Supplementary Information: Measuring the time that atoms spend in the excited state due to a photon they don't absorb

Josiah Sinclair,[1] Daniela Angulo,[1] Kyle Thompson,[1] Kent Bonsma-Fisher,[1,2] Aharon Brodutch,[1] and Aephraim M. Steinberg[1,3]

[1]*Department of Physics, and Centre for Quantum Information and Quantum Control, University of Toronto, 60 St. George Street, Toronto, Ontario, Canada M5S 1A7*

[2]*National Research Council of Canada, 100 Sussex Drive, Ottawa, Ontario, Canada K1A 0R6*

[3]*Canadian Institute For Advanced Research, 180 Dundas St. W., Toronto, Ontario, Canada, M5G 1Z8*

(Dated: August 23, 2021)

CONTENTS

I. CHARACTERIZING THE DEPENDENCE OF THE PER-PHOTON XPS ON PULSE LENGTH

As explained in the main text, short signal pulses were superior to long ones in terms of distinguishing the effect of a transmitted photon from the slowly-varying background in $\Delta\phi(t)$. However, the size of ϕ_0 and the signal-to-noise ratio depend on both the pulse duration and the bandwidth of the atomic nonlinearity (pulses that are too short interact weakly with atoms because of low spectral overlap with the atomic transition). We observed that for the shortest pulses we could conveniently produce (which had an rms width of approximately 10ns), no significant decrease of ϕ_0 had yet occurred, and we therefore carried out the experiment with these pulses. This duration was found to be short enough to easily separate the effect of a transmitted photon from the slowly-varying background.

In Fig. 1a, the XPS, low-pass filtered at 25 MHz, is plotted versus time for signal pulses of several different durations. Signal pulses are created using an single-pass AOM with a rise time of approximately 10 ns (rms width) modulated on and off by a square-wave. This figure is plotted against the duration of the square-wave pulse, but for signal pulses shorter than 50 ns, the finite switching time of the AOM becomes comparable to the pulse duration, and the signal pulses become highly Gaussian in shape. Consequently, for short pulses, the rms width of the pulse is a better measure of duration than the square-wave pulse length. During data collection, the signal power entering the AOM was held constant for different pulse lengths. However, for short pulses, the finite switching time of the AOM affects the peak power in the pulse. To account for this variation in peak power, signal pulses are measured separately using a calibrated avalanche-photodiode (calibration is done using a calibrated Thorlabs power meter) and the energy in each pulse is integrated to find the average number of photons contained in pulses of each duration. Next, the phase shift per photon, ϕ_0, was found by dividing the peak XPS by the number of photons in a pulse. In Fig. 1b, the XPS per photon is plotted as a function of pulse length. Once the pulses are longer than about 100 ns, the time-integrated phase shift per photon saturates, while the *duration* of the XPS continues to grow; therefore ϕ_0 falls off as one over the pulse duration, as can be seen in Fig. 1b. For pulses shorter than 100 ns, ϕ_0 was observed to be relatively constant. The shortest pulses used in this characterization, which we went on to use for the rest of the experiment, had a square-wave pulse duration of 30 ns, and their rms width was separately measured to be 10 ns.

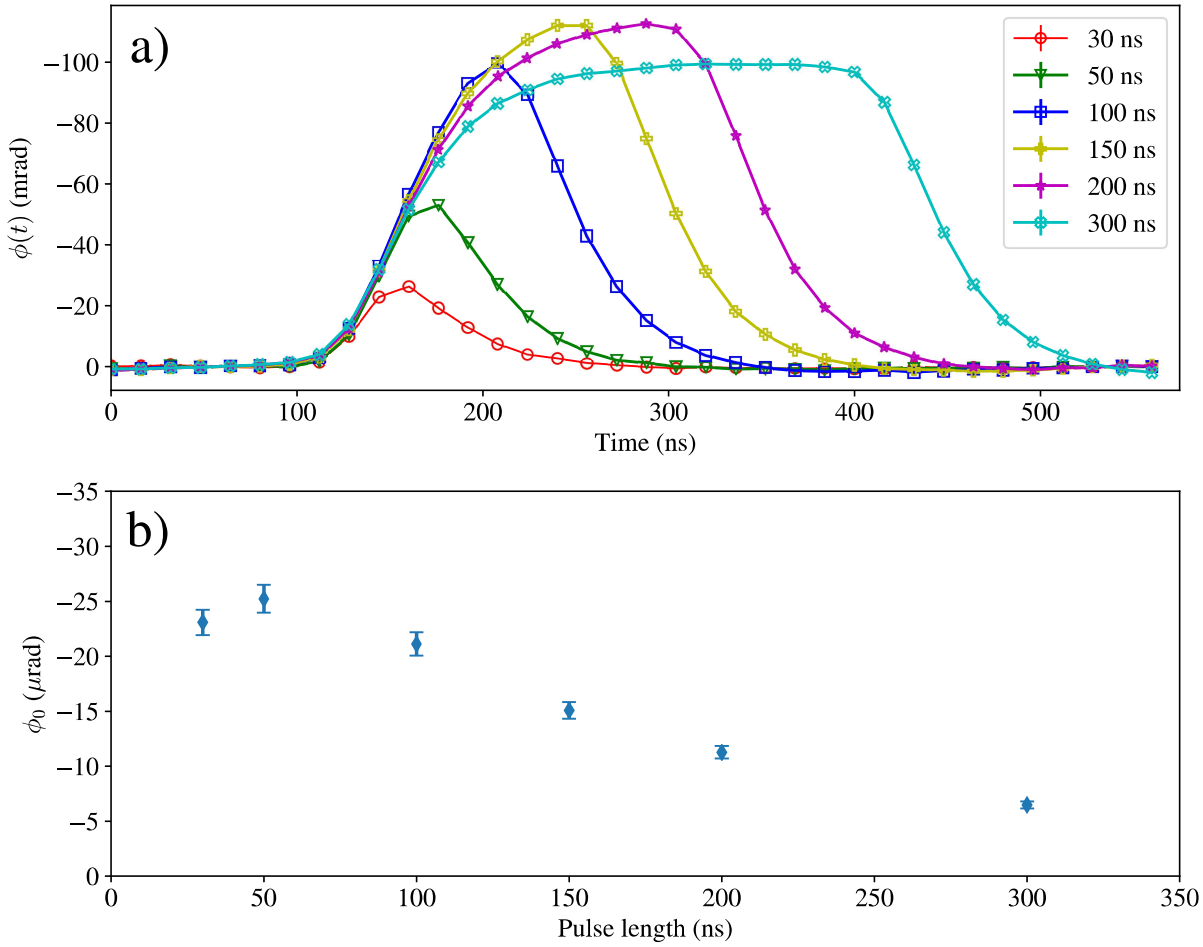

FIG. 1. (a) **XPS ($\phi(t)$) versus time for different pulse lengths.** (b) **Peak XPS divided by the number of photons in a pulse (ϕ_0) as a function of pulse length.** The probe detuning was $\Delta_p = -5.6$ MHz and the signal was on resonance. The signal pulse length used is given by the signal pulsing function generator setting. For pulses longer than 50 ns this approximates the actual length of the pulse well. For pulses shorter than 50 ns, the pulse shape becomes quite Gaussian as a result of the finite-switching speed of the AOM used to create signal pulses. The number of photons in each pulses takes into account the temporal shape of the pulse, which was measured separately.

II. CALIBRATING THE CROSS-PHASE SHIFT PER SIGNAL PHOTON

To calibrate the cross-phase shift (XPS) per photon, we vary the number of photons in a signal pulse between $0-3000$ and use a line of best fit to estimate the slope of the peak XPS (see Fig. 2). Measurements of the peak XPS above 2000 photons (indicated with circle markers and darker colour) were excluded from the fit because of the onset of atomic saturation (the critical photon number for saturation is around 10^4, given our beam waist of approximately 10^4 cross-sections and our pulse duration, which was on the order of τ_{sp}). The signal pulses used have an rms width of 10 ns, and the measurement is performed at two probe detunings, -5.6 MHz (orange data) and $+4.2$ MHz (blue data), resulting in $\phi_0 = -24.2 \pm 0.1$ μrad and $\phi_0 = 19.3 \pm 0.1$ μrad, respectively. The number of signal photons at the atom cloud is measured using a Thorlabs calibrated power meter and a fast avalanche-photodiode, which we need to measure the total number of photons in a 10 ns long Gaussian pulse. These data were obtained by taking four billion phase

measurements over approximately 30 minutes. The error bars shown in Fig. 2 are the standard error of the mean for the XPS measurement and the reduced Chi-Squared on the fit was 1.7, for positive probe detuning, and 2.3, for negative probe detuning.

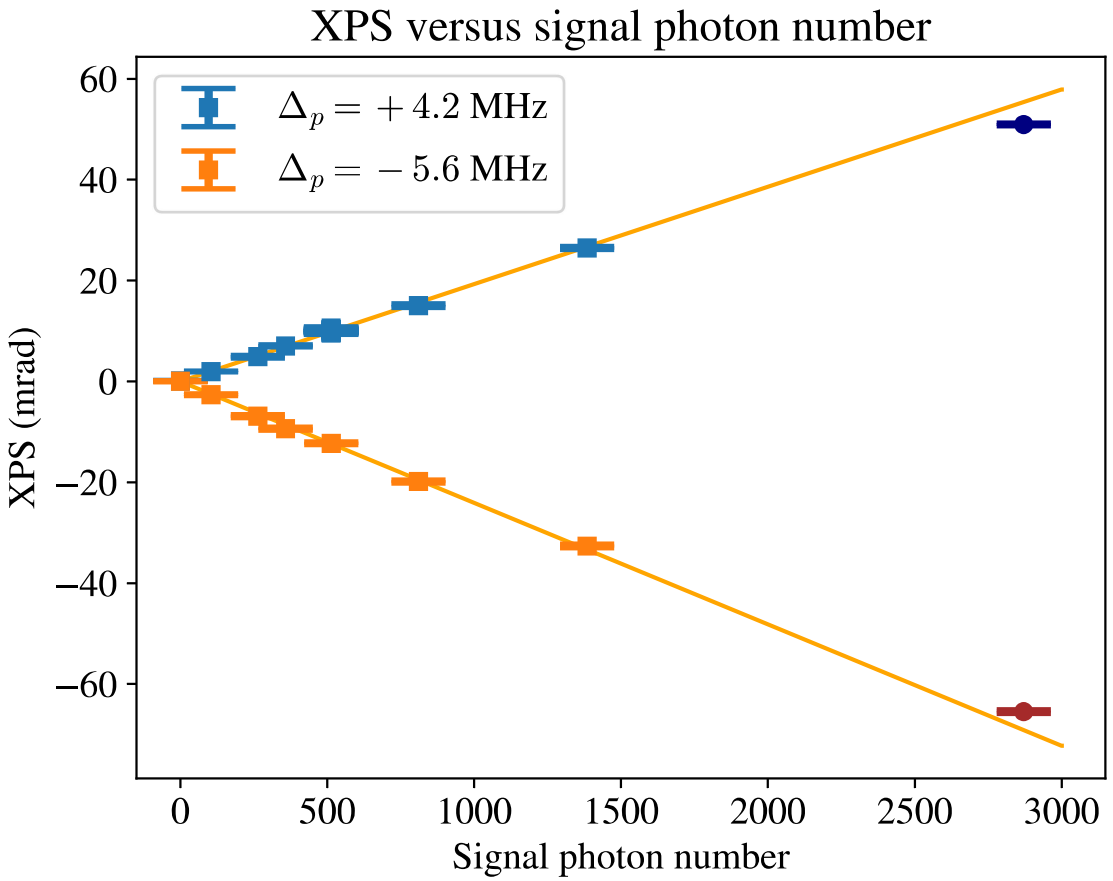

FIG. 2. **XPS as a function of signal photon number.** The probe detuning is Δ_p = +4.2 MHz and Δ_p = −5.6 MHz, and the signal pulse is on resonance with an rms width of 10 ns. The orange line is the line of best fit to measurements below 2000 photons, and ϕ_0 is determined from the fit to be -24.2 ± 0.1 μrad for Δ_p = −5.6 MHz, and 19.3 ± 0.1 μrad for Δ_p = +4.2 MHz.

III. ASSESSING THE TOTAL AMOUNT OF PROPORTIONAL NOISE IN THE SIGNAL

A possibility that we must contend with is that the signal pulses used in the experiment could have photon number (intensity) fluctuations larger than would be expected for a coherent state (whose photon number fluctuations are given by a Poissonian distribution). If there are large shot-to-shot fluctuations in the number of signal photons, then the photon number becomes classically correlated with whether or not the detector fires. Instead of a detection event adding a single transmitted photon, it indicates that there was a larger number of incident photons to begin with, some of which would be transmitted and some lost, both potentially leading to cross-phase shifts.

To understand the effect of intensity noise on the signal, consider a signal pulse that contains a number

of photons that is Gaussian-distributed

$$P(n) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(n-\bar{n})^2/2\sigma^2}, \tag{1}$$

where σ is the rms uncertainty in photon number in a signal pulse and $\bar{n}$ is the average number of photons in a signal pulse (which is assumed to be large enough that we can treat the photon number as approximately continuous). To ensure that $P(n < 0) \approx 0$, we assume $\bar{n} \gg \sigma$. According to Bayes's rule, the probability for the pulse to contain n photons given that one was detected is

$$P(n|c) = \frac{P(c|n) \cdot P(n)}{P(c)}. \tag{2}$$

The probability for a click given n photons is $P(c|n) = \eta n$, in the limit where the detector efficiency $\eta \ll 1/n$. Furthermore, the probability of getting a click is given by $P(c) = \eta\bar{n}$. For $P(c|n)$ ($P(c)$), we have approximated our single-photon counting module as an idealized linear detector. The a posteriori mean of the photon number conditioned on detection is

$$\langle n \rangle_c = \int_{-\infty}^{\infty} n \cdot P(n|c) \cdot dn, \tag{3}$$

which can be evaluated to be

$$\langle n \rangle_c = \frac{\sigma^2}{\bar{n}} + \bar{n}, \tag{4}$$

which is larger than the unconditioned mean, $\bar{n}$, by $\sigma^2/\bar{n}$ (note that if we consider the case of Poissonian noise, $\sigma = \sqrt{\bar{n}}$, we recover the result of +1 photons [1]). The implication is that the XPS difference will contain the effect of not just one transmitted photon, but $\sigma^2/\bar{n}$ photons. To work out the exact effect of these extra photons on the XPS difference would require one to know both ϕ_T, ϕ_L, and P_L and P_T. It suffices to say, however, that even if $\phi_T = 0$, the effect of intensity noise is to produce a XPS difference which scales with σ^2. We experimentally confirm this by using the white noise output on a function generator to add intensity noise to the signal. In Fig. 3a, $\Delta\phi(t)$ is plotted for three different values of added noise. These measurements were taken using 10 ns rms signal pulses with 34 photons and with a probe detuning of -5.6 MHz. The size of the peak $\Delta\phi(t)$ phase difference varies from about 1 mrad to about 150 μrad as the magnitude of the noise is varied. To distinguish the XPS difference (which could now be made quite large) from any remaining slowly-varying background, $\Delta\phi(t)$ was fitted to the following equation

$$y(t) = \delta\phi \cdot e^{-(t-t_0)/\tau_{\text{decay}}} \cdot \left(1 + \text{erf}\left[(t - t_0)/\tau_{\text{rise}}\right]\right) + \beta t^3 + \kappa t^2 + \delta t + \gamma, \tag{5}$$

where $\delta\phi$ is approximately the peak excursion of the post-selected XPS, τ_{rise}, τ_{decay}, and t_0 are the rise, fall, and peak times, and β, κ, δ, and γ are used to fit any slow-varying background. For fitting the post-selected XPS, the rise, fall and peak times were determined by a fit to the average XPS, and only the amplitude ($\delta\phi$) and the cubic parameters were left free. The amplitude of the post-selected XPS divided by the peak XPS ($\delta\phi/|\alpha|^2\phi_0$) is plotted as a function of the ratio of proportional noise to mean photon number ($\sigma/\bar{n}$) in Fig. 3b, and scales as $\sigma^2/\bar{n}^2$, as predicted by Eqn. 4.

Returning to our original concern: what if, even when we don't intentionally add any intensity noise, technical effects persist and introduce small extra fluctuations in photon number? Technical sources of noise are likely to be proportional, meaning that $\sigma = p\bar{n}$ (where p characterizes the magnitude of the proportional noise). Substituting into Eqn. 4, we find that the XPS difference will grow in proportion to the mean photon number, meaning that the effects of a small amount of proportional noise may only become significant for larger signal photon numbers. There are many possible sources of proportional noise—ranging from signal

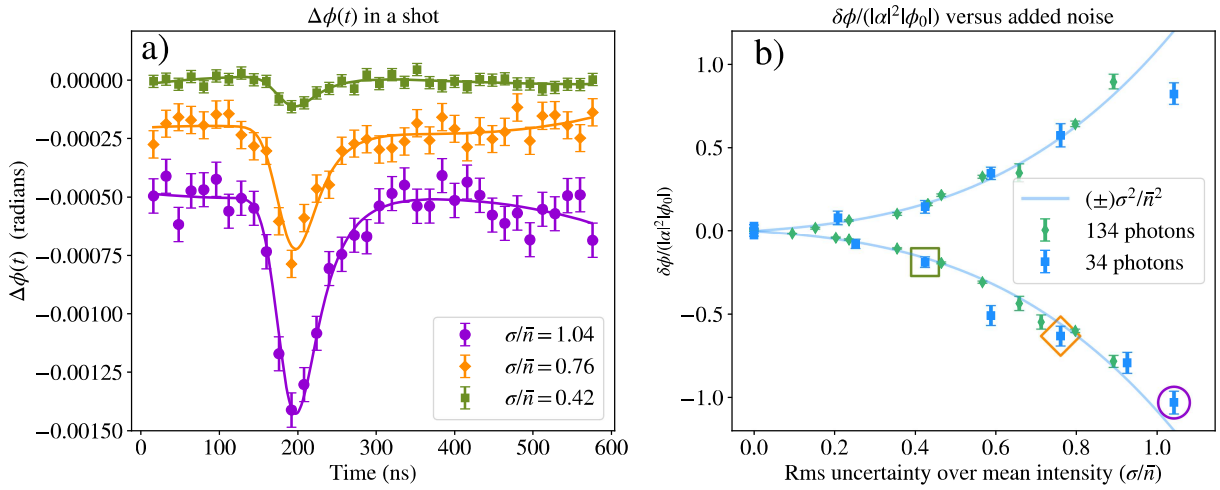

FIG. 3. **(a) Three measurements of $\Delta\phi(t)$ with different amounts of noise are plotted as a function of time**. Each measurement of $\Delta\phi(t)$ is fitted to Eqn. 5, and the resulting fit is plotted on top of the measurements. From the fit, $\delta\phi$ is extracted and used to create (b), where the **peak cross-phase shift difference ($\delta\phi$) divided by the absolute value of the average cross-phase shift ($|\phi_0||\alpha|^2$) is plotted versus the ratio of rms intensity noise to mean intensity** ($\sigma/\bar{n}$). The mean photon number and the fluctuations are directly measured on a fast APD with a 5 MHz bandwidth, matched to the specified bandwidth on the function generator's white noise output. Note that our data extend to values of $\sigma/\bar{n}$ even greater than 1, which cannot occur within our Gaussian approximation; for smaller values of noise, however, we believe the Gaussian approximation to be good. The ratio of $\delta\phi$ and $|\alpha|^2|\phi_0|$ was measured for signal pulses containing 34 and 134 photons at Δ_p = −5.6 MHz and Δ_p = +4.7 MHz and scales as $\sigma^2/\bar{n}^2$ for $\sigma/\bar{n} < 1$ as expected. The three values of $\delta\phi$ extracted from (a) are highlighted with an enclosing square, diamond, or circle.

intensity fluctuations to optical depth fluctuations or other—so rather than independently characterizing each potential source, we choose instead to assess the *total amount* of proportional noise in the experiment. We do this by varying the number of photons in the signal pulse and looking for a linearly changing XPS difference (as predicted by Eqn. 4).

In Fig. 4, $\delta\phi$ is plotted as a function of signal photon number. From the slope (gold line), the total amount of proportional noise in the experiment is estimated to be less than or equal to about 3%. This is substantially less than Poisson fluctuations for coherent states with average photon number of 34 (17%) or 134 (9%), so we conclude that the total effect of proportional noise in our experiment is quite insignificant and that there is very little error introduced by estimating ϕ_T from a weighted average of $\delta\phi$ measured with 34 and 134 photons, as we do in Fig. 3a in the main text. In general, however, the y-intercept of the linear fit to $\delta\phi$ versus signal photon number corresponds to ϕ_T, and this is how the values of ϕ_T reported in Fig. 3b and 3c in the main text are found.

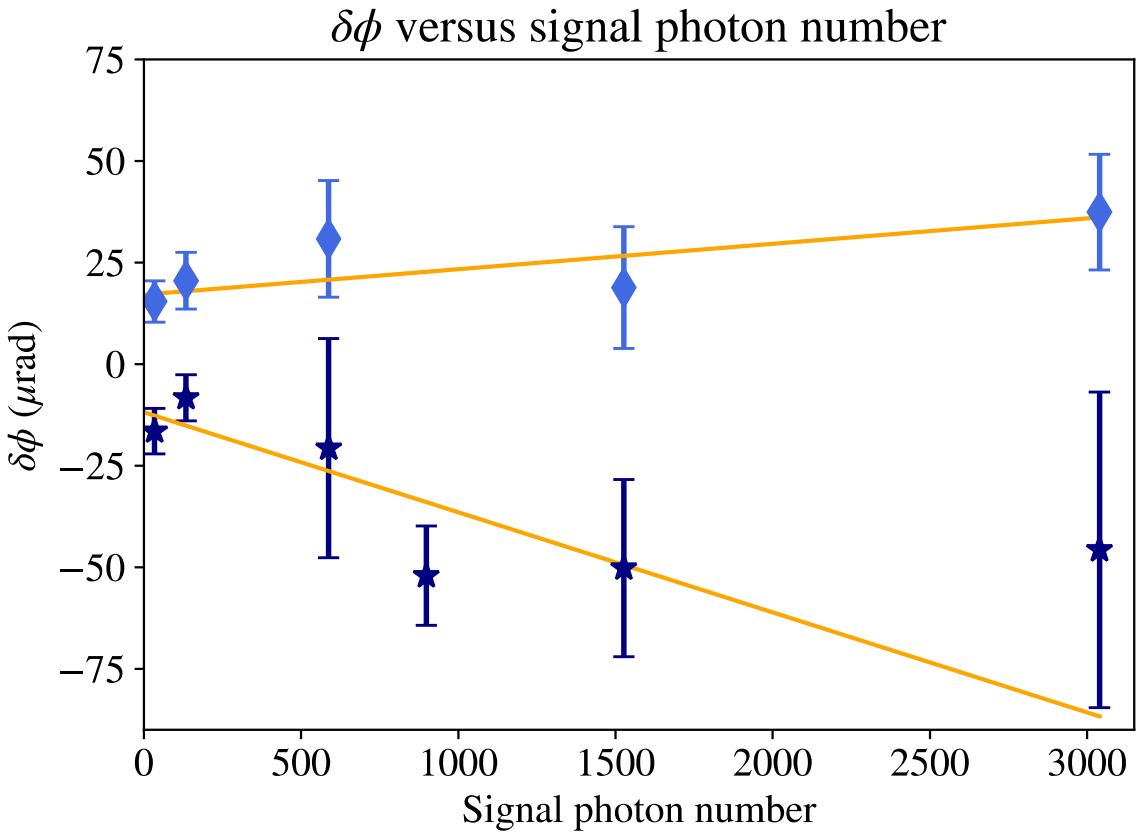

FIG. 4. **Peak cross-phase difference versus average number of signal photons per pulse** for $\Delta_p = -5.6$ MHz (light blue diamonds) and $\Delta_p = +4.7$ MHz (dark blue stars). For all measurements the signal is resonant, and the rms pulse length is 10 ns. ϕ_T is extracted from the y-intercept of the linear fit (gold line) for both positive and negative probe detunings. The error bars shown are the standard error of the mean.

IV. A BRIEF DESCRIPTION OF THE MINIMUM-COHERENT-EMISSION MODEL

As the atoms interact with light, there are three conceivable processes by which they can change energy level: absorption of incident photons, spontaneous emission into side modes (whose rate per excited atom, Γ, is known and fixed), and potentially, coherent emission into the forward mode (we work in a limit where stimulated emission is negligible). The minimum-coherent-emission model springs from the simple observation that, in certain situations, the Maxwell-Bloch equations predict that the number of excited atoms will decrease faster than can be explained by spontaneous emission. This is unambiguous evidence of coherent forward emission. The semiclassical equations describe the time rate of change of the excited-state population, setting a value for the *difference* between the absorption and coherent-emission rates, but cannot fix those two rates independently. Our minimum-coherent-emission model starts from the hypothesis that whenever the excited state population decreases faster than would be expected from spontaneous decay, this rate of change can be entirely attributed to coherent forward emission, and no absorption is occurring; and conversely that whenever the decrease is slower than the spontaneous emission rate, it is because of absorption, and no coherent emission is occurring. From these rates, one can define at each position along the propagation direction an average time for photons to be absorbed and an average time for the emission of

those photons which are coherently re-scattered into the forward mode. We use the difference between these two times to quantify how long such forward-scattered photons spent in the atoms. Combining this with our knowledge of the total number of such photons (from the time-integral of the emission rate discussed earlier), we can calculate the time an average transmitted photon causes atoms in each plane to spend in the excited state, and integrate over the medium to extract τ_T. Of course, we cannot rule out the possibility that this model severely underestimates the rate of coherent-forward emission in our experiment; there is no a priori reason to assume, as the model does, that absorption and coherent forward emission cannot be concurrent. However, the model's predictions are consistent with our experimental results and our expectations that $\phi_T \to 0$ as OD $\to 0$. Finally, the model itself raises interesting new questions about the sensitivity of ϕ_T to the details of the pulse shape and bandwidth, motivating further experiments that are currently underway.

V. CHECKING FOR SYSTEMATIC EFFECTS

The null results of checks for systematic effects are shown in the shaded area of Fig. 3b in the main text. As we ultimately measure a post-selected cross-phase shift that is 10^4 times smaller than the single-shot phase noise, it is imperative to check for minuscule systematic effects that could be mistaken for the effect of a transmitted photon. Examples of possible systematic effects include electrical cross-talk between BNC cables carrying information about the phase of the probe and signal photon detection, and various common-cause effects that introduce a spurious correlation between signal transmission and the phase of the probe. These correlations might be mediated by the atomic sample (e.g. electrical noise could affect the atomic transition frequency—for instance, by creating background magnetic fields—while the probe phase fluctuates in a correlated fashion), or they could be mediated by something else (e.g. electrical noise could perturb both the phase of the probe and the SPCM detection efficiency [2]). To rule out systematic effects such as, but not limited to, those listed above, we carry out measurements of $\Delta\phi(t)$ (and ϕ_T) with no atoms; with the signal far-detuned, to make the signal insensitive to the atomic medium; and with the probe far-detuned, to make the probe insensitive to the atomic medium. Combined, checks for systematic effects required 30 billion signal pulses and about 150 hours of data acquisition.

A. No atoms

To prove that the correlations we observe between the phase of the probe and the signal transmission are mediated by the atomic medium, we measured $\Delta\phi(t)$ (and ϕ_T) with no atoms present. To eliminate the atomic sample in a minimally invasive way (i.e. no electrical signals were changed), the magneto-optical trapping and repumper beams were physically blocked. The signal pulses used for this measurement contained, on average, 34 photons and had an rms length of 10 ns. To achieve a precision of < 2 μrad, the measurement required approximately 12 billion signal pulses (half a trillion phase measurements) and took about 60 hours. Periodically, the measurement would be paused, and the trapping and repumper beams turned back on to check that ϕ_0 had not changed. The rise and fall time used for fitting $\Delta\phi(t)$ was taken from the XPS measured in one of these periodic checks.

B. Signal far-detuned

Motivated by similar concerns as in the "no atoms" systematic check, and to rule out an admittedly somewhat conspiratorial scenario where some common-cause noise introduces a correlation between the phase of the probe and the likelihood of the SPCM firing 500 ns in the future, and perhaps mainly out of an abundance of caution, we measured $\Delta\phi(t)$ (and ϕ_T) with the signal far-detuned. This check was performed

twice, once with the probe at $\Delta_p = -5.6$ MHz and once with the probe at $\Delta_p = +4.7$ MHz. In both measurements, the signal pulses contained an average of 34 photons and had an rms length of 10 ns. For the measurement with negative probe detuning (labelled '−'), the signal was locked 60 MHz above resonance, approximately 60 MHz away from the nearest hyperfine transition. As the XPS due to 34 photons could not be measured at this detuning, the rise and fall times used for fitting were taken from the XPS in Fig. 2a in the main text. This measurement involved approximately 11 billion signal pulses and took about 45 hours, resulting in a precision of about 5 μrad. For the systematic check with positive probe detuning (labelled '+'), the signal was locked approximately 6 GHz below resonance to a Rubidium 87 hyperfine $F = 1 \rightarrow 2$ crossover. As before, the signal was too far detuned to observe an XPS. This time, the rise and fall times used for fitting were supplied by a measurement made right after the run was complete, with the signal set back on resonance, and with signal pulses containing approximately 800 photons. The second measurement of ϕ_T carried out with the signal far-detuned took slightly less time than the first and achieved a similar precision. It lasted for 25 hours and required 7 billion signal pulses.

C. Probe far-detuned

To rule out the possibility that a common-cause effect is introducing a correlation between the probe's phase and signal transmission by, for example, modulating the phase of the probe and the atomic transition frequency (affecting signal transmission), we detune the probe several GHz away from resonance and measure $\Delta\phi(t)$ (and ϕ_T). As expected, once the probe becomes insensitive to the atomic medium, ϕ_T becomes consistent with zero. As in the previous check, the signal pulses contained an average of 34 photons and had an rms length of 10 ns. To achieve a precision of approximately 4 μrad, the measurement required about 2 billion signal pulses and took about 13 hours. Signal detection rates were approximately 55% with atoms present, and 25% with atoms absent, and these detection rates were used to monitor the optical depth. As there is no measurable XPS at this detuning, ϕ_T was extracted by fitting $\Delta\phi(t)$ twice, first using the rise and fall times from the XPS in Fig. 2a in the main text, and then using the rise and fall times taken from Fig. 2b in the main text, and averaging.

D. Further checks

After these systematic checks, an important remaining question is: "is the measured value of ϕ_T truly due to absorption?" It was partially to confirm this that we performed measurements at positive and negative probe detuning, demonstrating the expected detuning dependence for an absorptive (saturation) nonlinearity. More crucial, however, is being certain that the signal was resonant and that other dispersive nonlinearities did not contribute to ϕ_T. The absence of a larger slow-varying background in $\Delta\phi(t)$, however, is clear evidence that our signal was resonant, as detuning the signal even a few MHz to the red or blue was observed to produce a ± 200 μrad oscillation in $\Delta\phi(t)$.

VI. DATA AVAILABILITY

The data that support the findings of this article form a data set of approximately 100 TB in size (containing approximately 10 trillion phase measurements). This entire data set, or specific portions of it, will

be made available upon reasonable request. For more information, contact AMS.

---

[1] Matin Hallaji, Amir Feizpour, Greg Dmochowski, Josiah Sinclair, and Aephraim M. Steinberg, "Weak-value amplification of the nonlinear effect of a single photon," Nature Physics **13**, 540–544 (2017).

[2] Although this would require a 500 ns delay matched improbably to our optical delay set by a long Thorlabs single-mode fibre.