

Observing the quantum behavior of light in an undergraduate laboratory

J. J. Thorn, M. S. Neel, V. W. Donato, G. S. Bergreen, R. E. Davies, and M. Beck^{a)}

Department of Physics, Whitman College, Walla Walla, Washington 99362

(Received 4 December 2003; accepted 15 March 2004)

While the classical, wavelike behavior of light (interference and diffraction) has been easily observed in undergraduate laboratories for many years, explicit observation of the quantum nature of light (i.e., photons) is much more difficult. For example, while well-known phenomena such as the photoelectric effect and Compton scattering strongly suggest the existence of photons, they are not definitive proof of their existence. Here we present an experiment, suitable for an undergraduate laboratory, that unequivocally demonstrates the quantum nature of light. Spontaneously downconverted light is incident on a beamsplitter and the outputs are monitored with single-photon counting detectors. We observe a near absence of coincidence counts between the two detectors—a result inconsistent with a classical wave model of light, but consistent with a quantum description in which individual photons are incident on the beamsplitter. More explicitly, we measured the degree of second-order coherence between the outputs to be $g^{(2)}(0) = 0.0177 \pm 0.0026$, which violates the classical inequality $g^{(2)}(0) \geq 1$ by 377 standard deviations. © 2004 American Association of Physics Teachers.

[DOI: 10.1119/1.1737397]

I. INTRODUCTION

Students often believe that the photoelectric effect, and Einstein's explanation of it, proves that light is made of photons. This is simply not true; while the photoelectric effect strongly suggests the existence of photons, it does not demand it.^{1,2} It was shown in the 1960s by Lamb and Scully that the photoelectric effect can be explained by assuming that the detector atoms are quantized, but that the field is not (i.e., by assuming light to be a classical wave). This explanation is based on the semiclassical model of photoelectric detection, which we will discuss further below.^{3,4}

How then does one prove that photons exist? Here, we will assume that proving photons exist is equivalent to observing an effect that requires a quantum mechanical description of the field; equivalently, we will say that photons exist if the results of an experiment cannot be explained using a classical wave theory of light. Ideally, an experiment to prove the existence of photons will also demonstrate that light has “granular” properties. While physicists may argue about which specific experiment was the first to conclusively demonstrate the existence of a field requiring a quantum mechanical (QM) description, one can be fairly certain that this experiment was carried out in the 1970s.^{5–7} While many such experiments have subsequently been performed, we know of very few that are well-suited for an undergraduate laboratory.^{8–10}

In 1986, Grangier, Roger, and Aspect performed an elegant experiment.^{11,12} Conceptually very simple, their approach was to examine correlations between photodetections at the transmission and reflection outputs of a 50/50 beamsplitter. To quote the experimenters, “a single photon can only be detected once!”¹¹ Hence, if a single quantum of light is incident on the beamsplitter (BS), it should be detected at the transmission output or at the reflection output, but not both: there should be no coincident detections between the two outputs. In fact, Grangier *et al.* measured fewer coincidences than predicted by a classical wave theory, violating a classical inequality by 13 standard deviations, and demonstrating that the field incident on the beamsplitter was well

described by a single-photon state.¹¹ As discussed below in more detail, the key challenge in such a measurement is to create a field that truly has a *single-photon* incident on the BS; a weak beam containing on *average* a single photon (or less) is not sufficient.

Here, we have repeated the experiment of Grangier *et al.*, adapting it for an undergraduate laboratory. We have taken advantage of over 15 years of technological advancements to obtain orders of magnitude increased count rates over those obtained by Grangier *et al.* The increased count rate in our experiment allows us to violate a classical inequality by 146 standard deviations with only 5 min of counting time. Our experiment is well described by the QM description of a field in a single photon state incident on a beamsplitter.

II. HISTORY AND THEORY

A. Early measurements

As stated above, we are interested in examining correlations between the photocounts on two detectors situated at the output ports of a BS (Fig. 1). The first experiment to examine these correlations was carried out by Hanbury Brown and Twiss,^{13,14} who found a positive correlation between the outputs of the two detectors. It should be noted that in this first experiment, Hanbury Brown and Twiss were not counting individual photons, but were instead measuring correlations between continuous currents from photomultiplier tubes (PMTs).¹³ As such, this positive correlation indicated that when the current from one PMT increased, the current on the second tended to increase simultaneously. While the intent of Hanbury Brown and Twiss was to develop a new technique for measuring the angular diameter of stars,¹⁵ their work played an important role in creating the field of quantum optics.

A brief controversy arose when Brannen and Ferguson performed a similar experiment in which they observed no positive correlation, and then claimed “that if such a positive correlation did exist, it would call for a major revision of some fundamental concepts in quantum mechanics.”¹⁶ However, Purcell¹⁷ and Hanbury Brown and Twiss¹⁸ quickly




Fig. 1. Coincidence measurement. The incident (I) beam is split into transmitted (T) and reflected (R) beams at a 50/50 BS. Detections at T and R are examined to see whether or not they occur simultaneously.

noted that the experimental parameters used by Brannen and Ferguson precluded the observation of positive correlations. They also showed that positive correlations are not only allowed by quantum mechanics, but are a natural consequence of the tendency for photons (and other bosons) to “bunch” together. The first experiment to observe positive correlations using coincidence detection of individual photocounts (not just photocurrents) from PMTs was performed by Twiss, Little, and Hanbury Brown,¹⁹ who observed positive correlations of a few percent. This amount of correlation was consistent with that expected, given their experimental parameters.

B. Classical fields

By a classical field, we mean an electromagnetic wave that is perfectly described by Maxwell’s equations. For such a field, the correlations between the intensities of the transmitted I_T and reflected I_R beams are given by the *degree of second-order (temporal) coherence*, $g_{T,R}^{(2)}(\tau)$, which is a function of the time delay τ between the intensity measurements:²⁰

$$g_{T,R}^{(2)}(\tau) = \frac{\langle I_T(t+\tau)I_R(t) \rangle}{\langle I_T(t+\tau) \rangle \langle I_R(t) \rangle}. \quad (1)$$

If the light source is *stationary* (i.e., if its statistics do not change in time), then we can interpret the brackets as referring to ensemble averages rather than time averages. It is called the degree of second-order coherence because it involves correlations between intensities, whereas the degree of first-order coherence describes correlations between fields.

Of particular importance to the present work is the case of simultaneous ($\tau=0$) intensity measurements. In this case, and furthermore assuming a 50/50 BS in which the transmitted, reflected, and incident intensities are related by $I_T(t) = I_R(t) = \frac{1}{2}I_I(t)$, it is straightforward to see that

$$g_{T,R}^{(2)}(0) = g_{I,I}^{(2)}(0) = \frac{\langle [I_I(t)]^2 \rangle}{\langle I_I(t) \rangle^2} = g^{(2)}(0). \quad (2)$$

From the Cauchy–Schwartz inequality, it is straightforward to prove that $\langle I_I(t) \rangle^2 \leq \langle [I_I(t)]^2 \rangle$.^{20,21} Using this, we find that

$$g_{T,R}^{(2)}(0) = g^{(2)}(0) \geq 1 \quad (\text{classical fields}), \quad (3)$$

where we emphasize that this result has been derived using classical wave theory. In Eq. (3), equality with 1 is achieved if the input field is perfectly stable with no fluctuations, while for fluctuating fields the second-order coherence is greater than 1. For “chaotic” light (e.g., light from a thermal

source that is either collisionally or Doppler broadened), it can be shown that $g^{(2)}(0) = 2$.²⁰ In an ingenious set of experiments involving a “pseudothermal” light source (a laser whose phase was randomized by a rotating ground-glass slide), Arrechi *et al.* were able to measure fields with $g^{(2)}(0) = 1$ and $g^{(2)}(0) = 2$.²²

C. Semiclassical theory of photodetection

So far, we have been speaking about correlations between the intensities of the fields leaving the BS. In an experiment, however, one does not measure the intensity directly, but rather the photocurrent from a detector. It is then necessary to model the photodetection process. Since to this point we have been discussing classical fields, it is most appropriate to use a model that treats the field classically. The most rigorous theory of this sort is the semiclassical theory of photoelectric detection, in which the field is treated classically and the photodetector is treated quantum mechanically.²³ For the purposes of the discussion here, it is convenient to refer to the detector monitoring the transmitted (reflected) field as detector T (R).

In the semiclassical theory of photoelectric detection, it is found that the conversion of continuous electromagnetic radiation into discrete photoelectrons is a random process. The probability of obtaining a single photocount from a single detector (for example, detector T) within a short time window Δt is proportional to the average intensity of the field striking that detector, given as

$$P_T \Delta t = \eta_T \langle I_T(t) \rangle \Delta t, \quad (4)$$

where η_T is a constant that characterizes the detection efficiency of detector T . The joint probability of obtaining a photocount (within a time window Δt) at detector R , and then after a time τ obtaining a photocount at detector T (within a time window Δt), is given by

$$P_{TR}(\tau) \Delta t^2 = \eta_T \eta_R \langle I_T(t+\tau)I_R(t) \rangle \Delta t^2. \quad (5)$$

It is then easily seen that if one measures the probability of joint and individual photocounts at detectors T and R , one can determine the degree of second-order coherence from

$$g_{T,R}^{(2)}(\tau) = \frac{P_{TR}(\tau)}{P_T P_R}. \quad (6)$$

Again, we are most interested in simultaneous, $\tau=0$, detection of photocounts at detectors T and R , which occurs with probability $P_{TR}(0)$. Using Eq. (3), we find that for classical fields, the measured degree of second-order coherence must be greater than or equal to 1:

$$g_{T,R}^{(2)}(0) = \frac{P_{TR}(0)}{P_T P_R} = g^{(2)}(0) \geq 1 \quad (\text{classical fields}). \quad (7)$$

Here, we see that if the joint probability factorizes, $P_{TR}(0) = P_T P_R$, which occurs when the detections at T and R are completely uncorrelated, then $g^{(2)}(0)$ is minimized and is equal to 1.

We can summarize what we have learned about classical field statistics as follows. It is possible to measure the degree of second-order coherence between the fields leaving a beamsplitter $g^{(2)}(0)$ by measuring the probability of joint and individual photocounts at detectors T and R . The second-order coherence must satisfy the inequality $g^{(2)}(0) \geq 1$.

When the photocounts at T and R are completely uncorrelated, $g^{(2)}(0)=1$, which occurs when the input field to the beamsplitter is a perfectly stable wave. If the input field fluctuates, then $g^{(2)}(0)>1$, indicating positive correlations between the photocounts.

Since $g^{(2)}(0)$ cannot be less than 1, we are left with the conclusion that for classical fields the measured photocounts at T and R cannot be anticorrelated. This makes sense because a BS simply splits a classical input field into two identical copies. These output fields either fluctuate together (positive correlation) or do not fluctuate at all (no correlation). It is not possible for the intensity of one to decrease while the intensity of the other increases (anticorrelation).

D. Quantum fields

From the time of the original Hanbury Brown and Twiss experiment in 1956,¹³ the importance of a rigorous theory of photoelectric counting was recognized. The first attempts were the semiclassical theories discussed in the previous section. In the mid-1960s sophisticated fully QM theories, in which both the electromagnetic field and the detector atoms are treated quantum mechanically, were developed by Kelly and Kleiner,²⁴ Glauber,²⁵ and others (see Refs. 20 and 23, and the references therein). A QM field is not fully described by Maxwell's equations.

In the fully quantum theory, the correlations between the output fields from the BS in Fig. 1 are still described by the degree of second-order coherence $g_{T,R}^{(2)}(\tau)$, although now the electric fields (and corresponding intensities) are treated as QM operators, rather than as classical waves. Again, we are interested in simultaneous ($\tau=0$) detection of photons at the outputs; quantum mechanically $g_{T,R}^{(2)}(0)$ is defined as

$$g_{T,R}^{(2)}(0) = \frac{\langle : \hat{I}_T \hat{I}_R : \rangle}{\langle \hat{I}_T \rangle \langle \hat{I}_R \rangle}. \quad (8)$$

Here the colons indicate that the creation \hat{a}^\dagger and annihilation \hat{a} operators corresponding to the electric fields are to be placed in normal order, which means that all creation operators appear to the left of all annihilation operators. The intensity operator is proportional to the photon number operator for the field $\hat{n} = \hat{a}^\dagger \hat{a}$, so that

$$g_{T,R}^{(2)}(0) = \frac{\langle : \hat{n}_T \hat{n}_R : \rangle}{\langle \hat{n}_T \rangle \langle \hat{n}_R \rangle} = \frac{\langle \hat{a}_T^\dagger \hat{a}_R^\dagger \hat{a}_R \hat{a}_T \rangle}{\langle \hat{a}_T^\dagger \hat{a}_T \rangle \langle \hat{a}_R^\dagger \hat{a}_R \rangle}, \quad (9)$$

where we have explicitly placed the field operators in normal order.

The averages in Eqs. (8) and (9) are given by QM expectation values. The expectation value is computed using the field states at the detectors. These states can be derived from the input state to the BS.^{23,26} Alternatively, the operators for the reflected and transmitted fields can be written in terms of the operators for the input field \hat{a}_I , and the unoccupied (vacuum) field \hat{a}_V that enters the unused port of the BS (Fig. 2). For a 50/50 BS, and one particular choice of phase for the BS, it is straightforward to show that

$$\hat{a}_R = \frac{1}{\sqrt{2}}(\hat{a}_I + \hat{a}_V), \quad \hat{a}_T = \frac{1}{\sqrt{2}}(\hat{a}_I - \hat{a}_V). \quad (10)$$

Substituting the reciprocity relations Eq. (10) into Eq. (9), and using the fact that the unoccupied field mode is in a




Fig. 2. Field operators corresponding to BS input and output ports.

vacuum state, the second-order coherence can be rewritten as²⁰

$$g_{T,R}^{(2)}(0) = \frac{\langle \hat{n}_I(\hat{n}_I - 1) \rangle}{\langle \hat{n}_I \rangle^2} = g_{I,I}^{(2)}(0) = g^{(2)}(0), \quad (11)$$

where now the expectation value is computed using the QM state of the field incident on the BS. As in the classical case, the second-order coherence between the BS outputs is equal to the second-order coherence of the input.

Quantum mechanically, the measured correlations at the detectors are determined by the state of the field incident on the BS (the input state). The QM equivalent to a stable classical wave is a coherent state $|\alpha\rangle$, which is the eigenstate of the annihilation operator $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$.²⁰ If one evaluates the second-order coherence [Eq. (11)], assuming an input field in a coherent state, one finds $g^{(2)}(0)=1$, which is the same as the classical result for a stable classical wave. Evaluating Eq. (11) assuming an input field in a thermal state (which is an incoherent mixture described by a density operator) one finds $g^{(2)}(0)=2$.²⁰ Such a field is said to be “bunched,” because one interpretation of this result is that the photons tend to come in bunches; once they strike the BS, some are transmitted and others are reflected, leading to positive correlations between the output fields.

Thus, the quantum theory of photoelectric detection is in agreement with the classical theory described in Sec. II, as long as one uses the appropriate field states. However, there exist certain field states that are inherently QM in nature, and for which there is no classical wave theory counterpart. Such nonclassical fields are not in general constrained by the limits discussed in Sec. II C. An example of a nonclassical field state is one containing exactly one photon; this state is an eigenstate of the photon number operator, with eigenvalue 1: $\hat{n}|1\rangle = 1|1\rangle$. Evaluating Eq. (11) using an input field in a single-photon state yields $g^{(2)}(0)=0$, which violates the classical inequality $g^{(2)}(0) \geq 1$.

Theoretically predicting the existence of nonclassical fields, and generating them in the laboratory, however, are two very different matters. One of the first experiments to demonstrate the existence of a nonclassical field was performed by Kimble, Dagenais, and Mandel in 1977.⁷ They measured the light emitted by a single atom (“resonance




Fig. 3. Coincidence measurements with a gate. A source emits pairs of photons simultaneously, and the photons travel in opposite directions. Detection of the gate signal tells the T and R detectors when to expect a “proper” detection on the experiment side.

fluorescence”) and found $g^{(2)}(0) = 0.4 \leq 1$, proving that the field was “anti-bunched.” An anti-bunched field can be interpreted as one in which the photons do not clump together, and hence tend to arrive one at a time. When these individual photons strike the beamsplitter, they are either transmitted or reflected (but not both), leading to anticorrelations in the photocounts at the detectors.

Despite clearly demonstrating that the light emitted by a single atom is anti-bunched, this experiment was complicated by the difficulty of isolating the light coming from the atom from the background scattered light. This complication was due to the fact that the laser light used to excite the atom and the resonance fluorescence were both at the same frequency. To isolate the resonance fluorescence, it was necessary to use a detailed model of the experiment, and to correct for the expected contribution from the scattered laser light.

A conceptually much simpler demonstration of photon anti-bunching was performed by Grangier, Roger, and Aspect in 1986.¹¹ A schematic of their experiment is shown in Fig. 3. They circumvented the problem of background light by using a two-photon cascade in Ca. In this process, a Ca atom absorbs two photons, one each from two lasers operating at frequencies ν_{11} and ν_{12} , promoting it to an excited state. The Ca atom then decays by emitting two photons at different frequencies: one at frequency ν_1 by decaying to a short-lived intermediate level, and a second at frequency ν_2 by decaying to the ground state. All four frequencies are distinct and can be isolated using filters, greatly reducing the problem of scattered background from the intense laser beams. Furthermore, angular momentum conservation ensured that the two photons always were emitted in opposite directions. The detection of one photon at one detector ensured that there would be a photon heading in the opposite direction, so that the first photon could be used as a gate to tag the arrival of the second. Thus, when a gate photon was detected, it was known with high confidence that there was one (and only one) photon incident on the BS.

For this experiment, detections at T and R were condi-

tioned upon detections at the gate detector (detector G). With this conditioning, the measured degree of second-order coherence [Eq. (6)] is given by

$$g^{(2)}(0) = \frac{P_{GTR}}{P_{GT}P_{GR}}. \quad (12)$$

Here, $P_{GT}(P_{GR})$ is the probability of measuring simultaneous photocounts at detector $T(R)$ and detector G , and P_{GTR} is the probability of obtaining a threefold coincidence between detectors T , R , and G . The probabilities can be written as

$$P_{GTR} = \frac{N_{GTR}}{N_G}, \quad P_{GT} = \frac{N_{GT}}{N_G}, \quad P_{GR} = \frac{N_{GR}}{N_G}, \quad (13)$$

where, given a specified time window, $N_{GT}(N_{GR})$ is the number of simultaneous photocounts at detector $T(R)$ and detector G , N_{GTR} is the number of threefold coincidences, and N_G is the number of singles counts at detector G . By using Eq. (13), we can rewrite the experimentally determined degree of second-order coherence [Eq. (12)], as

$$g^{(2)}(0) = \frac{N_{GTR}N_G}{N_{GT}N_{GR}}. \quad (14)$$

In an experimental tour-de-force, Grangier *et al.* measured a second-order coherence of $g^{(2)}(0) = 0.18 \pm 0.06$, which violated the classical inequality $g^{(2)}(0) \geq 1$ by 13 standard deviations.¹¹ In a 5 h experiment, they measured a total of nine threefold coincidences, while a classical wave theory would have predicted greater than 50 threefold coincidences. If the state were a perfect one-photon state, Grangier *et al.* would have measured no threefold coincidences.

We have repeated the experiment of Grangier *et al.*; with advances in technology over the past 15+ years, however, a tour-de-force is no longer required. By using only readily available, off-the-shelf components, we were able to assemble this experiment in an undergraduate laboratory. In a typical run lasting less than 5 min, we measure $g^{(2)}(0) = 0.0188 \pm 0.0067$, where no corrections have been applied to the data to account for accidental coincidences. We have also determined that by accounting for the expected accidental coincidences (see Appendix A), the difference between our result and $g^{(2)}(0) = 0$ (i.e., that expected from a true single-photon incident on the BS) is fully explained by the accidental coincidences.

III. PARAMETRIC DOWNCONVERSION

The key to our ability to perform the experiment is our use of a parametric down-conversion source in place of the atomic Ca cascade source used by Grangier *et al.*¹¹ This method has the advantages of increased simplicity, reduced cost, and increased count rates (several orders of magnitude greater than those observed by Grangier *et al.*) In the process of spontaneous parametric downconversion, a single photon of one frequency is converted into two photons of lower frequency (by approximately a factor of 2). Although down-conversion is extremely inefficient (milliwatts of input power generate output beams that must be detected using photon counting), it is much more efficient than the Ca cascade.

The input is referred to as the pump (at angular frequency ω_p), while the two outputs are referred to as the signal and idler (at angular frequencies ω_s and ω_i). Energy conservation requires that

$$\hbar \omega_p = \hbar \omega_s + \hbar \omega_i, \quad \omega_p = \omega_s + \omega_i. \quad (15)$$

Momentum conservation is equivalent to the classical phase-matching condition, which requires that the wave vectors of the input and output fields satisfy

$$\vec{k}_p = \vec{k}_s + \vec{k}_i. \quad (16)$$

The frequencies and wave vectors are not independent of each other, and are related by the dispersion relation

$$k_p = \frac{n_p(\omega_p) \omega_p}{c}, \quad (17)$$

where $n_p(\omega_p)$ is the index of refraction of the pump wave at the pump frequency, and similarly for the signal and idler waves.

In Type-I downconversion, which is what we use in our experiments, the signal and idler beams are polarized parallel to each other, and their polarization is perpendicular to that of the pump; all polarizations are linear. By proper orientation of the pump beam wave vector \vec{k}_p with respect to the optic axis of the crystal, it is possible to satisfy the constraints imposed in Eqs. (15)–(17). Because only the relative angle between the pump, signal, and idler are important, the downconverted light is emitted into a cone surrounding the pump beam (see, for example, Ref. 9).

Typically, the frequencies of the signal and idler beam are chosen to be equal to each other, at half the frequency (twice the wavelength) of the pump. In order to separate the signal and idler, they are chosen to make a small angle (a few degrees) with the pump beam so that the signal comes out a few degrees from the pump, and the idler comes out a few degrees on the other side of the pump.

However, for a given crystal orientation, there is no unique solution to the constraints imposed in Eqs. (15)–(17). The sums of the frequencies and wave vectors are constrained, but not the individual frequencies and wave vectors. For instance, if the signal frequency is less than half the pump frequency by a certain amount, it is possible for energy to be conserved [Eq. (15)], if the idler frequency is an equal amount greater than half the pump frequency. In order for momentum to be conserved [Eq. (16)], the signal makes a slightly greater angle with respect to the pump, and the idler makes a slightly smaller angle. Thus, the light coming out of a downconversion crystal is emitted into a range of angles (several degrees), and wavelengths (on the order of 10s of nm, centered about twice the pump wavelength).

The similarity between the Ca cascade source used by Grangier *et al.*¹¹ and our downconversion source is that both sources produce two photons, one of which is used as a gate. In our experiment, we use the idler photons as a gate—the detection of an idler photon in one beam (using detector *G*) indicates that there is a signal photon present in the other. The signal beam is directed to a beamsplitter with two detectors at its outputs (detectors *T* and *R*). Just as in the experiment of Grangier *et al.*, we expect to see an absence of coincidences between the *T* and *R* detectors, conditioned on a detection at *G*. This absence is equivalent to an absence of threefold coincidences between *G*, *T*, and *R*. We can use Eq.




Fig. 4. Experimental apparatus. Major components include an ultraviolet laser, downconversion crystal (DCC), polarizing beamsplitter (PBS), single-photon counting modules (SPCM's), and gating, transmission-side, and reflection-side collection optics (*G*, *T*, *R*). Optical fibers direct the light from *G*, *T*, and *R* to their corresponding SPCMs.

(14) as a measure of the second-order coherence of the signal beam, and a result of $g^{(2)}(0) < 1$ is inconsistent with a classical wave description of our system.

IV. EXPERIMENT

We now describe the major components for our updated version of the experiment of Grangier *et al.* The layout of these components is presented in Fig. 4. In brief, a beam of ultraviolet laser light enters a nonlinear crystal where, via spontaneous parametric downconversion, some of the light is converted into IR light in two beams. Light from one of the IR beams (the idler) is used as a *gating beam* and passes directly from the crystal into a photodetector. Light from the other beam (the signal), which we shall call the *experiment beam*, is directed into a 50/50 BS and subsequently observed by photodetectors placed in both the transmission and reflection ports of the beamsplitter. A photodetection in the gating beam is used to signal that the experiment beam has been prepared in the proper single-photon state, and it is the light in the experiment beam whose second-order coherence is measured. Detections from the three detectors *G*, *T*, and *R* are registered by a series of time-to-amplitude converters and single-channel analyzers; coincidence statistics are then compiled and analyzed.

For a more detailed discussion, it is convenient to group components of the instrument into three categories: (i) light source, (ii) light detection, and (iii) coincidence-counting electronics; there also are some diagnostic instruments that make the experiments easier to perform. A list of major components, manufacturers, and part numbers is provided in Appendix C; all of the equipment is commercially available and relatively affordable; a complete parts list and further information is available on our website.²⁷

A. Light source

Our light source was designed to be sufficiently bright so that alignment can be done in real time; we obtain sufficient coincidences in a 100-ms counting window to use the ~ 10 -Hz rate of raw coincidence measurements to perform final alignment. Given this high brightness, data collection occurs over a few minutes (our experimental runs last from approximately 5 to 40 min).

The pump laser used in this work is a cw ultraviolet (409 nm), diode-pumped, frequency-doubled, solid-state laser; the pump is linearly polarized. It was chosen for its turnkey operation, high output power (20 mW), long advertised lifetime (10 000 h) and comparative value. Before entering the downconversion crystal, the pump passes through a zero-order, 400-nm half-wave plate, which allows us to adjust the pump polarization to maximize the downconversion rate (by rotating a half-wave plate in its mount, the direction of linear polarization also rotates.) Downconversion is accomplished in a 5×5 mm aperture, 3-mm-long beta-barium borate (BBO) crystal. It is cut for Type-I downconversion of 405-nm pump light, with a 810-nm signal and idler waves making angles of $\theta = 3^\circ$ with respect to the pump. Because the crystal is hygroscopic, the crystal faces have humidity-barrier, antireflective coatings. The crystal is mounted so that a small flow of dry nitrogen flows over it while in use on the optical table. When not in use, the crystal is stored in a desiccant jar.

When discussing the performance of the source, it is useful to talk in terms of the count rates R , measured in counts per second (cps); $R_G = N_G / \Delta T$, where ΔT is the measurement time, and similarly for other count rates. Our source regularly produces singles count rates in the signal and idler beams (e.g., R_G) of $\sim 110\,000$ cps, and total coincidence rates between the signal and idler beams of ~ 8800 cps (coincidence rates for the counters behind the BS, R_{GT} and R_{GR} are half this value.)

The downconverted light is vertically polarized. Instead of using an ordinary 50/50 BS, we use a combination of a half-wave plate and a polarizing beamsplitter (PBS). The half-wave plate is adjusted so that the light entering the PBS is polarized at 45° with respect to the polarization axis of the PBS; the light then splits equally between the two outputs. By rotating the half-wave plate, we can adjust the input polarization (and hence the splitting ratio), allowing us to fine tune the splitting to be as close to 50/50 as possible. We also can easily transmit or reflect 100% of the beam, which is useful during alignment.

B. Light detection

Our light collection optics are designed for ease of alignment and ambient light rejection. The use of fiber optic cables also makes the system very flexible. We highly recommend that anyone performing experiments with downconverted light consider using a similar fiber-based system.

The collection optics and detection systems for the three detectors (G , T , and R) are identical. Downconverted light is collected by a converging lens and focused into the end of a $62.5\text{-}\mu\text{m}$ -diameter, 1-m-long multimode fiber optic cable that has fiber-coupling (FC) connectors on both ends. The lens is a fiber-coupling lens (Thorlabs F220FC-B), and is pre-aligned to place the tip of the fiber cable in the focal plane of the lens, so that no alignment of the lens to the fiber is necessary. The other end of the fiber connects to a fiber-to-

fiber coupler, which couples light into a second, identical, fiber. This arrangement allows us the flexibility of swapping connections between the coupling lenses and different detectors, which is useful in setting up the coincidence counting electronics (detailed below). It also allows us to easily connect a fiber-coupled laser diode which shines light backward through the coupling lens onto the downconversion crystal for alignment purposes (detailed below). The second fiber carries the downconverted light into a light-tight enclosure which houses the optical filters and detectors. The only light entering this enclosure comes through the fibers, which greatly eliminates problems with stray light.

Light is coupled out of the second fiber with another fiber-coupling lens, passes through an RG 780 filter (which passes wavelengths longer than 780 nm), and is coupled with a third lens into a third fiber cable ($50\text{-}\mu\text{m}$ diameter, FC connectors, and an opaque jacket). We use a kinematic mount to align the output of one lens with the input of the other. We also surround the lenses and filter with beam tubes to further eliminate the possibility of collecting ambient light. The third fiber cable transports the light to the single-photon counting module (SPCM), which has its own FC connector which is pre-aligned to image the fiber tip onto the active area of the detector.

The SPCMs use an avalanche photodiode operated in Geiger mode to detect the light. They output a 30 ns, 4.5 V (into $50\ \Omega$) pulse when they detect a photon, with a 50 ns dead time between pulses. The SPCMs have a specified quantum efficiency of $\sim 50\%$ at 810 nm, and the model we used had dark count rates of ~ 250 cps. With this dark count rate and our 2.5-ns coincidence window, coincidences due to dark counts are negligible.

C. Coincidence counting electronics

As described above, we are interested in detecting coincidence counts between the outputs of different detectors. We use a coincidence window of 2.5 ns, and coincidences are determined using a combination of a time-to-amplitude converter (TAC) and a single-channel analyzer (SCA). Three such coincidence units are used (one each for GT , GR , and GTR coincidences), and their outputs are recorded by a counting board in our computer. We briefly describe how we use the TAC/SCA to determine twofold GT coincidences. (GR coincidences are determined in the same manner). Modification of the TAC/SCA configuration to obtain threefold GTR coincidences also is described.

A TAC operates by receiving two inputs, called START and STOP, and then outputting a pulse, the amplitude of which is proportional to the time interval between the rising edges of the START and STOP signals. The proportionality between the amplitude and the time interval is controlled by the gain of the TAC, and we typically use a value of 0.2 V/ns. To measure GT coincidences, the START input comes from the output of detector G and the STOP input comes from detector T (see Fig. 5). To ensure that the START pulse precedes the STOP pulse, we insert an extra length of coaxial cable, corresponding to a delay of ~ 6 ns, between T and the STOP input. Thus, if detectors G and T record simultaneous detections, the delay between START and STOP signals is 6 ns, and the output from the TAC is 1.2 V.

The SCA operates by receiving an input pulse, and then outputting a pulse (with an amplitude of 5 V) *only* if the amplitude of the input pulse falls within a certain voltage window. The width of the window is adjustable, as is the




Fig. 5. Coincidence counting electronics. TACs and SCAs are used to identify GT , GR , and GTR coincident detections. Here ST , SP , and GT refer to START, STOP, and START GATE inputs, respectively. The outputs go to six input channels on the counter.

lower level of the window. The input to the SCA is the output from the TAC. Using the values for the TAC output above, a coincidence window of 2.5 ns centered about 6 ns corresponds to a voltage range of 0.95–1.45 V, and our SCA is configured to output a pulse if the amplitude of the input pulse lies within this range. The only trick to configuring the TAC/SCA setup is in properly setting the SCA window to maximize true coincidences and reject false coincidences. This procedure is described in Appendix B.

In order to measure the threefold GTR coincidences, we use T as the START input and R as the STOP input, and configure the TAC/SCA as described above to register TR coincidences. To ensure that these TR coincidences also are coincident with a detection at G , we operate the TAC in “start gate coincidence” mode, and feed the G signal to the START GATE input of the TAC. If an output pulse from G is not present at the START GATE when the pulse from T arrives at START, then the timing circuitry in the TAC is disabled, and no output is produced.

There is one last trick used in setting up this threefold coincidence unit. The technique for setting the SCA window described in Appendix B relies on observing coincidences between the detectors measuring the START and STOP input; however, we expect an *absence* of coincidences between T and R . In order to obtain coincidences between these detectors so that we can set the window, we switch the fiber optic cables so that the idler (gate) beam is fed into the detector that ordinarily measures the R output. Now, we have coincident photons entering the two detectors, so that we can set the window as described in Appendix B. The delays are all set by the coaxial (electrical) cables between the detectors and the coincidence units. Because all of the fiber cables have the same length, the optical delays are the same, and switching the fiber cables back after the window is set does not affect the timing.

We measure a total of six photocounts in each data acquisition interval: singles counts from each of the three detectors, N_G , N_R , and N_T , as well as the coincidence counts N_{GR} , N_{GT} , and N_{GTR} . We use a counting board that plugs into a PCI slot in our computer, and it simultaneously records these counts on six different channels. A LABVIEW program reads the data from the board, computes the second-order coherence [Eq. (14)], and saves the data.

D. Optical alignment

Although requiring some care, we have found the setup and alignment process to be sufficiently straightforward that two undergraduates having some familiarity with the experiment were able to start from a bare optical table and complete the process with minimal supervision over the course of one or two days. The alignment is robust once it has been completed. For example, we remove the downconversion crystal when it is not in use; reinserting the crystal and tweaking-up the alignment takes only a few minutes. When starting from scratch, the major components are first affixed to the optical table in rough alignment as illustrated in Fig. 4; although at first we are interested solely in obtaining coincidences between the idler and signal beams, so that the half-wave plate, PBS, and detector R are not used. The pump beam is aligned level to the table using the two mirrors, and the electrical connections are completed.

The first component to be aligned is the collection optics (i.e., fiber optic cable/lens assembly) for the G detector. The collection lens is mounted in a kinematic mount that allows for horizontal, vertical, and angular adjustment, and the center height of the lens is adjustable using a post holder and post. The height is initially adjusted so that the center of the lens is at the same height as the pump beam. Light from a fiber-coupled 780 nm laser diode is coupled (via the fiber-to-fiber coupler) backward through the fiber cable, and out through the lens. The lens is placed so that angle of this beam is set to be 3° off of the pump beam, and the mount is adjusted so that the laser shines back onto the center of the downconversion crystal. The alignment laser is now removed and the fiber cable is connected to the detector. By monitoring the count rate from the G detector, the polarization of the pump and the horizontal and vertical tilt of the downconversion crystal are adjusted to maximize the count rate. Once this adjustment has been accomplished, the kinematic mount controlling the alignment of the two lenses surrounding the RG780 filter also is adjusted to maximize the count rate. Now, the horizontal placement of the collection lens (and hence the angle between the collection lens and the pump beam) is carefully adjusted to maximize the count rate. The alignment of the downconversion crystal and placement of the collection beam is then iterated to maximize the count rate on the G detector. As stated above, we typically obtain $\sim 100\,000$ cps on this beam.

Next, the T detector is aligned in nearly the same way. At first the goal is not to painstakingly align this detector for maximum counts, but simply to get enough counts so that the coincidence window between the G and T detectors can be set as described in Appendix B. Once this alignment is set, the alignment of the T collection optics is adjusted to maximize the *coincidence* rate between G and T , not the raw count rate on T . We easily obtain a coincidence rate of over 7000 cps, and frequently achieve a rate of ~ 8800 cps. Once the alignment of the T optics has been accomplished, the alignment laser is shone backward through this optics, and adjustable iris diaphragms are aligned with the beam in between the downconversion crystal and the T optics. These irises serve to identify the beam path and assist in aligning the R detector.

The half-wave plate and PBS are now inserted, and the alignment laser is shone backward through the collection optics of the R detector. The optics are adjusted so that the light goes back through the irises and onto the crystal. The R detector is connected and when counts are achieved, the win-

Table I. Measurements of $g^{(2)}(0)$, the degree of second-order coherence.

Trial (Total acquisition time, min)	Integration time per point (s)	Number of points	$g^{(2)}(0)$	Standard deviation of $g^{(2)}(0)$
1 (~5)	2.7	110	0.0188	0.0067
2 (~10)	5.4	108	0.0180	0.0041
3 (~20)	11.7	103	0.0191	0.0035
4 (~40)	23.4	100	0.0177	0.0026

dow on the *GR* coincidence unit is set, and the *GR* coincidence count is maximized. Lastly, the *GTR* coincidence unit is configured using the procedure described above.

V. RESULTS

One of the primary advantages of the apparatus described in this paper is the ability to acquire good counting statistics in time periods reasonable for an undergraduate laboratory. In Table I, we present the results of four experimental runs. In each of these runs, we performed ~ 100 measurements of $g^{(2)}(0)$, while in each run we changed the integration time for each measurement. These results are clearly inconsistent with a classical wave theory, which predicts $g^{(2)}(0) \geq 1$. Even for counting times of less than 5 min, we obtain a value of $g^{(2)}(0)$ that is lower than the classical lower limit by 146 standard deviations. Increasing the counting time does not affect the measured value of $g^{(2)}(0)$ (to within the statistical error of our measurement), but increasing the counting time does decrease the statistical error. Our best results is $g^{(2)}(0) = 0.0177 \pm 0.0026$, which violates the classical inequality by 377 standard deviations.

If a truly single-photon state were incident on the BS, QM would predict that $g^{(2)}(0) = 0$. Why don't we see this? A consequence of defining a "coincidence" with a finite time window is an expected nonzero anticorrelation parameter. This is because there is the possibility that uncorrelated photons from different downconversion events may hit the *T* and *R* detectors within our finite coincidence window; these are "accidental" coincidences. As the count rates and coincidence window increase, so do the number of accidental coincidences. In Appendix A, we analyze the effect of these accidental coincidences on our measurements of $g^{(2)}(0)$. For our experimental parameters, when we account for accidental coincidences, we calculate an expected value of the second-order coherence $g^{(2)}(0) = 0.0164$, which is what we observe to within our statistical error.

As a final check on the instrument and method, the measurement was repeated, but with an extra length of coaxial cable (corresponding to a delay of 6 ns) placed after the *R* detector. In this case we are not measuring true coincidences, but instead coincidences between measurements made at time t at detector *R*, and time $t + 6$ ns for detector *T*. This means we do not measure the quantity $g^{(2)}(\tau = 0)$, but instead we measure $g^{(2)}(\tau = 6 \text{ ns})$. Under such circumstances, *GTR* coincidences are not excluded because we expect the detections at *T* and *R* to be due to uncorrelated downconversion events. Indeed, we would expect to obtain a measured

value $g^{(2)}(\tau = 6 \text{ ns}) \geq 1$. In two different experimental runs we obtained measured values for $g^{(2)}(\tau = 6 \text{ ns})$ in the range 2–3.

VI. CONCLUSIONS

We have performed an experiment whose results cannot be explained using a classical wave description for light. The results are consistent with a quantum mechanical description in which a field in a single-photon state is incident on a beamsplitter, and as such we take this experiment as proof of the existence of photons. The experiment is conceptually simple, and is suitable for an undergraduate laboratory.

While we would not describe the cost of this experiment as inexpensive (total cost of $\sim \$40\,000$), the cost is not prohibitive; a more detailed discussion of this cost is presented in Appendix C, along with the parts list. Furthermore, the equipment is extremely versatile and can be used for a number of other experiments. By adding approximately $\$2,500$ in components, we have extended the work described here to demonstrate that (i) single photons interfere with themselves as they pass through the two arms of an interferometer, and (ii) that the frequencies of the signal and idler beams generated in our experiment are highly correlated. These experiments will be described in a future publication. With other small additional purchases, it will be possible to perform tests of Bell's inequalities^{9,10} and to demonstrate two-photon interference.^{28,29} Thus, for less than $\$50\,000$, one could implement five experiments suitable for undergraduates that demonstrate interesting features of quantum mechanics. While the total cost is not inexpensive, it is most certainly cost effective.

ACKNOWLEDGMENTS

We wish to acknowledge several helpful discussions with David Branning. We also acknowledge financial support from the National Science Foundation and from Whitman College.

Note added in proof. We have recently become aware that similar experiments have been carried out by a group at Colgate University.³⁰

APPENDIX A: ACCIDENTAL COINCIDENCES

The time interval defining a coincidence is determined by the windowing of the SCAs. Specifically, the SCAs are configured with finite time windows of $\Delta t = 2.5 \text{ ns}$, giving the term "coincidence" the meaning "within 2.5 ns". A consequence of this finite window is a finite probability, proportional to Δt , of registering *GTR* coincidences that have no relation to the coincidences of interest; these are accidental coincidences.

For example, suppose we obtain a valid coincidence between detectors *G* and *T*, which occurs with probability $P_{GT} = R_{GT}/R_G$, where we have written Eq. (13) in terms of the count rates. Within Δt of this coincidence, there is a random chance that the *R* detector also will measure a count, leading to an accidental threefold coincidence. If the time interval Δt is small enough, then the probability of this random *R* detection occurring can be approximated by $P'_R \approx R_R \Delta t$, where the prime indicates that this is an accidental event, occurring within a specific time window. Similarly, a valid *GR* coincidence and a chance detection at *T* also will

yield an accidental threefold coincidence. The probability of the accidental coincidences can then be written as

$$P'_{GTR} = P_{GT}P'_R + P_{GR}P'_T = P_{GT}R_R\Delta t + P_{GR}R_T\Delta t. \quad (18)$$

Here, we have ignored the probability that the accidental threefold coincidences may be due to pure chance detections at all three detectors, because for our count rates and coincidence window, this probability is negligible.

We can now calculate the effect of these accidental coincidences on the second-order coherence. Substituting Eq. (18) into Eq. (12) yields

$$\begin{aligned} g^{(2)}(0) &= \frac{P_{GTR}}{P_{GT}P_{GR}} \\ &= \frac{P_{GT}R_R\Delta t + P_{GR}R_T\Delta t}{P_{GT}P_{GR}} \\ &= \frac{R_R\Delta t}{P_{GR}} + \frac{R_T\Delta t}{P_{GT}} = R_G\Delta t \left(\frac{R_R}{R_{GR}} + \frac{R_T}{R_{GT}} \right). \end{aligned} \quad (19)$$

Using the average count rates obtained from the data collected during trial 4 of Table I, we calculate the contribution to $g^{(2)}(0)$ from the accidental coincidences to be 0.0164.

APPENDIX B: SETTING THE SINGLE-CHANNEL-ANALYZER WINDOW

One technique for setting the voltage window of the SCA is to simply set the window width to some value, and to slowly adjust the lower level of the window while monitoring the SCA output. The goal is to maximize the coincidence rate. The window is then adjusted to be just wide enough so that a further increase in width does not significantly increase the count rate. Adjustments of the width and lower level of the window can be iterated to optimize the count rate.

An easier way to set the SCA window is to use a multi-channel analyzer (MCA). Our MCA is on a PCI card that plugs into our computer and comes with its own software. We use it as a diagnostic tool for setting the window, but do not use it in our experiments to determine $g^{(2)}(0)$. An MCA histograms voltage pulses of varying amplitude. The histogram is displayed in real time, with an update rate of a few hertz, so that one can watch the histogram build over time. The input to the MCA is the output from the TAC, so that the histogram can be interpreted as measuring time intervals instead of voltages. As stated above, coincidence counts at G and T are separated at the TAC by 6 ns, and with the coincidence rates in our experiment, we easily see a peak in the histogram generated by the MCA centered at this 6-ns time delay. The wings of this peak extend outward to a width of approximately 2.5 ns, which is the reason we chose this value for our coincidence window. (This width is due almost entirely to the properties of the SPCMs, as the time interval between the photon pairs produced in our experiment is certainly much less than this.³¹) Uncorrelated photodetections (arising from G seeing a photon from one pair-production event and T seeing a photon from a different pair-production event) contribute a uniform background that the coincidence peak sits on top of.

Simply looking at the output of the TAC on the MCA displays the coincidence peak, but yields no information about the window of the SCA, which is what we are really interested in. In order to set the SCA window, we throw a switch on the TAC/SCA unit, which causes the SCA to win-

now the output of the TAC. In this mode of operation, if the output amplitude of the TAC falls within the voltage window set by the SCA, the TAC operates normally and outputs a voltage proportional to the time difference between the START and STOP pulses. However, if the output of the TAC falls outside of the SCA window, the TAC output is inhibited and there is no TAC output. Thus, if the SCA window is not properly set, no peak appears in the MCA histogram. We thus set the SCA window by simply monitoring the MCA histogram and adjusting the SCA controls until only the coincidence peak is seen and the uncorrelated background is eliminated.

APPENDIX C: PARTS LIST AND COST OF EXPERIMENT

Here, we list the major components for this experiment. A detailed list of all the components can be found on our website.²⁷

Pump Laser: Edmund Industrial Optics, (<http://www.edmundoptics.com/>). Diode-pumped, frequency-doubled, solid-state laser (405–410 nm); model NT55-872; \$5,800.

Downconversion Crystal: Cleveland Crystals, (<http://www.clevelandcrystals.com/>). Beta-barium borate (BBO) crystal, 3 mm long, for converting a cw 405-nm input to an 810-nm output, 3° cone angle on signal and idler, XH0503 housing with a 5-mm aperture, Humidity-barrier antireflective coatings on the crystal faces, nitrogen purge connections, and no windows; \$2,160.

Single-Photon Counting Modules: Pacer Components, (<http://www.pacer.co.uk/>). Single-photon counting module (Perkin Elmer model SPCM-AQR-13-FC), dark count less than 250 cps, FC fiber connector; quantity 3; \$4,300/each.

Counting Electronics: ORTEC, (<http://www.ortec-online.com/>). TAC/SCA model 567; nuclear instruments modular (NIM) plug-in module; quantity 3; \$1,656/each. These modules plug into a NIM crate with associated power supply, which we already had available to us. If a NIM crate is needed, the ORTEC model 4001A/4002D @\$2,500 should be suitable. MCA model TRUMP-PCI-2K (diagnostic for setting up SCA window); PCI plug-in card with software; \$2,370.

Counter: National Instruments, (<http://www.ni.com/>). 8-channel counter/timer model PCI-6602; plug-in card. (Note that the optional BNC-2121 connector block and SH68-68-D1 shielded cable greatly simplify connecting to the counter). Total cost with options, \$1,000.

Alignment Laser and Power Supply: Thorlabs, (<http://www.thorlabs.com/>). 785-nm laser coupled to a single-mode fiber with FC connector; model LPS-4224-785-FC; \$400. ILX Lightwave, (<http://www.ilxlightwave.com/>). Current source model LDX-3412; \$930.

The total cost of the experiment is ~\$40 000. This cost includes all of the equipment necessary to carry out the experiments, with the exception of a computer, LABVIEW software, and the 3×5 foot optical table. The experiment does not require a full optical table—an optical breadboard would be sufficient, and it should be possible to fit everything on a 3×4 ft breadboard. Below, we discuss a few possibilities for reducing the cost of the experiment.

Approximately \$2,500 of the cost is for standard optical components: mirrors, kinematic mounts, posts, etc. A laboratory with a stock of such components could save much of

this cost. The fiber-coupled alignment laser and power supply are nothing special—any available laser coupled into a fiber would suffice.

Another opportunity for reducing cost is in the counting electronics. Many laboratories have some of the necessary electronics as part of existing nuclear physics experiments. Thus, it may be possible to save on the TAC/SCA cost and the cost of a NIM crate. It also is possible to build coincidence counting electronics from integrated circuits,⁹ eliminating nearly \$7,400 in cost.

After our apparatus was assembled, Perkin Elmer introduced the SPCM-AQ4C, which consists of four fiber-coupled photon-counting modules in one unit. These modules have a larger dark count rate (500 cps), but that should have little or no effect on the experiments described here. The cost of this unit is \$9,000, which is significantly cheaper than purchasing three separate counters. If we were building a new system, we would use this unit.

We do not recommend replacing the avalanche photodiode-based photon-counting modules with photomultiplier tubes (PMTs). Quantum efficiencies of most PMTs above 800 nm are about 100 times smaller than avalanche photodiodes, meaning that the count rates would be 100 times lower. PMTs with GaAs photocathodes have efficiencies that are only 5 or 6 times lower than avalanche photodiodes, which is not too bad. However, PMT-based systems further require the use of cooled housings, high-voltage (HV) power supplies, discriminators, and possibly high speed amplifiers; this increases the cost and complexity of PMT-based systems. Even PMT-based photon-counting modules (which incorporate the housing, HV supply, and discriminator) operate significantly better with external temperature control circuitry, which makes their cost higher than the avalanche diode systems.

If a high-compliance voltage laser diode current source is already available, stand-alone 30 mW blue laser diodes are available directly from Nichia, (www.nichia.com), for \$2,000 each. It would be advisable to have a temperature-controlled mount for this diode. If a laboratory already has an Ar-ion laser, it could be used. The bluest line with significant power (10 s of mW) is typically at 458 nm, which places the downconversion at 916 nm. There is a slight reduction in quantum efficiency of the avalanche diodes at this wavelength, but we envision no significant obstacles to operating at this wavelength.

^aElectronic mail: beckmk@whitman.edu

¹R. Q. Stanley, “Question #45: What (if anything) does the photoelectric effect teach us?,” *Am. J. Phys.* **64**, 839 (1996).

²P. W. Milonni, “Answer to Question #45: What (if anything) does the photoelectric effect teach us?,” *Am. J. Phys.* **65**, 11–12 (1997).

³W. E. Lamb, Jr. and M. O. Scully, “The photoelectric effect without photons,” in *Polarization, Matière et Rayonnement* (Presses University de France, Paris, 1969).

⁴L. Mandel, “The case for and against semiclassical radiation theory,” in *Progress in Optics*, edited by E. Wolf (North-Holland, Amsterdam, 1976), Vol. XIII, pp. 27–68.

⁵D. C. Burnham and D. L. Weinberg, “Observation of simultaneity in para-

metric production of optical photon pairs,” *Phys. Rev. Lett.* **25**, 84–87 (1970).

⁶J. F. Clauser, “Experimental distinction between the quantum and classical field-theoretic predictions for the photoelectric effect,” *Phys. Rev. D* **9**, 853–860 (1974).

⁷H. J. Kimble, M. Dagenais, and L. Mandel, “Photon antibunching in resonance fluorescence,” *Phys. Rev. Lett.* **39**, 691–695 (1977).

⁸A. C. Funk and M. Beck, “Sub-Poissonian photocurrent statistics: Theory and undergraduate experiment,” *Am. J. Phys.* **65**, 492–500 (1997).

⁹D. Dehlinger and M. W. Mitchell, “Entangled photon apparatus for the undergraduate laboratory,” *Am. J. Phys.* **70**, 898–902 (2002).

¹⁰D. Dehlinger and M. W. Mitchell, “Entangled photons, nonlocality, and Bell inequalities in the undergraduate laboratory,” *Am. J. Phys.* **70**, 903–910 (2002).

¹¹P. Grangier, G. Roger, and A. Aspect, “Experimental evidence for a photon anticorrelation effect on a beam splitter: A new light on single-photon interferences,” *Europhys. Lett.* **1**, 173–179 (1986).

¹²G. Greenstein and A. G. Zajonc, *The Quantum Challenge, Modern Research on the Foundations of Quantum Theory* (Jones and Bartlett, Sudbury, MA, 1997).

¹³R. Hanbury Brown and R. Q. Twiss, “Correlation between photons in two coherent beams of light,” *Nature (London)* **177**, 27–29 (1956).

¹⁴An excellent compilation of reprints, including many of the articles referenced here, is L. Mandel and E. Wolf, *Selected Papers on Coherence and Fluctuations of Light (1850–1966)* (SPIE, Bellingham, WA, 1990).

¹⁵R. Hanbury Brown and R. Q. Twiss, “A test of a new type of stellar interferometer on Sirius,” *Nature (London)* **178**, 1046–1048 (1956).

¹⁶E. Brannen and H. I. S. Ferguson, “The question of correlation between photons in coherent light rays,” *Nature (London)* **178**, 481–482 (1956).

¹⁷E. M. Purcell, “The question of correlation between photons in coherent light rays,” *Nature (London)* **178**, 1449–1450 (1956).

¹⁸R. Hanbury Brown and R. Q. Twiss, “Interferometry of the intensity fluctuations in light I. Basic theory: the correlation between photons in coherent beams of radiation,” *Proc. R. Soc. London, Ser. A* **242**, 300–324 (1957).

¹⁹R. Q. Twiss, A. G. Little, and R. Hanbury Brown, “Correlation between photons in coherent beams of light, detected by a coincidence counting technique,” *Nature (London)* **180**, 324–326 (1957).

²⁰R. Loudon, *The Quantum Theory of Light*, 3rd ed. (Clarendon, Oxford, 2000).

²¹A way to motivate the validity of this inequality is to note that the variance of the intensity must be a positive number. Since the variance can be written as $\Delta I^2 = \langle I^2 \rangle - \langle I \rangle^2 \geq 0$, it must be true that $\langle I^2 \rangle \geq \langle I \rangle^2$.

²²F. T. Arecchi, E. Gatti, and A. Sona, “Time distribution of photons from coherent and Gaussian sources,” *Phys. Lett.* **20**, 27–29 (1966).

²³L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge U.P., Cambridge, 1995).

²⁴P. L. Kelly and W. H. Kleiner, “Theory of electromagnetic field measurement and photoelectron counting,” *Phys. Rev.* **136**, A316–A334 (1964).

²⁵R. J. Glauber, “Optical coherence and photon statistics,” in *Quantum Optics and Electronics*, edited by C. DeWitt-Moretz, A. Blandin, and C. Cohen-Tannoudji (Gordon and Breach, New York, 1965), pp. 63–185.

²⁶Z. Y. Ou, C. K. Hong, and L. Mandel, “Relation between input and output states for a beamsplitter,” *Opt. Commun.* **63**, 118–122 (1987).

²⁷<http://www.whitman.edu/~beckmk/QM/>

²⁸C. K. Hong, Z. Y. Ou, and L. Mandel, “Measurement of subpicosecond time intervals between two photons by interference,” *Phys. Rev. Lett.* **59**, 2044–2046 (1987).

²⁹C. H. Holbrow, E. Galvez, and M. E. Parks, “Photon quantum mechanics and beamsplitters,” *Am. J. Phys.* **70**, 260–265 (2002).

³⁰E. J. Galvez *et al.*, personal communication, 2004.

³¹S. Friberg, C. K. Hong, and L. Mandel, “Measurement of time delays in the parametric production of photon pairs,” *Phys. Rev. Lett.* **54**, 2011–2013 (1985).