AMERICAN
Scientist

March–April 2010

www.americanscientist.org

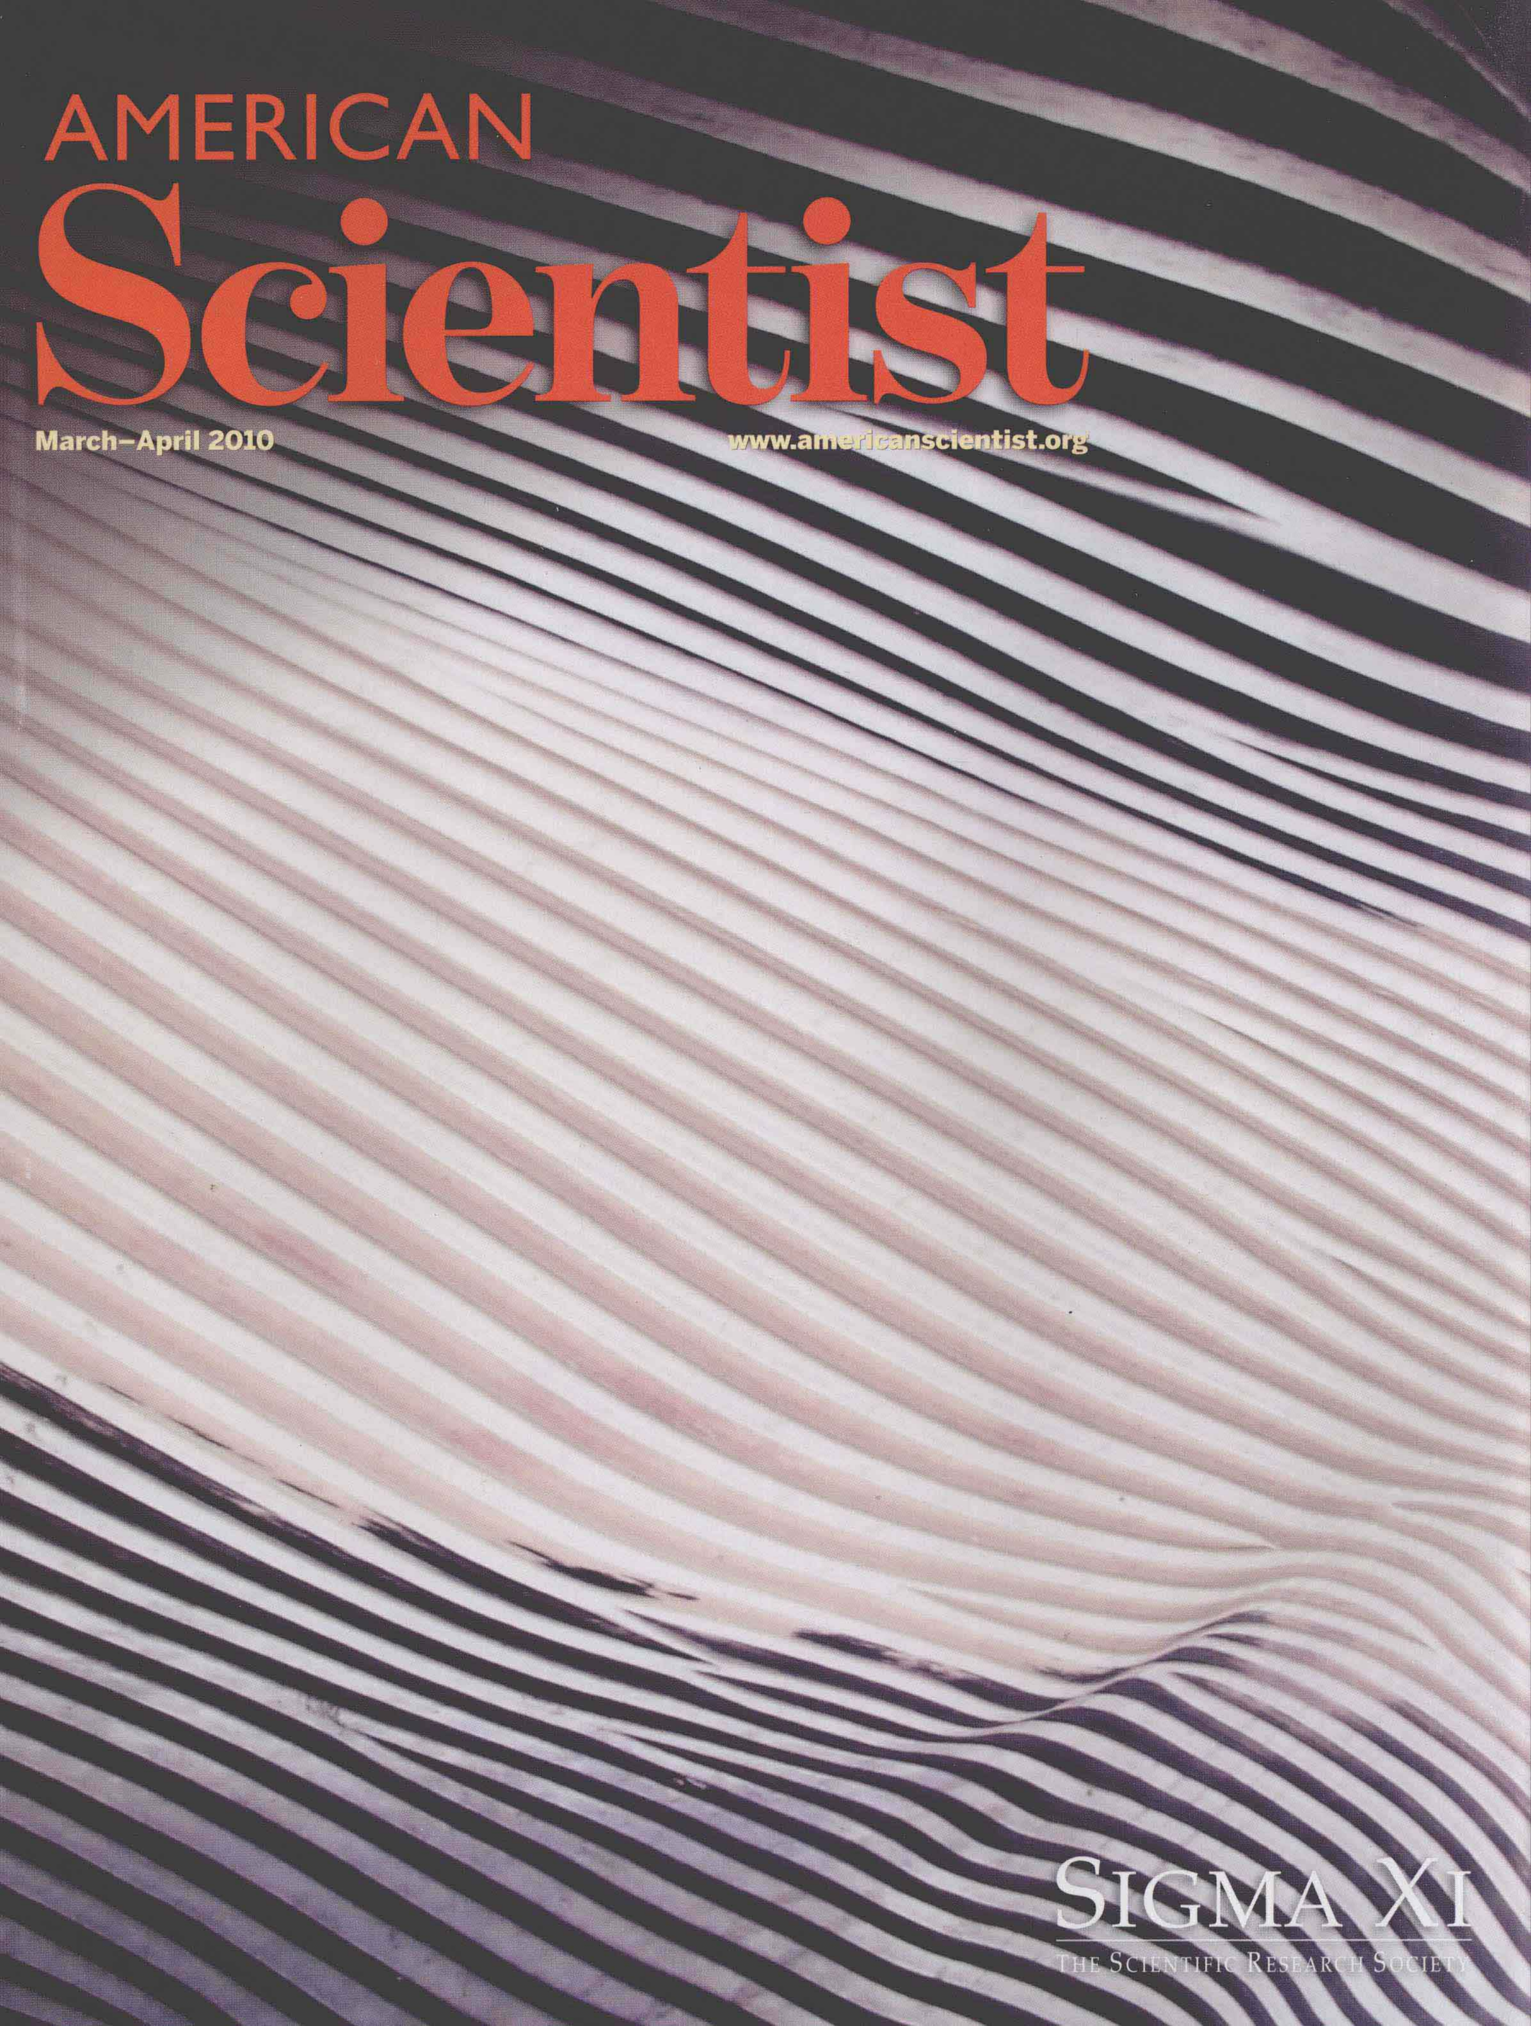

SIGMA XI
THE SCIENTIFIC RESEARCH SOCIETY

Congratulations

2010 Sigma Xi Award Winners

William Procter Prize for Scientific Achievement

The award includes a $5,000 Grant-in-Aid of Research given to a young scholar selected by Dr. Spivey.

Michael Spivey, professor of cognitive science, University of California-Merced

John P. McGovern Science and Society Award

The McGovern Award was established in 1984 to honor an individual who has made an outstanding contribution to science and society.

Barbara Gastel, professor of veterinary integrative biosciences and of humanities in medicine and biotechnology, Texas A&M University

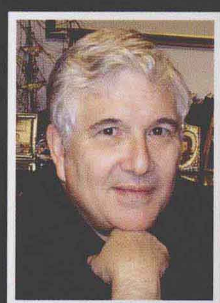

Walston Chubb Award for Innovation

The Chubb Award honors and promotes creativity among scientists and engineers.

Howard Moskowitz, an expert on sensory psychology and its commercial application, is president and CEO of Moskowitz Jacobs Inc.

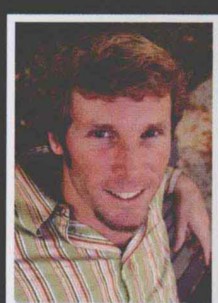

Young Investigator Award

Sigma Xi members within ten years of their highest earned degree are eligible for this award.

Kevin Gurney, assistant professor of Earth and atmospheric science, Purdue University

Sigma Xi prize lectures will be highlights of the 2010 Sigma Xi Annual Meeting and International Research Conference on November 11-14 at Raleigh Convention Center, Raleigh, North Carolina

www.sigmaxi.org

SIGMA XI

THE SCIENTIFIC RESEARCH SOCIETY

AMERICAN
Scientist

Volume 98 • Number 2 • March–April 2010

Departments

Scientists' Bookshelf

From Sigma Xi

Feature Articles

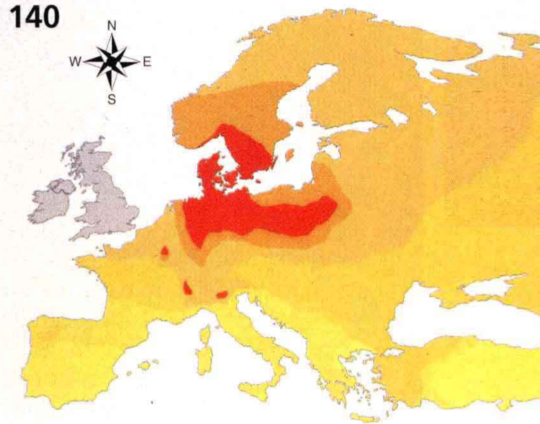

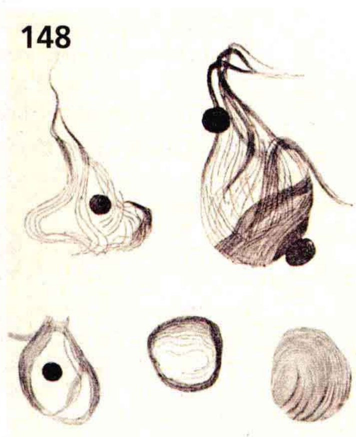

The Cover

The accordionlike blubber on a blue whale's underside extends from mouth to bellybutton (on the cover). *The structure, found only in the family of baleen whales called rorquals, is made from firm ridges* (left) *connected by deep furrows of delicate elastic tissue, and can stretch to more than twice its original length. Thus the whale's oral cavity can expand to enormous size and hold many tens of tonnes of water and krill; the whale then filters out the water with its baleen while retaining its tiny shrimplike prey. Exactly how rorquals engulf such quantities of water has long been obscured by ocean depths, but as Jeremy A. Goldbogen recounts in "The Ultimate Mouthful: Lunge Feeding in Rorqual Whales" (pages 124–131), electronic devices are aiding researchers in understanding the complex biomechanics behind how these enormous animals eat. (Cover image and image at left courtesy of Nick Pyenson.)*

Methodology in Our Madness

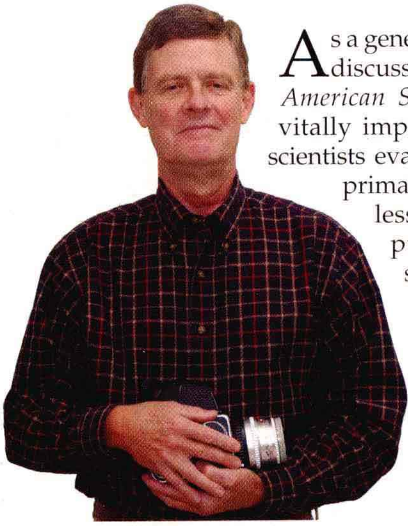

As a general rule, we avoid much discussion of methodology in *American Scientist*. Although it's vitally important information for scientists evaluating other scientists' primary work, it usually has less value in a secondary publication, where the science reported has already been through the peer-review filter. Sometimes, though, it's just too interesting for us to push aside.

Jeremy Goldbogen's piece "The Ultimate Mouthful: Lunge Feeding in Rorqual Whales" (pp. 124–131) is a great example. For decades, what baleen whales are up to when they dive as much as 300 meters deep in search of krill has been shrouded in mystery. Theoretical studies could speculate on the biomechanics of rorqual feeding, but it took the development of temporary tags and infrared cameras to actually ride along to dinner with the largest of marine mammals. What Goldbogen and his colleagues found is truly extraordinary, but I won't spoil it by revealing more than that we're talking school-bus scale here.

Once you've digested the biomechanics of whale feeding, you need only turn the page to find an example of technology where methodology *is* the message. Henrik Wann Jensen and Tomas Akenine-Möller are experts in making dancing pixels as convincing as possible. In "The Race for Real-time Photorealism" (pp. 132–139) they describe approaches to representing images in ways that prove convincing to the human eye yet remain computationally practical. To be honest, much of the driving force behind photorealistic rendering has been to satisfy the ravenous consumers who support the video gaming industry, and science has been the happy secondary beneficiary. We'll take it.

Sometimes, though, the methodology confounds expectations—at least mine. As a photographer, I came late to the digital revolution. And like so many reluctant adopters, I became an enthusiastic (some might add "over-" to the previous word) proponent. Which is why I was so caught by surprise when I read the caption for "Sightings" (pp. 156–157). When Fabiano Ventura set out to duplicate scenes of glaciers captured by Vittorio Sella 100 years ago, digital was not his medium. Instead, he used a 4 x 5 view camera and film—exactly the same medium Sella used. This allowed him to duplicate the geometries of the originals for exact comparisons showing glacial changes. Then he scanned the film and stitched the digital images together to produce panoramas impossible in digital alone. The result certainly throws down the gauntlet to the photorealists mentioned above.

Recently Sigma Xi hosted *ScienceOnline2010*, a conference on electronic media and the communication of science held here each year since 2008. Over that time I've watched the conference grow from a gathering of bloggers to a remarkably diverse meeting of minds on everything from podcasting to social networking to citizen science to, well, blogging. Many of this year's sessions were video recorded, most of which should be available on YouTube by the time you read this. For starters, you might look for a talk by our own Elsa Youngsteadt and her Public Radio International counterpart, Rhitu Chaterjee, on "The World Science" podcast.—*David Schoonmaker*

American Scientist

David Schoonmaker *Editor*
Morgan Ryan *Managing Editor*
Fenella Saunders *Senior Editor*
Catherine Clabby *Associate Editor*
Mia Smith *Editorial Associate*
Barbara J. Aulicino *Art Director*
Tom Dunne *Assistant Art Director*
Brian Hayes *Senior Writer*
Christopher Brodie *Contributing Editor*
Rosalind Reid *Consulting Editor*
Elsa Youngsteadt *Contributing Editor*

Scientists' Bookshelf

Flora Taylor *Editor*
Anna Lena Phillips *Assistant Editor*

American Scientist Online

Greg Ross *Managing Editor*

www.americanscientist.org

Jerome F. Baker *Publisher*
Katie Lord *Associate Publisher*
Jennifer Dorff *Marketing Manager*
Eric Tolliver *Marketing Associate*

Advertising Sales

Kate Miller *Advertising Manager*
advertising@amsci.org • 800-282-0444

Editorial and Subscription Correspondence

American Scientist
P.O. Box 13975
Research Triangle Park, NC 27709
919-549-0097 • 919-549-0090 fax
editors@amscionline.org • subs@amsci.org

Published by Sigma Xi
The Scientific Research Society

Howard Ceri *President*
Richard L. Meyer *Treasurer*
Joseph A. Whittaker *President-Elect*
James F. Baur *Immediate Past President*
Jerome F. Baker *Executive Director*

Publications Committee

A. F. Spilhaus, Jr. *Chair*
Jerome F. Baker, Howard Ceri, Lawrence M. Kushner, David Schoonmaker

PRINTED IN USA

Sigma Xi, The Scientific Research Society *was founded in 1886 as an honor society for scientists and engineers. The goals of the Society are to foster interaction among science, technology and society; to encourage appreciation and support of original work in science and technology; and to honor scientific research accomplishments.*

Understand the Material

To the Editors:

Heather Patisaul's feature article "Assessing Risks from Bisphenol A," (January–February) nicely illustrates the difficulties in trying to assess human health effects from low levels of chemicals in our environment. The author would have benefited, however, from collaboration with a materials scientist.

Early on she notes that BPA is a common ingredient in many hard plastics. BPA is a monomer for polycarbonate (66 percent) and epoxy (30 percent). Those polymers constitute only about 3 percent of U.S. plastics production of over 100 billion pounds in 2008. BPA is hardly a common ingredient.

By the end of the article, Patisaul said, "DDT undoubtedly saved lives, and likely still does. No such case can be made for BPA. It is time to develop a clear and comprehensive strategy for assessing the potential public health consequences of endocrine disruptors such as BPA that may contribute only economic value." To understand the public health consequences and develop a clear strategy, one must understand the materials involved, how they are used and the routes and levels of exposure to compounds of concern. Why are epoxy coatings used for certain cans? They reduce the likelihood of botulism. Someone behind impact- or bullet-resistant windows might value the protection they give.

Technologies exist largely because of the underlying materials. Consider CDs. Polycarbonate is a lightweight, high-impact, heat-resistant, intrinsically flame-retardant plastic. The first three qualities are why BPA was used in baby bottles. If better materials are available, great. But let's not throw CDs, electrical appliances and bullet-resistant windows out with the baby bottles.

Gordon L. Nelson
Dean, College of Science
Florida Institute of Technology

Dr. Patisaul responds:

Information about BPA production levels comes from the "NTP-CERHR Monograph on the Potential Human Reproductive and Developmental Effects of Bisphenol A," at http://cerhr.niehs.nih.gov/chemicals/bisphenol/bisphenol.pdf.

As for routes of exposure, you don't need solvents to get BPA to migrate. As it turns out, you don't even need heat. A Harvard University research group reported in June that consumption of cold beverages from polycarbonate bottles containing BPA raises human urine levels of BPA by 69 percent.

Exposure to BPA is low but that does not mean it is innocuous. That type of "the dose makes the poison" thinking may not apply to endocrine disruptors because their dose responses appear to be non-monotonic in many cases. Given that, if you don't need it in food containers, why not pull it out and be on the safe side?

Another View of Hydrogen Sulfide

To the Editors:

I enjoyed Roger P. Smith's article "A Short History of Hydrogen Sulfide" (January–February). Many may not know that the gas plays productive roles in several geologic settings. First, hydrogen sulfide occurs as a minor constituent in most natural gas deposits and must be removed. It is then oxidized to elemental sulfur, a process that produces virtually the sole source of sulfur in North America. Previously, "biogenic" sulfur had to be mined by the Frasch process primarily in Gulf Coast salt domes.

Hydrogen sulfide also plays an important role in forming metallic-sulfide ores of zinc, lead and copper. The bearing fluids of these base metals must encounter a source of hydrogen sulfide (either biogenic or magmatic) along their flow path to precipitate metal-sulfide minerals. Alternatively, volcanic hydrogen sulfide helps form ores of gold and silver, as the precious metals form stable aqueous complexes with hydrogen sulfide, greatly enhancing their solubility in ore fluids. Precious metals often precipitate when the hydrogen sulfide is destroyed by a number of processes, including boiling, which puts hydrogen sulfide into the vapor phase. Hydrogen sulfide can also be oxidized by certain bacteria to make sulfuric acid, which is thought to be important in cave formation. Volcanic hydrogen sulfide from hot springs at mid-ocean ridges becomes the basis for complex biological communities where chemosynthetic (chemical producing) bacteria use it and carbon dioxide. Finally, we use hydrogen sulfide produced by stimulating sulfate-reducing bacteria to remediate groundwater contaminated by metals, arsenic and radionuclides (US Patent 5,833,855). With hydrogen sulfide, one should consider its good, bad and smelly aspects!

Jim Saunders
Auburn University

An Apollonian Opportunity

To the Editors:

In Dana Mackenzie's interesting column "A Tisket, a Tasket, an Apollonian Gasket" (January–February), Peter Sarnak remarked on the present-day inability of mathematics to prove or explain certain conjectures, including his own. Those conjectures concern the number series of the bends in Apollonian Gaskets. "The necessary mathematics has not been invented yet," Sarnak said.

It is interesting to remember something stated more than 200 years ago by Carl Friedrich Gauss. In his one-page proof of the long-unproven Wilson's prime number theorem, first published by Edward Waring, Gauss noted that "neither of them was able to prove the theorem, and Waring confessed that the demonstration seemed more difficult because no notation can be devised to express a prime number. But in our opinion truths of this kind should be drawn from notions rather than from notations."

Sarnak seems to have ignored Gauss's advice. That, unwittingly, may dissuade those who might otherwise attempt to prove those unsolved theorems.

Bernard H. Soffer
Pacific Palisades, CA

American Scientist (ISSN 0003-0996) is published bimonthly by Sigma Xi, The Scientific Research Society, P.O. Box 13975, Research Triangle Park, NC 27709 (919-549-0097). Newsstand single copy $4.95. Back issues $6.95 per copy for 1st class mailing. U.S. subscriptions: one year $28, two years $50, three years $70. Canadian subscriptions: one year $36; other foreign subscriptions: one year $43. U.S. institutional rate: $70; Canadian $78; other foreign $85. Second-class postage paid at Durham, NC, and additional mailing office. Postmaster: Send change of address form 3579 to Sigma Xi, P.O. Box 13975, Research Triangle Park, NC 27709. Canadian publications mail agreement no. 40040263. Return undeliverable Canadian addresses to P. O. Box 503, RPO West Beaver Creek, Richmond Hill, Ontario L4B 4R6.

To the Editors:

Considering that the geometry of the circle involves irrational numbers such as pi and square roots, I was struck by the seemingly infinite array of integers in the Apollonian gaskets described by Dana Mackenzie. One view of this is that each pair of mutually tangent circles has two infinite series of tangent circles spiraling into crevices between them. There are an infinite number of these mutually tangent pairs, each with a pair of infinite series. Some series appear more than once and some are part of other series.

I have found a linear relation that is somewhat different than Mackenzie's by choosing two tangent circles (say with curvatures a and b) from any four mutually tangent circles. Of the remaining two circles, call the curvature of the larger d_0 and the curvature of the smaller d_1. Thus d_0 is the starting term in a series of curvatures and d_1 is the second term. Other curvatures are determined by the linear formula obtained by subtracting Descartes's equation written for a, b, d_{n-2}, d_{n-1} from that for a, b, d_{n-1}, d_n. The resulting equation, $d_n = 2(a + b + d_{n-1}) - d_{n-2}$, can be used to determine the successive values of d_n by a process of iteration. Because Descartes's equation is a quadratic, the difference of the differences between consecutive terms is a constant and equal to $2(a + b)$ in each series. This seems to apply to the irrational roots of Descartes's equation, also.

Ronald Csuha
New York, NY

Dr. Mackenzie responds:

I see no conflict between Sarnak's quote and Gauss's admonition. Sarnak would certainly agree that new notions, not new notations, are needed to prove his "local-to-global principle" for Apollonian packings.

I am glad to report that Elena Fuchs (Sarnak's student) and Jean Bourgain have proven a "positive density theorem." This relates to the question Ron Graham asked, about whether even 1 percent of the numbers that could occur in an Apollonian gasket actually do occur. Fuchs has shown that the answer is yes, provided 1 percent is replaced by a sufficiently small (but positive) number. Interestingly, her approach was to use carefully selected subsets of the Apollonian gasket, an approach not too dissimilar from what Ronald Csuha proposes. Instead of the sequence of all circles tangent to two fixed circles, she looks at the somewhat more complicated sequence of circles tangent to a single fixed circle. The preprint will be posted at the open-access site http://arXiv.org.

How to Write to *American Scientist*

Brief letters commenting on articles that have appeared in the magazine are welcomed. The editors reserve the right to edit submissions. Please include a fax number or e-mail address if possible. Address: Letters to the Editors, American Scientist, P.O. Box 13975, Research Triangle Park, NC 27709 or editors@amscionline.org.

Illustration Credits

Macroscope
Page 103 Morgan Ryan and Barbara Aulicino
Page 104 Barbara Aulicino

Computing Science
Pages 199, 200 Tom Dunne

Engineering
Page 113 Tom Dunne

Marginalia
Pages 118, 119 Barbara Aulicino

The Ultimate Mouthful: Lunge Feeding in Rorqual Whales
Figures 3, 4 *(bottom)*, 6, 9, 10 Tom Dunne

The Race for Real-time Photorealism
Figures 3, 5, 6 Tom Dunne
Figure 4 Morgan Ryan

Gene-Culture Coevolution and Human Diet
Figures 2, 7, 8 Barbara Aulicino

The Origins of Alzheimer's Disease
Figure 7 *(left)* Barbara Aulicino

Just-as-good Medicine

David M. Kent

The rabbi's eulogy for Sheldon Kravitz solved a minor mystery for my father: what was behind the odd shape of the juice cups he had been drinking from after morning services for the last few years? Adding a bit of levity while praising his thrift and resourcefulness, the rabbi told of how Sheldon purchased, for pennies on the dollar, hundreds of urine specimen cups from Job Lot, that legendary collection of pushcarts in lower Manhattan carrying surplus goods—leftovers, overproduced or discontinued products, unclaimed cargo. At the risk of perpetuating a pernicious cultural stereotype, for men of my father's generation like Sheldon, raised during the Great Depression, bargain hunting was a contact sport and Job Lot was a beloved arena. My father, too, would respond to the extreme bargains there with ecstatic automatisms of purchasing behavior and come home with all manner of consumer refuse, including, and to my profound dismay, sneakers that bore (at best) a superficial resemblance to the suede Pumas worn and endorsed by my basketball idol, the incomparably smooth Walt "Clyde" Frazier. My father would insist that such items were "just as good" as the name brands. But we, of course, knew what "just as good" really meant.

In fairness to my father and his friends, from a utilitarian perspective (decidedly not the perspective of preadolescents), maximizing the overall good of the family involves economic trade-offs. Money saved from something "just as good" can be reallocated toward items that bring greater benefit than the value sacrificed. Indeed, these types of cost-versus-quality trade-offs are ubiquitous in our economy, and are especially useful when resources are tightly constrained. Those following the long march to health-care reform know that one of the few things beyond argument is that the old approach is unsustainable and threatens to bankrupt the country. Perhaps a little belt tightening and bargain hunting of this sort might make our health-care dollars stretch farther.

Less expensive, lower-quality innovations abound in every economic sector—except medicine

David M. Kent is an associate professor of medicine at Tufts Medical Center and the associate director of the Clinical and Translational Science Program at the Sackler School for Graduate Biomedical Sciences of Tufts University. Address: Tufts Medical Center, 800 Washington St. #63, Boston, MA 02111. Email: dkent1@tuftsmedicalcenter.org

The Cost-effectiveness Plane

To help maximize the overall benefits in health care under a utilitarian framework and conditions of constrained resources, health economists use an analytic tool called cost-effectiveness analysis (CEA) that quantifies the added expenditure necessary to obtain a unit of health benefit (typically measured in quality-adjusted life years or QALYs, pronounced "kwallies"). The most common application of CEA is to examine the value of medical innovations compared to the standard of care routinely available, since new technologies are an important cause of the increase in health-care costs.

If the "unit cost" for a QALY of benefit (that is, the cost-effectiveness ratio) is less than some threshold (conventionally $50,000 or $100,000 per QALY), then adoption of the innovation is deemed "incrementally cost-effective," since the benefit obtained compares favorably to that obtainable at similar cost using accepted medical technologies (such as dialysis, which has a cost-effectiveness ratio variously estimated at between $50,000 and $80,000 per QALY). Above the ratio, they are deemed not to be cost-effective. That is, the (relatively small) incremental benefits of the intervention do not justify the (relatively large) incremental costs.

Comparisons between alternative approaches in cost-effectiveness analyses can usefully be depicted on a cost-effectiveness plane, shown in the figure opposite. Most studied medical innovations fall into the northeast quadrant of this plane; that is, they increase both costs and health benefits. Within this quadrant, the acceptability threshold would be represented by a line of constant slope, indicating the "willingness to pay" (WTP) for a QALY, separating nominally cost-effective therapies from cost-ineffective therapies.

Of course, if all innovation in health care fell into this northeast quadrant, innovation could only increase the costs of care. That is, even so-called cost-effective health-care innovations would always cost more money than the alternatives they replaced. This is often a point of confusion, sometimes purposeful, as when our political leaders claim that "preventative medicine" is highly cost-effective and would therefore save money. In fact, while most *recommended* preventative services are cost-effective (meaning the value of their benefits in terms of QALYs gained justifies the costs in terms of dollars spent), only very rarely are preventative services actually cost-saving, even when all the "downstream" avoided medical expenses are folded into the analysis. Indeed, new "cost-effective" innovations are one of the principal reasons that health-care costs continue to soar.

In fact, only innovations that fall south of the equator in the cost-effectiveness plane are actually cost-saving. When those innovations are also superior to the alternative, or standard of care, they are considered "dominant" (that is, cost decreasing *and* quality improving); adoption of these southeast quadrant innovations should not be controversial. However, as health-care costs continue to rise, cost-saving innovations may be increasingly attractive even when they do not improve care, particularly in a weak economy. While some innovations in the southwest quadrant would clearly be unattractive because they are substantially worse than the available standard of care or offer only trivial cost savings, what about innovations that offer substantial cost saving and are genuinely almost as good as the standard? In a 2004 article in *Medical Decision Making*, fellow researchers and I described innovation that is greatly cost saving but only slightly quality reducing as "decrementally" cost-effective. In such cases, the savings could potentially increase the overall good despite the sacrificed benefit. Indeed, if "much cheaper, almost as good" products are attractive in other economic sectors because they permit the reallocation of saved resources to items of more value than the benefits sacrificed, why not in medical care as well?

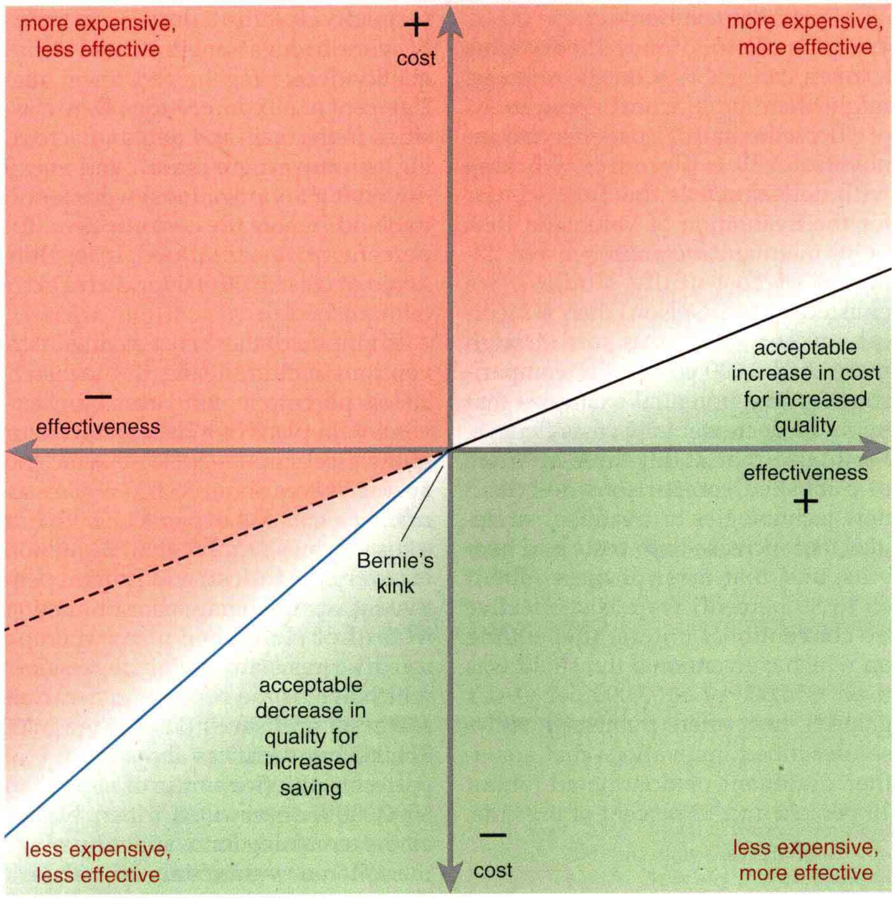

Medical innovations fall into one of four quadrants on the cost-effectiveness plane, based on how they compare with existing standards of care. For example, the top left quadrant represents innovative treatments that are more expensive and less effective—an off-putting combination; bottom right represents less expensive, more effective treatments—an easy decision. In between, the decision process is not so obvious. The diagonal lines represent thresholds for the acceptability of cost-effectiveness tradeoffs. Above the diagonals (in the red regions), the balance of cost and effectiveness is rejected. Of special interest is "Bernie's kink" at the origin, which reveals how medical markets actually behave. People prove to be unwilling to surrender quality using the same formula they would use to accept increased cost.

Bernie's Kink

> Men generally fix their affections more on what they are possessed of, than on what they never enjoyed: For this reason, it would be greater cruelty to dispossess a man of any thing than not to give it [to] him.—David Hume, *A Treatise on Human Nature*

Theoretically, perfectly rational economic agents seeking to maximize their welfare would be similarly willing to relinquish QALYs obtained from some routinely available standard-of-care for a new "much cheaper, almost as good" therapy, if the savings could be reallocated to an item of equal or higher value than what was sacrificed. Put another way, the selling price (often referred to as willingness to accept, or WTA) and the buying price (willing to pay, WTP) of a QALY should be similar, and the societal threshold for accepting or rejecting a technology should be symmetric and pass through the origin of the cost-effectiveness plane as a straight line. However, as David Hume anticipated, a reproducible observation is that consumers' willingness to accept monetary compensation to forgo something they have is typically greater, and often much greater, than their stated willingness to pay for the same benefit. Several explanations exist, including the so-called "endowment effect," the psychological principle that people value items that they already have simply because they already have them.

A 2002 review of 20 studies by the late Bernie O'Brien and his colleagues at McMaster University found that the ratio of individuals' WTA to WTP was always greater than 1 and ranged from 1.9 to 6.4 for two scenarios specifically related to health care. They suggested that rather than a symmetric accept-reject threshold on the cost-effectiveness plane, societal thresholds should reflect the WTA-WTP gap seen in individual preferences, which would be captured by a downward "kink" (subsequently known as "Bernie's kink") in the threshold as it passed through the origin, indicating that a QALY's selling price in the southwest would always be higher than a QALY's buying price in the northeast.

Thus, there may be an inherent cognitive bias against relinquishing the gains of health-care interventions that have already been accepted, and the cost savings from decrementally cost-effective innovation may need to be substantially greater than conventionally used thresholds suggest.

Bargain Hunting

Whereas all this fancy theory plus a token can get you on the subway, might there be practical applications of "decrementally" cost-effective innovation? To explore this, working with colleagues at the Tufts Center for the Evaluation of Value and Risk (who maintain a comprehensive database of cost-utility studies), we enlisted Aaron Nelson, then a medical student, to help us sort through more than 2,000 cost-utility comparisons for any potential examples that might be decrementally cost-effective. We found that about three-quarters of published comparisons described new technologies or treatment strategies that increase both costs and benefits, and that most of these (about 65 to 80 percent) were cost-effective by conventional criteria (depending on which conventional threshold was used, $50,000 or $100,000 per QALY gained). Less often, published analyses described innovations that are either dominant or dominated (about 10 percent and 15 percent of the time, respectively), but only very rarely were innovations both cost- and quality-decreasing. Indeed, fewer than 2 percent of all comparisons were classified in the cost- and quality-decreasing "southwest quadrant", and only 9 (involving 8 innovations) were found to be decrementally cost-effective (0.4 percent of the total)—that is, they saved at least $100,000 for each QALY relinquished.

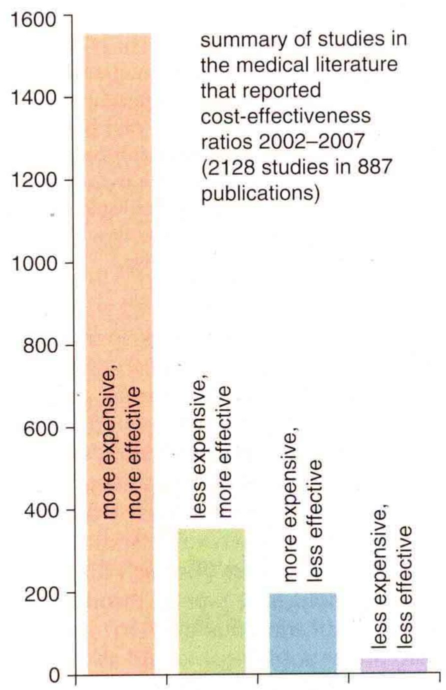

A survey of more than 2,000 medical studies that reported cost-effectiveness ratios highlights a striking difference between medical and other consumer markets. In the hurly-burly of retail markets, producing "nearly as good" products for less money is a major competitive strategy; in the medical literature, that type of innovation ("less expensive, less effective") is hardly represented at all *(purple bar)*.

Examples of these cost-saving interventions include using the catheter-based percutaneous coronary intervention in place of bypass surgery for multivessel coronary disease, which on average saves about $5,000 while sacrificing a half day of perfect health (for a cost-savings of more than $3 million for every QALY lost) and using repetitive transcranial magnetic stimulation instead of electroconvulsive therapy for drug-resistant major depression, which avoids the need for general anaesthesia and saves on average over $11,000 but sacrifices about a week of perfect health (for a ratio of more than $500,000 for every QALY lost). Nearly all the remaining innovations involved the tailored withholding of standard therapy, including watchful waiting for selected patients with inguinal hernia, withholding mediastinoscopy for selected patients with lung cancer, and abbreviated physiotherapy or psychotherapy for patients with neck pain or deliberate self-harm, respectively. Finally, the cost-saving innovations included the sterilization and reuse of dialysate, the chemical bath used in dialysis to draw fluids and toxins out of the bloodstream—a degree of thrift even the late Sheldon Kravitz would have to admire.

That decrementally cost-effective innovations are so rarely described in the health-care literature suggests that medicine is distinct from most other markets, in which cost-decreasing, quality-reducing products are continuously being introduced—think IKEA, Walmart and the Tata car. Several reasons may explain this "medical exceptionalism." First, there is fundamentally a lack of incentives both for physicians to control costs, especially under a fee-for-service regime, and for patients to demand less expensive treatment when insurance shields them from the direct costs of care. Second, medical "bargains" frequently come with health risks, and trading health for money strikes some as vulgar, regardless of ratio. The inherent ethical unease that decrementally cost-effective innovations can elicit poses a serious public relations and marketing challenge.

However, consumers have been comfortable with many decrementally cost-effective options outside of health care that pose similar health risks. For example, automobile manufacturers produce many vehicles that lack certain safety features (for example, side-impact airbags), because some consumers are willing to forgo those options to reduce the purchase price. Why not in health care?

Lowering Health Costs: Buy Less Stuff

Even by the standards of political rhetoric, it strains credulity when politicians suggest that the declared goals of health-care reform—increasing access, improving quality and controlling costs—are somehow mutually reinforcing. I'm no Peter Orszag, the über-wonk overseeing President Obama's Office of Management and Budget, but if my father taught me anything it was that saving money rarely involves buying more and better stuff. Plain talk about ways to cut costs are buried in rhetoric about rooting out inefficiencies and various prevarications about savings from investing in (that is, spending on) more preventative medicine, health information technology, and comparative effectiveness research about what therapies work best for which patients. While these goals may all be worthwhile, and there is much of little or no value in the current system (including the immense amount of money spent to maintain our Byzantine for-profit insurance system), ultimately we simply do not have the resources to give away an expensive commodity like health care in quantities that people want, subject to no budgetary constraints.

It is beyond dispute that some mechanisms for the controlled distribution of these expensive goods and services are required. In most markets, prices play this role, and many feel that the fundamental problem in health care is that many consumers are shielded from the costs of their care. A system based largely on prices (that is, price rationing) may control costs better than our current system, but it would of course mean that those with the most money have first dibs on scarce health-care resources, and there might be little left over for those without means. (There are other reasons too why most consumers can't be expected to comparison shop for emergency coronary angioplasty or for charged-particle radiosurgery for their glioblastoma the same way they might for gasoline, underwear and cling peaches). It is a fantasy to believe that price rationing alone can provide an acceptable mechanism for the controlled distribution of medical services, and some other means are thus also needed. Perhaps we should take it as a sign of the robustness of our democracy that this rather technical issue of the proper mix and variety of price and non-price rationing has somehow managed to plunge our national conversation about health-care reform into a Jerry Springer–style shouting match, except without the civility.

But regardless of the mix, expanding coverage to the uninsured, caring for our aging baby boomers, and accommodating new, effective technologies—while still feeding, clothing, housing, and educating ourselves, and catching an occasional movie—will require our system of distribution of health services to be more cost-sensitive, and will almost certainly mean the adoption of some decrementally cost-effective strategies for saving money. For example, Canadian-style delays for expensive diagnostic or surgical procedures certainly pose real, albeit small, medical risks, and would fall into this southwest category. Getting insured Americans to accept such new risks may be difficult, but slightly quality-reducing (that is, risk-increasing) cost-saving strategies have already been widely adopted within the American system, even if not studied or widely acknowledged. The gradual increase in the "hassle factor" in accessing medical care is one covert way that the industry has found to limit the distribution of services. More overt examples of rationing already adopted include aggressively shortening hospital stays and limiting formulary options (which sometimes require patients to change from a medicine they have been tolerating well to another in the same class). Despite the fact that doctors regularly (although sometimes disingenuously) deploy patter informing patients that the hospital is a dangerous place to stay and that the formulary medication is "just as good" as the one they've been taking, these strategies are certainly associated with small but real risks. Even a preadolescent quickly learns the true meaning of "just as good"; perhaps a more mature citizenry can also come to appreciate some of the upside of having "just as good" alternatives.

Bibliography

Orszag, P. R., and P. Ellis. 2007. The challenge of rising health care costs—a view from the Congressional Budget Office. *New England Journal of Medicine* 357:1793–1795.

Cohen, J. T., P. J. Neumann and M. C. Weinstein. 2008. Does preventive care save money? Health economics and the presidential candidates. *New England Journal of Medicine* 358:661–663.

Kent, D. M., A. M. Fendrick, and K. M. Langa. 2004. New and dis-improved: On the evaluation and use of less effective, less expensive medical interventions. *Medical Decision Making* 24:281–286.

O'Brien, B. J., K. Gertsen, A. R. Willan and L. A. Faulkner. 2002. Is there a kink in consumers' threshold value for cost-effectiveness in health care? *Health Economics* 11:175–180.

Nelson, A. L., J. T. Cohen, D. Greenberg and D. M. Kent. 2009. "Much Cheaper, Almost as Good": Decrementally Cost Effective Medical Innovation, *Annals of Internal Medicine* 151:662–667.

Avoiding a Digital Dark Age

Kurt D. Bollacker

When I was a boy, I discovered a magnetic reel-to-reel audio tape recorder that my father had used to create "audio letters" to my mother while he was serving in the Vietnam War. To my delight (and his horror), I could listen to many of the old tapes he had made a decade before. Even better, I could make recordings myself and listen to them. However, all of my father's tapes were decaying to some degree—flaking, stretching and breaking when played. It was clear that these tapes would not last forever, so I copied a few of them to new cassette tapes. While playing back the cassettes, I noticed that some of the sound quality was lost in the copying process. I wondered how many times I could make a copy before there was nothing left but a murky hiss.

A decade later in the 1980s I was in high school making backups of the hard drive of my PC onto 5-¼-inch floppy disks. I thought that because digital copies were "perfect," and I could make perfect copies of perfect copies, I couldn't lose my data, except by accident. I continued to believe that until years later in college, when I tried to restore my backup of 70 floppy disks onto a new PC. To my dismay, I discovered that I had lost the floppy disk containing the backup program itself, and thus could not restore my data. Some investigation revealed that the company that made the software had long since gone out of business. Requests on electronic bulletin board systems and searches on Usenet turned up nothing useful. Although all of the data on them may have survived, my disks were useless because of the proprietary encoding scheme used by my backup program.

Data longevity depends on both the storage medium and the ability to decipher the information

The Dead Sea scrolls, made out of still-readable parchment and papyrus, are believed to have been created more than 2,000 years ago. Yet my barely 10-year-old digital floppy disks were essentially lost. I was furious! How had the shiny new world of digital data, which I had been taught was so superior to the old "analog" world, failed me? I wondered: Had I had simply misplaced my faith, or was I missing something?

Over the course of the 20th century and into the 21st, an increasing proportion of the information we create and use has been in the form of digital data. Many (most?) of us have given up writing messages on paper, instead adopting electronic formats, and have exchanged film-based photographic cameras for digital ones. Will those precious family photographs and letters—that is, email messages—created today survive for future generations, or will they suffer a sad fate like my backup floppy disks? It seems unavoidable that most of the data in our future will be digital, so it behooves us to understand how to manage and preserve digital data so we can avoid what some have called the "digital dark age." This is the idea—or fear!—that if we cannot learn to explicitly save our digital data, we will lose that data and, with it, the record that future generations might use to remember and understand us.

Save Our Bits!

The general problem of data preservation is twofold. The first matter is preservation of the data itself: The physical media on which data are written must be preserved, and this media must continue to accurately hold the data that are entrusted to it. This problem is the same for analog and digital media, but unless we are careful, digital media can be more fragile.

The second part of the equation is the comprehensibility of the data. Even if the storage medium survives perfectly, it will be of no use unless we can read and understand the data on it. With most analog technologies such as photographic prints and paper text documents, one can look directly at the medium to access the information. With all digital media, a machine and software are required to read and translate the data into a human-observable and comprehensible form. If the machine or software is lost, the data are likely to be unavailable or, effectively, lost as well.

Preservation

Unlike the many venerable institutions that have for centuries refined their techniques for preserving analog data on clay, stone, ceramic or paper, we have no corresponding reservoir of historical wisdom to teach us how to save our digital data. That does not mean there is nothing to learn from the past, only that we must work a little harder to find it. We can start by briefly looking at the historical trends and advances in data representation in human history. We can also turn to nature for a few important lessons.

The earliest known human records are millennia-old physical scrapings on whatever hard materials were available. This medium was often stone, dried clay, bone, bamboo strips or even tortoise shells. These substances were very durable—indeed, some specimens have

Over the past two decades, Kurt D. Bollacker has romped through the fields of artificial intelligence, digital libraries, linguistics, databases and electrocardiology. He currently is the digital research director of the Long Now Foundation and gets his hands dirty as a freelance data miner and builder of collaborative knowledge-creation tools. He also works on the Rosetta Project. He received his Ph.D. in computer engineering from the University of Texas at Austin in 1998. Email: kurt@longnow.org

survived for more than 5,000 years. However, stone tablets were heavy and bulky, and thus not very practical.

Possibly the first big advance in data representation was the invention of papyrus in Egypt about 5,500 years ago. Paper was lighter and easier to make, and it took up considerably less space. It worked so well that paper and its variants, such as parchment and vellum, served as the primary repositories for most of the world's information until the advent of the technological revolution of the 20th century.

Technology brought us photographic film, analog phonographic records, magnetic tapes and disks, optical recording, and a myriad of exotic, experimental and often short-lived data media. These technologies were able to represent data for which paper cannot easily be used (video, for example). The successful ones were also usually smaller, faster, cheaper and easier to use for their intended applications. In the last half of the 20th century, a large part of this advancement included a transition from analog to digital representations of data.

Even a brief investigation into a small sampling of information-storage media technologies throughout history quickly uncovers much dispute regarding how long a single piece of each type of media might survive. Such uncertainty cannot be settled without a time machine, but we can make reasonable guesses based on several sources of varying reliability. If we look at the time of invention, the estimated lifespan of a single piece of each type of media and the encoding method (analog or digital) for each type of data storage *(see the table, above right)*, we can see that new media types tend to have shorter lifespans than older ones, and digital types have shorter lifespans than analog ones. Why are these new media types less durable? Shouldn't technology be getting better rather than worse? This mystery clamors for a little investigation.

To better understand the nature of and differences between analog and digital data encoding, let us use the example of magnetic tape, because it is one of the oldest media that has been used in both analog and digital domains. First, let's look at the relationship between information density and data-loss risk. A standard 90-minute analog compact cassette is 0.00381 meters wide by about 129 meters long, and a typical digital audio tape (DAT) is 0.004 meters wide by 60 meters long. For audio encodings of similar quality (such as 16 bit, 44.1 kilohertz for digital, or 47.6 millimeters per second for analog), the DAT can record 500 minutes of stereo audio data per square meter of recordable surface, whereas the analog cassette can record 184 minutes per square meter. This means the DAT holds data about 2.7 times more densely than the cassette. The second table *(below)* gives this comparison for several common consumer audio-recording media types. Furthermore, disk technologies tend to hold data more densely than tapes, so it is no surprise that magnetic tape has all but disappeared from the consumer marketplace.

However, enhanced recording density is a double-edged sword. Assume that for each medium a square millimeter of surface is completely corrupted. Common sense tells us that media that hold more data in this square millimeter would experience more actual data loss; thus for a given amount of lost physical medium, more data will be lost from digital formats. There is a way to design digital encoding with a lower data density so as to avoid this problem, but it is not often used. Why? Cost and efficiency: It is usually cheaper to store data on digital media because of the increased density.

type of medium	data medium	approximate year of invention	ideal expected lifetime of medium
analog	clay/stone tablet	8000 BC	>4,000 years
analog	pigment on paper	3500 BC	>2,000 years
analog	oil painting	600	centuries
analog	silver halide black and white photographic film	1820	>100 years
analog	modern color photographic film	1860	decades
analog	phonograph record	1877	>120 years
analog/digital	magnetic tape	1928	decades
analog/digital	magnetic disk	1950	3–20 years
analog/digital	polycarbonate optical WORM disk	1990	5–20 years

When we compare the different data-storage media that have appeared over the course of human history, a trend emerges: Digital data types are expected to have shorter lifetimes than analog ones.

type of medium	audio data medium	recording capacity (minutes per square meter)
analog	6.35 millimeter wide 190.5 millimeters per second reel-to-reel magnetic tape	13.8
analog	33-1/3 RPM vinyl album	411
analog	90-minute audio cassette	184
digital	compact disk (CD)	8,060
digital	60-meter digital audio tape (DAT)	500
digital	2 terabyte 89-millimeter hard drive	4,680,000

As technology has advanced, the density of data storage on analog and, subsequently, digital recording media has tended to increase. The downside of packing in data, however, is that more of the information will be lost if a portion of the recording medium becomes damaged.

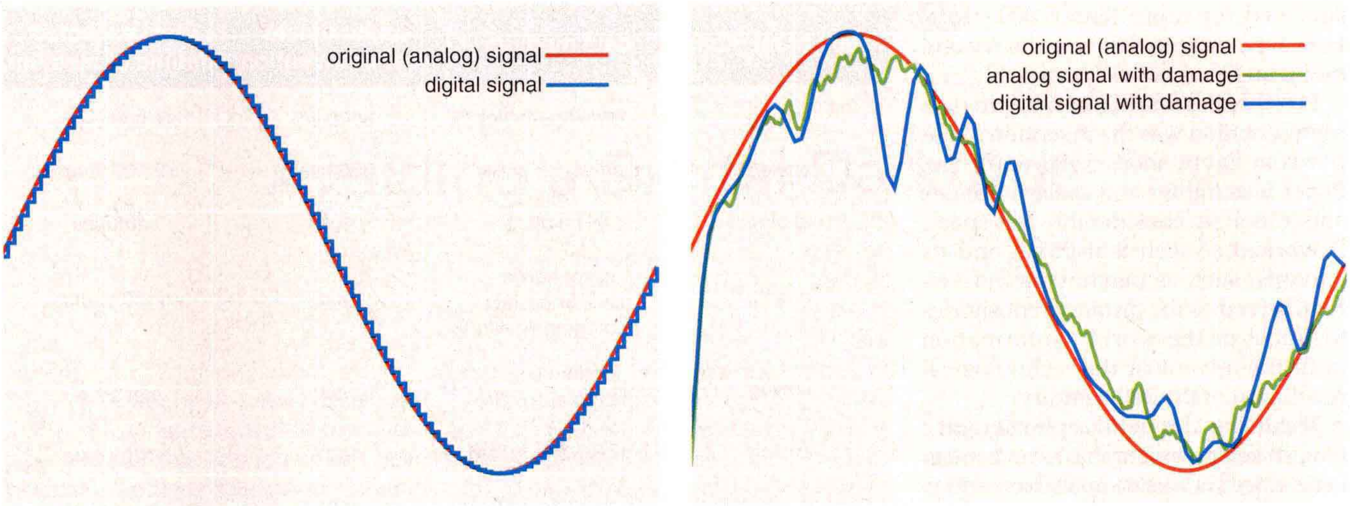

A simple audio tone is represented as a sine wave in an analog signal, and as a similar wave but with an approximated stepped shape in a digital signal *(left)*. If the data receive simulated damage, the analog signal output is more resistant to damage than the digital one, which has wilder swings and higher error peaks *(right)*. This result is largely because in a digital recording, all bits do not have the same worth, so damage causes random output error.

A possibly more important difference between digital and analog media comes from the intrinsic techniques that comprise their data representations. Analog is simply that—a physical analog of the data recorded. In the case of analog audio recordings on tape, the amplitude of the audio signal is represented as an amplitude in the magnetization of a point on the tape. If the tape is damaged, we hear a distortion, or "noise," in the signal as it is played back. In general, the worse the damage, the worse the noise, but it is a smooth transition known as *graceful degradation.* This is a common property of a system that exhibits *fault tolerance,* so that partial failure of a system does not mean total failure.

Unlike in the analog world, digital data representations do not inherently degrade gracefully, because digital encoding methods represent data as a string of binary digits ("bits"). In all digital symbol number systems, some digits are worth more than others. A common digital encoding mechanism, pulse code modulation (PCM), represents the total amplitude value of an audio signal as a binary number, so damage to a random bit causes an unpredictable amount of actual damage to the signal.

Let's use software to concoct a simulated experiment that demonstrates this difference. We will compare analog and PCM encoding responses to random damage to a theoretically perfect audiotape and playback system. The first graph in the third figure *(above)* shows analog and PCM representations of a single audio tone, represented as a simple sine wave. In our perfect system, the original audio source signal is identical to the analog encoding. The PCM encoding has a stepped shape showing what is known as *quantization error,* which results from turning a continuous analog signal into a discrete digital signal. This class of error is usually imperceptible in a well-designed system, so we will ignore it for now.

For our comparison, we then randomly damage one-eighth of the simulated perfect tape so that the damaged parts have a random amplitude re-

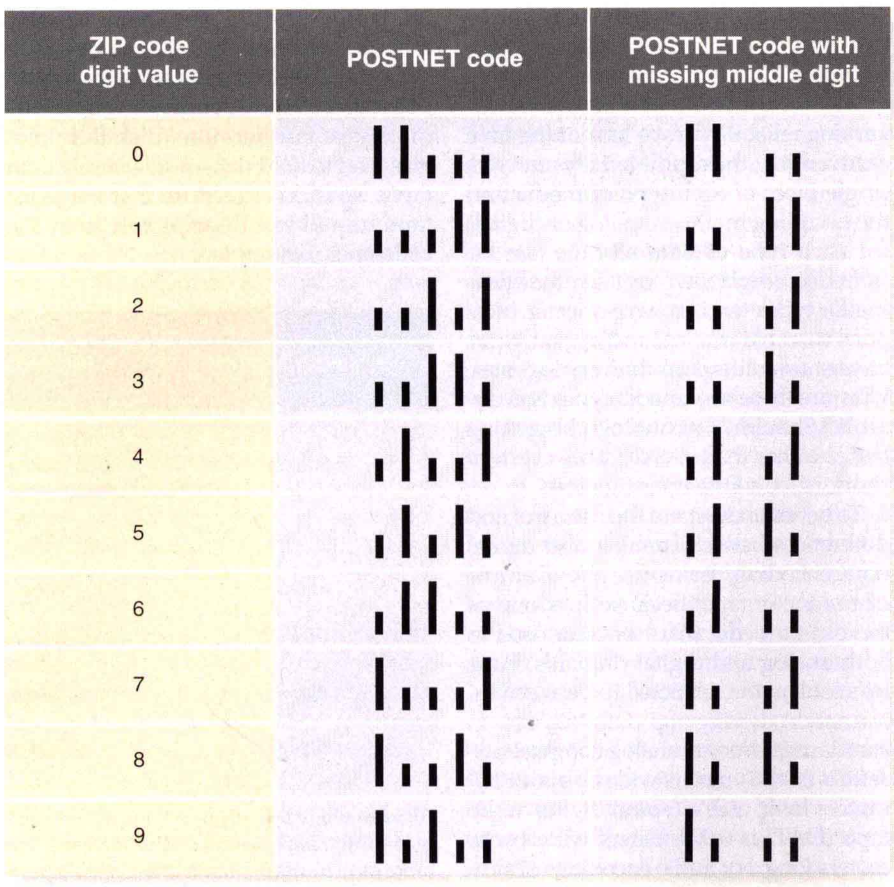

The U.S. Postal Service uses an encoding scheme for ZIP code numbers called POSTNET that uses an error-correcting code. Each decimal digit is represented as five bars. If, say, the middle bar disappears, each number is still distinguishable from all the others.

The Phaistos Disk, housed at the Heraklion Archaeological Museum in Crete, is well preserved and all its data are visible, but the information is essentially lost because the language in which it is written has been forgotten. (Photograph courtesy of Wikimedia Commons.)

sponse. The second graph in the third figure *(facing page, top)* shows the effect of the damage on the analog and digital encoding schemes. We use a common device called a *low-pass filter* to help minimize the effect of the damage on our simulated output. Comparing the original undamaged audio signal to the reconstructions of the damaged analog and digital signals shows that, although both the analog and digital recordings are distorted, the digital recording has wilder swings and higher error peaks than the analog one.

But digital media are supposed to be better, so what's wrong here? The answer is that analog data-encoding techniques are intrinsically more robust in cases of media damage than are naive digital-encoding schemes because of their inherent redundancy—there's more to them, because they're continuous signals. That does not mean digital encodings are worse; rather, it's just that we have to do more work to build a better system. Luckily, that is not too hard. A very common way to do this is to use a binary-number representation that does not mind if a few bits are missing or broken.

One important example where this technique is used is known as an error correcting code (ECC). A commonly used ECC is the U.S. Postal Service's POSTNET (Postal Numeric Encoding Technique), which represents ZIP codes on the front of posted envelopes. In this scheme, each decimal digit is represented as five binary digits, shown as long or short printed bars *(facing page, bottom)*. If any single bar for any decimal digit were missing or incorrect, the representation would still not be confused with that of any other digit. For example, in the rightmost column of the table, the middle bar for each number has been erased, yet none of the numbers is mistakable for any of the others.

Although there are limits to any specific ECC, in general, any digital-encoding scheme can be made as robust as desired against random errors by choosing an appropriate ECC. This is a basic result from the field of information theory, pioneered by Claude Shannon in the middle of the 20th century. However, whichever ECC we choose, there is an economic tradeoff: More redundancy usually means less efficiency.

Nature can also serve as a guide to the preservation of digital data. The digital data represented in the DNA of living creatures is copied into descendents, with only very rare errors when they reproduce. Bad copies (with destructive mutations) do not tend to survive. Similarly, we can copy digital data from medium to medium with very little or no error over a large number of generations. We can use easy and effective techniques to see whether a copy has errors, and if so, we can make another copy. For instance, a common error-catching program is called a *checksum function*: The algorithm breaks the data into binary numbers of arbitrary length and then adds them in some fashion to create a total, which can be compared to the total in the copied data. If the totals don't match, there was likely an accidental error in copying. Error-free copying is not possible with analog data: Each generation of copies is worse than the one before, as I learned from my father's reel-to-reel audiotapes.

Because any single piece of digital media tends to have a relatively short lifetime, we will have to make copies far more often than has been historically required of analog media. Like species in nature, a copy of data that is more easily "reproduced" before it dies makes the data more likely to survive. This notion of *data promiscuousness* is helpful in thinking about preserving our own data. As an example, compare storage on a typical PC hard drive to that of a magnetic tape. Typically, hard drives are installed in a PC and used frequently until they die or are replaced. Tapes are usually written to only a few times (often as a backup, ironically) and then placed on a shelf. If a hard drive starts to fail, the user is likely to notice and can quickly make a copy. If a tape on a shelf starts to die, there is no easy way for the user to know, so very often the data on the tape perishes silently, likely to the future disappointment of the user.

Comprehensibility

In the 1960s, NASA launched *Lunar Orbiter 1*, which took breathtaking, famous photographs of the Earth juxtaposed with the Moon. In their rush to get astronauts to the Moon, NASA engineers created a mountain of magnetic tapes containing these important digital images and other space-mission-related data. However, only a specific, rare model of tape drive made for the U.S. military could read these tapes, and at the time (the 1970s to 1980s), NASA had no interest in keeping even one compatible drive in good repair. A heroic NASA archivist kept several donated broken tape drives in her garage for two decades until she was able to gain enough public interest to find experts to repair the drives and help her recover these images.

Contrast this with the opposite problem of the analog Phaistos Disk *(above left)*, which was created some 3,500 years ago and is still in excellent physical condition. All of the data it stores (about 1,300 bits) have been preserved and are easily visible to the human eye. However, this disk shares one unfortunate characteristic with my set of 20-year-old floppy disks: No one can decipher the data on either one. The language in which the Phaistos disk was written has long since been forgotten, just like the software to read my floppies is equally irretrievable.

These two examples demonstrate digital data preservation's other challenge—comprehensibility. In order to survive, digital data must be understandable by both the machine reading them and the software interpreting them. Luckily, the short lifetime of digital media has forced us to gain some experience in solving this problem—the silver lining of the dark clouds of a looming potential digital dark age. There are at least two effective approaches: choosing data representation technologies wisely and creating mechanisms to reach backward in time from the future.

Make Good Choices …

In order to make sure digital data can be understood in the future, ideally we should choose representations for our data for which compatible hardware and software are likely to survive as well. Like species in nature, digital formats that are able to adapt to new environments and threats will tend to survive. Nature cannot predict the future, but the mechanism of mutation creates different species with different traits, and the fittest prevail.

Because we also can't predict the future to know the best data-representation choices, we try to do as nature does. We can copy our digital data into as many different media, formats and encodings as possible and hope that some survive.

Another way to make good choices is to simply follow the pack. A famous example comes from the 1970s, when two competing standards for home video recording existed: Betamax and VHS. Although Betamax, by many technical measures, was a superior standard and was introduced first, the companies supporting VHS had better business and marketing strategies and eventually won the standards war. Betamax mostly fell into disuse by the late 1980s; VHS survived until the mid-2000s. Thus if a format or media standard is in more common use, it may be a better choice than one that is rare.

... Or Fake It!

Once we've thrown the dice on our data-representation choices, is there anything else we can do? We can hope we will not be stuck for decades, like our NASA archivist, or left with a perfectly readable but incomprehensible Phaistos disk. But what if our scattershot strategy of data representation fails, and we can't read or understand our data with modern hardware and software? A very common approach is to fake it!

If we have old digital media for which no compatible hardware still exists, modern devices sometimes can be substituted. For example, cheap and ubiquitous optical scanners have been commonly used to read old 80-column IBM punchcards. This output solves half of the problem, leaving us with the task of finding hardware to run the software and interpret the data that we are again able to read.

In the late 1950s IBM introduced the IBM 709 computer as a replacement for the older model IBM 704. The many technical improvements in the 709 made it unable to directly run software written for the 704. Because customers did not want either to lose their investment in the old software or to forgo new technological advances, IBM sold what they called an *emulator* module for the 709, which allowed it to pretend to be a 704 for the purposes of running the old software. Emulation is now a common technique used to run old software on new hardware. It does, however, have a problem of recursion—what happens when there is no longer compatible hardware to run the emulator itself? Emulators can by layered like Matryoshka dolls, one running inside another running inside another.

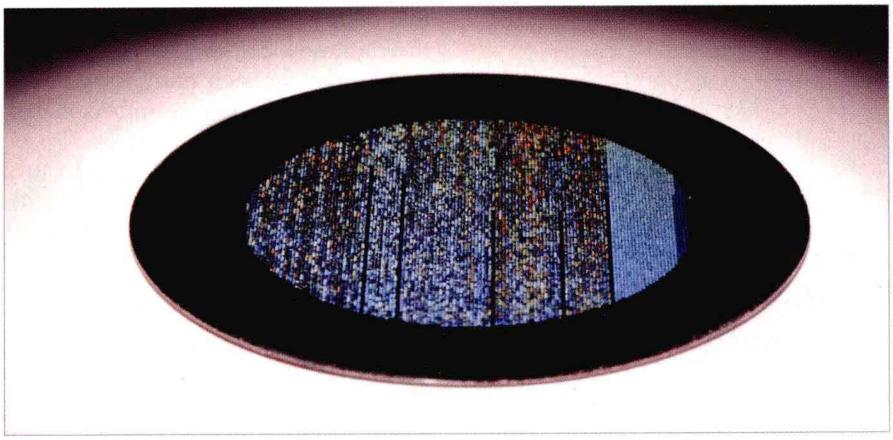

The Rosetta Project aims to preserve all of the world's written languages with a metal disk that could last up to 2,000 years. The disk records miniaturized versions of more than 13,000 pages of text and images, etched onto the surface using techniques similar to computer-chip lithography. (Photograph by Spencer Lowell, courtesy of the Long Now Foundation, www.longnow.org.)

Being Practical

Given all of this varied advice, what can we do to save our personal digital data? First and foremost, make regular backup copies onto easily copied media (such as hard drives) and place these copies in different locations. Try reading documents, photos and other media whenever upgrading software or hardware, and convert them to new formats as needed. Lastly, if possible, print out highly important items and store them safely—there seems to be no getting away from occasionally reverting to this "outdated" media type. None of these steps will guarantee the data's survival, but not taking them almost guarantees that the data will be lost, sooner or later. This process does seem to involve a lot more effort than my grandparents went to when shoving photos into a shoebox in the attic decades ago, but perhaps this is one of the costs for the miracles of our digital age.

If all this seems like too much work, there is one last possibility. We could revert our digital data back to an analog form and use traditional media-preservation techniques. An extreme example of this is demonstrated by the Rosetta Project, a scholarly endeavor to preserve parallel texts of all of the world's written languages. The project has created a metal disk *(above)* on which miniaturized versions of more than 13,000 pages of text and images have been etched using techniques similar to computer-chip lithography. It is expected that this disk could last up to 2,000 years because, physically, the disk has more in common with a stone tablet than a modern hard drive. Although this approach should work for some important data, it is much more expensive to use in the short term than almost any practical digital solution and is less capable in some cases (for example, it's not good for audio or video). Perhaps it is better thought of as a cautionary example of what our future might look like if we are not able to make the digital world in which we find ourselves remain successful over time.

Bibliography

Balistier, Thomas. 2000. *The Phaistos Disc: An Account of Its Unsolved Mystery.* New York: Springer-Verlag.

Besen, Stanley M., and Joseph Farrell. 1994. Choosing how to compete: Strategies and tactics in standardization. *Journal of Economic Perspectives* 8:117–131.

Camras, Marvin. 1988. *Magnetic Recording Handbook.* New York: Van Nostrand Reinhold Co.

The IBM 709 Data-Processing System. www-03.ibm.com/ibm/history/exhibits/mainframe/mainframe_PP709.html

Koops, Matthias. 1800. *Historical Account of the Substances Which Have Been Used to Describe Events, and to Convey Ideas, from the Earliest Date, to the Invention of Paper.* London: T. Burton.

Pohlmann, Ken C. 1985. *Principles of Digital Audio,* 2nd ed. Carmel, Indiana: Sams/Prentice-Hall Computer Publishing.

The Rosetta Project. www.rosettaproject.org

United States Postal Service, Domestic Mail Manual 708.4—Special Standards, Technical Specifications, Barcoding Standards for Letters and Flats.