, $F_Q\leqslant F_Q^{\mathrm{pur}}$, and with $n_{\mathrm{abs}}=(1-T)\bar n_{\mathrm{tot}}$,

$$\boxed{\frac{F_Q(T)}{n_{\mathrm{abs}}}\leqslant\frac{1}{T(1-T)^2}} \tag{26}$$

for every fixed multipass arrangement. ∎

Equation (25) says that $\bar n_{\mathrm{tot}}$ is the right way to count the light an object sees, and Eq. (26) then has a direct reading: a K-pass scheme is never more informative than the same incident flux spent on independent single-pass probes, since a single pass saturates $\bar n/[T(1-T)]$. Equality is attainable—trivially for $K=1$, and also for multipass arrangements whose optics route the photon away after its first encounter, which are single pass in disguise; numerical optimization of the intervening optics reaches the bound at $K=2$ and $K=3$. For arrangements in which the photon genuinely returns to the object the inequality is strict, and often badly so. Ten passes with no intervening optics at $T=0.5$ reach only 4.9% of the bound; five passes at $T=0.05$ reach 1.4×10^{-4} of it. Repeatedly interrogating an object does not merely fail to help, it discards most of the information the light carries.

A word is needed on multipass microscopy, which reports the opposite. Juffmann *et al.* measure a 4.8 ± 0.8 dB variance reduction for transmission measurements at constant damage [26]—a factor of three, in the same currency as Eq. (26). There is no conflict, and the reason is instructive. Their probe is classical light, and a coherent state is a poor probe of loss: its Fisher information per incident photon is $1/T$, against the $[T(1-T)]^{-1}$ of Eq. (26), so a classical single-pass measurement begins a factor $(1-T)^{-1}$ below the bound—a factor of

one hundred at the $T \simeq 0.99$ typical of microscopy. Multipassing classical light climbs inside that gap. For m passes with shot-noise-limited detection of the transmitted intensity, the Fisher information per unit dose is $m^2T^{m-2}(1-T)/(1-T^m)$, which at $T = 0.99$ rises from 1% of the bound at $m = 1$ to 65% at $m = 159$, and never crosses it: maximizing over $T \in (0, 1)$ and $m \leqslant 500$ gives 0.99 of the bound. A threefold gain over a classical baseline is therefore exactly what should be expected. It is a gain towards the single-photon limit, not past it.

The contrast with phase estimation is worth drawing, because it is the reason multipassing has a good reputation. Yu *et al.* analyze multipass interferometry in this same per-dose currency and find that for phase the sequential strategies do beat the parallel ones [24]. The asymmetry is not accidental. A phase is imprinted coherently and accumulates over passes, so the signal grows as m while the dose grows as m; each encounter with a lossy object, by contrast, is an independent Bernoulli trial, and independent trials do not compound.

Two features of the bound are worth noting. It is stated in terms of n_{abs}, the mean number of photons the object absorbs, which is a property of the apparatus rather than of a chosen dilation, so Eq. (26) does not depend on how the environment is modeled—only F_Q^{pur} would. And T is throughout the transmissivity of one elementary encounter; a device that presents several thicknesses per pass is described by redefining the encounter, not by re-using Eq. (26) with the same T.

Fisher-information analysis of lossy interferometers is a mature subject, but loss there is almost always a nuisance that degrades phase estimation [28]; here loss is the parameter being estimated. The route taken here—dilate, evaluate the quantum Fisher information of the purification, then invoke monotonicity under the partial trace—is Nair's [29], and Eq. (25) is his $\tilde{K}_\phi = 4\,\mathrm{diag}(N_k)$ written for one object met K times rather than K objects met once. The underlying constant is the Monras-Paris limit [30], which Fock states saturate [31]. Nair's scope is the ancilla-assisted parallel strategy, and he closes by noting that sequential strategies remain to be examined; Eqs. (25)–(26) settle the fixed sequential case at finite K, where the per-pass amplitudes v_k do the bookkeeping and the answer is that reusing a photon buys nothing. Protocols reported to beat the limit, such as the time-reversed scheme of Wang *et al.* [32], in fact reach it and stop. Massar, Mitchison, and Pironio obtained an inequality of this kind for a broader class by other means [15]. General arguments that adaptive strategies do not asymptotically outperform parallel ones [33,34] point the same way, but they count channel uses rather than incident energy and are proved for phase-type parameters.

Three points about scope. First, the gap between F_Q and F_Q^{pur} is the information locked in the environment: which pass absorbed the photon. That record is not available—an absorbed photon is simply a missing count—and for the Zeno chain the loss is complete in a strong sense. We find the classical Fisher information of the $\{H, V, \mathrm{lost}\}$ measurement to equal the quantum Fisher information of the reduced state at every N and T tested, so that measurement is optimal and the shortfall in Fig. 2(b) cannot be recovered by measuring better.

Second, the bound is on the quantum Fisher information, so no choice of measurement escapes it. Facchi *et al.* reached a comparable conclusion for standard and Zeno tomography, but only for binomial statistics [18]. Equation (26) drops that restriction and covers any fixed multipass arrangement.

Third, the three assumptions do real work. The optics W_j must leave the environment modes e_k alone. This is physically obvious, since one cannot redirect light the object has already absorbed, but it is what makes the induction close. The object must be memoryless, so that each pass meets fresh vacuum; an absorber interrogated faster than its own coherence time can reemit, and the identity fails. And the object must be nondispersive: a Kramers-Kronig phase on $\sqrt{T}$ builds up over passes and carries the problem outside the theorem. Anyone hoping to beat the exchange rate must look outside this class; breaking one of the three assumptions is necessary, not sufficient.

What Eq. (26) does not cover is multiphoton probes in sequential arrangements, and adaptive schemes in which the optics W_j are chosen conditionally on intermediate measurement outcomes. Both lie inside the broader class treated by Massar, Mitchison, and Pironio [15]; whether the short argument given here extends to them we have not determined. Generic asymptotic arguments that adaptive strategies do not outperform parallel ones [33,34] point the same way but count channel uses rather than incident energy, and so do not by themselves imply Eq. (26).

VI. CONCLUSION

Put in estimation-theoretic terms, the supersensitivity claim splits cleanly in two. Measuring how transparent an object is, the Elitzur-Vaidman interferometer carries half the Fisher information of direct probing, Eq. (6), and exactly the same Fisher information per absorbed photon, Eq. (7). That exchange rate also bounds every fixed multipass arrangement of a single photon, Eq. (26): no such scheme is more informative than the same incident flux spent on independent single-pass probes, and one that genuinely revisits the object is strictly worse. On this measure there is no interaction-free advantage, and no rearrangement of the interferometer will produce one.

Deciding whether an object is present at all is a different matter. There the discrimination rate per absorbed photon grows as $N \ln N$ in the cycle number, Eq. (14), and the advantage is real, quantifiable, and limited only by the loss per cycle through Eq. (18). It does not require an opaque object, only that the competing hypothesis be empty space. That widens the class of samples it serves, and at the same time sharpens the boundary: between two partial transparencies the rate does not grow at all.

So interaction-free sensing is the right tool wherever the sample, and not the photon budget, is the scarce resource: fragile biological specimens [35], photosensitive materials, x-ray interrogation [10]. It is the wrong tool wherever the goal is a more precise number. Massar, Mitchison, and Pironio reached that conclusion two decades ago in a broader setting [15]; Eqs. (25)–(26) recover it for single-photon multipass schemes by a short argument, and say when it is tight.

Two questions remain open. We do not know whether the identity extends to multiphoton sequential and adaptive strategies; the inequality is known there [15], the equality is not. And Facchi *et al.* report a real per-absorbed-particle advantage for Zeno tomography when the gray levels across a sample are unevenly distributed [18]. That is a claim about a prior over many pixels, not about estimating one parameter at fixed T, so the two results are probably answers to different questions—but showing that would sharpen the boundary further. A word finally on testing Eq. (18), because we are not able to propose a clean protocol and would rather say so. Measuring $\mathcal{R}_{\max}$ directly requires separating the photons the object absorbs from those the optics lose, and at $\epsilon = 10^{-2}$ the former are only a few percent of all missing counts. The natural evasion is to test the position of the maximum rather than its height, since an argmax is invariant under any N-independent multiplicative calibration error. That evasion fails on closer inspection. The obvious calibration-free estimator, $H_{\text{present}}/(\text{total}_{\text{absent}} - \text{total}_{\text{present}})$, has its $(1-\epsilon)^N$ factors cancel by Eq. (19) and returns the lossless ratio, which is monotone in N and has no maximum at all: at $\epsilon = 10^{-2}$ it runs 15.7, 31.9, 64.3, 129, 259, 518 over $N = 40$–1280 while the true $\mathcal{R}_\epsilon$ peaks at 25.7 near $N = 160$. Recovering $n^{\epsilon}_{\text{abs}}$ from the counts requires the model one is trying to test. Nor is the argmax as robust as the invariance argument suggests: additive N-dependent backgrounds are not removed by it, and a per-cycle retardance error is unitary, hence invisible to a calibration based on the total count rate, while placing a term growing as N^2 in the H port. We therefore leave the identification of a calibration-free observable that exhibits the predicted maximum as an open problem, and note that it is the main obstacle between Eq. (18) and an experiment.

ACKNOWLEDGMENT

The author thanks Norbert Sauer for valuable discussions.

DATA AVAILABILITY

There are no publicly available research data or software supporting this paper. Requests for further information or data should be sent to the author.

AI USE DISCLOSURE

This work was developed with the assistance of Claude (Anthropic, model CLAUDE-OPUS-5) for literature search and reference retrieval, algebraic and numerical derivation of the results in Secs. II–IV, the induction proof of the Lemma in Sec. V, and generation of the figures. All analytic results were checked by independent numerical computation, and all references were verified against publisher records. The author has verified the derivations reported here and is solely responsible for the content.

[1] A. C. Elitzur and L. Vaidman, Quantum mechanical interaction-free measurements, Found. Phys. **23**, 987 (1993).

[2] P. Kwiat, H. Weinfurter, T. Herzog, A. Zeilinger, and M. A. Kasevich, Interaction-free measurement, Phys. Rev. Lett. **74**, 4763 (1995).

[3] P. G. Kwiat, A. G. White, J. R. Mitchell, O. Nairz, G. Weihs, H. Weinfurter, and A. Zeilinger, High-efficiency quantum interrogation measurements via the quantum Zeno effect, Phys. Rev. Lett. **83**, 4725 (1999).

[4] L. Vaidman, The meaning of the interaction-free measurements, Found. Phys. **33**, 491 (2003).

[5] L. Vaidman, Analysis of counterfactuality of counterfactual communication protocols, Phys. Rev. A **99**, 052127 (2019).

[6] A. G. White, J. R. Mitchell, O. Nairz, and P. G. Kwiat, "Interaction-free" imaging, Phys. Rev. A **58**, 605 (1998).

[7] Y. Yang, H. Liang, X. Xu, L. Zhang, S. Zhu, and X.-S. Ma, Interaction-free, single-pixel quantum imaging with undetected photons, npj Quantum Inf. **9**, 2 (2023).

[8] J. R. Hance and J. Rarity, Counterfactual ghost imaging, npj Quantum Inf. **7**, 88 (2021).

[9] Y. Chen, Y.-J. Cai, X.-T. Li, K. Huang, J.-M. Liu, and E. Wu, Interaction-free quantum spectroscopy, Adv. Photon. Res. **2**, 2000206 (2021).

[10] R. Cohen, S. Shwartz, and E. Cohen, Loss-resilient, x-ray interaction-free measurements, Opt. Quantum **3**, 7 (2025).

[11] S. Dogra, J. J. McCord, and G. S. Paraoanu, Coherent interaction-free detection of microwave pulses with a superconducting circuit, Nat. Commun. **13**, 7528 (2022).

[12] J. J. McCord, S. Dogra, and G. S. Paraoanu, Coherent interaction-free detection of noise, Phys. Rev. A **110**, 032404 (2024).

[13] J. P. Dowling, Quantum optical metrology—the lowdown on high-N00N states, Contemp. Phys. **49**, 125 (2008).

[14] N. R. Miller, S. Ramelow, and W. N. Plick, Versatile supersensitive metrology using induced coherence, Quantum **5**, 458 (2021).

[15] S. Massar, G. Mitchison, and S. Pironio, Minimal absorption measurements, Phys. Rev. A **64**, 062303 (2001).

[16] G. Mitchison and S. Massar, Absorption-free discrimination between semitransparent objects, Phys. Rev. A **63**, 032105 (2001).

[17] G. Mitchison, S. Massar, and S. Pironio, Minimum number of photons needed to distinguish two transparencies, Phys. Rev. A **65**, 022110 (2002).

[18] P. Facchi, Z. Hradil, G. Krenn, S. Pascazio, and J. Řeháček, Quantum Zeno tomography, Phys. Rev. A **66**, 012110 (2002).

[19] S. Thomas, C. Kohstall, P. Kruit, and P. Hommelhoff, Semitransparency in interaction-free measurements, Phys. Rev. A **90**, 053840 (2014).

[20] P. M. Anisimov, D. J. Lum, S. Blane McCracken, H. Lee, and J. P. Dowling, An invisible quantum tripwire, New J. Phys. **12**, 083012 (2010).

[21] Y. Zhou and M.-H. Yung, Interaction-free measurement as quantum channel discrimination, Phys. Rev. A **96**, 062129 (2017).

[22] J. R. Hance, T. Matsushita, and H. F. Hofmann, Counterfactuality, back-action, and information gain in multi-path interferometers, Quantum Sci. Technol. **9**, 045015 (2024).

[23] J. R. Hance and J. Rarity, Interaction-free polarimetry of a polarising object, in *Quantum 2.0 Conference and Exhibition*, Technical Digest Series (Optica Publishing Group, 2022).

[24] Q. Yu, B. Wang, K. Zheng, M. Mi, H. Li, and L. Zhang, Dose-efficient quantum phase estimation in lossy optical interferometry, Optica **13**, 1130 (2026).

[25] P. Kruit, R. G. Hobbs, C.-S. Kim, Y. Yang, V. R. Manfrinato, J. Hammer, S. Thomas, P. Weber, B. Klopfer, C. Kohstall, T. Juffmann, M. A. Kasevich, P. Hommelhoff, and K. K. Berggren, Designs for a quantum electron microscope, Ultramicroscopy **164**, 31 (2016).

[26] T. Juffmann, B. B. Klopfer, T. L. I. Frankort, P. Haslinger, and M. A. Kasevich, Multi-pass microscopy, Nat. Commun. **7**, 12858 (2016).

[27] C. F. Wildfeuer, A. J. Pearlman, J. Chen, J. Fan, A. Migdall, and J. P. Dowling, Resolution and sensitivity of a Fabry-Perot interferometer with a photon-number-resolving detector, Phys. Rev. A **80**, 043822 (2009).

[28] T.-W. Lee, S. D. Huver, H. Lee, L. Kaplan, S. B. McCracken, C. Min, D. B. Uskov, C. F. Wildfeuer, G. Veronis, and J. P. Dowling, Optimization of quantum interferometric metrological sensors in the presence of photon loss, Phys. Rev. A **80**, 063803 (2009).

[29] R. Nair, Quantum-limited loss sensing: Multiparameter estimation and Bures distance between loss channels, Phys. Rev. Lett. **121**, 230801 (2018).

[30] A. Monras and M. G. A. Paris, Optimal quantum estimation of loss in bosonic channels, Phys. Rev. Lett. **98**, 160401 (2007).

[31] G. Adesso, F. Dell'Anno, S. De Siena, F. Illuminati, and L. A. M. Souza, Optimal estimation of losses at the ultimate quantum limit with non-Gaussian states, Phys. Rev. A **79**, 040305(R) (2009).

[32] J. Wang, R. L. de Matos Filho, G. S. Agarwal, and L. Davidovich, Quantum advantage of time-reversed ancilla-based metrology of absorption parameters, Phys. Rev. Res. **6**, 013034 (2024).

[33] R. Demkowicz-Dobrzański and L. Maccone, Using entanglement against noise in quantum metrology, Phys. Rev. Lett. **113**, 250801 (2014).

[34] S. Pirandola and C. Lupo, Ultimate precision of adaptive noise estimation, Phys. Rev. Lett. **118**, 100502 (2017).

[35] M. A. Taylor and W. P. Bowen, Quantum metrology and its application in biology, Phys. Rep. **615**, 1 (2016).