The Fifth Annual Arrow Lecture in honor of Kenneth J. Arrow

Moral Hazard in Health Insurance: Developments since Arrow (1963)

Table of Contents

Introduction to the Fifth Lecture

JOSEPH STIGLITZ

The Arrow lecture series is given in honor one of Columbia's most distinguished graduates, Ken Arrow. We were enthusiastic about having the lecture series, because we decided to have one lecture every year featuring what happened to a contribution that he had made 30, 40, or 50 years ago. The lecture has covered such a span of topics.

The first lecture followed up on his paper on learning by doing. The second one followed up on his book on social choice and individual values. The third was on financial markets, and the fourth was on work on global warming.

The fifth, on which this book is closely based, picks up on basic work that Ken did in 1963, where he developed concepts of moral hazard in the context of health insurance. He solved the problem then, but unfortunately the country did not fully take the analysis on board, and here we are many years later still talking about health insurance.

We were really pleased that Amy Finkelstein was willing to give the fifth annual lecture, because it is a topic that is of enormous current interest and a topic in which there has actually been a lot of good economic research. Her talk and the ensuing discussions were able to bring together some of these policy issues with the basic economic research.

Amy Finkelstein is Professor of Economics in MIT, Co-Director of the Public Economics Program at the National Bureau of Economic Research and Co-Editor of the Journal of Public Economics, a journal which I help found many years ago. She researches public finance and health economics, primarily focusing on market failures and government intervention in insurance markets and the impact of public policy under the healthcare sector.

Following Amy's presentation, Jonathan Gruber talked about her work. He is also a Professor of Economics in MIT and Co-Director of the Healthcare Program at the NBER. He has been involved in health economics for a very long time, and has won a number of awards for his work. He has also served as Deputy Assistant Secretary for Economic Policy at the Treasury Department. He has also been described as one of the key architects of Obamacare, for his work on the Massachusetts Healthcare Reform, which was the prototype for Obamacare.

One of the wonderful aspects of the Arrow Lecture is that Ken always graciously agrees to comment on the talk, and that is always an exciting moment, because it both reflects how he thought about the problem many years ago as well as how thinking about these debates have evolved over these subsequent decades. His commentary is included in this volume, as is a

transcript of the short discussion that took place after the lecture.

I want to acknowledge Columbia University Press, the Program for Economic Research, and the Committee on Global Thought for putting together this program, and Laura Morrison, Sasha de Vogel and Robin Stephenson. I also want to especially thank Ken for honoring us again with his attendance, and Amy and Jon for sharing their work and thoughts.

Fifth Annual Arrow Lecture

Moral Hazard in Health Insurance: Developments since Arrow (1963)

AMY FINKELSTEIN

I was honored to give the Fifth Annual Arrow Lecture in April 2012. When I was first asked to give the lecture in honor of Ken Arrow and on the topic of the economics of healthcare, I thought that this would be simple. Ken wrote a seminal 1963 paper entitled "Uncertainty and the Welfare Economics of Medical Care." My plan was to talk about how this paper has shaped and influenced the subsequent field of health economics, a field which did not exist when he wrote the paper. It really is striking to re-read the paper because you quickly realize that virtually everything that has happened in health economics in the last half century can trace its origins back to that paper. But that's also when I realized that I had a real problem: I only had 50 minutes to talk about the wide-ranging impact of the paper…. and even I can't talk quite that fast!

So what I decided to do instead was just to focus on just one aspect of Arrow's original paper, which relates to a topic I have been researching for the past several years: the economics of moral hazard.

To put this topic in perspective, there are essentially two central facts of the U.S. healthcare system. The first is the lack of health insurance by approximately 46 million people, or about 15% of the population. The second is that healthcare spending is a large and growing share of our economy. In 1960, only 5% of GDP was spent on healthcare; in 2012, that share was 17%. That has real implications for the public sector, because in 2012spending on healthcare was 20% of the federal budget, and a substantial share of State budgets. Moreover, the major driver of projected federal spending growth is from public health insurance, specifically Medicare and Medicaid (CBO, 2012).

Not surprisingly, the recent healthcare reform—the Affordable Care Act of 2010—tried to adjust both these issues. It focused on trying to cover the uninsured and to rein in the growth of healthcare spending. But there is a fundamental tension in these goals, in that by itself accomplishing the first goal exacerbates the second problem. All else equal, insuring the uninsured is likely to increase healthcare spending. Why is that? It has to do with the moral hazard effects of health insurance, which is the subject of this essay.

To my knowledge, Arrow (1963) was the first academic reference to moral hazard and health insurance. Arrow defined moral hazard in health insurance as the concept that "medical

insurance increases the demand for medical care" (p.961). At the end of the introduction in that article, he lays down a gauntlet for subsequent researchers. He says, and I really like this quotation:

> "The discussion is not designed to definitive, but provocative. In particular, I have been chary about drawing policy inferences; to a considerable extent, they depend on future research, for which the present paper is intended to provide a framework." (page 948).

And indeed it did so. What I want to explore here is how we have risen to this challenge that he laid down in this article now almost 50 years ago. In particular, I will talk about two separate empirical challenges posed by his discussion of moral hazard. The first is an existence question: is the idea of moral hazard, which is an interesting theory, empirically relevant? Is the demand for medical care really price-sensitive? And the second relates to the challenge he poses about drawing policy inferences: how to estimate the likely impact of alternative health insurance policies or contracts are both the level and the growth of healthcare spending?

Is demand for medical care really price sensitive?

What do we mean when we say moral hazard and health insurance? Arrow defined it as medical insurance increasing the demand for medical care. Now why would that happen?

There are essentially two ideas in the literature. The first, due to Ehrlich and Becker (1972), has come to be known as what's called *ex ante* moral hazard. This is the idea that if I have health insurance and it will pay my medical bills when I get sick, then I have less incentive to invest in maintaining my health, because when I get sick, the financial consequences will now be borne by someone else; therefore, I am going to eat, drink and be merry. I might smoke more, drink more, exercise less, et cetera.

The second idea, which was formalized by Mark Pauly (1968) has come to be known as *ex post* moral hazard. For this idea, let's forget that health insurance may affect my investments in my health; let's just take my health as given. Ex post moral hazard is the idea that, at a given level of health, I am going to choose to consume more medical care because the price of that medical care is lower. This is basically about a demand curve and the price sensitivity of demand for medical care.

For the most part, the literature has tended to adopt the second definition. Moral hazard and health insurance has come to mean price sensitivity of demand for medical care, rather than an impact of health insurance on investment in one's health. For the remainder of this essay, I'll focus on the price sensitivity of demand.

The first question is whether demand for medical care is in fact sensitive to the price of

care. On the one hand, this certainly seems like a natural property of any demand curve: if you lower the price, people buy more of it. There is an alternative view, however, that medical care is not determined by price, but by needs. One summary of this is what I will call the rhetorical case against the notion of moral hazard and health insurance, and it was put forward in a relatively recent New Yorker article called "The Moral Hazard Myth" by Malcolm Gladwell in 2005:

> "The moral-hazard argument makes sense… only if we consume health care in the same way that we consume other consumer goods, and to economists like [John] Nyman this assumption is plainly absurd. We go to the doctor grudgingly, only because we're sick. "Moral hazard is overblown," the Princeton economist Uwe Reinhardt says. "You always hear that the demand for health care is unlimited. This is just not true…. Do people really like to go to the doctor? Do they check into the hospital instead of playing golf?" (Gladwell, 2005)

If that is the rhetorical case, what exactly is the evidence? Ultimately it is an empirical question of whether when you give people health insurance, they consume more medical care and thus increase healthcare spending. Empirically, it is somewhat challenging to get at this. You do not want to just compare people with and without insurance and look at whether people with insurance spend more than people without insurance, because people with insurance are different than those without insurance. They have different income, different employment status, and different health. They are different in ways that are likely correlated with the demand for healthcare; basic adverse selection theory suggests that individuals who are less healthy will be more likely to purchase health insurance (Akerlof 1970, Rothschild and Stiglitz 1976).

In fact if you just do these cross-tabulations in the data, you will quickly discover what looks like moral hazard: people with health insurance spend more on healthcare than people without it. You'll also quickly discover that it looks like health insurance kills people, because people with health insurance have higher mortality than people without health insurance. So that might give you some pause in what you can infer about the impact of health insurance from these types of observational comparisons.

The ideal solution to test the null hypothesis that Malcolm Gladwell and others have put forward that there is no moral hazard or no price sensitivity of the demand for medical care would be a randomized control trial in which you randomly assign different insurance across individuals. Then you do not have the selection problem and this allows you to actually infer what the effect is of giving insurance to one group and not another.

The Oregon Medicaid Experiment

Remarkably there have been two randomized control trials in the United States on health insurance. (You can think that that's remarkable either that there have been "only two" randomized evaluations of such an important topic, or that it's remarkable that we have had any randomized trials).

One of the experiments was the RAND Health Insurance Experiment from the 1970s. The other, is the Oregon Health Insurance Experiment which I have been leading with Kate Baicker at the Harvard School of Public Health. It is based on a Medicaid lottery conducted in Oregon in 2008 and our analysis of it is ongoing. We of course have more detail available for those who are interested (see e.g. Finkelstein et al. (2012) and also our study website: www.nber.org/oregon).

Let me describe the Oregon experiment, and then you can judge the evidence for yourself. Medicaid is public health insurance for the indigent. In Oregon, they have an expansion program to cover people who are financially, but not categorically eligible for Medicaid. Essentially, these are low income people, below 100% of the federal poverty line. This is really poor—less than $10,000 for a single person in annual income. They are uninsured but they are able-bodied, so they're not getting on Medicaid because they are on welfare or they're disabled, and they are not otherwise legible for Medicaid.

Oregon had money to cover some of the low income, uninsured adults in this group, but not all of those that are eligible. They only had money to cover 10,000 people, so they had to think about how best to do this. They decided it was not fair to award their limited slots on a first come first served basis, because that would privilege people who are better connected.

Federal law actually prohibits you from awarding Medicaid on the basis of health status, so they decided the only fair and equitable thing to do was run a lottery. They ran a big public relations campaign, they asked interested individuals to sign up. About 90,000 low income adults signed up, and they randomly drew about 30,000 names to be eligible to apply for Medicaid.

We have results now from the first year of this experiment, which is ongoing. This is just one of many pieces of data, to give you a sense. Figure 1 looks at the probability of hospital admission over about a 16-month period after the lottery. The dark blue shows you the probability of hospital admission for the controls, those who did not win the lottery. (We try to avoid referring to them as "losers"). They have about a 7% admission rate over a 16 month period. And the light blue shows the implied increase in the probability of going to the hospital due to Medicaid. You can see that it's about two percentage points or 30% higher.

Figure 1: Impact of Medicaid on Hospital Admissions, Evidence from the Oregon Health Insurance Experiment

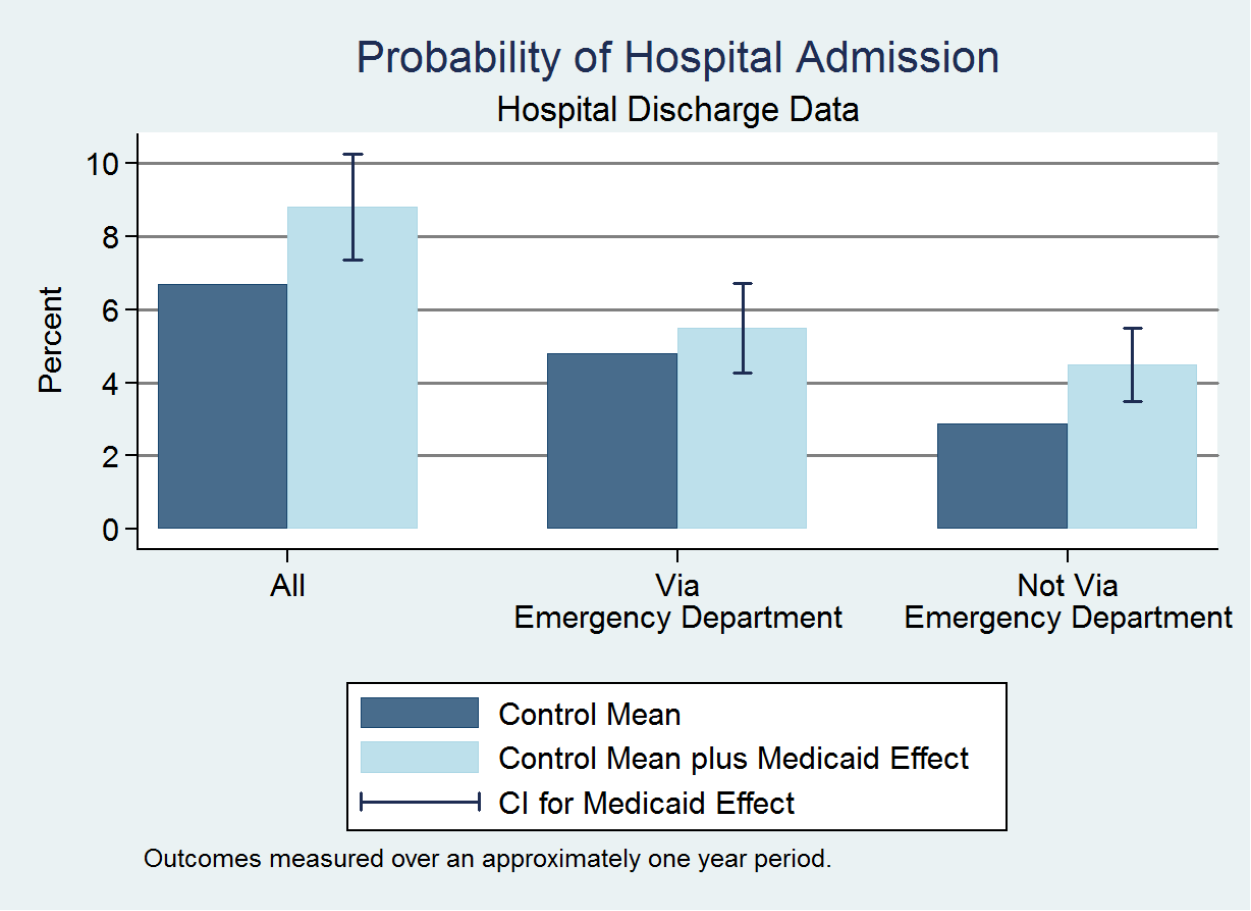

Source: Finkelstein et al. (2012)

That is the empirical evidence from a randomized control trial that belies this notion that demand for medical care is not price sensitive. That because no one wants to go to the hospital, so they're not going to do it even if it's less expensive.

There is an interesting effect if you look at the other two bars, where we broke up hospital admissions into those that come through the emergency department and those that do not. Perhaps not surprisingly, what you see is that almost all of the increase in hospital admissions is admissions coming not through the emergency department; these are presumably ones that are less serious or less time sensitive, and therefore, perhaps not surprisingly, more price sensitive. It is a fairly intuitive result.

There is a whole bunch of other evidence we have from just this first year. We see that Medicaid not only increases hospital admissions, as I just showed you, but also the probability of taking of prescription drugs and of going to the doctor. If you figure out the cost of all of these things, we estimate that Medicaid increases annual medical spending by about 25%. People on Medicaid, when they are randomly assigned to it, are spending 25% more, which is about $750 a year for this population, than those who are uninsured. I think this is fairly compelling and definitive evidence rejecting the null of no spending effective health insurance, or no moral hazard.

As an aside, we do not find, at least in its first year, any evidence of *ex ante* moral hazard, where someone might say, now that I have health insurance I don't have to take as good care of myself. For example, we did not find any change in smoking behavior. Now, this may be because the theoretical possibility that when I have health insurance I'm less concerned about my health is not actually operative, or it might be operative and counter-balanced by the fact that people with health insurance go to the doctor more and the doctor is urging you not to smoke. That is something that remains to be sorted out, but it is consistent with the intuition that there is not much of this *ex ante* effect.

The RAND Health Insurance Experiment

I also want to briefly mention the other health insurance experiment, the famous RAND Health Insurance Experiment from the 1970s. Here they randomly assigned about 6,000 individuals to plans with different levels of consumer cost-sharing. Unlike in the Oregon experiment, where we compared the impact of public insurance relative to having no insurance, here everyone in the experiment is getting private insurance; but the private insurance in different experimental arms differs in terms of how much consumer cost-sharing it requires. The RAND experiment was led by Joseph Newhouse, actually one of our collaborators on the Oregon experiment, and it summarized very nicely in his 1993 book, *Free for All* (Newhouse et al., 1993). I also wrote a shorter summary piece on the RAND with some collaborators that is designed to be accessible to advanced undergraduates (Aron-Dine et al., 2013).

In the RAND experiment, families were randomly assigned to private insurance plans with different amounts of consumer cost-sharing. The most comprehensive coverage was called the free care plan. There was zero cost-sharing, so zero what the consumer had to pay out-of-pocket, and then there were a bunch of positive consumer cost-sharing plans where the consumer had to pay something out-of-pocket, depending on what you were assigned, maybe you were to pay 25% of your cost out-of-pocket, or 50%, or 95%.

One thing I want to note, and I'll come back to this later, is that all of the plans had low out-of-pocket maximums or stop losses, so that even if you had to pay 95% out-of-pocket, you had to pay that only up to some maximum out-of-pocket amount which was about $1,000 in 1970 dollars, after which you then had insurance cover everything. Even the people in the least generous plan still had catastrophic coverage. And in fact about a third of families were hitting the stop loss, so it was quite binding.

Once again, there was compelling evidence of moral hazard—that spending is lower when consumers cost sharing is higher, and when consumers have to pay a higher share of the cost out of the pocket, they spend less. To give just one number: if you look at people assigned to the plan with 95% consumer cost sharing - so they have to pay 95% of their medical costs up to the stop loss - you find that their annual medical spending is almost two-fifths less than the annual spending for those assigned to the free care plan. Again, this is compelling evidence against this null hypothesis of no moral hazard effects.

How will alternative health insurance policies affect health care spending?

Implications

So what does this mean? Thus far what I have hopefully convinced you of is that contrary to what uninformed readers of the *New Yorker* might believe, moral hazard and health insurance is not a myth. There is compelling empirical evidence from randomized controlled trials that when you give people health insurance, they consume more medical care, and when you give them more comprehensive insurance, they consume even more medical care.

But key challenges remain for policy. This goes back to what Arrow talked about in his original article, which is what are the policy implications of moral hazard in health insurance? In other words, one question that comes up is how do you translate these experimental treatment effects – such as the differences in spending across plans that we see in the RAND Experiment – into economic objects of interest? In particular we are always very interested in how you can take these estimates from the RAND Health Insurance Experiment of what people spent in the specific plans they were randomly assigned to, and use those to forecast out of sample the likely effects of different plans or different policies that we might be considering today.

So I want to devote the second half of this essay to talking about how to translate the economics of moral hazard into policy. I want to discuss several key conceptual challenges that immediately confront you if you try to predict the spending effects of alternative health insurance policies or health insurance regimes, some of which I have been working on for the last few years.

It's often helpful just to have a specific example to ground discussion and focus ideas, so I am going to talk about how to think about the impact on spending of high deductible health insurance plans. High deductible insurance plans were encouraged by the 2003 Health Savings Accounts Act which encouraged people through tax subsidies to buy plans that are very high-deductible. During the deductible, you pay 100% out-of-pocket. Qualifying plans often have deductibles of say $3,000; but then you get catastrophic coverage for expenditures beyond that deductible. You are insured against really large expenditures, but you first have to pay that amount out-of-pocket yourself.

Again going back to Arrow's paper, perhaps not coincidentally, this design is actually the theoretically optimal design of a contract when you have risk-averse individuals and concerns about moral hazard. In fact, a goal of this legislation was to get people into high deductible plans to reduce the level and growth of health spending, while still providing catastrophic coverage that is so valuable to risk-averse individuals.

There are at least three questions that come up when you want to think about how a high deductible health insurance plan is going to affect health spending.

I. Which price matters to the consumer? Is it the deductible that they face at the start of the year or their anticipation that by the end of the year they may have spent past the deductible? Which are they going to respond to in their spending behavior?

II. Who is going to select the high-deductible plans, and how is that going to affect the impact of these plans on spending?

III. Why would health insurance affect spending growth, which is really the key problem, as opposed to just the level of health spending?

What price matters to consumers?

So thinking first about the first issue, what price matters to consumers? Typical plans - including those in the RAND Health Insurance Experiment and also in the real world, both then and now - are highly non-linear. Instead of facing a linear relationship between your total medical spending and your out-of-pocket costs, you more face something like what's shown by the red line in figure 2. Initially, you are in the deductible range and so it's about a 45 degree line there, because you are paying 100% out-of-pocket, so your out-of-pocket spending is rising one for one with your total spending.

Figure 2: Non-linear health insurance contracts: What price matters to consumers? (to come)

At some point you spend past your deductible, and into the coinsurance arm. Say you

face a 20% marginal price in the co-insurance arm, so your out-of-pocket spending is rising one-fifth for every dollar you spend. Then at some point you hit the stop loss—that's the horizontal line—so your out-of-pocket spending does not move with your total spending and medical care is now free.

At the beginning of the year when these plans reset and individuals are faced with a deductible, do they think that the marginal price of going to the doctor is the full price—the 100% that they have to pay out-of-pocket—or do they think about the fact that by the end of the year maybe they'll have had a bunch of illnesses anyway and be past that deductible? Which price do they respond to? That's obviously going to have implications for how much of a spending reduction you will get for introducing a high-deductible plan.

If people react as if the price is the spot or sticker price of 100%, they will reduce spending a lot more than if they forecast that by the end of the year they may face a much lower effective marginal price.

In particular, if individuals are myopic, they might respond to the introduction of a deductible as if their price has increased dramatically to 100%. On the other hand, if you have fully rational forward-looking and not liquidity-constrained individuals, they may recognize that the spot price should not affect their consumption decision, only the end-of-year price.

Imagine someone with $3,000 deductible, but who has a chronic condition that regularly costs over $10,000 a year. At the beginning of the year if they have a headache and they're considering whether to go to the doctor to get this unusual headache looked at, they should not treat the current doctor visit as if they bear the full cost out of pocket. Rather they should think, over the course of the year I'm naturally going to spend past the deductible anyway and face a marginal price much lower, perhaps zero, of going to the doctor. So the effective marginal price of the current visit is quite low. The question is, do they do that?

This is something I've tried to get at in a recent paper with Aviva Aron-Dine, a graduate student at MIT, and my collaborators, Liran Einav and Mark Cullen at Stanford (Aron-Dine et al., 2012), where we asked the question: how important is forward-looking behavior in moral hazard?

Like the question of whether moral hazard exists in health insurance, the question of how forward-looking people are in their medical consumption is ultimately an empirical question. In this paper, we take advantage of fact that with employer-provided health insurance, the annual deductible resets on January 1st, regardless of when the employee is hired. Whether you're hired in March or in October your deductible is going to reset in January.

As a result, if you imagine someone who is hired into a deductible plan in March and

someone else who is hired in October, they both face the same initial spot price of medical care of one. They have to pay fully for the first doctor visit, but they face very different expected end-of-year prices, because the guy hired in February or March has the whole year to have health shocks accumulate and push him past the deductible, while the guy hired in October or November - or December 31st in the extreme - is very unlikely to get past the deductible.

Our basic idea therefore is to use the variation across employees in hiring date to look at a group of people who face the same initial or spot price for medical care but different future prices; those hired later in the year face higher end-of-year expected prices ("future prices") because they have less time to spend past the deductible. We therefore compare the initial medical utilization of people hired into the same deductible plan at different months in the year. Of course, we have to take into account that people hired in February may be different than people hired in October and the onset of illness can be different in February than in October. We therefore look at patterns of initial medical use for those hired into a no deductible plan, which does not have this feature, as a control for seasonal variation in hires or medical spending.

What we find is a rejection the null of complete myopia, or the hypothesis that individuals respond only to the spot price. We find that individuals in a deductible plan hired earlier in the year, who face the same spot price but a lower expected end-of-year price than individuals in the deductible plan hired later in the year, initially use more medical care. In other words, they are more likely to go to the doctor in the first month or two than individuals hired later in the year.

That suggests that individuals are not completely myopic. They do not just respond to the sticker price, so there is some degree of forward-looking behavior, but how much and how important is it? To answer this, we write down and calibrate a simple model. The results suggest that individuals are not fully forward-looking, so they respond both to the spot price and to the future price. What are the implications? With the amount of forward-looking behavior we estimate, we find that the spending reduction of moving from a no deductible to a high deductible plan is a lot lower, about 25 to 50 percent, as you would predict if you felt that individuals were fully myopic.

The sign of that result is intuitive, and we don't need to estimate the model to know it. As I described, when individuals are fully myopic the deductible has the most bite, because they keep thinking as long as they are in the deductible range that's their effective price. So if individuals are not fully myopic, we know the spending reduction from introducing a high deductible plan will be lower. But the magnitude is also of interest and we would not have known it without this estimation.

I want to give you some context about why it is challenging going from the experiment to these policy implications. The RAND investigators in their original work were aware of all these issues and tried to translate the differences in spending across plans in their experiment into a prediction for how spending would respond in other plans. The result was their now-famous estimate of the price elasticity of demand of -0.2 which is widely used by the Congressional Budget Office and by many others. The RAND investigators computed this elasticity under the assumption that individuals were fully myopic. And if in fact that's not the case, the spending effects of some of these plans may be a lot lower than we thought.

Who Selects High Deductible Plans?

A second and related issue in trying to forecast the likely spending reduction from introducing a high deductible plan is: who is going to select it? In general we tend to think of selection and moral hazard as distinct issues. The adverse selection literature has focused on the fact that individuals may differ in their (privately known) underlying health, and that that is going to affect their demand for health insurance. The moral hazard literature has tended to ignore heterogeneity across individuals and just focus on average price sensitivity or the average slope of the demand curve.

But what if we start with the observation that people may differ not only in their health but also in their price sensitivity of demand? Then you have to think about who's going to select these high deductible plans: are they going to be the guys who are more or less price sensitive? That gets to this notion of what we call selection on moral hazard. This is the focus of another paper with my collaborators Liran Einav, Steve Ryan, Paul Schrimpf and Mark Cullen (Einav et al., 2013).

Here, we thought about a model in which there are three reasons that people might demand health insurance. One is the traditional adverse selection channel in which people have private information about their risk type (in this case, their health). Another is what we've called selection on moral hazard; the person knows that he has a high price sensitivity of demand, and thinks, if you make healthcare cheap I'm more likely to use a lot of it, so this affects his demand for a plan with low consumer cost sharing. And then the third is just risk aversion; more risk-averse people will always demand more insurance.

An analogy from our collaborator Mark Cullen helped us think about the difference between traditional selection and selection on moral hazard. Think for a moment of the context of all-you-can-eat restaurants.

In the context of an all-you-can-eat restaurant, traditional selection is that people with big appetites are more likely to go to all-you-can-eat restaurants. Selection on moral hazard is the idea that even if you have an average appetite, if you know that when food is free on the

margin - the marginal price of an additional entrée is 0 at an all-you-can-eat restaurant - you're going to consume a lot more than you usually do when things are priced a la carte. So people who tend to eat a lot more when the price of food is lower also find an all-you-can-eat restaurant appealing. Selection on moral hazard is thus selection on the slope, or on the price sensitivity of demand, rather than "traditional" selection on the intercept, or the level of demand.

We looked into this using data from Alcoa, where we had employee health insurance options, choices and medical claims. We found evidence of selection on moral hazard. The particular sign was that individuals who are higher moral hazard types—in other words, individuals who increased their medical consumption more when it was subsidized more —are more likely to choose more coverage. Thus not only do sicker individuals seek more medical coverage (traditional adverse selection), but also individuals with a higher price sensitivity of demand seek more coverage (selection on moral hazard). Quantitatively, in our particular setting, selection on moral hazard turned out to be quite important. It was almost as important as traditional adverse selection, at least in our particular context, in determining people's plan choices.

So why do we care? Again, let's remember the goal, which is to translate experimental estimates of the impact on spending from randomly assigning people to different plans to a forecast, for example, of what happens when we offer these tax-subsidized high deductible accounts. The key point is in most real world settings, including the health savings accounts one, people get to choose whether to be in a high deductible plan or not; we don't randomly assign people to the plan. The traditional approach to forecasting how health spending is going to respond to introducing high deductible plans is to estimate what percentage of the population will take up these plans given how they're priced, and then applying the average estimate of moral hazard that we get from experimental variation, such as the experimental variation in the RAND.

Remember that that experimental variation was key for being able to get the causal effect of one's health insurance plan on one's medical spending. But, here, the very feature that's solving your causal inference problem, namely random assignment, is shutting down the possibility of selection, which can be very important for the spending effects in a context where people can choose their plans. In particular, consider for example that you are trying to predict the spending reduction associated with offering a $3,000 per family deductible plan. If we do the traditional approach and assume that the people who are buying those plans are random with respect to their moral hazard type, we would use our average moral hazard estimate and predict a decrease of about $350 per employee spending, if we go from the current plan to a high deductible plan.

However, if you recognize based on our estimates that the people who are going to select the high deductible plan with less coverage are precisely the least price sensitive people—in other words, the low moral hazard types—the spending reduction you're going to get by these people moving into a high deductible plan is going to be much lower, because those who choose the high deductible plans are less responsive than average to consumer cost-sharing. Depending on how you've priced the plan, the utilization effect (or spending decrease) can be lower by a factor of two or three. This is going to have really important implications for how introducing a plan when you actually give people choice, as we usually do, is going to reduce health spending.

Spending Growth

The third and final thing I want to talk about is really, in some sense, the key issue in healthcare spending, which is not the level of healthcare spending, but the growth of healthcare spending. Thus far what I have discussed is the likely effect of introducing a high deductible plan on the level of spending. I have suggested that you may get much less of a spending reduction than you would have conventionally assumed, both because people are forward-looking in their medical consumption decisions, so they don't respond fully to just the current deductible, and because of the issue of selection on moral hazard, where people who choose the high deductible plan may be those who respond least to the deductible.

But the real story in the healthcare sector is not the level of spending so much, as it is growth, especially relative to GDP. Figure 3 shows the growth in health expenditure as a share of GDP. Health expenditure as a share of GDP has grown from about 5% in 1960 to about 18% in 2010.

Figure 3: National Health Expenditures as a Percentage of GDP, 1960-2010.

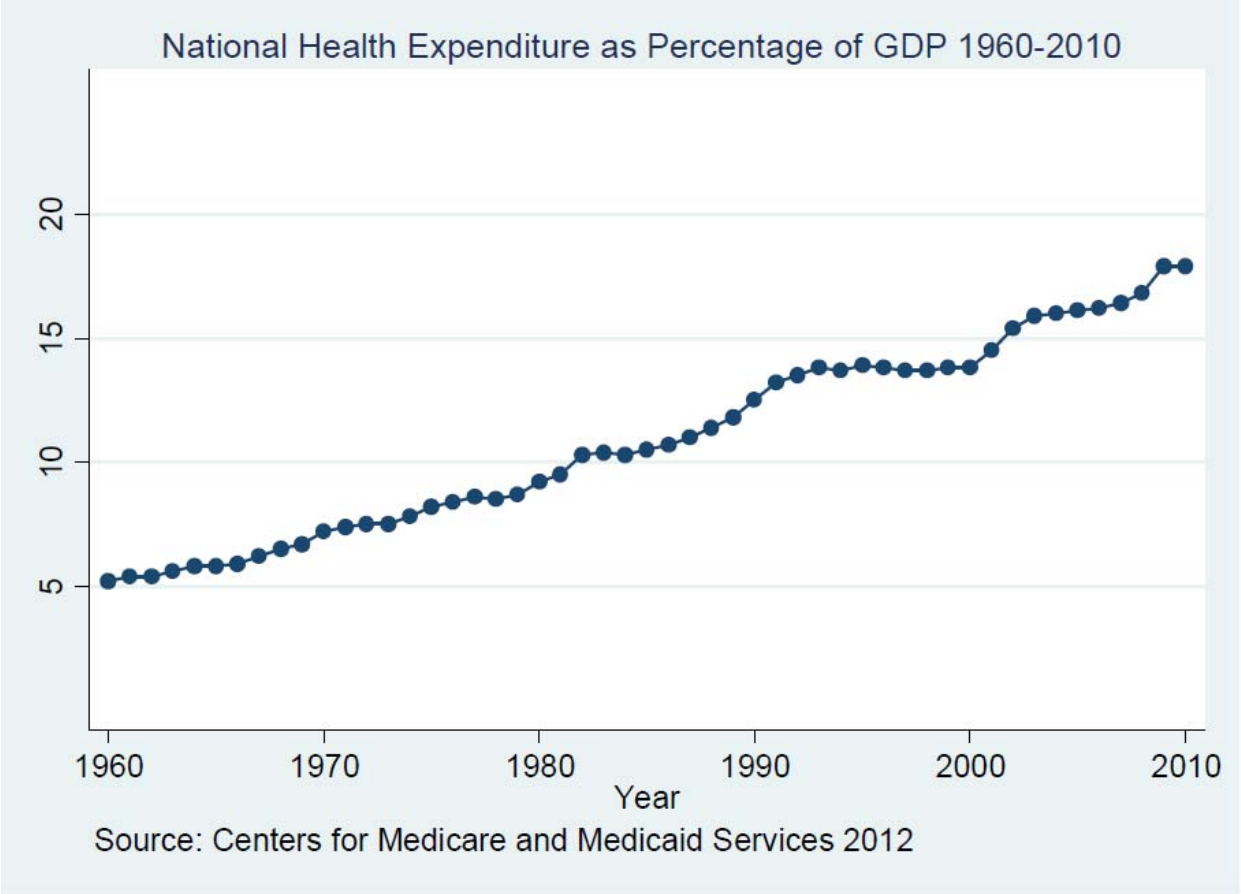

What is behind the growth of health spending? This is an area of rare consensus among economists. Joseph Newhouse and others have done work on this, looking at the roles of, for example, the aging of the population or other demographic factors (Newhouse 1992). While that all plays some role, there is widespread consensus, that technological change in medicine is the driving force behind the growth in health spending. But this just kicks the can down the road: what then drives technological change in medicine? And the particular question I'm interested in is: what role does health insurance plan in affecting technological change in medicine?

Another way of putting it is to note that, thus far, all of the analyses I have been describing have asked: how does my health insurance affect my medical spending, assuming that nothing else around me changes? The healthcare system is held constant in these analyses; this is what we call partial equilibrium analysis. In fact, however, major expansions of health insurance coverage incentives to develop and adopt new technologies, and therefore, have general equilibrium effects on the healthcare sector that are very different, perhaps bigger or perhaps smaller, than these partial equilibrium effects we've been talking about so far.

This is something I studied in an earlier paper looking at the introduction of Medicare in 1965 (Finkelstein 2007). Prior to the current healthcare reform, the introduction of Medicare was the single largest expansion of health insurance coverage. It provided health insurance to virtually all Americans aged 65 and older. To give you a sense of magnitude, prior to Medicare about three-quarters of the elderly were uninsured. And the elderly were about 10% of the population. Medicare therefore provided insurance coverage where there was none before to about 7.5% of the U.S. population, which is roughly similar to the share of the uninsured in Massachusetts prior to the 2006 Massachusetts reform. The Affordable Care Act could cover about 11% of the population, so the expansion is roughly similar in magnitude.

To estimate the impact of introducing Medicare, I used the fact that prior to Medicare, rates of health insurance coverage were very different across different regions of the country. On average about a quarter of the elderly had private health insurance coverage prior to Medicare, but the rates varied a lot across region. So for example, if you look at the percent of the elderly who didn't have insurance, who would be newly insured by Medicare, in New England Medicare increases the fraction with insurance by about 50%, but in the East South Central United States, by 90%. I looked before and after the introduction of Medicare in areas where Medicare had more or less of an impact on insurance coverage and looked at differential changes in the growth of health spending.

I found enormous spending effects. By 1970, only five years after Medicare came in, I estimated that hospital spending is almost 40% higher than it would have been absent the introduction of Medicare. That's total hospital spending, not just hospital spending on the elderly. One back of the envelope extrapolation result based on this is that the spread of all insurance, public and private, which occurred between 1950 and 1990, may be able to explain half of the six-fold growth in the real per capita spending over this time period.

These results suggest that the spread of insurance played a very big role in driving healthcare spending growth over the second half of the 20th century. This is contrary to what you would get if you use, say, the RAND estimates, which are partial equilibrium estimates, which suggest that the spread of health insurance is explaining maybe only a tenth of the growth of healthcare spending over that period (Newhouse 1992).

What is the difference? One difference is that RAND is a real randomized experiment. You can be pretty sure of those numbers, whereas I had to do my best to compare these different areas of the United States and control for different trends. There's always a possibility that that's part of what's going on, but I actually think - and what I have evidence for in the paper – is that the analysis of the introduction of Medicare is capturing general equilibrium – i.e. system-wide – effects that the RAND health insurance experiment cannot. In RAND, the sample is 6,000 people across the United States, so you are getting the effect of someone

newly having insurance on their health care use, *holding constant* they healthcare environment; the doctor they see and the hospital they go to are not doing anything different because a few more people have health insurance. However, when you suddenly have 7.5% of the population newly covered - and it is the highest spending part of the population (the elderly) - that is increasing aggregate demand for healthcare. This can provide incentives to providers to change the overall practice of medicine.

In fact I find evidence that the introduction of Medicare encouraged the adoption of new medical technologies. As I noted earlier, there's widespread consensus that technological change in medicine is behind the growth in healthcare spending. Now we find that when you have large-scale insurance changes that lead to a big aggregate increase in demand, hospitals have an incentive to adopt new medical technologies. People will use these new technologies, because they are not paying for them out-of-pocket, and they presumably would not have used them as much when they had to pay for them out-of-pocket.

You can trace that line of reasoning back further, not just to how insurance affects the adoption of new medical technologies, but to how it affects the development of those technologies in the first instance. This is harder to do in the case of Medicare, because it is a national or global innovation market, but other work that I've done (Finkelstein 2004) and that Daron Acemoglu and Josh Linn have done Acemoglu and Linn 2004) looking at pharmaceutical innovation, we have found that when you increase the expected size of the market for a particular drug or a vaccine by requiring that people get it, subsidizing it or by having a growing demographic base for it, you see increased new clinical trials and increased new drug approvals for these drugs and vaccines

It therefore looks like insurance, by increasing demand because it lowers the price of medical care, encourages both the adoption of new medical technologies—as we saw in the case of Medicare—and also, further down the pipeline, the innovation and development of these new technologies. Another way of saying this is that we usually think of insurance as providing insurance against current expenditure risk, but part of what insurance may be doing is changing the nature of expenditure risk that we will have to insure in subsequent generations.

Conclusions and Further Thoughts

In 1963, Ken Arrow proposed the concept of moral hazard in health insurance, the idea that health insurance may increase the demand for medical care. That creates a fundamental tension for health policy that is trying to both cover the uninsured and simultaneously reduce the level and growth of health spending.

The challenge posed in Arrow's paper to subsequent generations of economists was whether we could verify that this theoretical notion of moral hazard and health insurance actually existed, and whether we could quantify its magnitude and explore its nature and implications.

How have we risen to this challenge? There is compelling evidence from randomized trials that health insurance affects medical spending. Those who say otherwise are ignoring the evidence at their own peril. This is a fact of life. It may not be what we wished for, but we have to think about it, grapple with it and think about its implications for policy.

Some of the implications have to do with what price the individual considers: is it the spot price is it the end-of-year price? I discussed this in the context of high-deductible health insurance accounts but you can also think of this in terms of Medicare Part D, the prescription drug program for the elderly. You can consider the famous "doughnut hole," which is a range of spending in which suddenly your price goes back up to 100%. How the elderly respond to a policy with a doughnut hole depends on whether they are thinking about it at the beginning of the year or not. Additionally, issues of selection on your moral hazard type and how health insurance may affect the development of new technologies are important for the full "moral hazard" analysis.

I think we have made some progress, but naturally important questions remain. Let just mention a few of them, although I am sure there are many more. One is, where does moral hazard come from? How much of it comes from the patient and how much of it comes from the doctor? Clearly the decision to go to the doctor initially is primarily driven by the patient, but once there, how much of the decision of what care to get and how to adjust care based on your insurance coverage is driven by the patient or the doctor. That has implications for designing contracts to reduce excess moral hazard.

Another key issue is the welfare cost of moral hazard. The price of medical care does not reflect its true marginal social cost. As just one example, consider prescription drugs, which due to the patent system, have prices that are much higher than their marginal social cost of production, which is approximately zero. It is not obvious what the welfare cost is health insurance inducing increased consumption of an already developed product; it may be that the real welfare consequences are via health insurance's effect on the incentive to develop it in the first place. These are important issues to think more about.

References

Acemoglu, Daron and Joshua Linn. 2004. "Market Size and Innovation: Theory and Evidence from the Pharmaceutical Industry." *Quarterly Journal of Economics* 119(3): 1049-1090.

Akerlof, George. 1970. The market for "lemons": quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84: 488-500.

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, 2012. "Moral Hazard in Health Insurance: How Important is Forward Looking Behavior?" NBER Working Paper No. 17802.

Aron-Dine, Aviva, Liran Einav and Amy Finkelstein. 2013. "The RAND Health Insurance Experiment, Three Decades Later." *Journal of Economic Perspectives*, 27(1): 197-222.

Arrow, Kenneth, 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53(5), 941-973.

CBO. 2012. Long Term Budget Outlook
http://www.cbo.gov/sites/default/files/cbofiles/attachments/06-05-Long-Term_Budget_Outlook.pdf

CMS. 2012. National Health Expenditure Data. http://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical.html

Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf and Mark Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review* 103(1).

Finkelstein, Amy. 2004. "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry." *Quarterly Journal of Economics* 119(2): 527-564.

Finkelstein, Am. 2007. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare." *Quarterly Journal of Economics* 122(3), 1-37.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joe Newhouse, Heidi Allen, Katherine Baicker, and the Oregon Health Study Group, 2012. "The Oregon Health Insurance Experiment: Evidence from the first year." *Quarterly Journal of Economics*, 127 (3): 1057-1106.

Gladwell, Malcolm. 2005. "The Moral-Hazard Myth." *New Yorker* August 29.

Newhouse, Joseph P., 1992. "Medical Care Costs: How Much Welfare Loss?" *Journal of Economic Perspectives* 6(3), 3-21.

Newhouse, Joseph P., and the Insurance Experiment Group, 1993. *Free for All*. Cambridge: Harvard University Press.

Rothschild, Michael and Joseph Stiglitz. 1976. Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Quarterly Journal of Economics* 90: 630-49.

Discussion

JONATHAN GRUBER

I was thrilled to be involved in the fifth annual Arrow lecture and to honor my colleague, Amy Finkelstein. I remember diligently working to recruit Amy to come to graduate school at MIT when she was on the Council of Economic Advisors in the late 1990s. And I would like to think that Amy's MIT education is responsible for some of her enormous success since that time, but to be honest, we're probably just riding on her coattails.

And these coat tails are indeed impressive. Amy has gone on to become one of the best health economists in the country as well as one of the best overall applied micro economists in the country.

In dozens of innovative articles that we just got a sample of here, Amy has managed to really carry on this great MIT tradition, starting early on with Robert Solow and Paul Samuelson, and working through Paul Joskow and many others in our department, of combining rich theoretical insights with compelling empirical work to really bring together the whole package.

After many years of watching Amy's career skyrocket, I was fortunate enough a few years ago to actually be able to co-author on a project with Amy, the study of the randomized health insurance trial in Oregon. I could see firsthand how she works, and I guess I have to use the term 'co-author' quite loosely. The image you should have in mind is of someone sort of hanging on to a rocket as it zooms of in the stratosphere. I'm the hanger, and she's the rocket; basically that's sort of what co-authoring with Amy is like. I received literally dozens of emails a day, so fast that I couldn't respond before I get another one, and eventually I basically just dropped out. She was kind enough to continue to keep my name on the paper, but I'm still shuffling through the emails years later. It's really just a pleasure and an enormous privilege to get to work with Amy on a project.

Amy's lecture focused on one of the key health information symmetries. It's sort of ironic because in many ways we think about Ken Arrow's 1963 article, the first thought that comes to mind is not the moral hazard contribution, but the adverse selection contribution and the notion of insurance market failure through adverse selection.

What is also ironic is that Amy could have done an entire lecture on the adverse selection issues. She has many years of contributions prior to the seminal empirical article on adverse selection in the last couple of decades, including her work with Kathleen McGarry

showing how we think about adverse selection, and the presence of both virtuous and vicious selection. It is a real compliment to Amy that she could actually take on a second key element of Ken's article and talk about moral hazard.

I would just like to make four complimentary points to what Amy talked about, really coming back to what she touched on in the end.

The first is that she only focused on one kind of moral hazard, what she called the *ex post* moral hazard. Now she also mentions also *ex ante* moral hazard, or the notion that people might not take care of themselves because they are insured. I think this is kind of one of those things economists like to spend a lot of time on and does not really matter very much. We actually did investigate it in the Oregon experiment. We looked at when people had health insurance, whether they would for example smoke more or exercise less. So far, we do not have the evidence that that is the case, so I think this is really kind of a cute point more than an important point.

But there is another kind of moral hazard that Amy just mentioned briefly at the end that I think may be even more important than the one she discussed, which is provider-side moral hazard (to be clear, this is distinct from the general equilibrium effects Amy talked about in her Medicare work). This is the issue that is really summarized best in that quote that says that having a doctor tell you how much medical care to get is kind of like having a butcher tell you how much red meat to eat. We have this broken fee-for-service healthcare system where physicians are paid and providers are paid based on how much care they deliver, not on how healthy they make you.

This is really the focus of most discussions of healthcare cost control today. While many of us are fascinated by individual-side moral hazard, in fact most of the discussion that dominates policy circles is not about that; it is about the provider-side moral hazard and how can we move from our existing fee-for-service system towards things, for example, like accountable care organizations, where providers are brought together under one global reimbursement system and reimburse some aggregate amount based on enrolling you, not actually how they treat you.

Now there is a huge amount of evidence that provider-side moral hazard matters. Probably the most compelling example was also from a big change in Medicare in 1983. It was probably the biggest change since they introduced the program in 1965. Medicare led the way in moving from a fee-for-service reimbursement system to a prospective reimbursement system, where they paid hospitals, not based on what the hospitals did to you, but based on what was wrong with you. When they did so, there was an enormous change almost overnight. The average length of stay at the elderly hospitals fell by 15%, within one year. And yet there

was no effect on elderly health. That basically we are delivering all this excess care that was not improving the health of the elderly and we could get rid of that, save a lot of money and lower utilizations.

This is a really important part of moral hazard. This is not in any way in dispute with what Amy said, but it does raise a critical issue she talked about at the end, which is the complementarities between the patient-side moral hazard and the provider-side moral hazard. These communities are talking past each other, because if you talk to the provider-side moral hazard folks, they say, "Oh, that patient-side stuff is not important, it doesn't really matter. Co-pay is deductible, that's all just regressive and unimportant, we've just got to focus on how we reimburse providers." And then you have other folks—I'm more on the other side—saying it's really critical that we get consumers to be more cost-sensitive and to pay attention to the cost of medical care. Both sides are clearly right, and what neither side has really tackled is to think about how they interact. In particular, how do we interact the provider side incentives with the patient side incentives? How do we, for example, set up patients in plans which charge you more if you go to providers that are on fee-for-service not on capitated reimbursement? How do you put these two things together? That is a really critical issue to think about going forward.

The second issue is the interpretation of the evidence from RAND in Oregon. Amy mentioned that what these studies do not do is measure moral hazard. What they measure is the effect of price on medical care utilization, and that is because of two components. As economists, we think of them as income and substitution effects. More generally, think of the fact that on the one hand, the price of medical care is changing. On the other hand, when you have insurance, you have just been given a gift of money should you get sick, and you may choose to spend some of that money on your treatment. That is not a distortion of the system; that's just how you choose to spend your money. This is what we call liquidity effect.

That is probably what is partly going on when we get people health insurance and they spend more. People would always like to spend medical care, but they are constrained from doing so. By giving them money when they are sick, we allow them to afford that expensive operation they could not otherwise. That is not a distortion in our system; that is just getting rid of the problem of liquidity constraints that people face in our system.

The problem is these have very different policy implications, whether it's the true moral hazard liquidity effect, and we don't really have great evidence to separate them. In fact in other contexts and the context of unemployment insurance, Raj Chetty at Harvard and his colleagues have done very exciting work, showing that a lot of what we thought was moral hazard and unemployment insurance is actually just liquidity effects. We thought, gee, there's this evidence that if you give people higher unemployment insurance benefits, they stay

unemployed longer. People like myself have been running around saying, well, that means unemployment insurance is too generous, and when we give people these unemployment benefits, they sit on their couch. Chetty and others did a great study. They looked at Austria, where they do not give unemployment insurance benefits; they just give severance payments. You were not paid based on how long you stayed employed, but just got a check when you left your firm. And you know what? People stayed out of the labor force just as long; they still just sat around on their couch. The point is that people seem to value sitting around on their couch and that basically the liquidity benefits of these grants were mattering as much as the moral hazard effects. Really thinking about that more in the context of health insurance is critical.

In fact, we know the ideal experiment to run. You would like to take two people; to one you give health insurance, and to the other, you just give a check should they get sick and say use it any way you want. Then you ask if they behave differently. If it is moral hazard, they will behave differently, and if it is just liquidity effect, they will behave the same. I do not think we will ever get to run that randomized trial, but I'm sure Amy is clever enough to figure out a way to tease out that test in the real world.

My third comment is that if you read Amy's piece and think about these moral hazard effects, it seems pretty obvious that we ought to have patient-side cost-sharing. We got to have patients bearing some of the consequences of medical decisions.

There is of course a flipside of that, which are affordability concerns. Why are many people opposed to having patients face cost-sharing for their healthcare? It is really affordability concerns: they are worried that people cannot afford these costs. We have out-of-pocket limits to make sure it is affordable. If you take a very poor person, someone at the poverty line, individual with an income of $13,000 a year, they cannot afford a $2,000 deductible. They probably cannot afford a $1,000 deductible. This is not something that is really affordable for them.

Years ago, Martin Feldstein proposed the plan for which he chose the unfortunate acronym MRI, for Major Risk Insurance. That was sort of an unfortunate acronym given medical developments in later years. In his Major Risk Insurance plan, people would pay 50% of the cost of their healthcare until they spent 10% of their income. I believe this may be Feldstein's only progressive policy proposal. Basically what it said was that people would be protected from being bankrupted by medical costs, but be price sensitive on the margin.

Whether that is a good idea or not actually depends on another aspect, income heterogeneity, which has not really been studied. In the kind of moral hazard Amy discusses, whether a plan is good or not depends on whether poor people are more or less sensitive? If you protect the poor, but they are really super-price sensitive, then that might have a higher

welfare cost and you might want to not quite protect them as much we otherwise would. We have to understand not just the overall moral hazard effects that Amy mentioned, but we also have to think about how that affects differences on the income distribution. Think about setting up the optimal kind of major risk insurance system.

Finally, the trade-off between setting the right marginal incentives for consuming healthcare and certainty is another difficulty in translating Amy's work to policy. Economists have always said that the right way to set up a health insurance plan is to have what we call co-insurance, which is to say that you would pay X percent of your bill, 20% of your bill up to some maximum. People in the real world hate co-insurance. They like co-payments, and we ask them why. They say because it's too uncertain for people going to the doctor: they don't know what they're going to pay.

Part of that is the lack of price transparency. We do not know what 20% of the doctor's bill is going to be, because it is not posted. Part of it is true uncertainty, because you do not know what that doctor is going to do you or which tests you're going to get, et cetera. Subjecting people to co-insurance subjects them to some uncertainty. We know from the pioneering work of Ken Arrow and others that people do not like this uncertainty and that they would pay to insure against it.

That actually raises an interesting question of the right way to incentive people. The next step in Amy's research agenda can ask what is the right way to incentive people. On the one hand, you want them to face the price on the margin. You want co-insurance because that stops them from getting the really expensive treatment. On the other hand, that introduces more uncertainty into what they're going to have to pay out-of-pocket. A really interesting question, which is straight along the line of Amy's research agenda, is to understand those trade-offs between co-insurance, which economists find appealing, and co-payments, which people in the real world find appealing. How we actually resolve that trade-off going forward, I think, is really significant.

These are just a few ideas. Amy's work is uncriticizable because it is so excellent, so rather than trying to do the impossible, I thought I would offer a few suggestions where she could go next. I think Amy is really perfectly positioned to do that. I want to close by saying how thrilled I am, how wonderful I think it is that she was selected to honor Ken Arrow's work. She is just a terrific example of watching someone's work continually evolve and improve, and it's just exciting to have a front row seat for that ongoing process. Thank you very much.

Commentary

KENNETH ARROW

I was very happy to once again visit by graduate alma mater in 2012 for the fifth annual Arrow lecture.

It is a very different world in many ways. In the university, what was taught in the economics departments and the issues that we were concerned with have evolved, although I must say the Great Depression does seem to have a rather contemporary echo. Some things change and some things go around in cycles.

We have heard a wonderful speech that drew together evidence, and used the ability to develop experiments and to exploit natural experiments with a creative emphasis on the theoretical understanding. Amy's presentation has been a great pleasure to all of us. We have read the papers before, but this puts the whole thing into perspective that is already making a contribution to our study of our health problems.

Let me take a tack which was referred to by Amy and was spoken about somewhat more at length about by Jon, which I will call the non-economic side. In fact, part of this is in my original paper.

Let me start by talking about what moral hazard is and a little bit of its history. What I meant by moral hazard, is basically an asymmetry of information or knowledge between two parties. (Moral hazard is just one kind of asymmetry.) Let me just say a word on the origin of this project. The Ford Foundation then was a very large grantor to research in the social sciences, and one of the officers there was Victor Fuchs, who many of you know is one of the most distinguished health economists in the country and the world. I had met him, but we weren't close. He called me with the following project which he had gotten approval from the Director of Economics and Research Lloyd Reynolds, of Yale. He wanted to take three issues which involved public and private policy, and have each one studied by two people. One person, who had already worked in the field, would give some kind of summary, and the second person would preferably be a theorist who had not worked in the field at all, who would see what theory could do. He wanted me to be the theorist for health.

I have a tendency, which has worked out quite well for me in many ways, to take a dare. This is probably the best example, although *Social Choice and Individual Values* actually arose to some extent in the same way. When I'm challenged to do something which I am not doing and I'm not continuing research, I think, well, this is more exciting, I'll learn something. I thought, after all, health is a serious issue, and is theory is any good at all—which is a point which can be debated—then this would be a good place to apply it and see what we have to say.

I very diligently read the various books out there, what literature existed at the time. There were health economists even then, like C. Rufus Rorem, Harold Somers and Anne Somers. I read a lot of descriptive material. Of course, apart from social choice, my big interest had been the economics of uncertainty and risk-bearing. I was trying to fit it into a general equilibrium with uncertainty concepts.

In insurance at that point, there was a fair amount of hospitalization insurance. It was just the beginning of insurance for what they called major medical risks, by which is meant risks outside of the hospitalization area. HMOs and things like that did not exist at the time. There were a few clinic groups, but they were a very small part of the story. These were all incipient, although as you saw from the graphs, the health problem looked pretty big at the time, although it looks trivial by modern standards. It was probably on the order of 6% of national income. It was a significant problem, but not overwhelming.

It seemed pretty obvious that if you have risks, you ought to have insurance; it didn't seem that complicated. Why can't we have much more extensive health? I began trying to fill in this question. I thought, well, it's difficult for the insurance companies, because what exactly is the insurer insuring against? For example, with automobile insurance, when you have a smashed windshield, that's it. There's no choice involved: you have to replace the windshield (though even then there are problems of moral hazard, but they are not the dominant concern). If the procedure being done is pre-assigned, given the diagnosis—like replacing a windshield—then there is no moral hazard. But when you're sick, maybe your diagnosis turns out not to be what you thought it was, and it may be more expensive or less expensive. There are a lot of decisions that have to be made along the way, some of which are just questions of comfort, some of which are really medically significant. The insurer is not in a good position to monitor what's going on.

Then, I began to realize there was another problem: the relationship between the physician and the patient. How does the patient know the physician is putting forth full effort to on a case? Unfortunately, I know some people who have had serious consequences from that, so it's not imaginary.

I began to realize that these were the same issues; the light bulb went off, just like in the cartoon strips. I had spent several months trying to do a diligent job reading all this stuff, and it was kind of turning out to be kind of a workmanlike thing, where I said how each aspect was related to theory, and suddenly I got the idea there was a general principle here. The general principle was that in these relations there was asymmetric information. A doctor's function in life by definition is to know things you don't know. The insurance company is also at the mercy of the physician and patient.

Of course that raises an issue with relation to the concept of moral hazard. The adverse selection was another example. Now, I had had some slight exposure to the actuarial world. Namely, at one point in my life, not knowing what I was going to do for a living when money had become a problem and my family was badly hit by depression, I was concerned with some secure occupation. Academic life was sort of another world I didn't know about, so I was thinking of becoming an actuary. In the reading about it, I had come across words like moral hazard and adverse selection. I would not have thought of them as part of my working knowledge, but somehow, though this was probably 10 or 20 years earlier, those concepts came to me suddenly, and I said, "Oh, this is what the insurance company people are always talking about."

Now the question is, how do you deal with asymmetric information? This is a defect. In other words, you had a system of ideal future markets, with conditional uncertainties. But if there is asymmetric information, you and I cannot make a contract condition on something if only I know whether it happened or not. We have an immediate explanation of market failure. What does this create? If you have prices for everything, ethical considerations do not really enter the story, but as soon as you have this asymmetric information and the failure of prices, you have ethical problems, which true market system does not have.

There have some kind of a social benefit to ensuring truth-telling. Of course there is a vast literature on mechanism design. Personally, though this is a very important literature and it has its role, there are all sorts of reasons why it is not an adequate substitute for market failure due to asymmetric information. In some cases, it will be. Even in the financial world, we have seen so much of this, it is very difficult to devise incentives. In fact, we have some evidence the incentives they have supplied haven't functioned as well as might be desired, if I put it delicately, and it is not clear to me there is any way of completely correcting that.

One of the points that Amy raised briefly and that Jon talked about a little more extensively was that if you look at the system, the decision to buy and to use medical care is by no means a decision of the patient only. I rely on the doctor just as I rely on an attorney. There is a reason why professions have a special role. They are marked by strong inequalities in knowledge; basically what the doctor's services supply is not skill, though that matters, but knowledge. You can argue about why you have specialization, where a person is acquiring knowledge that can be used over and over again. Socially, it pays to have some people have knowledge and use that knowledge over and over again with new patients or clients, rather than everybody being educated in the medical field.

That creates, automatically, an asymmetry of information in the whole system, including physicians, hospitals and HMOs. One of the immediate reactions of insurers when they started gradually going into the health business, was they would start checking on the doctors. They

had doctors saying we should do one thing and the insurance company rejecting it, which has gotten so notorious. The insurance companies or HMOs are using their knowledge too.

You begin to get a very complicated kind of information system, in which the patient's role becomes very minimal. This worries me, particularly the most striking question of the differences in the moral hazard response. What exactly is the source of that? Is it that you are getting different doctors, for example, or different localities with different conditions?

Medical ethics has always been considered a very important matter. There was a whole school of thought, especially around the University of Chicago, that iargues that restrictions on medical procedures and medicines, and all these rules and ethics are essentially devices for improving monopoly power. Milton Friedman was a great advocate of that point of view. He went so far as to argue that we should not even license doctors and that we should have free entry into the medical profession.

But I think this misses the point very considerably. We find, for example, medical associations putting out information designed to prevent you from getting ill, which presumably could hardly be regarded as a profit maximization procedure. In my original article, I instance the fact that doctors and hospitals do not advertise as example to signal that they are not maximizing profit, although unfortunately now that particular story is quite wrong. When I pick up the papers and I see very respectable hospitals, like Sloan-Kettering, explaining why they are the best place to go for cancer or whatnot, I'm still a little shocked, but I suppose younger people are not. We rely on medical ethics, and we have since Hippocrates. Similar things happen in other professions

Let me mention some evidence. I am not an expert and some of these facts—well, some of these alleged facts may not be facts. I should say one more thing,

With introduction of new medications and devices, there is always the matter of who decides to use them. One of my students was studying that and what he found was that it was very hard to locate who made these decisions. Every HMO would say, it's up to the doctor, and if the doctor thinks it's necessary, we buy it. But it is clear that some HMOs are very resistant to that. This is a subtle thing which is not caught by the market system. We know for example, that medical costs vary greatly by region as the Dartmouth group has emphasized, with apparently no particular relation to outcomes, as studies have claimed.

If you take international comparisons, we have enormous differences in medical costs. Before and after it was pretty much the same allocation problem; it's not at all clear why there should be this difference. After all, Canada does not face different issues than the United States, so why is health care in Canada so much cheaper? As I said, if insurance raises prices, it should have much better care coverage than we do. You would expect to think it worked the

opposite way, and you would expect Canada to be more expensive than the United States. If you look at the French or Israeli system, all these seem to involve giving much better coverage and yet somehow getting by with much lower costs. Nobody is even close to the United States.

It is widely held, at least I believe probably correctly, that it always hard to prove these things: if you take different HMOs, their costs are quite different. Locally in California, everybody says Kaiser gives much better value than others, and yet Kaiser has not spread in spite of efforts, which is interesting. The workings of the market should imply that Kaiser and Kaiser-like plans should spread, but they have not spread nearly as rapidly as what you might expect.

Economists always hate to invoke the concept of culture. It has been around at least since the Max Weber's identification of emergence of capitalism with religion in 1905. In particular, if we go to our organizational behavior colleagues at business schools, we find they are always talking about corporate cultures.

If you have an organization with people working together, they develop ways of working which are hard to change. Newcomers come in, and they are going to be adapted to the culture rather than thinking of ways to change it. The fact that other firms are doing better does not necessarily mean that in the long run, you will eliminate the inefficient firms—especially in the medical field, where competition for various reasons is not so severe. We see certainly see differences even in firms that are much more exposed to the market. We clearly need to develop is a culture of accountability. Obviously there is room for accountability and accountable care organizations, which have a culture of cost-effectiveness.

How this is going to be achieved, I must say, is a mystery. Playing with the contract reimbursement systems and going to capitation instead of fee-for-services matter a great deal, but a lot more has to be achieved by organizational claims. Obviously, since organizations perpetuate themselves, they need some kind of outside leadership, but not, I think, central control. Something like the British National Institute of Clinical Excellence, which gives statements about the value of new drugs and with some idea of their costs, could play some important role.

Now, how this is going to be studied experimentally, I'm afraid I'm going to leave to those much more qualified than myself to discuss. Thank you.

Commentary

JOSEPH STIGLITZ

Let me begin by addressing the most fundamental issue that we are all interested in: the design of a healthcare system, including an insurance system, that would increase our welfare. When we think about this, one of the fundamental problems in the analysis is that in most other areas of economics, we begin with the presumption that markets are efficient, while this is an area where we should begin with the presumption that markets are not efficient. That really changes the analysis in a fundamental way.

One of the reasons that markets are not efficient has to do with the very general theorem that Bruce Greenwald and I proved a number of years ago, namely that whenever there are information asymmetries, whether adverse selection or moral hazard kind, or incomplete markets, then the economy is not efficient. (Those are assumptions that are always true but are particularly true in the healthcare market.) It is not constrained Pareto efficient taking into account the cost of information and the cost of setting up markets. There are always interventions that could make some individuals better off without making others worse off, or make everybody better off.

The kind of doctrine that you hear on one side of the political spectrum—that if we just led it to the markets everything would be right—has no intellectual foundation. There really is a need for government, and that is why the issue of how we design an appropriate set of institutions is absolutely key.

The second observation I want to make really echoes what Ken was saying: our current system is not working. To me, the issue is not so much that we are spending so much on healthcare. I do not have a view about the right amount that we ought to be spending on healthcare. People like to live longer, especially as they get near this boundary value conditions, so it is understandable that people would like want to spend a lot of money on healthcare. It is not that we spend 17% of our GDP on healthcare that is the problem; it is that we get so little for how much we spend. The real problem is the inefficiency of our healthcare system. The fact, as Ken pointed out, that we spend much more than any other advanced industrial country, not only much more as a percentage of our GDP, yet our health statistics are worse than other advanced industrial countries. In fact, in some dimensions in infant mortality, our health statistics compared with some developing countries; they are worse than Cuba, and they are an embarrassment.

Some of that has to do with the first issue that Amy brought up, the problem of access. Large fractions of America do not get access to adequate healthcare. I am going to come back

to that in a second.

One of the questions that we ought to be focusing on, which is a little bit outside of Amy's essay itself but which is part of the welfare analysis, is why is our healthcare system so inefficient? The point that Ken raised is absolutely right, the people are not basically less healthy. It is not as if only sick people immigrated to America. In fact historically, we always talk about the selection effect going the other way: it was not the sick, the ill and the genetically defective that came here. It is not as if the weather here on average is that much worse than other places; there are places that are better, I know, but it's not the worst. There is no inherent reason why we should spend so much to get so little. It is also clear, as Ken pointed out, that there are information asymmetries on both sides of the Atlantic and the Pacific, so it is not as if this is uniquely American problem. Moral hazard, adverse selection, problems of doctors exist everywhere, and the way we have solved those problems is not as good in some respects as the ways others have.

Part of the reason has to do with old-fashioned rent-seeking. We have created some institutions that are very good at rent-seeking and do a better job at it than those in other countries. We have public sector that has less control, so it is not a surprise then that we wind up having an expensive healthcare system.

Emblematic of that is the provision in the 2003 Medicare Drug Benefit where we said the U.S. government, the largest purchaser of drugs, could not negotiate with the drug companies. That provision is estimated to cost is excess of $50 billion a year. That is just a small example, but there are others.

We had a Medicare program called Medicare Advantage. Medicare Advantage was basically identical to Medicare, but it was in the private sector and it cost about 20% more. It did not do anything better. Why did we have it? Well, we had it because the health insurance industry lobbied to have it and they made money. Transaction costs are a bad thing for most of us, but for the health insurance industry they are a good thing, because that is what they live off.

Ken raised another significant point, that practices vary across the country and between countries. There are many practices that are not cost-effective, so you see in some places procedures like tonsillectomies are expensive with no evident benefit.

One of the good things that was in the Affordable Care Act was a provision for trying to do a more scientific study of these practices, and hopefully to create norms where you do not do costly things that do not yield any benefits. This is at least the beginning of a more scientific discussion of medicine.

The other set of remarks I want to make has to do with the development of economic theory. I was very much influenced in my work by reading Ken's little book that he wrote, his Yrjö Jahnsson lecture, called *Aspects of the Theory of Risk-Bearing*, as well as his paper on moral hazard. These were really seminal works that began to help us think about problems of the economics of information asymmetries.

The underlying problem in all of these examples is that, in the case of moral hazard, we cannot observe the output and we can observe the input. That is a real problem. We do not know what the doctor is doing, what he should do or what the consequences are, because even if he did the right thing, we could have a disease that will have bad outcome. So we can't judge whether the doctor has done the right thing by whether we survive. He may have done the right thing and we die anyway, or we may survive although he did the wrong thing. It is really difficult to have pay-by-performance. That is one of the inherent difficulties.

The models that I studied and that Ken studied and others studied some 30 years ago were always very simple. We looked at a single example of asymmetry of information, a single moral hazard, a single adverse selection. What the discussion this evening has highlighted is that most interesting problems involve multiple information problems. The doctor does not know the illnesses of the patient, the patient does not know whether the doctor is doing the right thing and the insurance company does not know about either of those two. Actually there are three moral hazard problems.

But in insurance market, there is an additional problem of information asymmetries. The insurance company does not know what my health condition is and I do not know whether the insurance company is going to pay me when I actually need them to pay, so there is a problem of payment. What is striking to me is that in the ensuing years, there has been remarkably little research done on a theoretical level, of trying to analyze models where there are these multiple information asymmetries.

The reason is very understandable, because the problems get very complex very quickly. One way of thinking about it is that you have incentive compatibility constraints, you have one set and then you multiply that but you have another set and you try to analyze how these interact simultaneously, and then you have adverse selection. They get very, very difficult.

I have done some of that and the only thing I can say about is, as far as I know, nobody has read it, because it is so difficult that even I struggle reading it after I write it. The problem is that the insight is not so clear, but what is clear is that the problems that we face at the practical level all involve these multiple levels of information asymmetry. In one way or another, eventually we're going to have to cut through and find what the right approximations and simplifications are that allow us to analyze it. It is great PhD thesis material for those

students who like to have complex problems that even their thesis advisors will not be able to read.

Further Considerations

AUDIENCE QUESTION: It seems to me that cross-country evidence argues against moral hazard in the sense that countries that have more comprehensive coverage have lower costs in terms of proportion of GDP spent on health. Is, then, the problem the moral hazard with not with consumers of healthcare, but with providers of healthcare? Presumably the consumers of healthcare have similar behavior with similar incentives, whereas what is different in the system are the incentives for the providers of healthcare. Is that where the problem lies and that's where research should be concentrated?

AUDIENCE QUESTION: There was no discussion about provider-side moral hazard. Actually patients only control, in their demand, about 10% of all costs. Most costs are in the hands of specialists in hospitals who have multiple profits, gains and incentives from ordering more tests and services for their professional wellbeing and their income. In your studies, did you have a control to examine the degree to which the increase in demand in cost, which you're attributing to individuals, is accounted for by providers?

AUDIENCE QUESTION: How does preventative care come into some of this modeling around increased use and potential reductions in future spending?

STIGLITZ: Thank you. Can I just add one question? There are two additional points I wanted to make. One of them goes to the welfare issue. Normally we think of technological advances as a good thing. There is a little bit of a hint you are saying that technological advances were driving costs, and since there is this paranoia about cost, there is the suggestion that it was a bad thing.

Let me give you my answer. Some of the technology changes are actually good things, that they are extending lifetime and they are enhancing welfare, but that in fact, the cost that we pay for them are not the marginal cost. They are a large ranks associated with it, so that our system of producing technology is not efficient, and our system of using it is not efficient. For instance, in some countries you can only use more expensive drugs if you can show that they have a benefit. In Australia, they have list of drugs that you can use, whereas we do not have those kinds of control mechanisms. In a way, while technology can have a potential for increasing welfare, the way we have used technology and the way our market distorts the production of technology means sometimes it is not welfare-enhancing.

The second point was really related to the last question. You have talked about your RAND study and a couple of the others, including the Oregon study. You talk about the usage hospital admissions now, but that does not talk about long-term healthcare costs. If some of these people are going to the hospital today when they would not have gone, they might have wound up with much worse diseases; that is to say that their problems would have

accumulated. It is not obvious that more usage is going to lead to higher healthcare costs.

FINKELSTEIN: There were really three sets of questions. One had to do with moral hazard on the provider side and the role of providers versus patients.

Jon didn't mention, but I'll mention that he is not only now my colleague at MIT, he was originally my thesis advisor at MIT. I think his discussion in some sense embodied what the ideal thesis advisor does. They encourage you, and they tell you what you're doing is wonderful. Then as you are walking out the door, they make a side comment that basically very gently suggests you've entirely missed the point, but you leave feeling really pepped up and encouraged. Three days later you realize, now I understand what I need to be working on.

I think his comments and the first two questions that came from the floor do really emphasize that one of the key place we need to shine the flashlight now is thinking about incentives on the provider side and not just the issue that I was alluding to of how much of the patient response is influenced by the physician. We need to think about how much of what the physician is doing is a response to their own financial incentives. That is an area where we need a lot more work.

The second question that came up was on preventive care and future healthcare spending. We did look at this in Oregon, and the results from the first year show that use of preventative care increased, such as use of mammograms, prostrate screens, checking blood cholesterol levels, et cetera. It is not clear what the long-term benefits of these are. For example, in the course of our study the recommendations on mammograms changed, so we did not know how to code recommended preventive care—recommended as of when?

The question Joe raises about the long-term future spending is something we would like to follow. We cannot follow it in the experiment, because the experiment itself lasted only two years, but we are getting physical measures of what happened to your blood pressure, blood cholesterol, and whether you were taking blood pressure medication, and so on. We can use those to try to forecast what the long-run health spending would be.

Jon may have a different view, but I have to say I am pessimistic that this will actually, in the long run, reduce spending. Then sometimes I feel that setting that up as an objective kind of misses the point. It is okay to spend money on your health: if health increases with health spending, then it is not necessarily a bad thing. It would be nice if there is a free lunch, and you could actually save money and improve people's health at the same time.

I am sort of pessimistic on that. That again relates to Joe's last point, which is a very important point to conclude on. You said you talked about technological change driving healthcare spending, which suggests it's a bad thing. You can also put up a graph that shows

that survival, particularly for people 65 plus has increased dramatically over the last 40 or 50 years, and that's in large part due to the technological change. I think you're absolutely right that the question is not whether we need more or less technological change, but how we align incentives to get the socially valuable type of technological change.

I'm a little less sanguine than you about the lists, in the sense there are certainly some things that we think almost everyone should do and some things that almost everyone should not, but there is probably a lot of heterogeneity both in the treatment benefits and the preference benefits, so it's hard to know how to work that out.

ARROW: I have said to several people that the real way to control costs is to abolish the National Institutes of Health. Someone said they would save us $50 billion a year, and I said that's not the point. It's all these expensive medications that result from this research.

Obviously we don't believe that. There is this rather heroic calculation by Holland Jones [phonetic] arguing that the idea that we should be spending 30% of our GNP on health is by no means absurd. The numbers are sort of made up, but they are not totally incredible numbers, that's the point. What is creating the health problem and why it is a special problem, is that it is a matter of public interest and there are taxes are involved. It would be less complicated if somehow we could do all this and insure privately. There is an interesting ethical dilemma here: we don't like to see people die if the medical care is available and the only problem is they can't paying for it, even if it's their fault for not having insurance.

Let me give a parallel. A number of years in the Sierra Club bulletin there was a letter that had to do with the fact that the National Park Service was proposing a rule that every climber had to have a helmet. This fellow, who had been a famous mountain climber and was then a manufacturer of very high quality climbing equipment, said he objected to it on grounds of freedom. He said he understood the reason for it, that if a person is trapped or injured but has his head intact, he can cooperate with the rescuers and make the rescue much easier if he has a helmet. He said, therefore, the only proposition is that if you go climbing without a helmet, you understand that you are not going to be rescued.

Now it shocked me. There is somehow a moral judgment here, that we can just let people die; therefore the financing becomes through the tax system, because of unequal distribution. You could think of, although it's difficult, an adequate insurance system where the ability to buy insurance would be related to income.

Let's for the moment imagine we can overcome all these moral hazard issues and we can have really good insurances really serves the proper function. Nevertheless, you would have the problem that after all health costs are by no means proportional to income. If you take the average health cost is at about $8,000 per person, it is obvious that an insurance adequate

to cover that would not be sustainable, so we have Medicaid. In other words, if you take an economic point of view, you can say that what you have is a lot of taxation, and taxation is distortionary. The costs are really higher than the money costs because there is a classic welfare loss due to excess burden due to taxation.

It seems to me that is why all these special problems arise, because we saying we're going to pay for it by taxes some way or another. Taxes mean a bigger role for the government, which we know is somewhat unpopular in some quarters, but even from a purely economic point of view, it does involve welfare losses, which have to be offset against what we think are the gains.

Interestingly nobody is going to really repeal a lot of this stuff. The marginal group of uninsured is only 50 million, that's why we can talk about throwing them to the dogs. We're not going to destroy their Medicare, and in fact there is no great opposition to Medicaid. Maybe there is opposition on the margin, but it's surprising to me that given the diversity of opinion, this is accepted. That is why it becomes a big issue, I think; it might not have if it was just a question of increased market share.

STIGLITZ: Just two comments: one, we've talked about the incentives of the providers, and one of the interesting aspects of that is the peculiarity of fee-for-provider, where you have an incentive to sell more. The other one of course is the HMOs, where they get per capitation, and they have an incentive to provide nothing. You hope that the reputation mechanism works, but we know that reputation mechanisms don't work very well, particularly when choices are limited. I mentioned that because it highlights the difficulty of solving the incentive problem; if it were easy we would have solved it by now. It was clear is that other countries have done a better job of this, and part of it may be sense of professional responsibilities, the kind of ethical things that Ken was talking about.

The second thing is that we have been focusing on the healthcare sector, but we have said actually relatively little about health. We ought to always remember that the healthcare sector is an input to something that we really care about, which is health, and it's not the most important input. Other things like nutrition, smoking and drinking probably have more important effects on average on a lot of aspects of our health, than does healthcare. The reason why this is partly an economic problem, of course, is that if we are going to subsidize the provision of medical care, it also raises question of our positions about taxing and subsidizing activities that affect health.

Erlich, Isaac, and Gary Becker, “Market Insurance, Self-Insurance and Self-Protection,” *The Journal of Political Economy,* vol. 80 no. 4 (Jul.-Aug., 1972), pp. 623-648.

Pauly, Mark, “The Economics of Moral Hazard: Comment,” *The American Economic Review,* vol. 58, no. 3, part 1, June 1968, pp. 531-537.