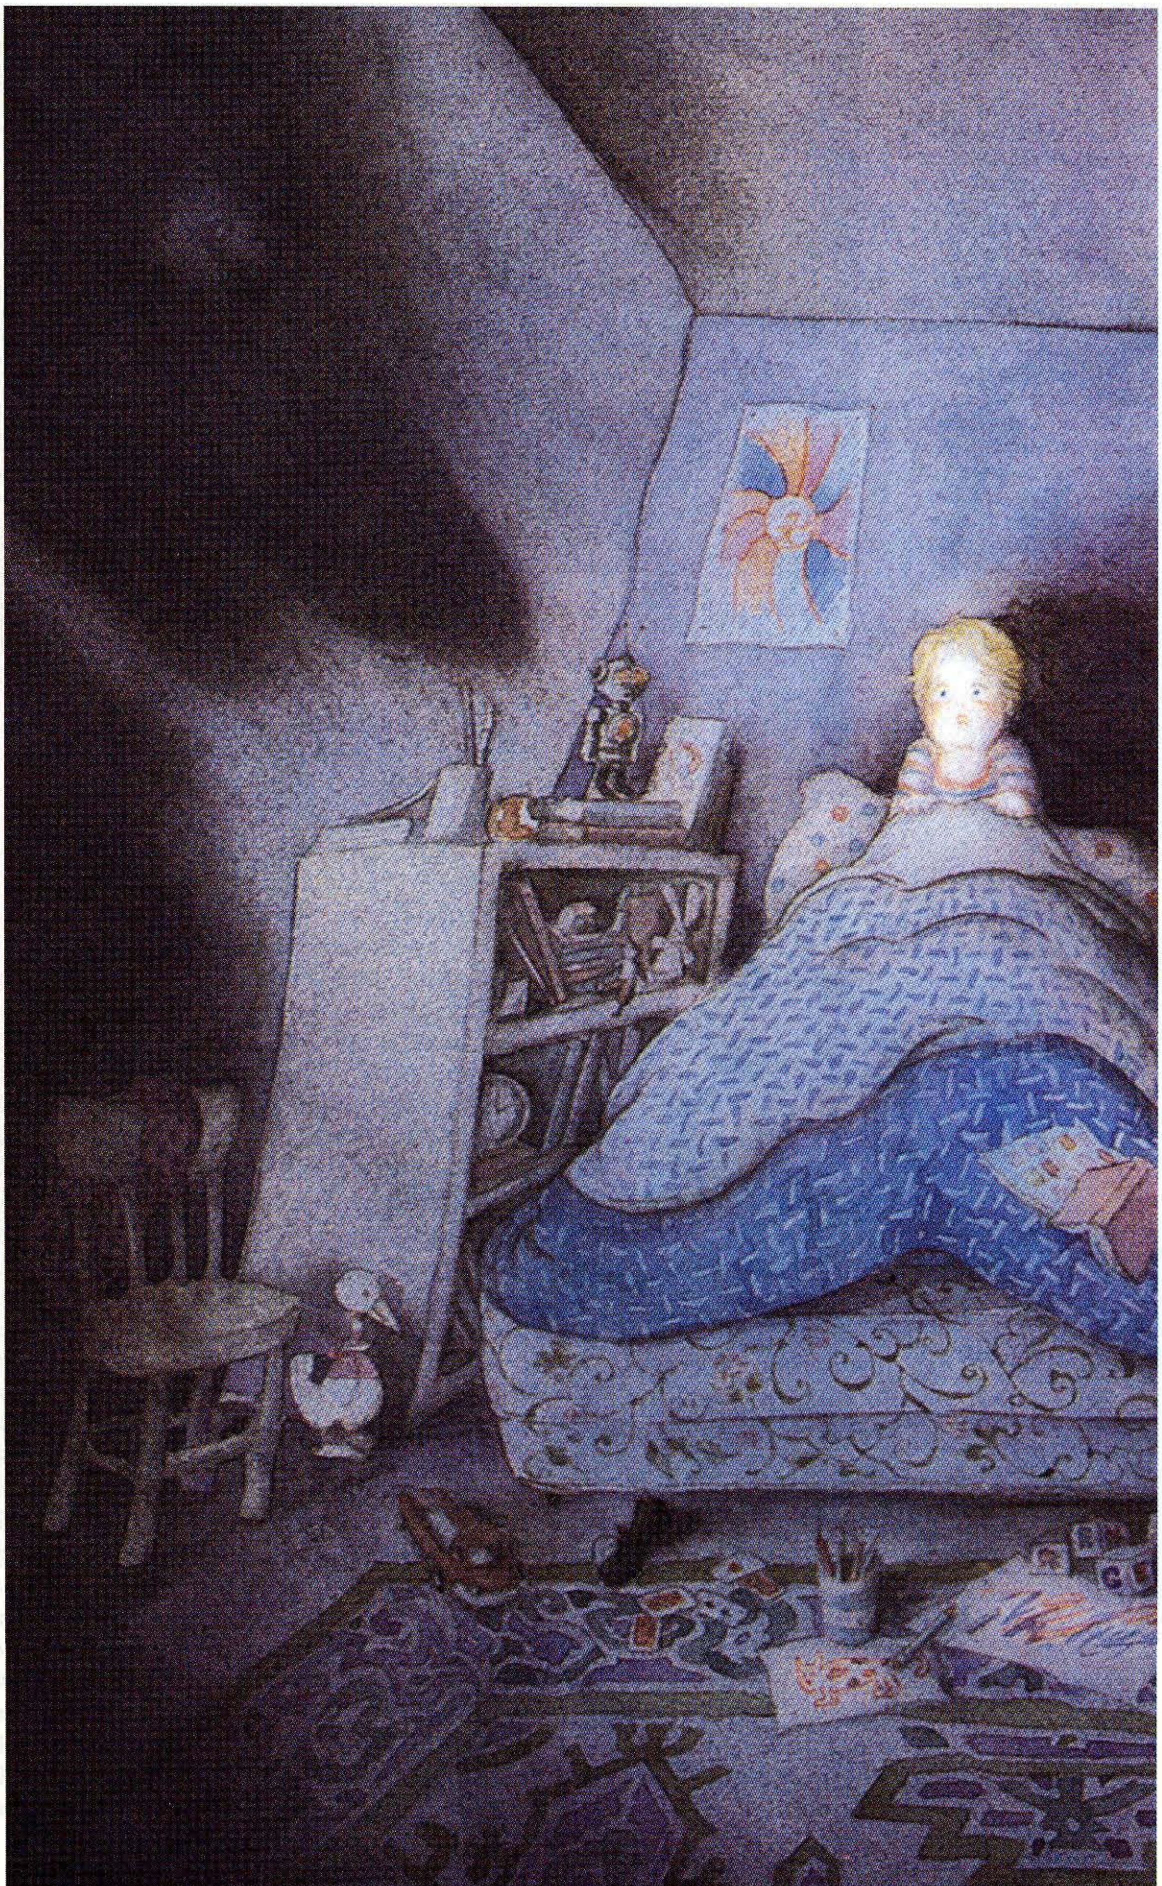

New Light on the Nature of Darkness

by H. Jeff Kimble

Dark is, in fact, an altogether more interesting character than is Light—at least light as most people understand these two characters.

In a book called *Alfi and the Dark* that my daughters Katie and Megan have generously lent me, the hero, a young man called Alfi, has a dialogue with the Dark. Although this is a child's story, it would be hard to find a better book in Millikan Library's physics collection to introduce my subject.

"Alfi was lying asleep in his bed
When he suddenly woke with a thought and he said,
'If I switch on the lights I'll be able to see
But where will the Dark go? Where will it be?'"

Alfi's question, in fact, is one of the central themes of my story. What is darkness? Where does it go in the presence of light? A long dialogue ensues between Alfi and the Dark, and each comes to understand the other somewhat better. Alfi learns that the Dark isn't such a happy fellow. Indeed,

"Dark felt so lonely. Dark felt so sad,
As he thought of the fun and the friends that Light had.
Wherever *he* went, people seemed to be scared.
He wanted a friend, just *someone* who cared."

In the end they become friends, and Dark reveals his secret to Alfi.

"Dark was so happy he laughed with delight.
'Now, I'll tell where I go when you switch on the light.
The answer is simple and you'll be amazed—
I NEVER GO ANYWHERE!' Alfi was dazed."

In the spirit of this book, my purpose here is to convey something about the modern view of darkness, and in the process to avoid Alfi's state of bewilderment at left.

Picture and text from the book *Alfi and the Dark* by Sally Miles and Errol Le Cane. Copyright 1988. Published by Chronicle Books, San Francisco.

The objectives are really twofold: First, to convince you that Dark is, in fact, an altogether more interesting character than is Light—at least light as most people understand these two characters; and secondly to tell you about the activities of the "Friends of Darkness"—that is, the graduate students and senior scientists in the quantum optics group here at Caltech. The experimental results I'll tell you about are really due to their hard work. In addition, I should note at the outset that the conceptual foundation for much of this research was laid by Caltech's Carlton Caves [PhD '79], visiting associate in physics; Kip Thorne [BS '62], Feynman Professor of Theoretical Physics; and Ron Drever, professor of physics; and by their collaborator (and frequent visitor to Caltech), Professor Vladimir Braginsky of Moscow State University.

Since light is fundamentally a wave phenomenon, we should get straight a few basic concepts about waves. Imagine that you're sitting on a raft in the ocean. As the waves pass, you will bob up and down. Instead of being waves in water, light is an oscillation of the electromagnetic field, so that if you were an electron immersed in the field—that is, if you had a charge—you would bob up and down as the light wave goes by. A raft in the ocean bobs every several seconds. By contrast, an electron bathed in red light oscillates with a frequency of 5×10^{14} cycles per second—roughly a million billion times per second. And while the distance between crests of ocean waves —the wavelength—might be a few meters (a dozen feet or so), red light's wavelength is only 6×10^{-7} meters or 6,000 Ångstroms (about

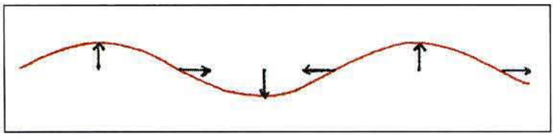

Above: How a rotating arrow can represent a wave.
Below: Changes in amplitude change the arrow's length.
Bottom: Changes in phase tilt the arrow.

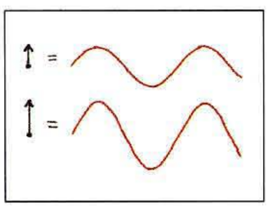

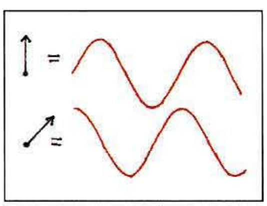

2/100,000ths of an inch, or roughly 100 times smaller than the thickness of this page). So visible light's wavelength is very short and its frequency is very high, which leads to some technical problems that make the experiments I'll be describing somewhat tricky to do.

Now, just as two variables, position and velocity, are required to describe the motion of a person, we require two variables to describe light—or, in general, any wave. These two variables are amplitude and phase, and they form the basis for our discussion of the physics of light. The amplitude of a wave is simply the height of the wave's crests, which translates into "brightness" for light. The phase specifies the time (or distance) between zero crossings—the points where the wave's amplitude is zero—and is thus related to the wave's frequency.

Now, with the objective of creating a more precise and powerful language to talk about light, let me get rid of the wave altogether and replace it by an arrow that rotates like the hand on a clock. The arrow's frequency of rotation represents the wave's fundamental frequency of oscillation. So, when the wave is at its peak, the arrow points to 12 o'clock, as shown above. As the wave's amplitude comes down to zero, the arrow rotates to 3 o'clock. At the wave's trough, the arrow is at 6 o'clock, and when the wave comes back through zero amplitude, the arrow reaches 9 o'clock. Now I don't really want to try rotating the arrow at the frequency of red light, so let's sit still—so that the arrow is stationary in our frame of reference—and assume that the world is instead rotating around us at 5×10^{14} times per second.

Translating our two variables for light into arrow language, we see that the arrow's length gives the light's amplitude, and the arrow's orientation gives the phase. If we increase the light's intensity, the arrow gets longer. If we instead change the light's phase, we tilt the arrow—that is, imagine two arrows spinning at the same rate, and hence fixed in our rotating frame of reference, but with one arrow tipped relative to the other; this is a difference in phase. Hence changes along the length of the arrow are amplitude changes, while deviations of the tip of the arrow perpendicular to its length (with its tail pinned down) are phase deviations.

Anyone who's tried to bodysurf knows that waves have certain irregularities. If you plotted the amplitude of successive ocean waves against their arrival time relative to the preceding wave, you'd find a spread of points clustering around the average wave amplitude and average time between waves. Likewise, light waves—even from a laser—are not perfectly regular, either. There are slight fluctuations in both amplitude and phase for any beam of light. Physicists call these fluctuations "noise" to indicate their random character, and to represent this noise, we "fuzz out" the tip of our arrow, so that its exact length (amplitude) and angle (phase) are now uncertain. I'll call this region of fuzziness a "noise blob." The larger the blob, the noisier the light.

I now want to turn to the fundamental rules and regulations specific to the electromagnetic field—that is, to light. What does physics have to say about the intrinsic amount of noise in a light wave? Nature's rule is simply that the product of the blob's noise in the amplitude dimension times the noise in the phase dimension has a minimum value set by Planck's constant. That is to say, because light is a quantum field, our noise blobs must have a minimum area. This is the Heisenberg uncertainty principle for light—the amplitude and phase of a beam of light cannot be precisely determined simultaneously, even in principle. (Heisenberg formulated the uncertainty principle to explain the quantum behavior of atoms and electrons; it is a direct and unavoidable consequence of the quantum theory.) Note that Planck's constant is a fundamental constant of nature—it sets the scale for the "graininess" of the atomic world. Light is quite remarkable in that, as we will see, this fundamental graininess leads to fluctuations in amplitude and phase that can have import, not only at the atomic level, but also in our macroscopic world. It is worth emphasizing that the fluctuations demanded by quantum mechanics are intrinsic and fundamentally unavoidable. Hence the slogan, "Quantum mechanics—It's not just a good idea; it's the law!"

So now that we know something about what light is, we can talk about what darkness is. In terms of our picture of light as an arrow with a quantum noise blob on its end, we simply shrink the arrow's length to zero, leaving only the blob. Thus zero isn't really zero; it is zero plus or minus the noise of the residual quantum blob, as set by Planck's constant and as demanded by the Heisenberg uncertainty principle. An electron still feels a noisy electromagnetic field when the lights go off. Alfi knows this noisy field as his friend, Dark. A physicist knows it as the quantum vacuum state. It's nothing. It's what is left when the arrow—the coherent amplitude of the quantum field—is turned off. But one cannot in principle turn off the quantum noise as well, so in fact there is something to nothing. Note that the vacuum noise blob is symmetric with respect to amplitude and phase fluctuations; any direction is equivalent to any other. Instead of dragging some cumbersome dimensions along, I'll assign the vacuum state's fluctuations a size of "one," in terms of some arbitrary unit. Thus darkness is really a circular quantum blob of radius one.

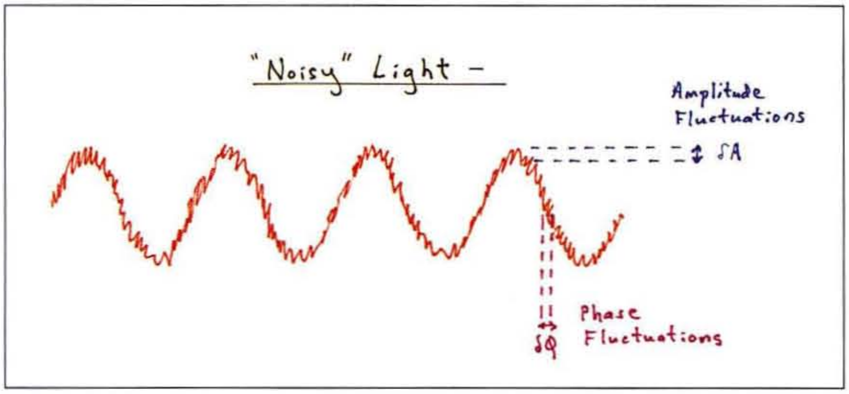

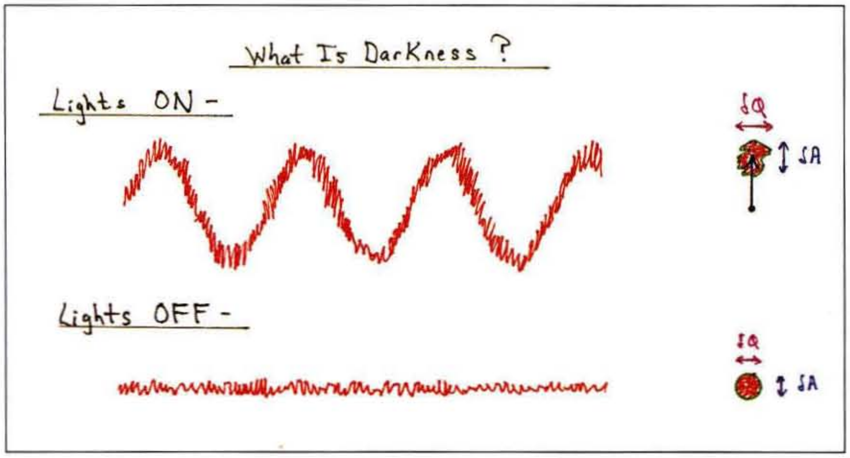

How noisy light becomes noisy darkness.
Top: A light wave has amplitude fluctuations, symbolized by δA, and phase fluctuations, symbolized by δΦ.
Bottom: When you turn off the wave or shrink the arrow's length to zero, the noise remains.

One cannot in principle turn off the quantum noise as well, so in fact there is something to nothing.

There is lots of evidence that these vacuum blobs are real. I'll mention two pieces, both of which have to do with the theory of quantum electrodynamics that Richard Feynman, Julian Schwinger, and Shinichiro Tomonaga pioneered in the late 1940s. The first piece of evidence is atomic. Consider the simplest atom—an electron orbiting a proton—to which nature inevitably adds a vacuum blob. (The vacuum fluctuations are everywhere!) Two funny things happen. One, the atom gets measurably bigger—by about one part in 100,000—because the electron is being jiggled by the fluctuations of the vacuum field. This is called the Lamb shift, named after Willis Lamb, Jr., who shared the Nobel Prize in physics in 1955 for the phenomenon's experimental discovery in the hydrogen atom. The other is that the atom spontaneously emits light because its otherwise stable excited state becomes unstable due to the inane and incessant noise of the vacuum. Sodium-vapor lamps glow orange-yellow because the sodium atoms in an excited state decay to the ground state, and that decay is caused by the vacuum jiggling the electron in a way that is perfectly calculable, and well-confirmed by experiment. The second piece of evidence is visible on a larger scale, and can be seen by holding two metal plates very close together. Even though there is nothing between the plates except the vacuum, one finds that the

In science as in Hollywood every successful story has a sequel, and I assure you that the story of the vacuum state has been very successful—in fact, a smash hit—through the past several decades.

plates actually attract each other with a tiny force—the Casimir force—resulting from the expulsion of the vacuum energy that's stored between those plates.

With the knowledge that darkness is the vacuum state packaged in quantum blobs, you might be tempted to think the story's over. It's not. In science as in Hollywood every successful story has a sequel, and I assure you that the story of the vacuum state has been very successful—in fact, a smash hit—through the past several decades. The sequel—Beyond the Vacuum Blob—is really what our research in the quantum optics group is about. Recall that our original vacuum state is a rather undistinguished circular blob. However, the only requirement of quantum mechanics is that a noise blob have a constant minimum area. The blob's dimensions along different directions don't have to be constant. For example, if we flatten the vacuum blob along the amplitude axis, the blob bulges out along the phase axis and the area remains constant. In other words, if the amplitude fluctuations get very small, the phase fluctuations must get very large. We call this a squeezed state. There are other ways to preserve the area. We could cut a hole out of the middle of the blob, and then stretch it into a thin ring, or annulus, of equal area, whose thickness in any direction is much less than the vacuum blob's diameter. The number of photons in this state is precisely known, but their phase is undetermined, as indicated by the circular symmetry of the blob. This is called the number state. And we need not restrict ourselves to the topology of only one blob—we can actually talk about two or more blobs at once. The laws of quantum mechanics demand that each blob have some minimum area, but the difference between the fluctuations each area represents—that is, between the "shapes" of the blobs—can be arbitrarily small. This is called the twin state because the blobs are, for our purposes, identical twins. The common ingredient in all these states is that some measurable aspect of the various blobs drops below one while some other aspect increases, keeping the area constant. (Remember, one is the size of the vacuum.) So, from the restricted viewpoint of the shrunken dimension only, we're making electromagnetic fields that have smaller fluctuations—less noise—than even the darkness of the vacuum state.

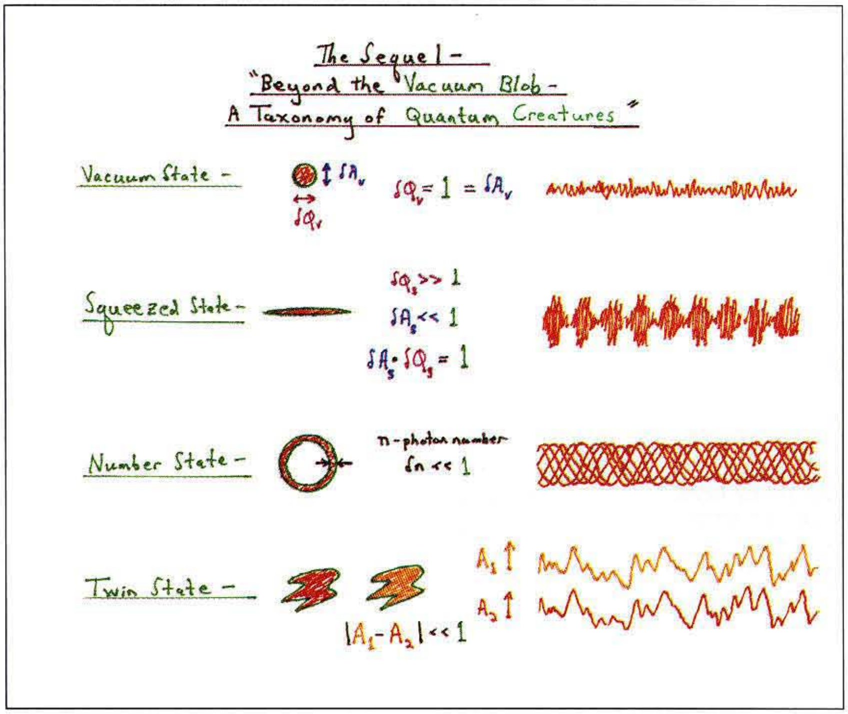

Each column contains (from left) the name of the blob, a portrait of its fuzzball in terms of amplitude and phase, its mathematical description, and the shape of its wave. A stands for amplitude, Φ is phase, n is the number of photons, and δ is the change or uncertainty in the variable.

Moving to a specific example, one can now ask, "How do we squeeze darkness?" We start with a vacuum state—and I should remind you, in this day of conservation, that the vacuum is an unlimited, inexhaustible natural resource. We

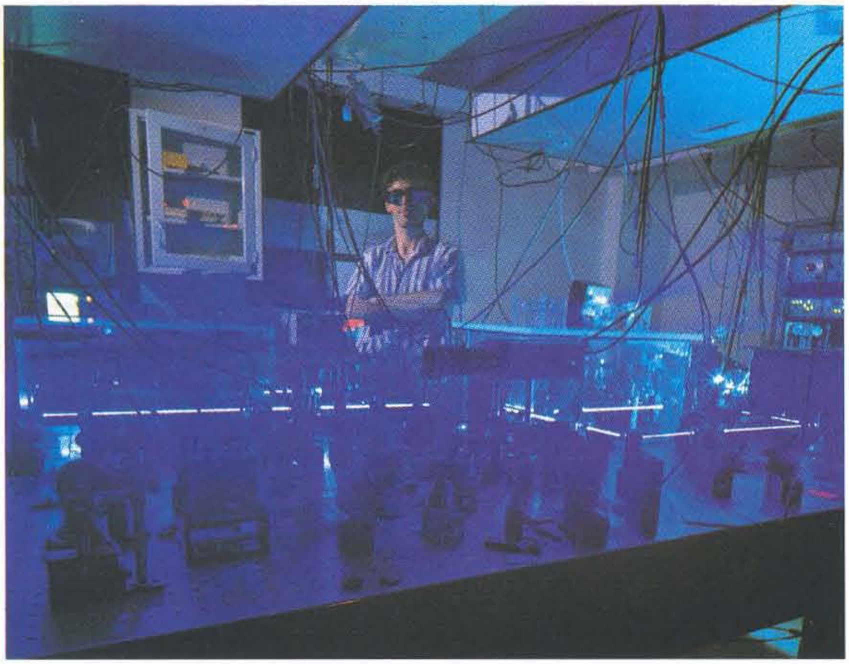

The light coming into our detector is four times darker than the darkness that the detector would see if it viewed empty space.

Grad student Nikos Georgiades and the darkness-squeezing factory. The blue lasers in front of him feed into a potassium niobate crystal, where the blue photons fission into squeezed red ones. The squeezed photons emerging from the crystal are actually in the infrared region of the spectrum, and can't be seen. The dark shapes in the foreground are a part of the interferometer that they live in.

send this vacuum state into our squeezing factory, which is the elaborate arrangement of lenses, prisms, and mirrors shown above. Of course, we have to be very careful in our choice of a squeezing machine. We have to somehow "squeeze" the vacuum without "touching" it—what I call a Platonic squeeze. We can't touch it directly because, after all, it's the vacuum, which is to say it's nothing at all. And once an apparatus has touched or interacted with the vacuum an unacceptable contamination usually results, because at the quantum level, macroscopic beings like graduate students are fairly shaky entities that impart their own uncorrelated fluctuations to the vacuum. There's no easy way to do Platonic squeezing in a satisfying manner, nor is there an easy way to explain it. The process that we use most is called "photon fission," in which a photon of blue light goes into a special "nonlinear" crystal and splits into two photons of red light. The law of conservation of energy must be obeyed, so the sum of the two red frequencies equals the blue frequency. This process doesn't occur to any significant degree in free space, but there are a variety of very interesting crystals, including the potassium niobate crystal that we use, that behave in unusual ways when illuminated. One of the seminal papers describing photon splitting was written by Amnon Yariv, Caltech's Myers Professor of Electrical Engineering and professor of applied physics, some 30 years ago.

Well, how do we squeeze without touching? There's a beam of blue light going into the crystal, but there's also a beam of red darkness, if you will—an initial vacuum state, pure and uncontaminated by the presence of red light—going into that same crystal. Into that red vacuum state, from the distant vantage point of the blue light, we take photons one by one from the blue beam and add them two by two to the red beam. As we do that, the initial vacuum for the red beam—its darkness, if you will—is turned into a squeezed state. And as we turn up the rate at which the photon pairs are added to the initial vacuum, the state is squeezed more and more into an ever thinner ellipse. (The process is mathematically identical to painting our circular vacuum blob on a rubber sheet and then stretching the sheet along one axis.) Surrounding the crystal is the actual apparatus for accomplishing this transformation. The apparatus is very complex—it looks like a kid went wild in a toy store and assembled the ultimate Lego set—because in essence we are trying to process the amplitude and phase fluctuations of a light wave (which is going up and down 5×10^{14} times per second) with a precision that is a small fraction of the size set by the vacuum blob. Therefore the entire apparatus, nonlinear crystal and all, is essentially a large interferometer whose arms are servo-controlled to keep the various waves in near-perfect alignment. In fact, the result of a lot of late-night effort, principally by associate scientist Eugene Polzik, is that we've been able to compress the vacuum state by a factor of four; that is, when measured along the squeezed dimension, the light coming into our detector is four times darker than the darkness that the detector would see if it viewed empty space.

Of course, the rules and regulations for quan-

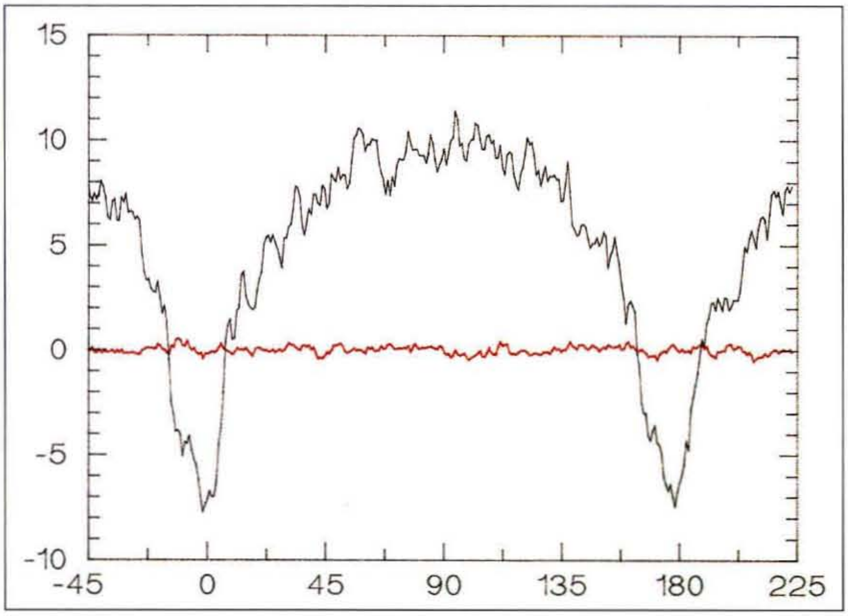

Reprinted from Wu, Xiao, and Kimble,
Journal of the Optical Society of America B, Volume 4, pp. 1465-75, October 1987.

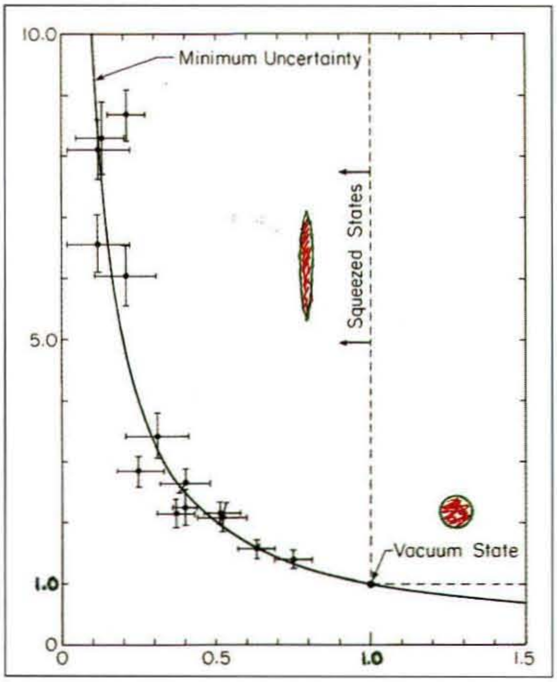

Above: A plot of noise versus the angle of tilt of the arrow for a squeezed fuzzball. The noise in decibels (vertical axis) is plotted logarithmically, so the noise of the unsqueezed vacuum is zero. The phase angle (horizontal axis) is plotted in degrees. The unsqueezed vacuum (red line) is equally noisy at all angles, whereas the squeezed fuzzball (black line) is much quieter than the vacuum at 0° and 180°, and noisier than the vacuum at 90°. Thus, to make a measurement using squeezed light, the detector would be locked at 0° or 180° in this case.
Right: The Heisenberg uncertainty principle for light. Uncertainty in phase ($\delta\Phi$) is plotted on the vertical axis; uncertainty in amplitude (δA) on the horizontal.

tum fuzzballs require that when we reduce the noise in one dimension, it must bulge out elsewhere. This is shown in the graph above, which plots the amount of noise as a function of angle in the two-dimensional space of amplitude and phase fluctuations. Thus, if we take the valleys in that graph as representing the short (squeezed) axis of the noise blob and the peaks as the long (bulging) axis, then a plot of one versus the other should be a hyperbola. (Remember, our unsqueezed vacuum fuzzball is one unit in radius, and the uncertainty principle sets a lower bound for the area.) And so, independent of how complicated the experiment is or how complicated the theory is, in the end the best that quantum mechanics lets us do is the hyperbola labeled "Minimum Uncertainty" in the graph above right, which agrees with our data reasonably well. Note that the data points have no adjustable parameters; we measure everything in absolute terms. To the right and above that figure's dashed lines, which mark where each dimension of the noise blob equals one, lies the land of classical physics. If all light behaved like that, you wouldn't be reading this. To the left and below these lines is the land of quantum darkness.

Now we're ready to think about making useful measurements with light. If I want to send a light wave to you, to talk to you on a fiber-optic telephone line, for example, what is the minimum modulation of the light—how much do I have to move the tip of the arrow—in order for you to notice any change? The classical answer is that the modulation can be arbitrarily small, because the position of the arrow's tip that represents the light wave is arbitrarily precisely defined. But this possibility is highly illegal because in quantum systems, the tip's exact location is no longer a defined quantity—it's just somewhere in a fuzzball of uncertainty. Heisenberg's uncertainty principle applies here to state that nature allows no naked arrows. I can represent this rule by impaling a quantum cabbage onto the point of the arrow. The laws of quantum mechanics say there have to be fluctuations—the arrow representing a perfectly smooth wave doesn't exist separately from the cabbage representing the quantum blob. In fact, to a physicist, a naked arrow is a much more heinous crime than is indecent exposure.

Since I can't remove the cabbage from the arrow, measurements involving a change of length of the arrow have to displace the arrow by an amount larger than the diameter of the cabbage—that is, of the vacuum fluctuations—in order to reliably discern any change at all. This displacement of the arrow by one diameter of the vacuum blob is the standard quantum limit for making measurements of the electromagnetic field. Over the history of the science of measurement, the standard quantum limit has stood as a seemingly impenetrable barrier, both conceptually and practically. And even making a measurement precise enough to approach the standard quantum limit in the first place is not trivial. However, in more recent times—over roughly the past 15 years—it has come to be appreciated that one can, in fact, do better than this limit. To do so, we squeeze our quantum cabbages into quantum cucumbers. Now a

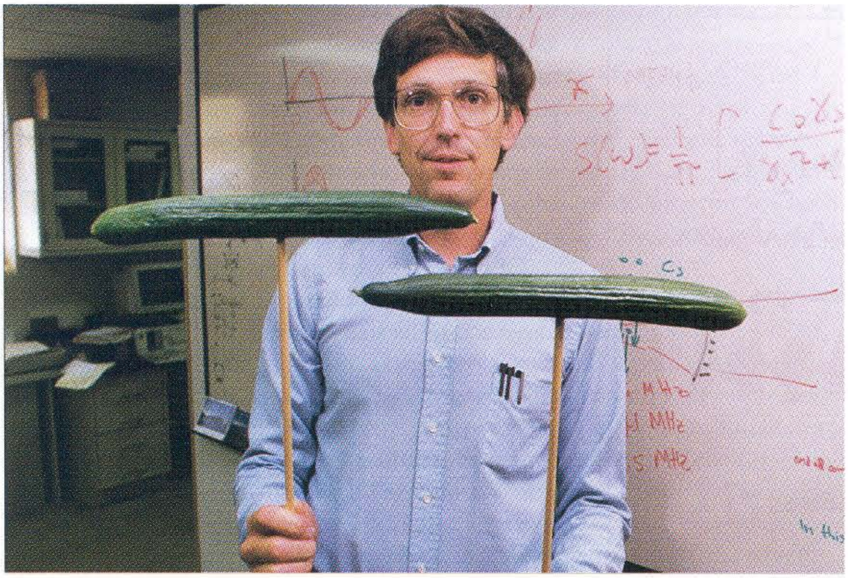

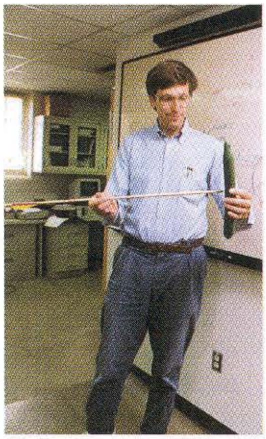

A quantum cabbage (right) has to be moved by roughly its diameter in order to be sure of displacing the tip of the arrow hidden within it. The same applies to a quantum cucumber (above), but it can be moved less, as its diameter is smaller. In terms of quantum fuzzballs (below), this means that a smaller δA is measurable. Are you *sure* Steve Martin got his start this way?

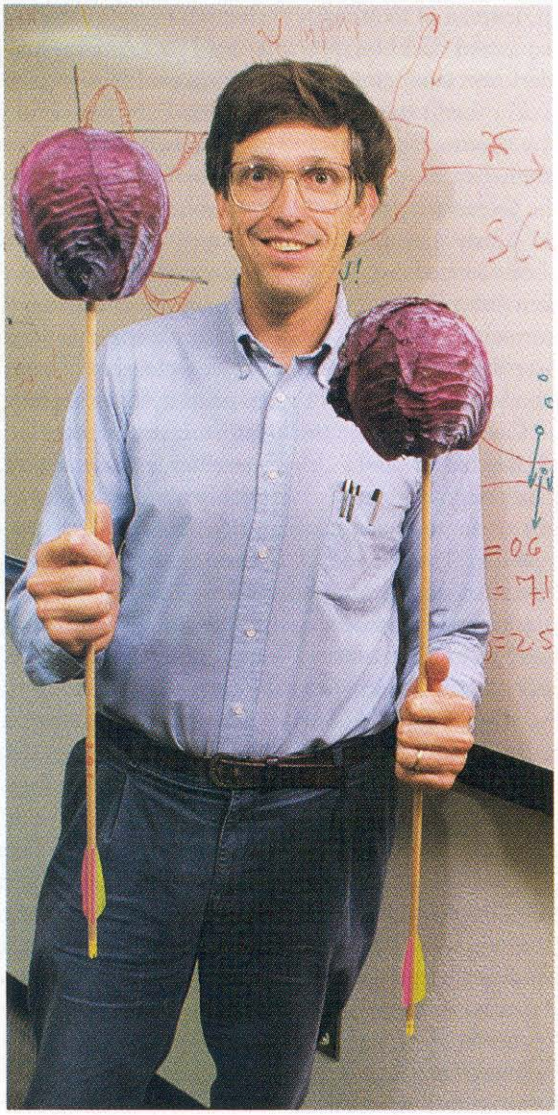

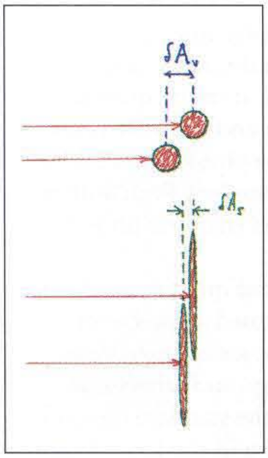

smaller displacement becomes discernible, because a cucumber is narrower than a cabbage—at least along its thin axis!—and we can make better measurements than had previously been thought possible.

It should be emphasized that, unlike cabbages, quantum entities are the same everywhere in the universe. While cabbages come in different sizes, the quantum fluctuations of the vacuum blob don't. Furthermore, these fluctuations are quite small. On a scale where a cabbage denotes a vacuum blob of radius one, the arrow's length, for even a laser of modest power, would be equal to the diameter of the earth.

These otherwise esoteric considerations of the quantum nature of light can be gainfully employed to detect a signal that couldn't otherwise be seen. Imagine that the quantum limit is a sea of fluctuations like the Pacific Ocean, and the signal we're looking for is the Hawaiian Islands. The Hawaiian Islands extend down to the ocean floor, but all we can see is what sticks up above the sea. If we aren't satisfied with this view, we could drain the ocean a bit. If we lower the ocean's level (the noise floor) by a factor of two, the Hawaiian Islands (the spectral peak we want to study) gets bigger relative to the noise by this same factor. That means we can see signals twice as small, or the same signals in half the time, as before. There is a caveat, of course, because this draining—which is really just a redistribution of quantum fluctuations—only happens along one axis. With the freedom to make this noise smaller comes the responsibility to make sure that we push the button that drains the ocean and

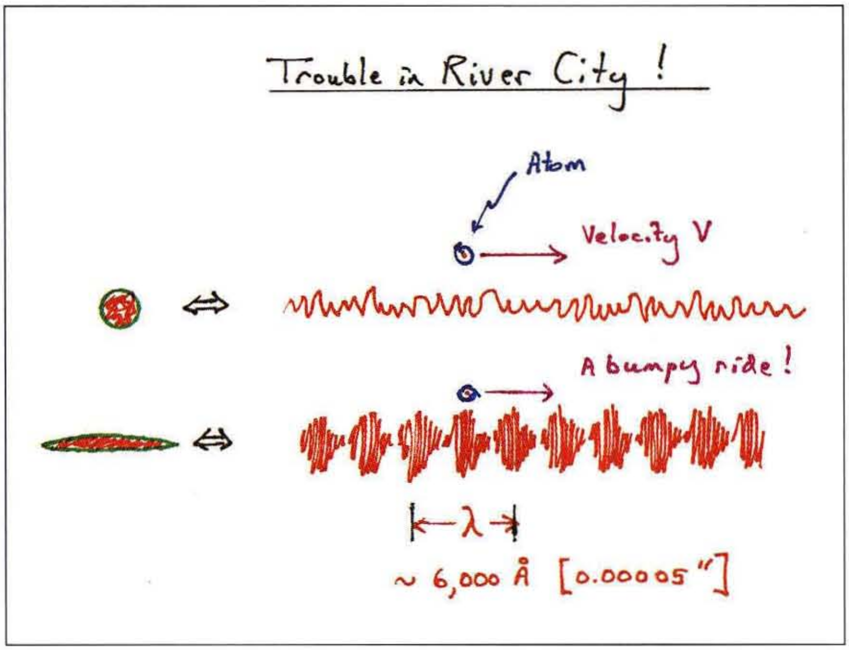

The problem of atomic motion: An atom traveling through an ordinary vacuum (top) has smooth sailing, but an atom moving through a squeezed vacuum (bottom) is in for a bumpy ride.

Electrons are reasonably intelligent. If one of them finds out that there's now a quiet dimension to its life where previously there was uniform noise in all directions (the usual vacuum state), it will try to live in the quiet dimension.

not the one that fills it. For if, instead of looking along the quiet dimension's squeezed darkness, we look at the squeezed antidarkness (the long dimension of our "quantum cucumber") the noise level goes up dramatically—by about a factor of ten in our most recent measurements, described below. That means the Pacific Ocean rises tenfold and the Hawaiian Islands vanish altogether.

One such experiment has been carried out here by Polzik and grad student John Carri. They were doing precision atomic spectroscopy—that is, detecting atoms by laser illumination. An atom has resonances—when you tickle it, it gets excited—so Polzik and Carri moved the laser's frequency around until they hit a resonance. The atoms—in this case cesium atoms in a vapor cell—absorbed light at the resonant frequency, betraying their presence. The particular technique we employed is called quantum-limited FM (frequency modulation) spectroscopy. It was pioneered by Drever here; John Hall, of the Joint Institute for Laboratory Astrophysics in Boulder, Colorado, and who was a Fairchild Scholar at Caltech in 1992; and a group at IBM. Our group added squeezed light to this technique, and has gone beyond the standard quantum limit by about a factor of two in spectroscopic sensitivity.

If we can use these "designer" fluctuations to probe atoms in new ways, can we also harness these funny fields to actually *change* the atoms? Remember, an atom coupled to a vacuum gives rise to the standard radiative processes—things like lasers, street lights, and interstellar nebulae. What if we instead couple the atom to a squeezed vacuum? Well, electrons are reasonably intelligent. If one of them finds out that there's now a quiet dimension to its life where previously there was uniform noise in all directions (the usual vacuum state), it will try to live in the quiet dimension. Indeed, there are stacks of theoretical papers indicating that atoms would behave in fundamentally different ways, if only we could couple them to the squeezed vacuum. Such coupling would affect all of traditional spectroscopy, as well as things such as how lasers work.

Of course, there are a few catches, at least one of which is the problem of atomic motion. Ordinary vacuum is structureless, so when an atom moves through it, the atom travels as though it's on a smooth, featureless road in North Dakota, as shown at left. On the other hand, if we use a squeezed vacuum, the atom's in for a bumpy ride—bouncing up and down over the spatially varying noise of the squeezed light. All the while, the atom's electron is trying to find the light's quiet dimension, which unfortunately changes every quarter of a wavelength—about every 1500 Ångstroms. One solution to this motion problem is to cool the atom's motion to almost absolute zero and to confine it to a distance much smaller than one-quarter of a wavelength. Graduate student Zhen Hu is doing such research in my group, trying to nail the atom down by using laser beams to build an atom trap. A trapped atom is also very cold, since temperature, on the atomic level, is really a measure of the atom's energy of motion. The photo opposite shows a cloud of cesium atoms cooled and trapped by laser beams. The cloud is about a millimeter—a twentieth of an inch or so—in diameter, and the atoms within it are cooled to within about 10^{-4} degrees Kelvin of absolute zero (–459° Fahrenheit). At the same time, we are working on ways to make clouds with fewer atoms, until we can eventually just trap a single atom. So we've almost got the atom nailed down, and once we do, we'll bathe it in the quantum quietness of squeezed light and see what happens. (Associate Professor of Astrophysics Ken Libbrecht [BS '80] and graduate student Phil Willems also have a laser cooling and trapping project on campus.)

Returning to the theme of quantum measurement, my group has performed a number of measurements over the past six or seven years—spectroscopy, interferometry, and others—at levels of precision beyond the standard quantum limit. But how far beyond will the laws of nature let us go? In terms of our previous analogy, we've lowered the ocean by about a factor of two, but where, actually, is the bottom? As far as I know, there's no totally satisfactory theoretical answer to this question. To find out, we need to optimize our measurement techniques over all possible quantum blobs—all shapes and states, not just the few I've told you about—and over all possible measurement strategies. That's a difficult thing to do. After all, we're using 19th-century techniques—for example, interferometry—and late-20th-century light. Nonetheless, some important theoretical progress has been made in recent years, notably by Carlton Caves and colleagues.

Apart from deep theoretical issues, there is a great deal of practical interest in manipulating the fundamental quantum fluctuations of light for such things as spectroscopy, quantitative analysis, and interferometry. Applications range from things on the scientific frontier, like the LIGO (Laser Interferometer Gravitational-Wave Observatory) program here, to more mundane things like the new aircraft-navigation systems, which use a laser gyroscope working near the standard quantum limit to sense rotation.

Top: A cloud of trapped cesium atoms.
Bottom: A pair of quantum tennis balls.

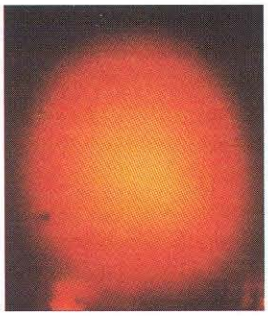

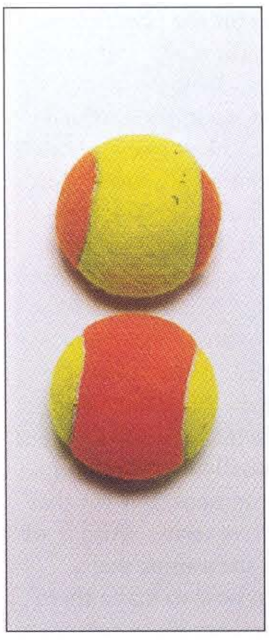

At this point, we might stop and ask, what does this all mean? What are these quantum blobs, really? This is, in fact, a very difficult question to answer. To avoid having to answer it myself, I will quote from *Dreams of a Final Theory*, by Stephen Weinberg, a Nobel laureate and one of the eminent scientists of this century. "A year or so ago, while Philip Candelas... and I were waiting for an elevator, our conversation turned to a young theorist who had been quite promising as a graduate student and who had then dropped out of sight. I asked Phil what had interfered with the ex-student's research. Phil shook his head sadly and said, 'He tried to understand quantum mechanics.'" Weinberg goes on to say, "I admit to some discomfort in working all my life in a theoretical framework that no one fully understands." The computational power of quantum mechanics is unquestioned. However, what it all "means" in any satisfactory sense is difficult to explain, even to oneself. Nonetheless, I'll try to illuminate some of the issues and conundrums in the following thought experiment.

Suppose I have a source that emits pairs of colored tennis balls, one to the right and one to the left, and detectors some distance away that catch the balls and register a reading of either red or green. The source always sends out correlated pairs of colored balls heading in opposite directions. Thus if I listed what each detector saw, the left detector would register a sequence of, say, red, red, green, red, green, and so on. And the right detector would register the opposite colors—green, green, red, green, red, and so forth. The question is, what inferences can I draw about the nature of these quantum blobs—here represented as tennis balls—as they propagate from source to detector, using the sequences recorded at the two detectors? For example, if I detect red at detector number one and green at detector number two, can I infer that a red blob actually traveled from the source to detector number one? Or, to paraphrase Einstein, "Do these blobs have any existence independent of one another?" Well, certainly they must. If one blob is just coming by me and the other is way over there in Andromeda, then surely nothing about what happens to this one can affect that one.

Unfortunately, or fortunately, this most sensible view of the nature of the physical world is not, in general, valid. The quantum world is indeed a strange place, with a large domain of exceptions to the rule of objective reality. It turns out that neither blob, for certain kinds of quantum systems, has a "color," where color is used metaphorically to refer to some property of the system in question, as for example its state of polarization. The "color" information is not carried by this blob or that blob, but rather resides in the correlations between the blobs. Either blob has the potential to be red or green—it's neither red nor green as it propagates, but somehow has the potential to be both colors at the same time. Hence physical properties for some microscopic quantum systems don't exist in the sense that I'd like to think that I exist. If you turn around, you can't see me, but I hope that I'm still here with an unchanged, definite set of properties. But for these quantum blobs, for these colored quantum tennis balls, color becomes well-determined—"exists," if you will—only when the blobs are detected. So a red click in my detector here, in some spooky way, means that the other blob must now be green, even if the detection events are light years apart. That's not a very comfortable thing, but that's the way it is. John Bell, who defined the limits of applicability of objective reality, called these correlations the irreducible nonlocal content of quantum mechanics. To paraphrase Bell, the speakable in quantum mechanics is the two detected sequences of reds and greens. The unspeakable, to which we are not allowed an answer in quantum mechanics, is the "real" color of one or the other blob as it propagates.

This is not a particularly comfortable situation, but is it refutable? No. A series of experiments by a number of groups, culminating in the work by Alain Aspect et al. in Paris, says that's the way nature is, like it or not. As for our own efforts in this regard, Zhe-Yu (Jeff) Ou—who has

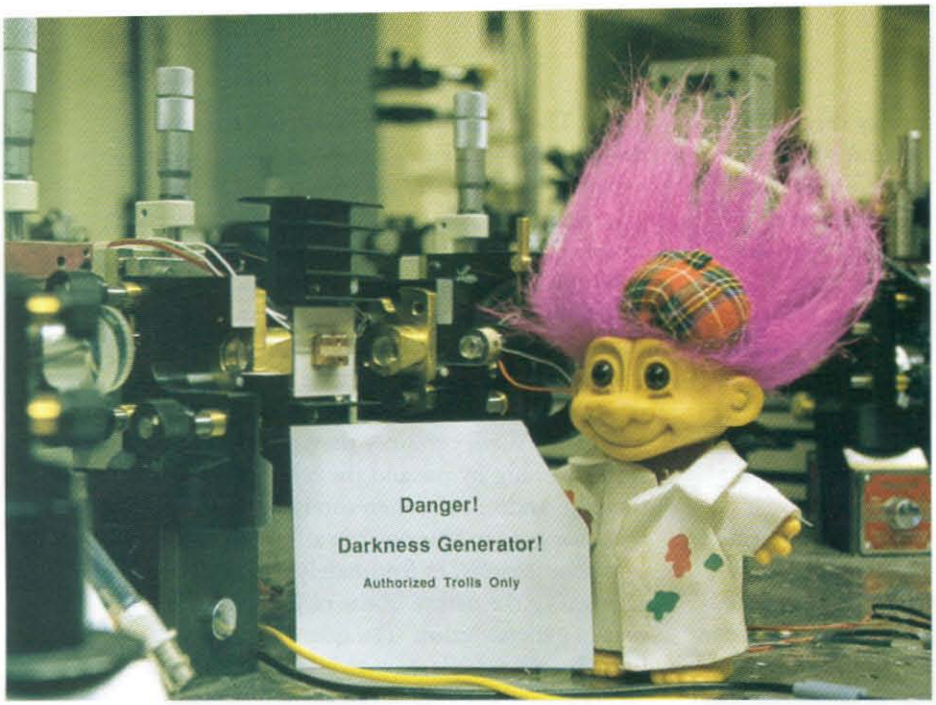

How darkness squeezing *really* works. (Troll courtesy of Megan Kimble.)

We'd like to find ... systems that are continually evolving and interacting with their environment, but yet that are not describable in objective terms. And if we can learn how to do this on the atomic scale, eventually we'd like to learn how to make them into macroscopic objects big enough to campaign for office.

since left for a faculty position at Purdue University in Indianapolis—and grad student Silvania Pereira have built an apparatus that makes correlated quanta in two spatially separated beams. They've carried out several experiments with this system over the last year and a half, but the one that I'll describe is related to quantum communication.

Imagine that I'm trying to send you a confidential message. Maybe it's about my bank account—a lot of such traffic is financial. Whatever it is, I don't want anybody to listen. Normally, my message would be encrypted in some code, such as the widely used Digital Encryption System (DES), that is nearly impossible to decode illicitly. Although such a code can be made extremely difficult to break in practice, nothing ensures that it cannot be broken in principle by some sufficiently clever person. One would like to protect these messages—not by my ingenuity or yours—but by the laws of quantum mechanics so that they are immune to interception in principle. So, by the process of photon fission, Ou and Pereira made two big, noisy quantum blobs that were arbitrarily large compared to the vacuum blob and that were quantum twins of each other. That is, their fluctuations in amplitude and phase were identical. Then, inside each one of these twins, we wrote a message so small that it was actually smaller than the vacuum level. The twins were then transmitted along different routes. Even if an eavesdropper detected one blob, the message was unrecoverable, because it was smaller than the standard quantum limit. Only when both blobs were detected and properly subtracted did the message emerge. Furthermore, if somebody did try to listen in, this interception would sound a burglar alarm, because detecting one blob destroys quantum correlations, and hence degrades the message to garbage.

Another experiment, which postdoc Olivier Carnal and grad students Robert Thompson and Quentin Turchette are pursuing, is difficult to describe, but the spirit is conveyed by comedian Robin Williams's line, "Reality—what a concept!" The issue is again the nature of reality, but now for a quantum system that's continuously interacting—being "measured," if you will—by its environment. Such "open" quantum systems are both driven by, and decay into, their surroundings, and are the basis for the phenomena that we know on a macroscopic scale. For any given open quantum system, there are many different measurements that we in the external world could choose to make: How many photons are coming out, and how are they distributed in time and space? What do their quantum fuzzballs look like? We could choose to ask a series of such questions by making a series of different measurements on the system. The $64 question —the sum of all questions—is whether there are systems whose "reality" is conditional upon the questions that we ask of them. We'd like to find such open quantum systems for which this is so—systems that are continually evolving and interacting with their environment, but yet that are not describable in objective terms. And if we can learn how to do this on the atomic scale, eventually we'd like to learn how to make them into macroscopic objects big enough to campaign

The Friends of Darkness, seen on the roof of Bridge Laboratory in the unaccustomed light of day. From left: Hu, Carnal, Thompson, Pereira, Georgiades, grad student Hideo Mabuchi, Turchette, and Willems.

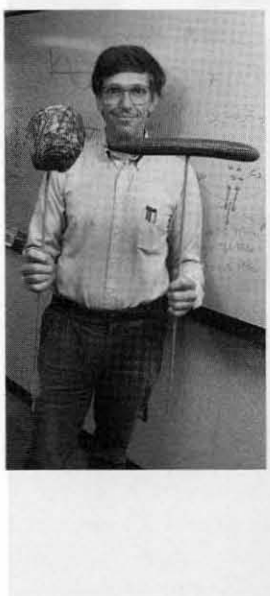

for office.

The particular system that Carnal, Thompson, and Turchette are looking at consists of a pair of parallel mirrors facing each other, some 300 millionths of a meter apart, with a stream of cesium atoms passing between them, like a single lane of cars between concrete dividers. The spacing between the mirrors is an exact multiple of a wavelength at which cesium atoms absorb and emit light, and is precisely controlled to within about 10^{-13} meters, or one-thousandth of the diameter of an atom. The cavity formed by the two mirrors serves as a very simple system whereby photons from a laser perpendicular to the path of the cesium atoms can be strongly coupled to them. That is, the interaction between the photon and the cesium atom is much stronger than the dissipative forces that normally cause the photon to lose its coherence. Thus, the excited atoms are in a nonequilibrium steady state, not unlike living beings—they take in energy, they move around and do things, and eventually they dissipate and die. The coupling is so strong that a mere 0.2 photons will evoke a nonlinear response from an atom, and a paltry 0.06 atoms significantly alters the photon's behavior. Hence the escape of a single quantum into the external environment can have a profound effect on the system, even though it contains hundreds of atoms and photons. The strong coupling means that quantum events occur at a faster rate than dissipative events, and the atom-photon system thus has enough time for at least the possibility of leading a life of manifestly quantum dynamics before the grim reaper of dissipation enters. It is to this type of system that we are currently turning our attention in a quest to explore the exquisite interplay of the birth and death of quantum states for driven open quantum systems.

Finally, then, let me come back to where we began—back to Alfi's question. If I turn on the lights, where does the darkness go? I hope I've given some sense of the answer to this seemingly simple question. We now know that darkness is the blob of noise representing fundamental fluctuations in the electromagnetic field. To produce light, we just put that blob on the end of an arrow. What Dark said to Alfi is precisely correct, "I never go anywhere." The dark is still there when we turn on the lights; it's just sitting on the end of an arrow that represents the basic coherent amplitude of the light. I couldn't have told my children about the nature of darkness any better. In fact, I use that book to tell them what I do in the laboratory. We've also seen that there are destinations beyond darkness. For example, I've told you about squeezed vacuum and some of its applications, and about twin states. In general, I've tried to convey a feeling for light that's even darker than the darkness of the vacuum, and about the activities of a group in a "mad pursuit" of the science of darkness. Finally, I would invite everyone to enjoy the darkness, much as Alfi can with his new-found understanding. □

H. Jeff Kimble received his BS from Abilene Christian University in 1971, and his MS and PhD from the University of Rochester in 1973 and 1978, respectively—all in physics. He came to Caltech as a professor of physics in 1989 from the University of Texas, where he was the Richardson Regents Professor of Physics. Kimble's PhD thesis research represented the first observation of a nonclassical state of light, and the research group he established at UT was one of the first to explore the field of squeezed light and related nonclassical phenomena. This article is adapted from Kimble's recent Watson lecture, which combined quantum physics and laser science with elements of Gallagher's vegetable-imperiling stand-up comedy and a tennis clinic. In fact, Kimble's research group had so many quantum and classical tennis balls lying around after the lecture that they recently held the First Annual Quantum Optics Tennis Tournament. (The Forces of Darkness beat the Forces of Light, 7-5, 6-2.) Kimble's daughters Megan and Katie are six and eight, respectively.