Exhibit 11

Risks And Benefits In Health Care: The View From Economics

A model that gives clarity to the discussion of risk and can be relevant to designing social institutions to deal with risk.

by Mark V. Pauly

ABSTRACT: This paper discusses the meaning of the term *risk* from the economic perspective. It argues that some consumer decisions about insurance and the use of medical care are consistent with the economic model, but many are not. When decisions are inconsistent, real-world democratic governments' ability to intervene is limited by politicians' desire to please voters. The choice of incomplete insurance coverage in private markets is often said to present a case for governmental intervention, but the choice of insurance design in the Medicare drug benefit shows that the political process also may fail to select insurance that is optimal from an economic viewpoint. [*Health Affairs* 26, no. 3 (2007): 653–662; 10.1377/hlthaff.26.3.653]

The language we have come to use in describing health, health care, and medical spending sometimes gets in the way of clear thinking and sometimes reflects (without being explicit) quite different ways of thinking. There is no better example of this than the use of the term *risk.* Policy discussions talk about "trade-offs between the risks and benefits" of medical interventions, "pooling risk" through insurance, and "analyzing risk" in clinical decisions, often as if there were consensus on the meaning of the term *risk* but in reality using that single short word in a variety of ways.

It is probably fair to say that the most rigorous and careful analysis of risk in its multiplicity of meanings has been the province of economics and the associated disciplines of decision analysis and actuarial science. But policymakers and voters do not necessarily use the term in the same way as the experts do. Some of the difference is attributable to confusion or imprecision, which (in my judgment) has adversely affected policy making as people argue in different languages. Because economic models and theories cannot capture all that is relevant to human preferences and behavior, some of the differences reflect substantive contrasts in how choices and policies are viewed—psychologically, subjectively, but validly.

Mark Pauly (pauly@wharton.upenn.edu) is the Bendheim Professor, Health Care Systems, at the Wharton School, University of Pennsylvania, in Philadelphia.

DOI 10.1377/hlthaff.26.3.653

In this paper I set out the economic views of risk, and I discuss where there is confusion and where we are uncertain—both because the analysis is complex and because a definitive model of actual or ideal behavior has yet to emerge. One obvious but important conclusion is that if people think and behave in ways different from the economic model, the outcome will, to some extent, be suboptimal from an economic perspective. What that would mean for practical public policy, however, is far from clear. I emphasize the approach most common in economics but most emphatically do not argue for its universal applicability, and I mention some other views.

The most fundamental characteristic of situations involving "risk" or "risks," in economics as in the dictionary, is the absence of certainty: when what will happen is not known beforehand for sure. In the achievement of health and the use of health care, risk is ubiquitous. I do not know what my health state will be in the future—what accident or illness might occur. If I am sick, I do not know for certain the outcome of treatments or medications; I hope for the best but realize that adverse side effects can happen. And from a financial perspective, I do not know what my medical bills will be.

Risks about health state, about the health outcomes of treatment, and about health care cost all usually bother middle-class people. It is not just that they do not like bad outcomes; rather, they do not like not knowing what the outcome will be. In a general sense, people prefer a surer thing to a "maybe"; they are "risk-averse." Uncertainty about health states and medical bills motivates the voluntary purchase of insurance that delivers money when bad events occur and takes money (in the form of premiums) when they do not. The risk of adverse treatment effects (which might also trigger medical bills) prompts demand for information about the trade-offs between these effects and benefits. Either situation also prompts a demand for regulation to ensure that insurance works properly and that trade-offs are appropriate.

There are three aspects of medical decisions under risk that are often misperceived: trade-offs in choices among procedures or medicines, opportunities for the voluntary purchase of market insurance, and proper decision making about both care and coverage.[1] In this paper I define *risk* or *risks* as necessarily involving uncertainty.

The Economic Model Of Actual And Ideal Decision Making

How do people generally make decisions in risky situations? How should they do so? The model that economists and decision theorists have traditionally used to answer these questions is the "expected utility" (EU) model.[2] It is important to recognize that like all theoretical models, this is at best a useful caricature describing general tendencies; it does not imply that literally everybody, every day, in every case uses this approach, that people even recognize that they think this way, or that it captures all of the influences on decisions. Its value is really comparative:

Does it do a better job of explaining or evaluating decisions in general than some other model, including the null hypothesis that people make choices at random?

There is considerable controversy about the EU hypothesis. It should not be regarded as the perfect truth, or the universal truth, but it is a place to start. Moreover, even when actual decisions are found not to fit this model, the judgment of whether that behavior is appropriate from a policy perspective is almost always made using this model as the normative benchmark.

Consider two different kinds of risky prospects. One has health status as the outcome (for example, future life-years); the other, money in the form of income minus health spending. We imagine that a typical person has a utility function that evaluates well-being depending on what actually happens: The outcome or "state of the world" with more life-years or more money left over after health expenses has more utility. The person also in some fashion attaches probabilities to each state; these estimates might come from statistics (50/50 for flipping a fair coin) or might just be a subjective guess.

The expected utility from a prospect of two outcome states is a weighted average of the utility in the two states, where the weights are the probabilities. So, for example, if one treatment would give a moderate level of life expectancy for sure while the other has the possibility of either complete cure or an unintended but serious adverse outcome (such as death), the person would decide what to do by comparing the sure utility from the first choice with the expected utility from the second and would select whichever is higher.

One other assumption that is usually made is that of diminishing marginal utility: The gain in utility from more life expectancy or more money to spend on things other than medical care is always positive but is larger if health is low or medical expenses are high. Given this assumption, we can conclude that people will prefer the certain prospect to the one with an equal expected or average value. For example, if under one treatment the chance of death (zero additional life-years) is 0.1 and that of normal life expectancy is 0.9, the person will prefer an alternative treatment that offers 90 percent of normal life expectancy with certainty. Uncertainty about treatment outcomes also translates into uncertainty about financial consequences. If the risky treatment has the higher expected utility, the person will prefer to fully insure the cost if the premium is 0.9 times the (probably low) cost of successful treatment plus 0.1 times the (probably very high) cost of treatment in case of an adverse outcome. How much more than this "average health" or "average cost" the person might be willing to experience and still prefer the same thing depends on how risk-averse the person is. There are differences on this preference parameter across people; some people are more willing to take chances for health or money than others are.

What does this model, highly simplified, tell us about choices and possibly about policy? As a general case, I discuss what are called "extreme events": low-probability, high-consequence events. Such an outcome might be the rare but se-

vere life-threatening side effect of a drug or device, or it might be the equally rare but spectacularly expensive treatment for an uncommon or uncommonly severe illness. Catastrophic treatment failures and catastrophic medical bills are things for which, in principle, correct decisions about risk are important, but they also are difficult for people to make and for policymakers to understand and control.

A Medical Treatment With Rare Side Effects

If the EU model is applied to a new treatment that has the potential for major improvement in life expectancy but also the chance of a fatal side effect, the person should choose the innovation, despite the chance of the catastrophic outcome, if the utility for the good outcome offsets the (negative) utility of the bad outcome by enough to make that choice better than the "moderately good" sure thing. But implementing this advice is difficult; knowing and understanding the probability of the bad effect and, even more, "appreciating" it can be a challenge. Even when probabilities are known statistically, people have problems understanding small probabilities and often set a near-zero probability to zero, thus ignoring it. Or sometimes they fixate on the adverse consequence and ignore the fact that it almost surely will not happen. The problem is even more severe when all anyone knows is that the bad outcome, although unlikely, could happen, but with what probability is anyone's guess. To the point, however, the decisionmaker will have to guess about some "subjective probability." The decisionmaker may, in a deviation from the EU model, attach an extra amount of disutility just because the probability is ambiguous.

It is obvious that there are values of benefits and probabilities that will make the innovative treatment the preferred one. Policy issues then are of two contrasting types: there may be concerns if the adverse outcome actually happens, but there may also be concerns that fixation on the adverse outcome causes people to choose the "safer" alternative, despite the very much greater average benefits from the riskier one.

In the first instance, it will be human nature, after the fact, to look for a cause on which to blame the bad outcome, even if its possibility and uncertainty were well known in advance. Of course, if such an explanation opens up a strategy to reduce the probability in a way that has more benefit than cost, this reaction will be salutary. But if not, it will be an understandable but undesirable waste of time. In the other case, excessive caution may lead to worse population health outcomes.

The policy goal here is the proper balancing of benefits and (true) risks. The approval process followed by the Food and Drug Administration (FDA) attempts to do this based on scientific evidence and scientific principles. But, according to the economic model, the former is necessarily incomplete, and the latter do not exist. The first part has some substance: Good information is better; better information (for example, from larger or longer trials) is best. The scientific way to compare risks and benefits is, to be charitable, quite unclear: How are risks and benefits to

be balanced, given the dependence of this balance, in the EU model, on the estimated probability of adverse outcomes, the dependence of the ideal choice on each person's utility function, and the likely interaction of probabilistic side effects and type of insurance? The utility function's form influences both the person's valuation of the adverse outcome and the degree of risk-aversion, neither of which are obvious subjects of scientific expertise at the FDA.

Criticisms of actual public policy or private decision patterns in this area can be interpreted as coming from the EU model. Some criticisms reflect the fact that citizens have different utilities for benefits, adverse outcomes, and the risk of both; there can be no judgment that is scientifically correct for all. Sellers of products, whether drug firms or health professionals, would be expected to take an optimistic view of net benefits of what they sell or do. Some consumer representatives think that the FDA has not been sufficiently appreciative of the possibility of adverse outcomes (a too-low subjective probability or a too-low valuation of the consequences), while others (primarily patient advocacy groups) think that it attaches excessively high values to adverse effects relative to the utility from positive health benefits made available as quickly as possible to those who have few alternatives. "Sufficiently" and "excessive" are the nonscientific, preference-related words here (even more elusive of meaning than "risk"). In the EU model, decisions cannot be based only on clinical or medical knowledge but depend in part on patients' preferences, about which the FDA has no great expertise. The real question is which (or whose) preferences should dominate.

Virtually any product will be too risky for the most risk-averse people, but any delay will be too long for many more willing to take a chance for a better average outcome. Some suggest that the regulatory process tends to overvalue avoiding adverse effects relative to delaying or preventing the emergence of beneficial outcomes.[3] The tendency to be politically cautious and to avoid the recriminations from bad outcomes (even ones based on a gamble with good prior odds) could further widen the deviation from the EU model.

But beyond predicting these general and unavoidable criticisms in a country in which people have different values, the EU model can do little more than offer a framework for classifying the things on which people may differ. It might suggest more explicit specification of that framework and more precise quantification of the parameters (of value and risk aversion) as a way of improving transparency and consistency.

Buying Insurance

The other application of EU theory is in choice of insurance coverage. A fundamental implication of the theory is that considering potential insurance coverage of two risky prospects with the same expected value, the person should attach more value to coverage of the lower-probability, higher-loss event. The intuitive reason: Sacrificing the premium to transfer dollars to the state in which the no-

insurance wealth would have been lower is to be preferred because then dollars will have higher marginal utility—"mean more to the unlucky person"—in that state than in the other. If we characterize these extreme events as those associated with catastrophic levels of expense, and if insurance comes with a premium that is close to the expected value, we can conclude that the person should prefer catastrophic coverage above all.

Do consumers seem to follow this model and its advice? To some extent they do. Inpatient hospitalization is a costly event, but it is a low-probability or rare event for "average risks" at any age. The proportion of such expense people have to pay out of pocket in the United States, even with a sizable uninsured minority, is only about 3 percent. By way of contrast, expense for dental care is lower on average but more likely (and also subject to some moral hazard). Here people have made choices that result in a much larger 44 percent paid out of pocket. Moreover, the great bulk of private insurance plans have good coverage of catastrophic costs, with upper limits in the hundreds of thousands or even millions of dollars.

In short, despite the complexity, confusion, and lack of confidence that are intrinsic to insurance choices, the pattern of coverage Americans choose seems quite consistent with the EU model. But there are different views than those based on this model, and those views sometimes lead to different perceptions and different choices. A good example, and one relevant to the specific policy question I discuss below, relates to efforts to design insurance coverage with lower premiums for lower-middle-income workers. The HR Policy Association, comprising benefits experts at major corporations, has designed some basic health insurance policies that can be offered at moderate premiums to currently uncovered workers (often new, temporary, or part-time workers). The hope is that these plans will be more attractive than being uninsured and will not require large employer contributions.

Focus groups were convened to provide information about what kinds of plans would be popular offerings. Consumers were asked, among a variety of plans with given premiums, what they most preferred. Catastrophic coverage (full coverage above a deductible) was not always the most preferred option. Based on this advice, the National Health Access Program offers a catastrophic option (major medical) but also other options that, after a modest deductible, cover the more likely ambulatory care services, while leaving the less likely (but much more expensive) hospital costs uncovered or moderately covered.[4] One rationale for a preference for the latter option, as suggested by research, is that people look at insurance as an investment, with the premium as the initial investment and benefits as the return.[5] In this view, a good investment would be one that a person would probably collect on (get a return from), whereas catastrophic coverage is not a good investment because it would rarely pay off.

This perspective is not consistent with the EU model, which actually views the best outcome as one where the person paid premiums but then was lucky enough almost never to get sick and so almost never to collect benefits, while still having

good protection should illness occur. More to the point, the expected utility under simple catastrophic coverage would be higher than that under the other kind of policy. It would therefore be common among analysts to suggest to policymakers that they attempt to correct these mistaken preferences, at least through guidance if not through regulation, and direct consumers toward the catastrophic plans that are better for them.

Example: Medicare Part D

Although the EU model says that people would prefer to use a given premium for catastrophic coverage, real people and real politics appear to have combined to frustrate the application of this advice in the public sector as well. In particular, the stylized form for the new middle-class Medicare drug benefit is not full (catastrophic) coverage above a deductible but rather is so-called doughnut coverage, with a modest deductible, then 75 percent coverage over a range of expenses, then 100 percent cost sharing (the "doughnut hole"), and finally a return to virtually complete catastrophic coverage.

Compared to traditional catastrophic coverage with a deductible, the politician-designers of Part D reduced coverage for people with high expenses (where, in theory, people would have gotten the most utility value from coverage) to offer rather generous coverage for people who happened to have low expenses (where, in theory, coverage should be less valuable), to provide most beneficiaries with a return on their premium. What was the attitude toward risk here, and what does it tell us about the relevance of the EU model? The simple answer to this question is that the comparative attractiveness of this design was the result of political catering to consumers' misperceptions.

The government did not play—or even try to play—the role of rational corrector of illogical consumer decisions. Rather, because of the need to please voters in a democracy, political leaders went along with and even endorsed a fundamentally incorrect view of insurance. This experience, in my mind, raises grave doubts about the general ability of democratically elected governments to intervene to prevent widespread mistakes about risk by consumers. The fundamental problem is that as long as consumers still think they are right, when they vote they will favor politicians who take positions that cater to their mistaken judgments. Government can limit things that most citizens agree are mistaken, but it cannot so easily limit things that most citizens are mistaken about.

Is there any way to explain the Part D design? One response is to surrender and accept citizens' preferences as given and legitimate: While the doughnut design does not fit the EU model, perhaps voters genuinely have different preferences about insurance—preferences that, as adults, they have a right to hold and to support politicians who hold them, too.

This is a logically consistent argument, but it is destructive to traditional policy analysis. It weakens the basis for judging large-scale choices of consumers as mis-

taken. If people can validly prefer insurance that exposes them to greater financial vulnerability, why not regard their other decisions, such as avoiding flu shots, as also fully rational reflections of special preferences involving some unusual likes and dislikes?

The argument that the Part D design is inconsistent with EU maximization is not completely bulletproof. In designing Part D, Congress set for itself a limit on total budgetary cost and the intent that beneficiary premiums should cover a quarter of this cost. If all Medicare beneficiaries had been of the same risk, the EU-maximizing design for the nonpoor would have been to buy catastrophic coverage with a deductible no larger than what the total premium (75 percent federal and 25 percent beneficiary) would buy.

But in reality, people are not of the same risk of outpatient drug expense. The existence of multiyear chronic illness treated with expensive drugs means that some people can confidently expect to get much more benefit than average from a given catastrophic policy, while others will expect much less. If the premium is the same for those high risks as for those who know themselves to be average or low risks, then low risks might decline catastrophic coverage even with a subsidy. One possibility—although not one advanced as part of the policy discussion—was that by offering some more up-front coverage that lower risks might expect to use, the plan might attract them.

An alternative and probably better way to prevent adverse selection would have been to provide catastrophic coverage for free. The budgeted subsidy would probably have been enough to offer a plan with a catastrophic threshold lower than that in the current plan. But perhaps the main rationale for the beneficiary premiums was itself political: to give the impression, despite the enormous subsidy, that the elderly were paying for their benefits and to maximize participation as evidence of support for the new benefit.

Whatever the precise political motivation, it is clear that the desire to offer insurance viewed as a profitable investment for many people shaped the design of this program. Finding the plan that maximized expected utility for the average beneficiary did not.

What Social Criteria?

Given the difficulty people have in understanding risk, much less dealing with it optimally, there might be a role for public policy. But it is far from clear how the objectives of that policy ought to be defined.

One possible definition is what we might call the "Man of the People" strategy. The policymaker should determine what strategy citizens would prefer—given their preferences for risk, their perception of the relevant probabilities, and their tastes and foibles—and then try to achieve that outcome. One problem with this approach is that citizens in unregulated markets could do this much as well as or better than government. The other problem is that these preferences might be

simply illogical, or wrong-headed, or harmful in the long run, and extending democracy to endorse reflection of citizens' mistakes might be hard to explain (much less to justify) as a policy goal.

Another perspective is frank paternalism, as advocated most recently by Richard Thaler and Cass Sunstein.[6] One problem here is that paternalists might well adopt a maximization of expected outcomes (for example, years of survival per capita) without regard to citizens' utilities and the values they imply for risk-aversion and the like.

The third approach is to force EU-maximizing outcomes on a reluctant public, based on the view that when outcomes are realized, this will maximize average utility after the fact (even if it might disappoint those unlucky enough to sustain no losses and therefore collect no benefits). A politician who is trusted might be able to pull this off, but it is a tall order.

Some Remedies

If citizens are confused by risk and have confused preferences that they try to foist on politicians, is policy condemned to an inefficient and unstable fate? Perhaps not. There might be strategies that even people with distorted or unusual preferences will support and yet represent a better outcome than simply doing what voters think they seem to want.

One example of such a strategy is the National Vaccine Injury Compensation Program. This program arose in part because different groups had very different views on the risks posed by pediatric vaccines. Some people believed that there was a high risk of side effects; others, that it was low. The strategy then was to levy what effectively was an excise tax on vaccine sales, use the proceeds to set up a trust fund, and pay claims for actual damages (on a no-fault basis). After an initial flurry of payments for previous injuries, the level of payments fell dramatically, and the trust-fund balance grew. Those who thought that adverse effects would be common turned out to be wrong. But the point is that those who thought side effects unlikely expected to get money back, and those who thought them likely felt that they were protected; both groups could agree with the proposal.

The second application is mutual insurance.[7] Suppose that consumers differ with actuaries or policymakers on their subjective judgment about the probability of loss. (If expected expenses vary across people, they must still all agree on who belongs in each expected expense category, but not on the expected loss in that category.) For example, many young people (especially males) now think that they will not need medical care, while actuaries see higher probabilities. These consumers will refuse to pay the actuarially based premiums; this group actually is the group most likely to decline to have insurance. But they might agree to contribute a payment to a young people's mutual insurance plan, in return for the promise that if they are right and are all low users, premiums will be refunded. Such a plan would agree to pay back "dividends" to policyholders if benefits pay-

outs are low; usually these payments are used to lower future premiums. Thus, young people who think that they rarely get sick should still be willing to join this plan as long as they know that the other insured people are like themselves and have the same risk as they—whatever that risk "really" is.

The EU economic model of private and public decision making under risk is far from perfect, either as a predictor of behavior or as a normative model. Yet it still seems to give some clarity in the discussion of risk and can be relevant to designing social institutions to deal with risk. Paradoxically, the more behavioral economics discovers cases where people act differently than the model would suggest, the more it might be called upon as a normative criterion for public policy. Even here there are limits in world of democratic public choice. Sometimes there are alternative designs that might not be as perfect as those from a perfect world of EU-maximizers but that can still be an improvement over muddling through with confused language and euphemistic slogans. The economic view might still help in dealing with risk in complex and complicated settings.

This paper was presented at "Confronting Dilemmas of Risk in Healthcare: Emerging Trends and Practical Approaches for Decision Makers," ECRI's Fourteenth Annual Conference on Healthcare Practice and Policy, 15 November 2006, in Plymouth Meeting, Pennsylvania.

NOTES

1. One source of error is in the language used to describe choices of medical therapies: "the trade-off between benefits and risks." The confusion arises because sometimes the word *risk* is really used as a synonym for adverse or undesired effects. Suppose a drug I might take for my allergies always and unequivocally in all patients does two things: It alleviates the symptoms of allergy, and it causes an unsightly rash. There is no uncertainty, no variation. The appropriate decision might still be ambiguous since it involves a trade-off between alleviated symptoms (for sure) and a red rash (for sure), and different people might value these outcomes differently. But there is no issue of risk.
2. C.E. Phelps, "Health Insurance," chap. 10 in *Health Economics*, 3d ed. (Boston: Addison Wesley, 2003), 318–365.
3. See S. Peltzman, "An Evaluation of Consumer Protection Legislation: The 1962 Drug Amendments," *Journal of Political Economy* 81, no. 5 (1973): 1049–1091, for a classic early statement of this view; see C.R. Sunstein, *Laws of Fear: Beyond the Precautionary Principle* (Cambridge, U.K., and New York: Cambridge University Press, 2005), for a more recent version.
4. National Health Access, "Limited Health Benefits" (table), http://www.nationalhealthaccess.com/ParticLM.aspx (accessed 9 February 2007).
5. H. Kunreuther, "Risk Analysis and Risk Management in an Uncertain World," *Risk Analysis* 22, no. 4 (2002): 655–664.
6. R.H. Thaler and C.R. Sunstein, "Libertarian Paternalism," *American Economic Review* 93, no. 2 (2003): 175–179.
7. M. Pauly, H. Kunreuther, and J. Vaupel, "Public Protection against Misperceived Risks: Insights from Positive Political Economy," *Public Choice* 43, no. 1 (1984): 45–64.