TIME, RISK, PRECOMMITMENT, AND ADVERSE SELECTION IN COMPETITIVE INSURANCE MARKETS

MARK V. PAULY

TIME, RISK, PRECOMMITMENT, AND ADVERSE SELECTION IN COMPETITIVE INSURANCE MARKETS

Abstract

This informal paper explores models of competitive insurance market equilibrium when individuals of initially similar apparent risk experience divergence in risk levels over time. The information structure is modeled in three alternative ways: all insurers and insureds know risk at any point in time, current insurer and insured know risk, and only the individual knows risk. Insurers always know the average risk. It is shown that some models lead to "backloading" of premiums in which initial premiums are less than initial period expected expense, and that other models lead to "frontloading" of premiums and policy provisions of "guaranteed renewability." Finally, it is shown that guaranteed renewability greatly reduces the possibility of adverse selection.

JEL Classification: D8, G2.

Mark V. Pauly
Department of Health Care Systems
The Wharton School
University of Pennsylvania
Room 208 Colonial Penn Center
3641 Locust Walk
Philadelphia, Penna. 190104
USA
pauly@wharton.upenn.edu

Introduction

Many of the risks an individual faces evolve and change over a lifetime. Not only do risks change, they change in different ways for different people, so that the distribution of risks becomes more diverse over time. That is, while individual risks of death, of illness, and of wealth reduction do differ to some extent at young ages (because of inheritance of wealth and because of observed aspects of inheritance of genetic makeup), for the most part it is the passage of time that is associated with the evolution of different risks for individuals changing in different ways or becoming gradually more apparent throughout life. Moreover, the knowledge of a risk is dependent to some extent on its longevity: old risks that remain with an individual or the individual's assets can become better known than risks reflecting changes that are recent or yet to occur.

The passage of time thus presents the conflicting trends of greater ability to measure existing risk differences and greater opportunity for risks to change. So how might or does insurance work for a set of agents who begin their (economic) lifetime with roughly similar perceived risks but gradually experience (as they age) greater divergence in the risk of future adverse events? Basic insurance theory tells us that competitive insurance markets can work efficiently in dealing with adverse events whose probabilities are known to all and stable over time. How do markets work when it is certain that risks will change according to a predictable pattern but it is uncertain which persons will experiences those changes at which time?

In this keynote paper, I informally review some recent research that develops the theory for how insurance markets develop methods to deal with risks that evolve over time. I specifically discuss guaranteed renewability as one such device, which not only deals with what Kenneth Arrow called "insurance with a long time perspective," but also goes much of the way toward solving the problem of adverse selection.[1]

Information Structures

In this paper I distinguish three different possible structures for risk information for insurance consumers and firms. One case, which I will save for later, is the situation in which at least some of the information is known only to the consumer and to none of the insurers; this is classic adverse selection. The second case is one in which all insurers know what consumers know about risks at every point in time; this is the case that leads to risk rating, in which premiums vary with risk and change when (perceived) risk changes. The third case is one in which the consumer and the consumer's "current insurer" (the one sold insurance to the consumer in the last period) have the same information about the consumer's risk, but other insurers have less complete information. I begin with a special version of this third case.

Models of Insurance With Constant But Initially Unknown Risk

The first set of models are ones in which risk "truly" differs even in the first period but in which information that establishes theses differences takes time to accumulate. We assume that, initially, the (representative) consumer and insurer do not know the consumer's particular risk level, but that time permits the accumulation of evidence, equally available to both, but hidden from others insurers, that leads to more precise estimate of individual risk. The classic example here is automobile insurance. Even if there are (intrinsically) "good drivers" and "bad drivers," newly licensed drivers will not be able to be categorized, and so would initially be assigned, by consumer and insurer alike, the average risk (which, of course, could be higher for new drivers than for experienced drivers). As time passes, however, events occur (or do not occur) which help to sharpen estimates. Accidents happen, and traffic citations are accumulated, which help to distinguish (though not perfectly) between drivers of different risk levels. A key assumption in these models is that at least some of this driving record (the frequency of claims) is known to the consumer and to the consumer's current insurer, but not to other potential insurers.

Given these assumptions, models generally indicate that insurers will underprice products initially in order to collect this private knowledge, but will eventually earn profits on those they identify as low risks.[2] (Other models indicate that the effect of insurance on premiums may be different for those consumers who initially represented themselves as low risk than for those who represented themselves as high risks; in this

discussion I ignore the possibility that the consumer has some influence over the initial classification.[3])

Those who turn out to be high risks will engage in adverse selection but will do so against other insurers whose information on the person's driving record is imperfect. Insurance premiums are "backloaded," in the sense that the ratio of the premium to current period expected expense rises over time, on average. Moreover, this equilibrium may be inefficient, because buyers will be faced at some point with premiums not properly reflective of risk.

Models of Insurance with Changing Risk

The bulk of my discussion will apply to a second case which has recently become a topic of discussion: a situation in which a consumer's risk level is perfectly known to all at any point in time, but which may change in the future in unpredictable ways. Health insurance is the most common example, in which the onset of a previously unsuspected high cost chronic condition changes the person's risk level in the future. This phenomenon also potentially applies to other insurance in which the risk depends on individual health, such as life insurance and nursing home (Long Term Care) insurance.

In this case, I will argue (in more detail) that the optimal insurance is *frontloaded*, with initial premiums above actuarially determined values. This argument is based on papers by Pauly Kunreuther and Hirth (hereafter PKH); Cochrane, and, more recently, for the specific case of term life insurance without formal guaranteed renewability, Hendel and Lizzeri.[4] I provide a simple example to illustrate the main idea of these papers. I further show that, in some circumstances, the equilibrium premiums are efficient, solving the problem of *both* "premium risk" *and* adverse selection. I also summarize more recent empirical evidence suggesting that actual markets may approximate this ideal rather well, despite some threats to equilibrium and efficiency. Finally, in addition to explaining in more detail what some recent research tells us (theoretically and empirically) about the changing risk model, I will offer some thoughts on what happens if the two models are combined.

A population is assumed to begin its life cycle with equal apparent risk, but events occur over time which in part measure the initial (but unknown) differences in risk and in part signal changes in risk. Will the final market equilibrium be one with frontloaded, backloaded, or actuarially based premiums? And can we say anything about its efficiency?

Guaranteed Renewability as a Solution to Insurance Market Problems

I first explain and summarize with a simple example some recent research which demonstrates a solution to the problem of risk variation. Suppose each of a set of similar persons is at risk of an adverse event per time period with probability p. Suppose there are two periods. If a loss occurs in period 1, those persons who incur the loss ("high risks") are assumed to have higher loss probability p´ in period 2; if a loss does not occur, the probability in period 2 remains at p for those who continue to be "low risks." (The onset of chronic illness would produce roughly this pattern.) Full coverage insurance (with benefits equal to the loss) is offered by competitive insurers, and every insurer knows each person's loss probability.

Ignoring administrative costs, the breakeven premium for any insurer in period 1 is p. In period 2, in the absence of anything else, insurers would charge p´ to those high risks who had suffered a loss in period 1 because of a continuing chronic condition, and p to those who had not. A uniform premium of $pp' + (1-p)p = \bar{p}$ would also cover losses in period 2. However, any insurer who tried to charge $\bar{p}$ would only be selected by the high risks; the low risks would defect to an insurer charging them just p.

Risk averse consumers would presumably prefer to avoid the risk of having to pay a high risk premium in period 2. An insurer could guarantee not to do so, still attract the lower risks, and (ignoring interest) just cover its cost by charging a premium of $p + p(p' - p)$ in period 1 and a uniform premium of p in period 2. That is, the insurer would provide *guaranteed renewability (GR) at nondiscriminatory premiums* for period 2. It is easy to see that such a schedule would be incentive compatible in the sense that (a) all consumers would be willing to pay the period 1 premium and (b) all consumers would be willing to pay the period 2 premium. Clause (a) is true because all insureds initially assume themselves to be the same risk and would

prefer to pay the two period premium to bearing the risk of losses in period 1 and uncertain premiums in

period 2. Clause (b) is true because both high and low risks would be willing to pay the low risk premium

in period 2. The insurer would break even since its per person premium collections over both

periods, $p + p(p' - p) + p$, would equal its expected losses over both periods in either case $(p + pp' + (1-p)p)$.

(That is, both revenue and cost equal $2p + pp^i - p^2$.)

This is a specific example of a set of incentive compatible premium schedules proposed by PKH and by

Cochrane.[5] (There are relatively minor differences between the two models in the extent of "lock in" to a

given insurer, depending on what one assumes about whether a high risk consumer is permitted to and is

able to recover the "excess" prepaid premium; there are administrative advantages and disadvantages of

either approach.)

Making GR Theoretically Bullet-Proof

The policy provision called "guaranteed renewability" (GR) means that the insurer agrees to renew a policy

at the next policy anniversary period at a premium which is non-discriminatory. Such a provision is

actually common and well known in insurance practice (if not to insurance theorists). Under this provision,

future premiums can change, but they must do so to the same extent for all in the (initial) rating class.

Administering this provision is in practice almost surely less costly than continual individual risk rating. It

means that the insurer need determine only the new (average) premium for the coverage purchased by

people in a risk class for the next time period, and does not need to assemble or analyze data on the claim

or risk levels of individual members of the class. In a practical sense, individual "re-underwriting" is costly,

but failure to re-underwrite is free of cost. The life insurer need not administer another medical

examination, the health insurer need not assemble data on individual health claims, and the long term care

insurer need not inquire about chronic conditions or frailty.

If some insurers engage in guaranteed renewability as a policy feature, is this behavior consistent with market equilibrium? The only conceptual treatment of this subject that I have found before the Cochrane and PKH papers was in Peter Diamond's 1991 Econometric Society Presidential Address.[6] After noting that guaranteed renewability was common in individual term life insurance (but failing to note that it was also common in individual health insurance), he says (correctly) that "the restriction on price changes coming from this contract clause will depend on the marketing strategy being followed by the company for the set of current insurees"—even under GR a company could stop selling insurance altogether, for example. But he also argues that "this guarantee could be of little value." His example is that of a life insurance company that initially charges the same rate to smokers and non-smokers. "Then some companies introduce new policies only available to nonsmokers. By evaluating new applicants, unhealthy non-smokers are excluded from buying these policies. All healthy non-smokers switch to this new class of policies," while the unhealthy non-smokers are now lumped together with the smokers in a higher premium policy. The problem with this example is that if the non-smokers were all equally healthy initially, the premium that would be charged to all non-smokers in each future period by their original insurer would be no higher than the premium a new insurer would propose to charge to those non-smokers who remained healthy. There would be no "new applicants." That is, under incentive-compatible GR, the premium charged to the healthy in any period makes renewed insurance purchase by them as or more attractive than switching to a new company or a new policy.

To be sure, in all periods (but the last) the premium charged to the healthy for insurance for next year would be higher than a premium charged by an insurer who just wanted to offer one year of coverage without GR. But risk averse people would prefer the higher one period premium with GR protection to the lower non-GR premium that buys one year of coverage but leaves the person exposed to the chance of a substantial jump in premium for the years following. This is the meaning of incentive compatibility.

The more serious threat to GR is whether a contract that delivers what is promised can be written and enforced. In one sense, any solution here is bedeviled by the usual problem of long term contracting in economics. Once payment has been made for some types of costly long term contract provision, it is

always to the advantage of the agent who would have to pay to try to renege on the contract—and it is usually impossible to write a contract so detailed that there is no possibility for the agent bearing the cost to shirk or welch on some part of the deal. The standard solution to this problem (despair being disproved by the existence in the real world of long term arrangements and contracts that seem to work) is to invoke some assumption of either repeat dealing or reputation that serves to force both agents to live up to the contract. Our empirical research in fact suggests that insurance brokers know which insurers play games at renewal time, and either advise their customers to avoid them or at least to pay a lower price.[7]

The other solution is a kind of second-best one: as long as the major or most important features can be specified in a contract, the outcome can be pretty good, even if not perfect. To be specific: it is important for buyers to be able to anticipate what the path of future premiums will be and to be assured that future premiums will change only if expected expenses truly do change.

For life insurance this is not a serious problem. Benefits are stated in nominal monetary terms, and there is usually good agreement on what death probabilities are represented in the mortality tables. In contrast, health insurance premiums in the future depend in part on how frequently people get sick, but even more on the cost of treating them, and that cost is affected by new technology, health worker wages, and the extent of insurance coverage itself. In the face of medical spending growth, Diamond argues that "premiums can rise enormously" and (he might have added) unpredictably. The question then is: to what is the future GR insurance premium indexed?

This is a serious problem. Suppose a company that sold GR health coverage raised its premiums (equally) for all of those in rating class, those who had contracted chronic conditions and those that had not. If it raised its premium by less than the increase in market-wide insurance cost, it would suffer losses. If it raises its premium by more than market wide insurance costs rose, it (or another insurer) could propose to charge a premium based more accurately on those costs for a new rating class that included only good risks that the good risks would find attractive. That would leave only the bad risks in the original pool, and the company could and would raise the premium it charges (non-discriminatorily) to those remaining in the

poll. In effect, it could pocket the excess premiums it collected in previous years, by driving off the good risks.[8] In doing so, it would not have violated any terms of the GR feature. Here again, repeated play of this strategy by an insurer would eventually bring recognition and retaliation, but it does appear that GR contracts are fragile. Such pricing mistakes, whether intentional or unintentional, make GR less valuable as a solution than the theoretical models would suggest. Moreover, consumers would have difficulty determining whether a given increase was or was not justified on the basis of cost, and thus difficulty detecting the kind of strategy just described. I will deal further with strategies that could and actually have been used to deal with this problem when I discuss the behavior of insurers in practice.

"Sure, it Works in Practice, But Can it Work in Theory?"

Despite the prominence in insurance economics of theoretical arguments about problems raised by adverse selection and single-period risk rating, guaranteed renewability (in some form or other) exists in many insurance markets. It continues to survive even when not required by law or regulation, and therefore (by the usual Darwinian model of positive economics) must be performing some useful function in some way. Our recent empirical research has developed some evidence on how guaranteed renewability works in actual individual health insurance markets in the US, which are relatively lightly regulated in most states, and so it is perhaps worthwhile to summarize that research.

While insurance with GR is theoretically "front-loaded," the risk of death or illness rises with age even for apparently healthy people. This means that, even with guaranteed renewability, premiums would be expected to rise with age—just not as steeply as in the absence of GR.

Our recent research has examined the way in which premiums for individual health insurance vary with age. One empirical discovery, which turned out to be highly consistent with the hypothesis of effective GR coverage, was the finding that, premiums rose with age at a much lower rate than did expected (average) expenses. That is, relative to a premium that would cover current period costs—younger buyers of

insurance were being "over-charged" much more than older buyers.[9] But this is exactly the pattern we would expect to prevail under GR!

Following this hint, we examined the actual path of premiums with age more directly. Using data on the probability of contracting a costly chronic condition at any age and on the persistence of high spending, we developed a model to predict or estimate the time path of premiums that would prevail under "optimal" GR—in which the premium at each age had to cover costs but to continue to appeal to the low risks.[10] Both interest rates and a forecast of the rate of growth of medical spending were included in the model. We then compared this theoretically optimal and equilibrium path with the actual path of premiums insurers charge. The results for males, shown in Figures 1 and 2, are striking. Figure 1 shows three alternative actuarially fair premium or average benefits paths for a standard policy. It shows the premium of "low risks" (those with no chronic conditions), "high risks" (those with one or more such conditions), and the average breakeven premium, all for a standardized policy. As expected, all premiums rise with age.

FIGURE 1
Male Risk-Rated Premium Schedules

Note: These estimates are given in 2003 dollars.

Next we solved for the incentive-compatible premium: that premium which (in any time period) covered both the costs incurred in that period by low risks and the expected discounted future premiums for those who become high risks in that period. (Both death and recovery are taken into account.) Finally, from the Community Tracking Survey (CTS), a large household survey of insurance purchasers, we calculated the actual premiums paid for the standardized policy as a function of age. These two premium paths are overlaid on the other paths in Figure 2. The comparisons are striking. First, actual premiums definitely deviate from the single period breakeven premium for both good risks and average risks. The margin of the actual premium over the age-rated average premium is high at younger ages, but eventually becomes negative. Second, and even more surprising, the optimal time path and the actual time path are fairly similar, as indicated in Figure 2. (The similarity still exists but is less precise for females, perhaps because the expenditures due to childbearing do not reflect an "unexpected" chronic condition for which GR funding is most appropriate.) In any case, the overall pattern is one in which reality, once observed, is highly consistent with the recently-developed theory.

FIGURE 2
Male Actual CTS Nongroup Premiums

Note: These values are given in 2003 dollars.

Modifications to the Market Equilibrium GR Model

The theoretically optimal equilibrium GR schedule charges the good risks at any age a premium that they are willing to pay, and at the same time collects enough to cover all future above-good-risk costs of those who become chronically ill. Here I consider some possible modifications of this basic model.

Level premiums. First, it is clearly possible to think of setting future premiums below the premiums good risks would be willing to pay, while at the same time having yet higher front-loading. The most extreme version of this would be to charge a single lump sum for lifetime coverage at some age, and zero thereafter. While that strategy would not seem reasonable, another similar approach apparently has attracted the attention of private insurers and insurance regulators in Germany: a requirement that plans charge "level premiums" that vary with the age at which the person first begins insurance with a particular firm but not thereafter. Compared to the optimal schedule at any age, this factor will yield higher initial premiums and lower ones later in life. Clearly it will be stable, but as with a lump sum premium the important question is whether this pattern is efficient. If we assume that capital markets are imperfect, so that young consumers would prefer to minimize the premiums they pay when wealth is low, this pattern is less than ideal. It also complicates the problem of calculating reserves by requiring larger future reserves.

Presumably this pattern exists in Germany because of the need to compete with the statutory health insurance that charges premiums based only on earnings (up to a limit); the "lock in" thus keeps people from defecting to that insurance as they age. What would emerge without the presence of such funding for quasi-public plans is an open question. My conjecture is that flat premium plans would not survive in competition with plans with GR whose premiums grow moderately with age.

Lock-in. One of the problems with front-loading is the mirror image of the protection the firm agrees to provide against increases in premiums: those who become high risks are locked in to the insurance policy and insurance firm they initially chose. It is important to note that the more numerous low risks are definitely *not* locked in; they can switch insurers or switch policies as they wish. It is only those who are high risks at the moment who are locked in. The possibility of lock in presumably motivates people to

search more intensively among insurers in making their initial purchase and to choose those with a reputation for treating high risks well. However, people may not want to make a binding commitment to a particular insurer if they become high risks. The data suggests that those with chronic conditions who survive often do not remain high expected expense for the rest of their lives; cancer survivors, for example, eventually become close-to-average risks. But people may well still be concerned with lock-in to individual insurance, just as they are supposedly concerned with lock-in to a job in the case of job-based insurance.[11]

There are ways to mitigate the lock in effect but they raise problems of their own. For example, the insurer could be required, as Cochrane has suggested, to pay a lump sum bonus, reflecting the expected value of future above-average premiums, to those high risks who wish to leave; they could then use this amount to cover their premiums at other insurers. However, following this procedure then requires determining precisely what the person's risk level is, which reinstitutes costly underwriting and is sure to be controversial. The PKH approach does not explicitly permit opting out, and therefore would require prospective buyers to put more effort into searching for good insurance from good firms in the first instance.

One issue here raises one of the paradoxes of economics. If competition, based on reputation, say, is vigorous, consumers who select a firm in the market will probably get a price close to cost and future performance that does not permit the insurer to take advantage of them. (More on this later on.) But for competition to be most vigorous, consumers should not only have an opportunity to choose initially but also an opportunity to leave firms that try to take advantage of them. More generally, the portability of initial excess premiums and the method of determining the amount of any exit payment to high risks is something that competitive markets could determine.

Apparently, in individual health insurance in the U.S., the problem of lock-in is not severe enough to make it worthwhile for any insurer (life or health) to follow the Cochrane model. However, it is instructive that

individual health insurance is much less likely to be a "preferred provider" model with limited but changing panels of providers (rather than indemnity (all-provider) insurance).

A key issue is whether there are good reasons why a person's demand for insurance (and therefore the specifics of a policy) should change. Lock-in is only a problem if there are likely to be reasons why one would want to change insurers or coverage. One reason for desired change should probably be ruled out: an increase in demand for coverage of a particular type of care after a person discovers a chronic condition that uses that type of care. This would ordinarily represent adverse selection.

At a more general level, the question is whether the demand for insurance would be expected to change when risk changes, given that premiums also change in response to risk. As Ehrlich and Becker point out, the demand for insurance should respond primarily to loading, not to loss probability; the only effect that the latter would have is as an income effect which should be minor.[12] Thus an optimal indemnity policy may depend on the person's extent of risk aversion and on the loading, but could and should be written in such a way that the form of the contract can remain the same even as risk changes. Proportional coinsurance, for example, displays this property: when risk increases expected out of pocket payment also increases (proportionately), even with no changes in explicit policy provisions.

Changes in a person's circumstances (income, location, family structure) are more plausible reasons for seeking a change in the form of insurance. Many insurers will permit changes in coverage and still continue GR for such explicit and documented alterations of characteristics. Clearly lock-in need not be absolute but must still be present to some extent; the real problem arises when it is hard to distinguish objectively between changes in the person's demand for insurance, given risk, and changes in risk.

The other reason why lock-in would be a problem is if the insurer changes even while the consumer's demand remains constant. As already noted, the danger of reneging on explicit or implicit long term contracts is always present, and the changing cost and quality of medical care made it harder to monitor what the insurer is providing. One would think that, if this were a serious problem in practice, safeguards

would be built into actual policies. The apparent absence of safeguards is consistent with the hypothesis that reputation provides a sufficient constraint on opportunistic behavior by insurers. It would also be consistent with the hypothesis that consumers are totally unaware of their vulnerability to lock-in—but then one would have to explain why they value GR protection that is empty. Empirical work seems to be the best way to make progress on these issues.

Time Period between Changes. While the theory envisions that GR provides a guarantee for a lifetime, there may be reasons to shorten the time period—without going all the way to the one-year interval typical of ordinary insurance plans. Term life insurance provides a model here, with term insurance available for periods of one, five, or ten years—over which interval not only renewal but premiums themselves are guaranteed. Buyers can choose the time period, and there is generally a tradeoff between length of commitment and premium.

Requiring a waiting period before a change in coverage can be made is equivalent to choosing the term of insurance. Single-year term insurance permits changes every year, but ten-year term insurance only every decade. The primary advantage of ten year term insurance over single year term insurance is that the latter does not specify what the premium will be for the next time period, only that evidence of insurability will not be required so that there will be no discrimination.

The longer the time period between a decision to increase, reduce, drop, or change insurance and the time at which that change becomes effective, the lower the chance of adverse selection by the buyer or re-underwriting by the seller. The veil of ignorance becomes more opaque, and decisions will be made in a less opportunistic fashion. But the longer the lag, the less closely coverage can be tailored to the person's current circumstances. (This raises again the issues associated with lock in.) Again there is a tradeoff, with greater ability to change insurance associated on average with a longer time period between changes.

It seems that markets could solve the problem of optimal waiting period for changes in coverage. It can offer insurances for which change can be rapid or slow, at different premiums, and let buyers choose what

they want. While a shorter waiting period will be associated with more adverse selection, there should not be adverse selection associated in making the choice of a shorter waiting period itself. Instead, the frequency of changes in external circumstances, the person's taste for bearing the cost of change, and the elasticity of the person's demand function (as a proxy for lost consumers' surplus from the wrong quantity) should influence this choice.

Guaranteed Renewability in Term Life and Long Term Care Insurance

It appears that the first and most common use of guaranteed renewability was in life insurance.[13] Under this provision, the person who has paid their premiums for a given amount of coverage for the current term is permitted to renew the same policy without the need to take a physical exam or provide other evidence on health. It is also even common (more so in whole life insurance) to buy a supplementary policy that allows the person to increase coverage at the anniversary date without providing evidence of insurability.

The premium for life insurance policies typically increases with age; the new premium is at the discretion of the insurer. The premiums for policies with longer terms either remain constant or increase with age according to a prespecified schedule. This is the case in which it is clearest the GR is really passive rather than active; it simply represents the absence of effort to determine the buyer's current-period risk level. Because the payout of life insurance is in monetary terms and because mortality tables are general knowledge, insurers set aside a portion of the premiums collected as a reserve against the GR promise.

As far as I am aware, the guaranteed renewability feature in life insurance works very well. Those who survive and have good health do not drop coverage in a future year even if, by that time, it will be apparent which of the original purchasers have a chronic condition which means that they are more likely to claim death benefits. The premium must be low enough to keep the good risks renewing from the initial firm and not volunteering for new medical examinations in order to get lower premiums from new firms.

Whole life insurance and long term care insurance share the property that the premium depends on the age at which the contract begins, and generally does not increase over time even though the probability of making a claim of either type surely does increase over time. As in German private health insurance, this feature implies even more front loading of the premium and even less likelihood of adverse selection by those who turn out to be good risks. In essence, the person prepays part of the future premium. In whole life a person who decides to drop out can claim some of this prepayment back—the so-called surrender value of the policy—although the frontloading of expenses (especially commissions) means that the surrender amount is initially low relative to the present value of the amounts used for prepayment.

Long term care insurance also has a similar prepayment feature, but usually does not offer a cash value if the policy terminates or lapses. Not only is the premium lower the younger the age at which the person starts to pay for long term care insurance, the difference between the "spot" premium and the present value of prefunded premiums at an earlier age is positive; the insurer is assuming adverse selection.

LTC insurers initially determined the amount of frontloading based on assumptions about the proportion of policies that would lapse. What they found was that the lapse rate was both lower than expected and skewed toward heavier users of long term care, so that the prefunding was inadequate.

Adverse Selection and GR

One of the aspects of insurance which has most intrigued and bothered economic theorists is the problem of adverse selection. *What I want to suggest here is that GR provides a solution to much of the problem of adverse selection*. (See Pauly and Herring (2003) for a more complete treatment.[14]) The solution is not quite complete in theory, but perhaps close enough in practice.

Suppose that, in contrast to the earlier model, no insurer can tell which of the insureds develops the chronic condition in period 1, so that there is informed asymmetry. However, the insurance has a GR provision. Would this situation be vulnerable to adverse selection?

Given the assumption of the Cochrane and PKH models, the answer is negative. Both of those models assume that the amount of insurance in the policy is given or fixed over time, and it is the premium for just this fixed amount of coverage which is "guaranteed." That is, neither model assumes that the individual obtains the privilege of buying more insurance at the average premium, or of recovering some of the excess premium if he subsequently decides to buy less. In effect, in these models the consumer's only choice is either to buy the predetermined level of coverage or none at all, at least from this insurer.

To see why the GR equilibrium is stable and is not subject to adverse selection, consider whether, in the spirit of the Rothschild-Stiglitz (RS) model, it is possible for a new insurer to profitably offer a different policy to those with the "GR-or-nothing" option.[15] Would low risks be able to be attracted away? The negative insurer has already been given: in the next period or periods the low risks expect to pay only their own expected expenses to the current insurer, so the new insurer has nothing to offer them.

In the RS model, the key reward to low risks from being separated from a pooling equilibrium is the chance to avoid transfers to high risks. But with GR, there are no transfers *ex post* from those who know themselves to be low risks to those who know themselves to be high risks. Thus, in this model, there is equilibrium, and no adverse selection, without even invoking the alternative concepts of equilibrium.

What about the people who know they have become high risks? If a new insurer proposed to charge the high risk premium, either for the current level of coverage or some different amount, the high risks would also not be attracted, since, having prepaid the difference between high and low-risk premia, they would expect to pay the latter in the future.

We do not need to require that the quantity be the "original" or "standard" quantity and zero. Insurers could be free to sell less generous policies; low risks would not be induced to reduce the quantity below their original level, since the premium was and is based on the low risk probability.

The one possible vehicle for adverse selection to occur would be if insurers were willing to sell more coverage than the initial or standard amount. The possibility here is that those consumers who (secretly) know they have become high risk might be more likely to buy more coverage because the premium has fallen well below one based on their risk level. But what premium would insurers charge for such additional coverage after the first period, given that they know that some of the potential buyers of additional coverage have now become high risk? The answer obviously depends on the reason they imagine that someone might want additional coverage. If the reason is always because the person has become high risk, insurers would then propose to change the high risk premium for such coverage, and high risks would have no special incentive to buy more coverage.

If the only change that was occurring for consumers was the change in loss probability, we know (as Ehrlich and Becker have shown) that change in probability, loading held constant, should not change the desired level of coverage (except for minor income effects).[16] However, there might be good reasons why some consumers would want to buy more coverage. They might have new dependents or a higher level of income. They might experience a change in the value of health that would affect the level of health insurance coverage they desire. Then some consumers seeking more coverage might be low risks. The insurer might be able to identify risk levels by discovering the reason for changing coverage. If that proved difficult, however, there could be adverse selection applied to *increments* in coverage, above the standard or average quantity chosen when insurance is first purchased.

Insurer Inside Information and Guaranteed Renewability

In this last section combine the model of guaranteed renewability, which both smooths risks over the lifetime and strongly inhibits adverse selection, with a model, especially valid for automobile insurance, in which the insurer from which coverage is bought learns through a consumer's claims history what that person's risk level is. However, other insurers, without access to as accurate a claims history, will be less adept at identifying specific risks. Will the current insurer try to exploit its informational advantages, and where will there be an equilibrium?

The insurer might plan to raise the premium it charges to especially good risks up to the higher premium they would be charged by outside insurers. One factor that would inhibit this behavior would be fear of a bad reputation, but the imprecision of what expected expenses are for coverages with necessary "repair costs," like auto and health, gives an insurer some leeway. If reputation fails to discipline adequately, however, insurers could expect to earn profits on those they know to be better risks than the market judges; in turn, this would provide a motivation for offering low prices initially in order to "buy" this information, as Kunreuther and Pauly have suggested.[17] Such "backloading" of profits and administrative costs would offset the "frontloading" inherent in GR.

It seems that the question is ultimately an empirical one of which influence is stronger. We know that, on balance, individual health insurance is "frontloaded." The elasticity of premiums with respect to changes in expected expense due to age is substantially less than unity, in the range of 0.2 to 0.45.[18] I am less familiar with such information on automobile insurance. But generally it does not seem that backloading or initial period discounts are very common.

Conclusion

In ways largely unrecognized by insurance theorists, some insurance markets have set up institutional structures that appear to deal in a good (though not perfect) way with both risk rating/time consistency and adverse selection. There is a problem in insurance markets if people wait to buy insurance until risk differentiation occurs or becomes apparent. GR tends to inhibit such behavior, offering a strong incentive to buy insurance *before* events occur which increase risk, but customer ignorance may mean that some consumers still wait to seek coverage until they become high risk. The political sensitivity to such high risks then paying high premiums or being refused coverage at reasonable premiums frequently motivates what in my view are ill-considered, unnecessary, inefficient, and inequitable efforts at *ex post* risk pooling. Regulations forbidding insurers from risk rating those who do delay destroys the incentive to take GR coverage; in this sense, "community rating" is a humanitarian temptation that needs to be resisted in favor

of more efficient ways to deal with high risks. To be sure, differences in risk level at the start of (adult) life that come from genetics or misfortune in childhood *do* deserve to be averaged or pooled across everyone in the community. But the fraction of the 18-year-old population with a chronic condition is less than 3 percent, so this cannot be a problem that requires heavy-handed interaction in the insurance transactions for the great majority of consumers.

The main message to me from the ideas in this paper is that insurance theorists may have devoted excessive attention to the theory of adverse selection, relative to the real life relevance of the problem in markets that actually exist. Most of the examples of severe adverse selection we have come from situations in which community rating or some other deviation from risk rating are required by law. At least we need to recognize that markets that are more complex than those usually used in the theoretical models have developed ways to capture, at least in part, the efficiency that would be lost in market failure.

[1] Arrow KJ. Uncertainty and Welfare Economics of Medical Care. *American Economic Review* (December 1963); 53(5): 941-973.

[2] Kunreuther H, Pauly MV. Market Equilibrium With Private Knowledge: An Insurance Example. *Journal of Public Economics* (May, 1985); 26(3): 269-288

[3] Cooper RW, Hayes B. Multiperiod Insurance Contracts. *International Journal of Industrial Organization* (1987); 5: 221-232.

[4] Pauly MV, Kunreuther H, Hirth R. Guaranteed Renewability in Insurance. *Journal of Risk and Uncertainty* (March 1995); 10(2): 143-156; Cochrane J. Time-Consistent Health Insurance. *Journal of Political Economy* (June 1995); 103(3): 445-473; Hendel I, Lizzeri A. The Role of Commitment in Dynamic Contracts: Evidence From Life Insurance. *The Quarterly Journal of Economics* (February 2003); 118(1): 299-327.

[5] Pauly, Kunreuther, and Hirth, 1995; Cochrane, 1995

[6] Cochrane, 1995; Pauly, Kunreuther, Hirth, 1995; Diamond P. Organizing the Health Insurance Market. *Econometrica* (November 1992); 60(6): 1233-1254.

[7] Pauly MV, Rosenbloom J, Herring BJ, Szrek H, Rosenquist JN. Guaranteed Renewability in Individual and Group Health Insurance: Functioning and Future Prospects. Final Report on Robert Wood Johnson Foundation Grant 44065, Unpublished, The Wharton School, University of Pennsylvania, August 2003.

[8] If it were forbidden to discriminate its premium and/or a new contract, it could still get the risks to self-select by offering a Rothschild-Stiglitz policy. Rothschild M, Stiglitz JE. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics* (November 1976); 90(4): 630-649.

[9] Pauly MV, Herring B. *Pooling Health Insurance Risks* (Washington, D. C.: AEI Press, 1999).

[10] Herring B, Pauly MV. Incentive-Compatible Guaranteed Renewable Health Insurance. NBER Working Paper 9888, July 2003. http://www.nber.org/papers/w9888, September 10, 2003.

[11] Madrian, B. Employment-Based Health Insurance and Job Mobility: Is There Evidence of Job-Lock? *Quarterly Journal of Economics* (February 1994); 109(1): 27-54.

[12] Ehrlich I, Becker, GS. Market Insurance, Self-Insurance, and Self-Protection. *Journal of Political Economy* (July/August 1972); 80(4): 623-648.

[13] Harrington SE, Niehaus GR. *Risk Management and Insurance* (Boston: Irwin/McGraw-Hill, 1999).

[14] Pauly MV, Herring BJ. An Almost Optimal Solution to the Problem of Adverse Selection with Time-Dependent Probability. Unpublished Manuscript, The Wharton School, University of Pennsylvania, August 2003.

[15] Rothschild and Stiglitz, 1976.

[16] Ehrlich and Becker, 1972.

[17] Kunreuther and Pauly, 1985.

[18] Pauly and Herring, 1999.

CESifo Working Paper Series

(for full list see www.cesifo.de)