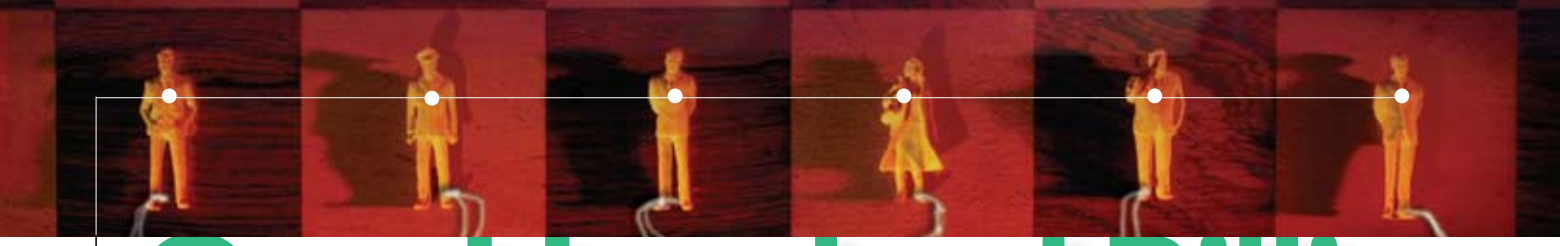

One Hundred Billion

BY TOM BAKOS

In 2005, University of Chicago economist Kevin Murphy was awarded a MacArthur Fellowship "genius" grant for his theory that the investment in basic health research and care results in orders-of-magnitude returns in economic value. But how does his theory stack up actuarially?

"THE VALUE OF HEALTH AND LONGEVITY," a paper by Kevin M. Murphy and Robert H. Topel published in June 2005, presents an economic framework for valuing improvements in human life expectancy and health. Certainly, a discussion of valuing anything is likely to get an actuary's attention. In fact, at least one among us (see Don Behan's web page at http://behan.ws/lifevalue.htm?Num=3) has already addressed this human life value issue. He takes a more traditional view, however, that's directed at placing a value on a single human life in terms of life expectancy and earning potential. This individualized approach is often used in connection with wrongful death or personal injury cases. The scope of Murphy and Topel's paper is much grander than that.

The authors assume humanity to be a sort of monolithic entity with a shared social mind, defined perhaps by an inherent majority consensus and a presumption that altruism rules individuals' pocketbooks. Such a society would make economic decisions regarding spending for life or health improvements based on the value those expenditures added to society as a whole. This seems to be consistent, generally, with the economists' view of the world or, at least, the economy they're evaluating.

For example, the authors calculate that a permanent 1 percent reduction in the mortality rate of cancer would improve the value of human life in the United States by almost $500 billion. A 100 percent cancer cure, they calculate, would be worth about $50 trillion. These numbers, which represent the economic value of the postulated reduction in cancer mortality or a cancer cure, set a limit on what they think Americans should be willing to pay for such a reduction or cure.

In this example, the authors go on to speculate that a "war on cancer" costing $100 billion (to develop and implement a new treatment) would be "worth it" if it had only a 20 percent chance of reducing cancer mortality by 1 percent. The analysis done is relative to the U.S. population, which consists of around 297 million people by the time you read this.

The Theory

On the whole, this seems like an interesting, effective, and objective decision-making process to apply in a society that has limited dollar resources. The theory would be to spend money on human health and longevity improvements only where it would do the most good and only if the value of the good done is greater than the cost. Of course, the example implies that there also has to be a willingness to accept some risk. After all, there is a four-out-of-five chance that spending the $100 billion won't help cancer mortality at all.

The workability of the theory also depends on individuals' willingness (note this is plural) to pay. It's always easier to spend somebody else's money or, more simply, not to have any strong objections at all if somebody else's money is being spent.

Governmental tax structures, to the extent they're progressive, have a Robin Hood flair to them. Solving society's health problems with tax dollars, therefore, won't require an equal contribution from everybody, and any single individual's "willingness to pay" may be affected by how much that particular individual might be expected to pay. The theory does address this issue, however, as we'll see later.

Even more to the point, if private or corporate dollars paid the $100 billion upfront research and development costs for a new drug that, say, reduced cancer mortality by 1 percent, these expenditures would ultimately be transferred to society through higher drug prices. Who would pay the higher drug prices?

Well, to some extent we all would, through higher insurance premiums, but that gets a little complicated. If a high-priced drug really could reduce the incidence of cancer, then cancer claim costs would be reduced, and the increased cost of the new drug and health insurance premiums would be mitigated.

Let's assume that the authors considered this and that the hypothetical $100 billion is a *net dollar cost* that has to be passed on to society, offset only by the improvement in the value of human life. Unless this new drug were a preventive as opposed to a curative treatment, its cost would be borne entirely by the smaller

PUNCHSTOCK/BONOTOM STUDIO

Gazillion Dollars

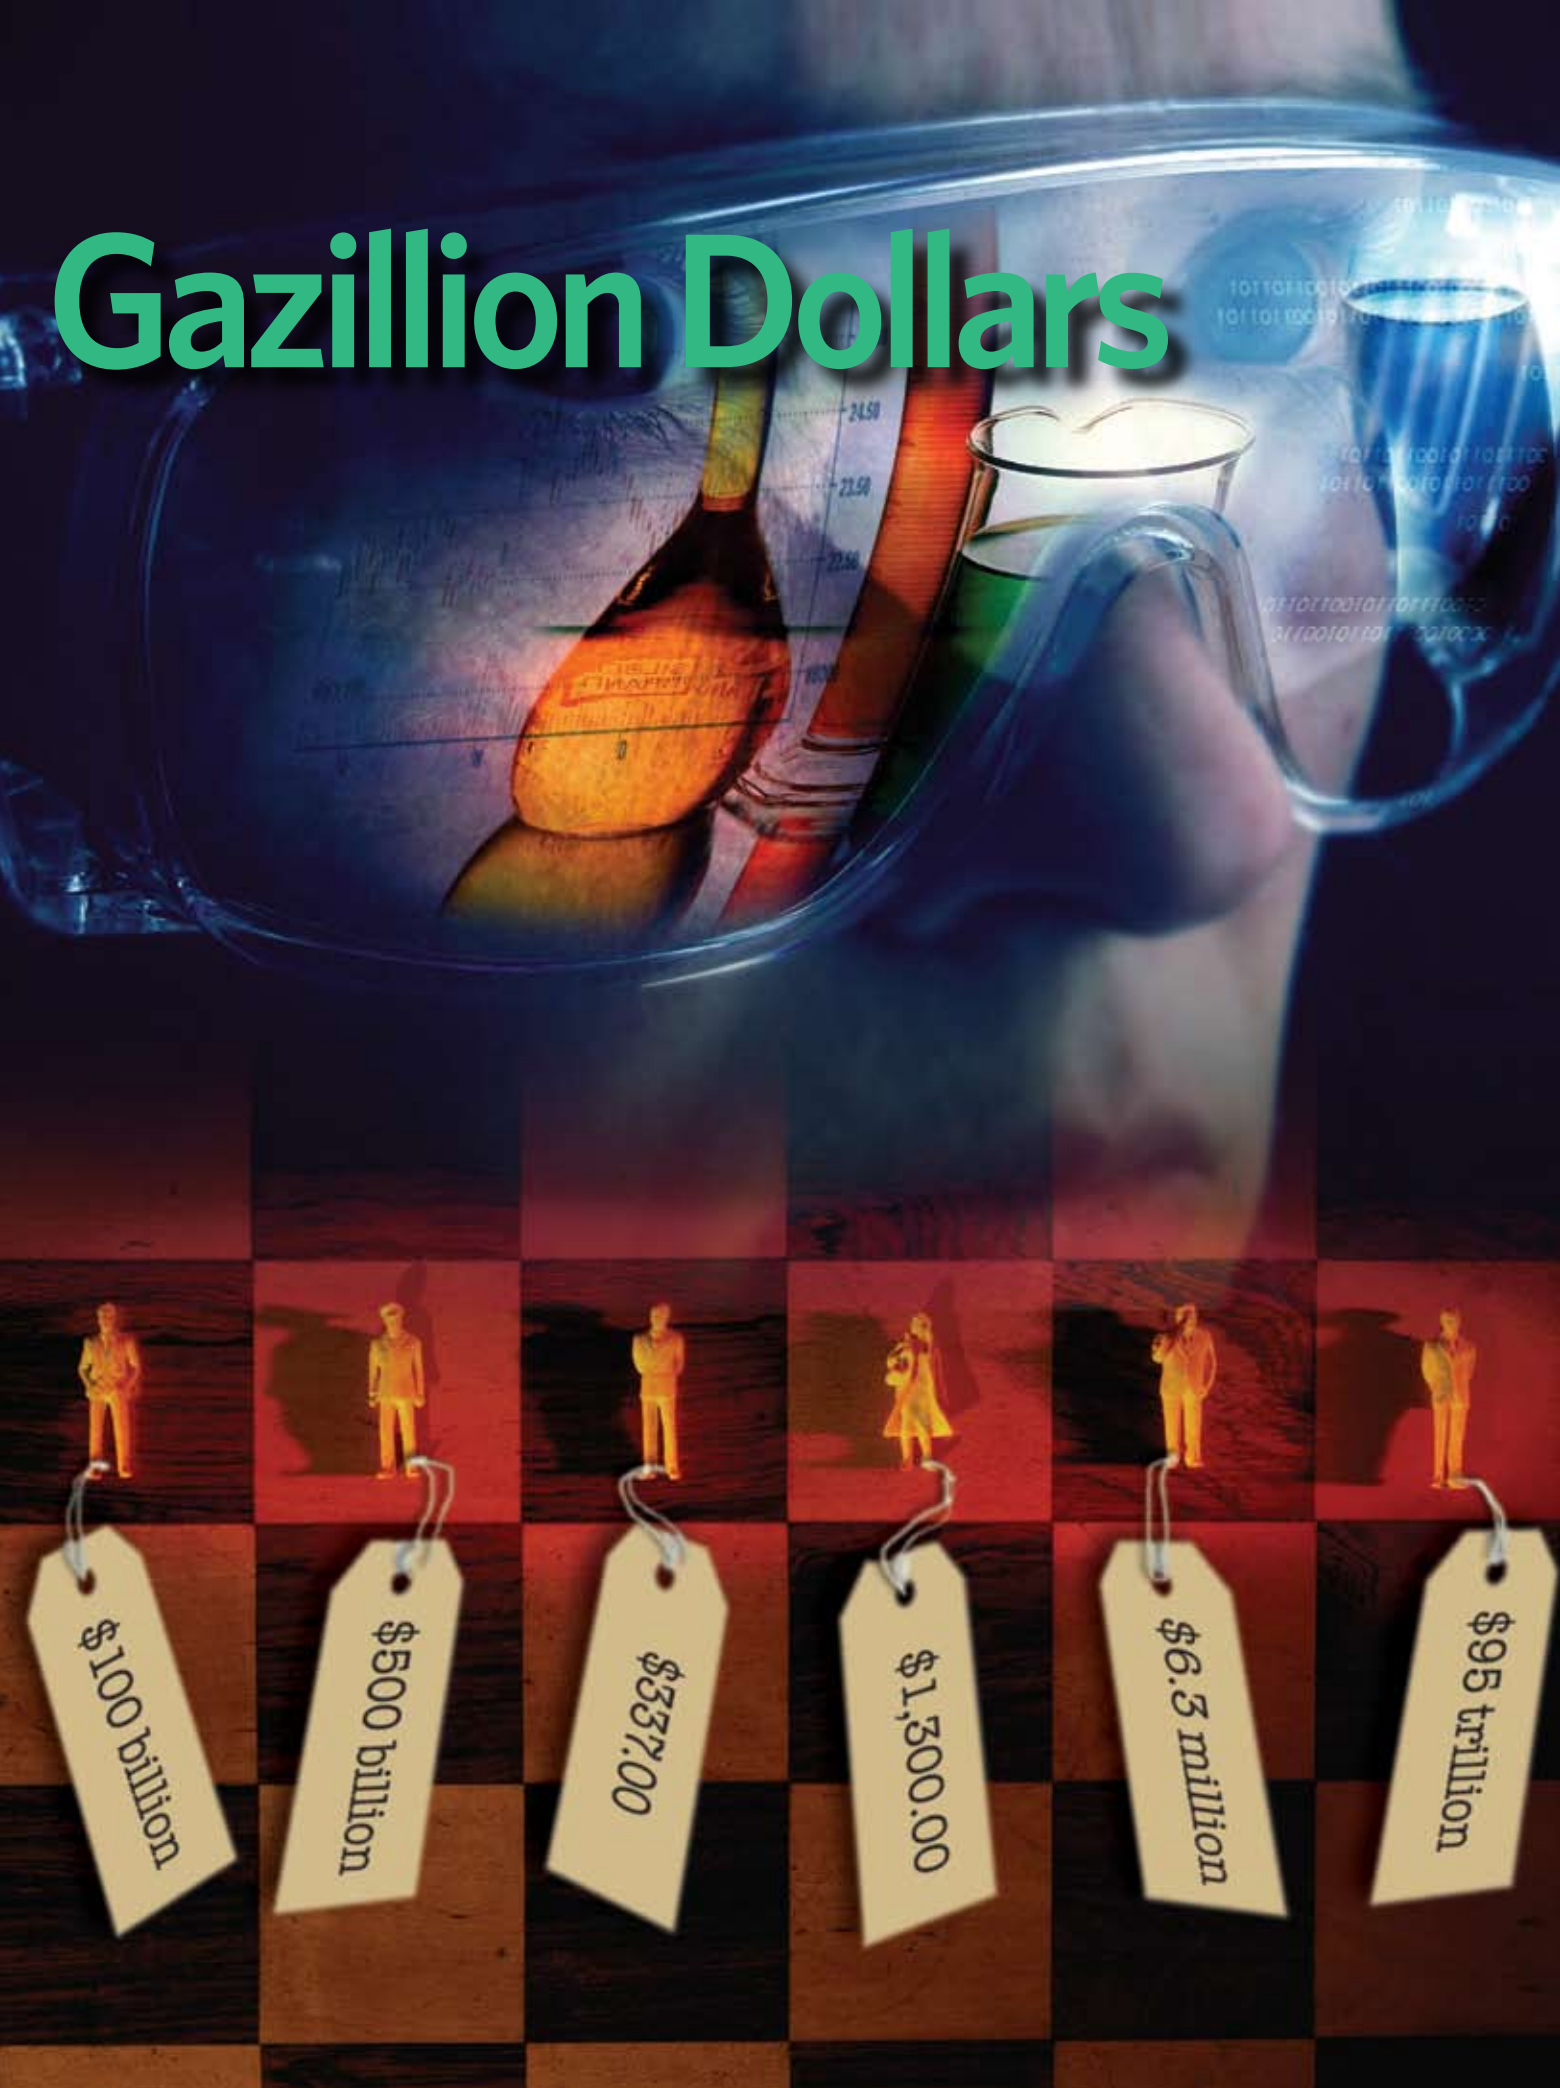

segment of the population that actually needed it—those with cancer. So the question should be: "What would those people be willing to pay?"

But ignore that question for the moment. In order to make "empirical headway" in their presentation, the authors assume utility to be *homothetic*. What does it mean for "utility" to be "homothetic"? Well, consider first that *utility*, generally, means being of practical value or use. In economics, utility is a measure of the happiness or satisfaction individuals gain from the ownership or use of a good or service. Assuming that utility is *homothetic* means that the demand for a good or service is related only to the relative prices of goods or services and not dependent on the consumer's ability to pay those prices. That is, differences in personal characteristics, such as tax brackets or ability to pay, aren't a factor in determining a willingness to pay.

So, let's look at this $100 billion cost on a homothetic utility basis. If it were assessed equally among all U.S. citizens, each would have to chip in about $337 in order to reduce cancer mortality by 1 percent. Now, it would be easy to say that the benefit in increased life value would be five times this, because the potential gain to society is $500 billion. But we would be forgetting about the risk. Since there is only a one-in-five chance that the per-life expenditure would improve cancer mortality at all, the expected return from this expenditure at the time it's made is really only the $337 spent. Of course, if the first try works, it would be like winning a lottery, and the value of human life would go up by $500 million at a cost of only $100 million.

But let's look more deeply into this cost issue. In the United States, the economic unit for a social expenditure to improve life or health is more likely to be a household, of which there are about 77 million. On a per-household basis, the cost for a 20 percent chance of a 1 percent drop in cancer mortality would be, roughly, almost $1,300. What impact would that kind of cost have on your willingness to pay?

Willingness to Pay

The concept of an individual's willingness to pay—or, more precisely, individuals' willingness to pay—is a key element of the authors' discussion about an economic basis for making investment decisions about life or health improvement in a society.

The economic approach they use looks at social rather than individual costs and benefits. As Spock would say (the Vulcan, not the doctor), "The needs of the many outweigh the needs of the few, or the one." In economic theory, to the extent that individuals are involved, they are "representative" individuals or Spock-like in their perceived duty to society. They are, in effect, merged entities representing the *fully informed* best interests of the population they serve. They're not uninformed, consensus individuals, 20 percent of whom believe the sun revolves around the Earth, and they're certainly not actual specific individuals with personal agendas. In insurance terms, economic theory, at least this one, has room for only one broad underwriting class.

However, this economic approach does recognize that different segments of a society may have different views of social spending on health care improvement. More generally, and in keeping with the representative individual model, as a society changes in its average mean values, it recognizes that its representative views on social spending for life and health improvement may also change.

For example, perhaps it's obvious that wealthier individuals are likely to be more willing to spend for improvements in life and health, not just because they have more to spend but because they value life more. They'd have more to lose if they didn't and more to gain if they did. And this easily translates into an assumption that the representative member of wealthier societies, standing for the overall interests of the wealthier society, is likely to be willing to spend more for life and health improvements.

In addition, it's recognized that spending for life and health improvements may affect segments of society unequally or that the timing of the improvements may be significantly offset from the expenditure. For example, the adverse effects of cardiovascular diseases are focused on individuals age 50 and older. While spending money to improve cardiovascular mortality will have a positive effect on everyone eventually, even 20-year-olds who will for the most part become 50-year-olds, such expenditures made now are likely to be of greater value to current 50-year-olds than to current 20-year-olds.

Value of Human Life

A willingness to pay is decided by a representative individual through a hypothetical comparison between the cost of a life or health improvement and the increase in the value of life or health that expenditure is expected to create. There are two types of improvement the authors recognize: a type-H improvement raises the *quality of life* without affecting mortality; and a type G *improves mortality* without affecting the quality of life.

Type-H improvements are, for example, new medicines that control the symptoms of a disease or condition (e.g., arthritis, migraine headaches) or prosthetics that make a disability easier to deal with but have no curative effect and, therefore, no impact on longevity. Indeed, the condition being addressed by a type-H improvement may, itself, not have a negative impact on life expectancy. Some health improvements can be both type H and type G, like blood pressure medication, which, if taken to reduce high blood pressure, will affect both the quality and the length of life.

Type-H improvements that increase only the quality of life are notoriously difficult to measure. The authors draw the conclusion from their numerical analysis that "living a bit better is like living a bit longer."

The type-G improvement reduces mortality, increasing life expectancy, and therefore increases the total number of years lived in a society. One method for valuing a type-G improvement considers the extra salary a worker would demand for a job with a higher risk of death. This is used to calculate the value of

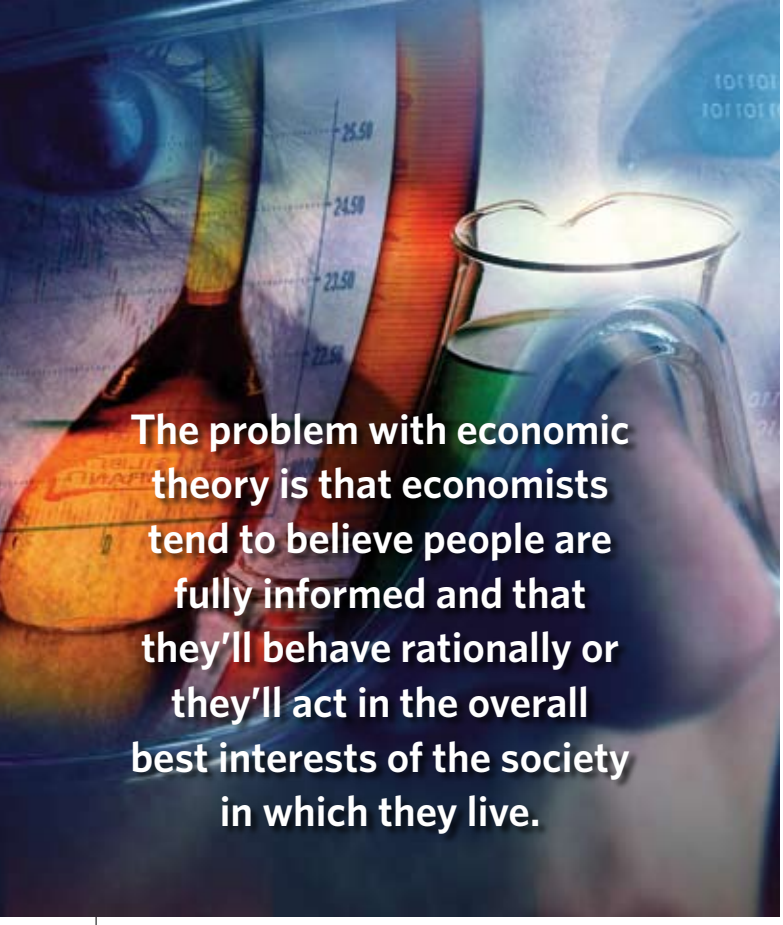

a "statistical life." A "statistical life" is one life that is either lost or being saved in one year.

The authors use, by way of example, a job that has an additional death rate of one per 10,000 per year. If a set of 10,000 of these workers demanded $500 more per year to offset the expectation that an additional one of them was likely to die from a work-related injury (compared with a similar job without the additional risk), then these workers would have implied a statistical life value of $5 million ($500 x 10,000).

This is a typical approach to finding the value of a statistical life. Of course, it assumes a fully informed workforce with access to data and that's capable of doing the math. Alternatively, one could examine what a group of people would be willing to pay for a mortality improvement that would save one life per year and use that to place a value on a statistical life.

For their purposes, Murphy and Topel use $6.3 million as the value of a statistical life, which is a value used since 1999 (adjusted for inflation) by the Environmental Protection Agency (EPA) in doing cost-benefit analyses of its rules and regulations. Of course, the value of a statistical life is subjective, and one might want to use a higher or lower value to estimate the value of improvements in mortality.

This value of a statistical life is a value related only to type-G or mortality improvements. The authors use it to calculate significant economic gains from mortality reductions achieved from 1970 to 2000, approximately equal to $95 trillion. During the same period, they estimate the increase in health expenditures was only about $35 trillion. Therefore, it appears that the health expenditures were well spent. And if a type-H (or quality-of-life) improvement value is factored in, the benefits would be even greater.

It's suggested that quality of life can be quantified using something along the lines of an inverse age rating technique. That is, by comparing mortality rates, one might find that a person age x in 2000 has mortality equivalent to a younger person age x-k in 1970. Therefore, he or she has experienced a gain in quality of life that can be quantified by valuing the (roughly) additional k more years of life. This, the authors estimate, would add about $1 million to the value of a statistical life—more for men than for women.

It's an approximate value that relies a great deal on averages, and improvements can certainly be made. But it's only one approach to valuing quality of life. With baby boomers now reaching retirement age in greater numbers and in better health than their grandparents and parents, it also seems like a somewhat reasonable approach until something better comes along.

Persistent Questions

OK. Great theory. But does it explain everything? If the numbers so compellingly show that we in the United States have received far more value than the cost we've paid for the improvements in life and health we've so far enjoyed, then why is there such a loud and persistent discussion on what has been characterized as the outrageous cost of health care today? Why do health care costs need to be controlled? Why is everybody looking for someone to blame? Why aren't we happy with the results instead of being critical of the costs? Shouldn't we be spending more so we can get more?

And how does one explain that some of us are willing to pay for utility, goods, and services, that reduce our life expectancy? For example, smoking is measurably bad for mortality; smoker mortality is about 2.5 times non-smoker mortality. Yet, although smokers are fewer in number, some people still smoke and tobacco products aren't cheap. One might blend this into economic theory by assuming that the quality of life type-H improvement smoking brings offsets the mortality type-G loss in value for individuals who smoke.

And, of course, hazardous avocations require a willing and able spender who again, I assume, sees non-economic quality-of-life value added that more than offsets the economic value lost in higher mortality.

Murphy and Topel's paper isn't unique in its approach. Willingness to pay appears to be a fairly standard economic method for valuing decisions that affect life and health. The EPA uses lives saved, and the value of a statistical life saved, as a standard benefit offset in evaluating the costs of its rules and regulations aimed at improving life and health expectancy. In automobile safety, cost-benefit analyses are used to justify or attack the cost of safety equipment.

The problem with economic theory is that economists tend to believe people are fully informed and that they'll behave rationally. Or they believe that there are representative individuals who will act in the overall best interests of the society in which they

live. Expense and value are viewed in the abstract, and economic theory seems to believe that it really is all about money.

Suppose some stunning medical breakthrough pointed the way to the eradication of all disease, providing all people with the ability to live healthy lives to ages in the 140s. But additional research and development were required to make this improvement a reality, and that would consume all the economic wealth in the world—which, by the way, is one hundred billion gazillion dollars. And let's further assume that the economic theory we've been discussing calculates that the gain in human life value for this expenditure would be only a gazillion dollars.

Well, our representative, acting in the best economic interests of the public, recognizes the dollar waste in funding further research. In good economic fashion, he vetoes any future development of the medical breakthrough. When this secret decision is leaked, the public outcry is filled with demands for explanations. The nature of the complaint is not that the public was misread and that people would be willing to pay for this medical breakthrough. The complaint is an entitlement argument that people have an inalienable right to life and the pursuit of happiness, regardless of the cost. When it comes to life and health decisions, utility is not price sensitive. It's demanded at any cost.

Or maybe, more precisely, health and life improvement is demanded regardless of the cost. The pursuit of happiness, an inalienable right, is, in part, the pursuit of the utility provided by good health and a long, healthy life. The objections to the cost associated with this pursuit aren't so much a willingness-to-pay issue as they are an ability-to-pay issue. Economics does encompass ability to pay as a criterion for structuring taxes and, implicitly, in consumer spending. One might want a lot but can get only what one can afford. But consumer spending for health care is different from consumer spending for a yacht. We can live without a yacht.

In the end, a good economic theory can be used to justify, through cost-benefit-type analyses, the level of health care made available in a society as a whole. But even good economic theory can't provide a rational and acceptable argument for how individual members of such a society should expect to benefit from the health care generally available. Sorry. We're still in the box on that one. ●

TOM BAKOS is a consulting actuary in Ridgway, Colo. He can be reached at Tbakos@BakosEnterprises.com.