STATEMENT OF MARK V. PAULY, PH.D.

Bendheim Professor, Professor of Health Care Systems, Insurance and Risk

Management, and Business and Public Policy in the Wharton School, and

Professor of Economics in the College of Arts and Sciences,

University of Pennsylvania

BEFORE THE JOINT ECONOMIC COMMITTEE

September 22, 2004

Thank you, Mr. Chairman, and members of the committee, for an opportunity to testify on adverse selection in health insurance and related issues.

Private health insurance would be far less controversial if we lived in a world where everyone was similar in terms of risk. Then insurers would charge similar premiums to everyone who put similar effort into shopping for a given policy, and would be equally eager to sell insurance to anyone. After the fact, those lucky enough to have low actual health expenses would have paid in more than they got back from insurance, but this redistribution from those who did not become sick to those who did would be something that everyone would agree before the fact was both fair and attractive, and all would be eager to buy insurance as long as the premium was not too much higher than expected benefits.

The world in which we do live, it is obvious, is different. It is one in which "risk" varies before the fact, in the sense that different consumers reasonably expect to collect different amounts in benefits from a given policy because they expect to get sick with different frequencies and severities.   Insurers can identify and measure some characteristics that they know predict above or below average benefits, characteristics such as age, location, and the presence of chronic conditions. Insurance markets can still function in such a world, but now either premiums or purchases will be different for different people.

What will happen depends crucially on whether insurers have and can use the same information that predicts benefits as consumers can use. If everyone has the same information, and the information does predict different risk levels, then insurance theory (Rothschild and Stiglitz, 1976) tells us that insurers will choose to charge below average premium to the lower risks and above average premiums to the higher risks. Someone who has four times the expected benefits from a given policy compared to someone else will be charged about four times the premium. At those premiums, insurers will be equally eager to sell to low and high risks. In insurance theory, this situation of *proportional risk rating* will be stable and probably will be one in which low risks are no less likely to buy insurance than high risks. (Some very high risks with low incomes may find that premiums are so high and expenses so near certain that they are just as well off not buying insurance they cannot afford and paying those expenses directly when and if they can.)

Insurance markets, the same theory also tells us, will be very different if insurers do not have equal information to that buyers have, or if insurers are not allowed to use the information they do have in setting premiums and bidding for business. In the extreme case in which insurers either cannot distinguish among risks or are not permitted to do so, they will be forced to charge the same premium to everyone who buys insurance. But if the insurance purchasers know their risk levels, their willingness to buy insurance at this premium will vary. Higher risks will be very enthusiastic about buying, since they can on average collect in

benefits more than they pay in premiums. But low risks may, in the limit, decide

not to buy insurance at all because it looks like a bad deal to them, or may at

least seek to buy less generous coverage than the high risks desire. This

situation of *community rating* will be one in which the low risks are less likely to

buy insurance than under risk rating. In the limiting case in which the low risks

bail out altogether, the so-called *death spiral,* the premium insurers end up

charging to the high risks will be the same as they would have charged under risk

rating; the effect of community rating will only be to drive out all of the low risks

(which is definitely not the same as no risk) from the insurance market, with

resulting adverse effects on access to care and financial stability. It is in this

sense that community rating can be inefficient compared to risk rating, since it

can make the low risks worse off and not make the high risks better off (Pauly,

1970). In the less extreme case in which some low risks might continue to buy,

the high risks could be better off but the low risks will still be worse off than they

would have been under risk rating. There will still be inefficiency compared to the

ideal because the low risks will choose less coverage than they would have

chosen if they had faced premiums reflective of the true cost of their coverage.

Whether there will be *cream skimming*, in which insurers are more eager to sell

to low risks than to high, depends on whether the adverse selection-community

rating is *essential (*caused by insurer inability to tell risks apart) or *inessential*

(caused by regulations or policies which forbid insurers from using information

they have to set lower premiums for lower risk and higher premiums for higher

risks). In the case of essential adverse selection, as in the case of risk rating, there should be no cream skimming because all potential purchasers look equally profitable to insurers. Insurers might want to cover only the low risks, but they cannot tell who is who.   In the less extreme case of regulation-required community rating, insurers will try to avoid selling to high risks they can identify, on whom (as a group) they are sure to lose money; there will be *cream skimming.*

For these kinds of reasons, some insurance analysts think risk rating is better than community rating. But many policymakers, and some other analysts, do not look at it that way. They do note the downside of community rating in terms of squeezing out the low risks, even to the extent of a death spiral in which at some point only the highest risks end up buying. (This should really be called a near-death spiral because at that point it will be profitable for some insurer to enter and offer a less generous plan at a much lower premium that can pull some of the lower risks back into the market; the market will rise, phoenix-like, only to go into another spiral.) But policymakers also find much not to like in risk rating, precisely because the higher premiums for higher risks may bite into their ability to consume other necessities for life if they have low income, and sometimes because observing higher income high risks paying more than higher income low risks still looks unfair, especially compared to a policymaker's dream world in which everyone pays a low premium. That this is impossible in a world of

competitive but unsubsidized insurance markets only marginally dampens their ardor.

The most obvious way to deal with these problems is to use regulation. Require insurers to charge similar premiums (or limit premiums for high risks), but forbid low risks from buying less generous policies. Then require insurers to sell policies to high risks they know will be causing losses, and, when there is enough political nerve, forbid insurers and the low risks from dropping out by mandating insurance purchasing. Measures short of this draconian one can still lead to bad adverse-selection type outcomes, especially when community-rating rules force insurers to ignore information they have and thus lead to *inessential adverse selection.* Then, when insurers respond to community rating regulations with cream skimming, one needs to write yet more regulations to require open enrollment and guaranteed issue. To avoid the death spiral, we move to a regulatory spiral. As with other kinds of health care regulation, how bad (or good) the regulatory outcome will be seems in practice to vary across states, depending on the characteristics of their potential insureds and the form and administration of the rules. In some states such rules seriously curtail the size of the insurance markets, while in others the main effect is only discontent among the low risks and the insurers who would like to sell to them.

The main novel point I want to make here is that recent research suggests that, in both theory and practice, there are ways alternative to regulation to get closer

to what policymakers want (or should want) when risk rating and adverse selection are possible. Compared to perfect regulation administered with perfect regulation, or even to the wise and prudent regulation that occasionally happens, these alternatives may still leave something to be desired. But compared to the kind of regulation we have had or can generally expect to have, they at least deserve equal billing and equal consideration. These alternatives may work better if some other government actions are curtailed and some modest regulation applied to encouraging the alternatives.

To be specific: one might suppose that, as is often the case, policymakers must choose between two undesirable outcomes—unfair risk rating or inefficient community rating—in order to deal reasonably well with risk variation. New developments in research (Pauly, Kunreuther, and Hirth, 1995; Cochrane, 1995) suggests that, in theory and in fact, in many circumstances realistic competitive insurance markets can avoid much of both bad situations, and that a relatively modest amount of public intervention can deal with the cases that fall through the remaining cracks. The fundamental reason for this market behavior is that potential insurance consumers also dislike the more negative aspects of either kind of behavior, and competitive insurers have developed methods to avoid them.   The fundamental reason for the political behavior is that some policymakers have already developed some well-tailored solutions that leave the market intact but rein in the worst cases.

The three kinds of "solutions" to which I want to draw your attention are (1) *guaranteed renewability at uniform premiums, (2) group insurance,* and *(3) high risk pools.* Because the first is much less well understood than the other two, I will discuss it in more detail, but I will also comment on the other two devices.

The great majority of people who are high risk today were not sicker than average at all times in their lives. Data shows what common sense tells us: even people who are in excellent health have higher medical expenses on average as they age, and some pick up chronic conditions. The age-related part of increasing risk is perfectly predictable; what is not predictable is the random onset of a chronic condition that makes a person high risk not only initially but for some time to come, possibly for life.

Most medical expenses for people under 65 are not related to chronic conditions; they come from the "bolt-from-the-blue" event of an accident, a stroke, or a complication of pregnancy that we know will happen on average but whose victim we cannot (and they cannot) predict well in advance. This is precisely the kind of low probability, high cost event for which insurance works extremely well as a device for substituting a smaller certain payment for an unexpected rare but large payment. Sometimes, however, what strikes unexpectedly is a condition from which the person is unlikely to recover rapidly; such random but then chronic conditions make future medical expenses higher for people who have them. If insurance premiums were proportionately risk rated

to the risk prevailing for the next year (the usual time period for health insurance), people who are well today and have no chronic conditions at the moment would face the chance of contracting such a condition with two bad financial outcomes. Not only is diagnosis usually associated with high immediate medical expenses, it would also be associated with a sudden and serious jump in premiums.

Risk averse people should want to have protection not only against high current period expenses but against the unexpected onset of a condition that might entail high lifetime premiums; they would seek protection against "the risk of becoming a high risk." In some real world health insurance markets such protection exists and was quite common even in the absence of regulatory rules. Specifically, most health insurance policies bought on an individual basis contained a provision also common in individual term life or disability insurance: guaranteed renewability at class average premiums. With this provision, the insurer promises not to single out insureds whose risk has increased more than average for high premiums when they renew their coverage. Instead, they are to be charged the same premiums as are charged to everyone else who was in the same initial (usually low) risk class as they and bought the same type of coverage. Administering such a guarantee is easy for an insurer: it promises to base its future premiums only on whatever information it collected about risk when it initially sold the coverage; it promises not to revisit the question of risk based on new data that might be obtained from the person or even based on the claims history data that the insurer has; it promises not to "re-underwrite." This

provision does not guarantee constant premiums; premiums can rise if expected medical expenses rise for everyone in the risk class (say, because of higher medical prices), and premiums may rise according to a schedule specified in advance as a function of perfectly predictable things, like growing older. But the person with coverage with this feature is protected against the bad luck of becoming riskier than average, and therefore will not pay a higher premium on becoming a high risk. This feature is not free, of course; policies that contain it must have higher initial premiums ("frontloading") than would premiums for a policy for which the insurer retained the right to increase premiums for people who contracted a chronic illness. But it is easy to see why rational, foresighted people would prefer the slightly more expensive but surer policy to the cheaper but riskier one.

Federal law now requires states to ensure guaranteed renewability for individual (but not group) insurance policies. But even before the spread of such state laws, industry observers estimated that about 80 percent of policies voluntarily (on the parts of both buyers and sellers) contained such provisions. (Pauly, Percy, and Herring, 1999) There is, however, considerable debate about how they work in practice, debate which is assisted by the absence of nationwide comprehensive data on practices in insurance markets, especially in the individual market, so that evidence tends to consist of anecdotes and problematic surmises. There certainly have been cases in which insurers were caught engaging in re-underwriting even when they were forbidden to do so, and a

number of state insurance departments have said that they would prohibit risk rating at renewal even in the absence of specific state law under their general authority to limit arbitrary and excessively discriminatory premiums (Patel and Pauly, 2003). Some insurers are said to have gotten around the requirement to continue to cover high risks by raising premiums for all insureds so that all drop out of the risk class, and then selectively re-enrolling only those low risks who have not been put off by this behavior. Insurance brokers and agents insist that they pay attention to this kind of behavior and steer customers who come to them for advice away from insurers who engage in semi-shady practices. We know that this feature does not work perfectly everywhere for everyone, but how well does it work on average?

Research has provided some data that is highly consistent with guaranteed renewability generally operating as the theory and the intent of the contractual provision suggests (Pauly and Herring, 1999; Pauly and Herring, 2001). This finding is striking enough that it deserves to be emphasized even beyond the issue of guaranteed renewability. To be specific*; there is very strong empirical evidence that the premiums higher risk insureds pay are much lower than would be consistent with proportional risk rating.* Stated slightly differently, while high risks do pay higher premiums than low risks, the increase in premium with risk is much less than proportional to the increase in risk.

This result has been obtained in a large number of studies using large nationwide data sets from different time periods. Depending on the measure used of risk, the "elasticity of premiums with respect to risk" in multivariate analysis of data ranges from about 20 percent to less than 50 percent; never higher. That is, a person whose risk is twice has high as average will pay a premium only 20 percent higher. Table 1 shows more intuitive evidence for this proposition. It uses data from the late 1980s before there was widespread premium regulation in the individual insurance market or requirements of guaranteed renewability, but when that feature was common nevertheless. The risk level for a person in the data set is characterized by the person's age, gender, location (to measure differences in medical cost), and pre-existing chronic conditions. Statistical models were used to relate the actual medical expenses, and the actual insurance benefits received for each person, to that person's values for these variables; the estimate of risk for that person is then the "predicted value" of their medical expenses (that is, the average medical expense for a large number of people with the same values for these characteristics as they). Those risk estimates were then used to select a sample of people with individual health insurance expected to have medical expenses in the top 10 percent of possible values of risk, and another sample of people in the bottom half of those values. As the first line of the table shows, the expected expenses, the actual average expenses, and the actual average insurance benefits were much higher for the high risks than the below-average risks. The average benefits for the high risks were 11 times greater (at $2054 per person) than for the lower risks (at $187).

The premiums were higher for the higher risks too, but the key point is that the premium for these very high risks (at $1150) were only 1.4 times greater than that of the low risks ($825); there was a substantial amount of averaging of risk in the premium structure.

While there are doubtless many causes for this phenomenon, one of them probably is guaranteed renewability. People with such provisions would not be paying premiums that were higher than average because they became higher risks. Of course, some people in the data were new purchasers of insurance whose premiums would be risk rated, but apparently by no means all. There is even stronger evidence. We looked at how premiums and risk varied with age for similar policies. Insurers certainly can determine a buyer's age, and they certainly can determine that, other things held constant, expected expenses and benefit payments will rise with age (especially for men). What we found, however, was that the premium paid by the average older man was only about 40 percent higher than that for the average younger man when the expected expenses differed by a factor of two to one. But this pattern of overpayment relative to expected expense for the younger people who would generally be the new buyers of insurance is exactly the frontloading that would be predicted to arise under guaranteed renewability (but that would be unstable in competitive insurance markets under proportional risk rating). We have further examined the path of premiums and benefits with age in this market and find that it corresponds rather well with the path that would be consistent with guaranteed

renewability. In doing this analysis, we adjusted for the fact that people often do not keep their individual coverage from a given firm but drop it because they have taken a job that carries coverage or because they switch insurers. Because the low risks have already prepaid their contribution to the high risks, their dropping out does not cause any problems for the ability of insurers to continue to maintain protection for higher risks. Some high risks do drop out as well but, as expected, at a much lower rate.

In our analysis of individual insurance data we found that only the locational and demographic variables were consistently related to higher premiums. The person's health status when they bought insurance (measured by the presence of a pre-existing chronic condition) was not statistically related to premiums, but the scarcity of observations on people with such conditions means that our estimates are themselves necessarily imprecise, Jack Hadley and James Reschovsky (2003), using a different risk measure (contemporaneous health status) and a more sophisticated but somewhat delicate statistical technique, did find that people in poorer health paid higher premiums, but even there the increase in premiums was much less than the increase in risk. I therefore conclude that individual insurance markets (even when they were unregulated) provided a substantial amount of protection against the adverse effects of risk rating to people who did what we want them to do—bought insurance before they

became high risks, and stuck with their insurance rather than becoming uninsured.

Risk rating can only occur if insurers can determine risk levels; under perfect risk rating, there can be no adverse selection. However, in a world in which buyers of insurance may sometimes know more than sellers, it is interesting to note that guaranteed renewability provides potentially important protection against adverse selection. If people buy this coverage early in life (as they should to take advantage of the provision), they are likely to be much more similar in risk levels than they will become later on. And since it is rational for the people who remain healthy to stay in their original policy where they have already made transfers to those in their cohort who became higher risks, it is less likely that they will drop out and start a death spiral. Finally, if those who remain lower risk do drop out or are lured away, because they have already prepaid their transfer to the high risks, the insurer does not need to raise premiums to the high risks.

We have investigated some of the other reasons why higher risks pay premiums that are less than proportional to their relative risk levels. There is evidence that higher risks search more intensively to find a premium that is low relative to the expected benefits; it makes more sense to check out many insurers (or use a broker to do so) when one is paying $400 a month for insurance than when one is 25 and paying less than $100 a month for insurance (Pauly, Herring, and Song, 2003). And it probably is true that some risk factors, like the decision on

the timing of the next child or the repair of an old football injury, is better known to the insured than to the insurer. But this phenomenon may be partially offset by the fact that insurers actually have more accurate data on risks than typical insurance consumers do.

Another feature of insurance that can protect against uncertain jumps in premiums and adverse selection is group insurance. The great bulk of Americans obtain their health insurance as group insurance related to their employment. Probably the main reason they do so does not have to do with any risk variation factors, but rather to the substantial tax subsidy to workers (*not* to employers) present in the exclusion of compensation received as health benefits from income and payroll taxation. But group insurance probably does have some features that deter the kind of behavior theory was earlier said to predict.

 Most simply (but not most obviously), group insurance offers a much better deal for your money for a given policy than does individual insurance. The difference between the premium one pays and the benefit one should expect on average to get in group insurance is lower than for individual insurance both because of economies of scale associated with group purchasing (especially lower selling and billing costs) and because of the tax subsidy. These features in effect may make insurance such a good deal for the wealthiest low risks (who get the biggest tax subsidies) that they will not be motivated to drop coverage and start a death spiral even if their premium is not properly tailored to their risk. As long as

a low risk's net premium is low enough after the tax benefits are taken into account, the fact that there is some cross subsidy to higher risks may not matter.

A more complicated issue is whether or not employment-based group insurance in some sense "pools risk" more than other arrangements. For large groups, there is no explicit individual underwriting, but the cost of that function is only a tiny fraction of any insurer's administrative cost. There can be variation in premiums with risk across small groups; a firm of three 25-year-olds in good health will pay much less than a firm of three 60 year olds who are out of shape. Moreover, the requirement that one be able to work to qualify for one's own employment based insurance serves to automatically screen out the highest risks and those unable to take a job because they are caring for a dependent with high risk. But the key determinant of access to insurance and net payments for insurance is the policies employers follow with regard to this benefit.

One thing that employers are motivated to do is to try to keep as many of their employees in the insurance plan as they can, because the premium, or even the availability of group insurance, depends on the participation level of workers in the firm. Let too many of them drop out, and the group insurance may not be offered by an outside insurer. Even self-insured employers (who cover the majority of workers nowadays) want to achieve economies of scale. Thus employers should want to avoid death spirals and widespread non-participation.

Probably most importantly, workers in group insurance almost never pay an explicit total premium that is related to their precise risk levels; they almost always all pay the same employee premium if they choose the same policy for the same-sized household unit. (There *is* explicit risk rating for the higher risk associated with having more people covered under a family policy relative to an individual policy). However, economists believe that workers pay for the bulk of their group insurance not through explicit premiums but through lower wages, and generally money wages are not explicitly adjusted based on an individual employee's risk level.

The evidence does, however, strongly suggest that worker wages are adjusted to some extent to reflect the different cost of insurance as a function of risk (Pauly and Herring, 1999; Sheiner, 1994). Wages vary by seniority, and more senior workers are usually older. What we found was that, other things equal, wages increased significantly less rapidly with seniority for workers who obtained job based insurance than for those who did not; we interpret this as the effect of higher insurance costs taking away some of what would have been the usual raise associated with more experience and seniority. Moreover, common sense tells us that an employer cannot take the typical $6000 "employer contribution" out of the wages of younger workers and still expect to compete to hire those workers with other firms that offer higher cash wages and no coverage.

There is no evidence that wages vary with health status given age and gender (though the lower wages of women could in part reflect their higher medical costs). But remember that with guaranteed renewability, premiums in individual insurance also need not vary with health status. Thus I would conclude that the amount of risk pooling in group insurance is at best only very modestly greater than in individual insurance on average. The difference would be greatest between a high risk person able to get a job at a firm that offers benefits and what that person would be charged as a new applicant for individual insurance. But the job with insurance is by no means assured to a high risk, and the typical buyer of individual insurance is renewing, not buying new, so this difference tends to average out to a small number if it is present at all.

The main virtue of group coverage in terms of risk variation is not risk pooling per se but rather that it discourages adverse selection. It does so in several ways. Most obviously, the range of insurance choices a person has within a firm is usually much smaller than the range of choices in individual insurance, and any opportunity to choose less generous coverage (whether it is a high deductible plan or a cost constraining HMO) offers a chance for low risks to separate themselves out. The downside of this advantage is less choice, but firms and their workforces are free to make this choice not to have many choices.

Equally if not more important is the fact that the worker who chooses to decline group insurance while remaining in the firm almost never recaptures the full

premium for that coverage. Instead the worker will get back any employee

premium and (in some firms) a small bonus for refusing coverage, but that

reward is almost always much less than the value of the insurance even to a low

risk. We do have a problem with more workers offered employment-based

coverage rejecting it, especially as the average explicit employee premium has

risen, but there are almost no cases where rejecting coverage to save the

employee premium would be rational behavior if the person thought that without

coverage they would have to pay for all of their medical care out of pocket.

(They might drop and expect to rely on family assistance or charity care, and the

still tiny fraction of people offered coverage who reject it may just be the minority

of any population who are irrational or unthinking.)

So there is very little total dropping out by lower risks, but do they inefficiently

drop back to less generous coverage? Not necessarily, because employers can

if they wish control adverse selection and risk rating. The simplest way to do this

is to offer only one plan. But even when employers offer several plans, the key

to controlling selection is to properly set the difference in employee premiums (or

in the contribution to spending accounts) across plans (Cutler and Reber, 1998;

Pauly and Herring, 2000). If employers foolishly make the premium much lower

for the less generous plans, then all but the highest risk will join them, leaving the

few remaining high risks in a more generous plan. But research shows

employers how to calibrate the premium difference to reflect the premium cost

reduction associated with the low risks (not the average and certainly not the

difference in expected benefits when the low risks have already sorted into the less generous plan). So employers who want to control adverse selection can do so to a considerable extent (though not perfectly), especially if they self insure all of the plans they offer. Things are somewhat more complex if multiple outside insurers are used and those insurers are not given the data they need to estimate the risk levels of the people who will choose their plans. Risk adjustment of the total premium the insurer gets combined with appropriate setting of premium differentials will prevent adverse selection if that is an employer goal.

Research (Pauly, Percy, Rosenbloom and Shih, 2000) suggests that some employers try to limit the choice of options and set the premiums to control adverse selection, while others take the view that any redistribution away from older workers in their health plan offering is probably offset from redistribution toward such workers in their pension plan or in other benefits, and that the total amount of redistribution (and inefficiency) is small. As long as the least generous plan offered is still a decent plan even for higher risks, there probably need be little policymaker concern about adverse selection in group insurance. Personally I would only be concerned about offering a health savings account type plan to very low income workers, or offering a very restrictive HMO to workers who would react strongly to limits on access, but I would not be much concerned in general.

How does the rate of takeup of insurance vary with risk level in group and individual insurance? Are higher risks more likely to have coverage than lower risks (which would be consistent with adverse selection), are they less likely (which would be consistent with very strong risk rating), or is coverage nearly universal and independent of risk (which would be ideal)?   Research on this subject is far from definitive. Studies that have looked at people in households where someone is a full time employee (and therefore potentially eligible for group insurance if the person chooses or is able to get a job at a firm offering coverage), the strongest and most consistent finding is that the size of firm in the industry or occupation of the worker is by far the most important predictor of having coverage (along with the size of the tax subsidy and therefore income) (Pauly and Herring, 2000). People who work in industries dominated by larger firms are much more likely to end up with coverage than those who work in small firm industries. The relationship of coverage to risk, given firm size, is less well understood. What we observe seems to depend on what measure of health risk we use. If we use chronic conditions as the measure, employed higher risks are more likely to have coverage than employed lower risks. If we use self reported health status, coverage may be less likely for high risks. Analysis of the late 1980s data showed that high risks were significantly less likely to have group coverage only if they were low income people working in small firms, but not otherwise (Pauly and Herring, 1999). There is little evidence that employers in general have difficulty in continuing to offer coverage to people who become high

risks, and no evidence at all that they have problems with people who have unexpected high expenditures.

It is much harder to determine how risk levels affect the likelihood of having coverage in the individual market because anyone can participate in that market, but most people do not do so and instead obtain group insurance. We have looked at people in households where no one is a full time employee—the household's income comes from self employment, part time work, or non-work sources. The relationship here depends even more on the measure of risk. Len Nichols and I (2002) found that if we measure risk by age, controlling for income, older people in "non-group" households were much more likely to have individual coverage, despite higher premiums, than younger people.   We also found that people with chronic conditions were more likely to have coverage, although the relationship was not as strong. On the other hand, when risk is measured by self reported health status, people who label their health as fair or poor are less likely to have individual coverage controlling for income; this is the opposite of adverse selection. One puzzle in the data is that many of those with insurance, who say that no one in their household works full time, still list themselves as having obtained group insurance coverage; there is no clean division of the population between those with access to group insurance and those who
must use the individual market.

Precisely for this reason one should be very cautious in trying to draw conclusions about the comparative performance of individual and group insurance markets. If I was forced to do so, I would conclude that there may be differences in the likelihood of obtaining individual insurance coverage by people who are very high risks when they seek coverage, but that if the group market does better, the differences are small, and are limited by the fact that many very high risks do not have access to employment-based insurance. It would be nearly impossible to provide those currently without a group option access to that option on the same terms as the current users. I think the differences in the extent to which net premiums do (or could) vary with risk are small, and any stronger relationship in the individual market is attributable to its small size and marginal or add-on character. For example, a person who had group insurance, who contracts a high risk condition, loses their job and insurance, and uses up their COBRA coverage, will be recorded as a high risk trying to buy new coverage in the individual market. But one could argue that placing the person in that situation is as much the fault of the link between tax subsides and group insurance which does not provide guaranteed renwability protection to individual workers as it is the fault of individual insurance.

Fortunately, there is a device available to pick up the pieces without requiring the imputation of blame: high risk pools. I do not intend to discuss the actual working of these pools in detail. Instead I want to point out that the concept of having a subsidized, decent though limited coverage policy available to high risks unable

to obtain or retain individual or group coverage makes great sense as a safety net. Since the number of high risks is by definition low, it avoids having to distort insurance markets for the great majority who are not high risks in order to make transfers to a few unlucky people. Some of the more anecdotal research shows that almost any risk can obtain individual coverage if they persist at searching long enough, but those who have already been rejected or quoted very high premiums perhaps ought to have another option than spending their time with insurance brokers. In idealized concept, a high risk pool ought to offer coverage at premiums somewhat higher than those charged for good risks but still at reasonable levels to people who have tried and failed to obtain coverage on their own. The financing of these pools should be generous enough to accommodate those who need to use them, and that financing should be raised by general revenue taxation, not by requiring insurers to contribute and thus raising premiums which drive more people out of regular insurance. The terms of coverage (premiums, type of coverage) should be only moderately attractive, because we want to preserve incentives to people to obtain voluntary coverage before they become high risk, rather than wait to pick up attractive subsidized high risk coverage when and if that happens. I am hopeful that it is possible to design a plan that walks this fine line and still preserves an opportunity for people to obtain coverage that will give them financial protection and access to care. Coordinating high risk pools with guaranteed renewability provisions would seem to be desirable.

To sum up: the important problems with private health insurance in the United States are not associated with the risk variation-risk segmentation issues that are so prominent in insurance theory and many policy discussions. Our problem is not the insurance is expensive and unattractive for high risks; it is that in some cases it is expensive and unattractive for all risks. It is true that the largest single segment of the uninsured population is low risk healthy twenty-somethings, and some adverse selection in group and individual insurance may modestly contribute to this. But I believe that a much larger contributor is the absence of generous subsidies and the absence of marketing efforts targeted at this group; there may actually be too little effort at cream skimming those low risks who remain uninsured.

This is especially the case for people who are discriminated against by being ineligible for generous tax subsidies when they buy insurance (the non-self-employed in the individual market) and those who could have access to products with lower across-the-board administrative costs but do not currently have such access. Finally, the key background issue of what if anything we want to do when premiums are rising not because of insurance market behavior but because medical care is becoming both more costly and yet much better should really be front and center in the policy debate.

TABLE 1

EXPENSES IN NONGROUP INDIVIDUAL COVERAGE, BY RISK

(EXPECTED EXPENSE)

	Bottom 50 %	Top 10%
Actual benefits	$187	$2054
Premiums	825	1150
Actual expenses (total)	555	3504

Source: Pauly and Herring (1999), based on 1987 NMES data.

REFERENCES

J Cochrane, " Time-consistent Health Insurance," *Journal of Political Economy*, 1995, 445-473.

D Cutler and S Reber, " Paying for Health Insurance: The Tradeoff between Competition and Adverse Selection," *Quarterly Journal of Economics,* 1998, 433-466

J Hadley and JReschovsky, unpublished paper, 2002.

L Nichols and M Pauly, "The Non-group Health Insurance Market: Short on Facts, Long on Opinions and Policy Disputes," *Health Affairs Web Exclusive,* October 23, 2002

V Patel and M Pauly, "Guaranteed Renewability and the Problem of Risk Variation in Individual Health Insurance Markets," *Health Affairs Web Exclusive,* September 8, 2002.

M Pauly, "The Welfare Economics of Community Rating," *Journal of Risk and Insurance,* 1970, 407-418.

M Pauly, H Kunreuther, and R Hirth , "Guaranteed Renewability in Insurance," *Journal of Risk and Uncertainty,* 1995, 143-156.

M Pauly, *Health Benefits at Work,* University of Michigan Press, 1997.

M Pauly and B Herring, *Pooling Health Insurance Risks,* AEI Press, 1999

M Pauly, A Percy, and B Herring, "Individual versus Job-based Insurance: Weighing the Pros and Cons," *Health Affairs,* 1999, 28-44.

M Pauly, A Percy, J Rosenbloom, and D Shih, "What Benefits Specialists Think about Medical Savings Account Options for Large Firms, *Benefits Quarterly,* 2000, 39-46.

M Pauly and B Herring, "An Efficient Employer Strategy for Dealing with Adverse Selection in Multiple-plan Offerings: An MSA Example," *Journal of Health Economics,* 2000, 513-528.

M Pauly and B Herring, "Expanding Health Insurance Through Tax Credits: Tradeoffs and Options," *Health Affairs,* 2001, 1-18.

M Pauly and B Herring, "Premium Variation in the Individual Health Insurance Market," *International Journal of Health Finance and Economics,* 2001, 43-58.

M Pauly and B Herring, "Incentive-compatible Guaranteed Renewable Health Insurance," NBER Working Paper 9888, July, 2003.

M Pauly, B Herring and D Song, "Health Insurance on the Internet and the Economics of Search," unpublished paper, 2003

M Rothschild and J Stiglitz, " Equilibrium in Competitive Insurance Markets," *Quarterly Journal of Economics,* 1976, 130-149.

L Sheiner, "Health Care Costs, Wages, and Aging: Assessing the Impact of Community Rating," Federal Reserve Board, December, 1994.