

**HALF FULL OR HALF EMPTY?
ECONOMIC AND BEHAVIORAL EXPLANATIONS FOR INSURANCE
DEMAND**

Mark V. Pauly
The Wharton School, University of Pennsylvania

Paper prepared for
“The Irrational Economist” Conference
and Book Writing in honor of Howard Kunreuther

December 4-5, 2008

Chapter to appear in
The Irrational Economist: Future Directions in Behavioral Economics and Risk Management
E. Michel-Kerjan and P. Slovic (eds.) *Public Affairs Press*. (forthcoming 2009)

**HALF FULL OR HALF EMPTY?
ECONOMIC AND BEHAVIORAL EXPLANATIONS FOR INSURANCE DEMAND**

Mark V. Pauly
The Wharton School, University of Pennsylvania

Paper prepared for conference in honor of Howard Kunreuther's seventieth birthday
The Irrational Economist
December 4-5, 2008

(DRAFT: REFERENCES INCOMPLETE)

Introduction.

Insurance is one of mankind's greatest inventions. I used to rank it just a little behind fractional reserve banking, but the recent turmoil in capital markets (with insurance not totally unaffected) suggests that the two financial instruments may now be even. But while most consumers buy some insurance and some investors make (or lose) fortunes at it, consumer behavior on its face often fails to fit the usual models of expected utility maximization. Faced with this phenomenon, Howard Kunreuther and I have over the years collaborated on research that tries to identify positive explanations for what happens and to suggest normative public policy strategies toward it. Our work largely deals with the broad question of what to do when there is a fork in the intellectual road: when confronted with behavior that appears to be inconsistent with a simple version of the traditional EU model, should one posit a different kinds of utility or choice function for consumers, or should one modify or amplify the traditional model by adding constraints on time, money, and information? We have benefitted from going a little distance down each path, contrasting what we have found, and trying to make a judgment about which answer is best when.

In this paper I will give some examples of insurance buyer behaviors that raise these kinds of problems—some from our own past work and some newer ones we have been considering—indicate the kinds of analysis that comes from either approach, and use these to discuss more generally the proper methodological approach in these circumstances. The reader should be forewarned that, although we sometimes come to a defensible conclusion, we often end up agreeing to disagree—in the sense that we lay out two alternative explanations that both seem reasonable, and then regard the best strategy as one of letting the intellectual audience make their own choice.

Deductible Aversion.

Suppose insurance is offered with a proportional administrative loading. A well known proposition in insurance theory then says that, without more, the optimal policy will be full coverage above a deductible (Arrow, 1963). The optimal deductible is determined by comparing the marginal reduction in the risk premium the consumer is willing to pay (dependent on his risk aversion and the distribution of losses) with the marginal reduction in the loading cost. Beginning with a zero deductible, the marginal reduction in the risk premium starts out at zero (because wealth is the same in all states with full coverage), but increases in absolute value as the deductible grows. At some point a higher deductible would increase risk by too much (relative to the reduction in premium); that is the optimal deductible.

In many circumstances, however, consumers seem to prefer policies with low or no deductibles: for auto insurance, health insurance, or homeowners insurance even comfortably middle class households choose plans with deductibles of only a few hundred dollars, and they do so even in the face of premium differentials that make the incremental premium for a lower deductible large relative to the expected value of additional claims. This preference for a low deductible policy by so many buyers seems too strong to square with reasonable assumptions about risk aversion: people generally seem to prefer low deductibles even in policies that cover risks that are small relative to their wealth. (They also seem to prefer these policies to such an extent that they choose them even when the loading on the coverage near the deductible is unusually high, that is, when the premium reduction associated with an increase in the deductible is quite high relative to the increment in expected out of pocket payments.) Indeed, when I show my students a numerical example where the increase in the deductible is smaller than the reduction in premium (so their guaranteed premium saving is more than the worst that could happen to them under the higher deductible), some still say they prefer the low deductible policy.

How to explain this? One strategy is to posit a different (or at least an enriched) utility function in which “peace of mind” and “freedom from regret” have positive values. The only real way to test such hypotheses is to have some further basis for identifying characteristics of individuals who are more or less likely to have such preferences. (One conversation Howard and I have had over the years is the desirability of taking full insurance coverage for a rental car if your wife is along, our own wives excepted, of course.)

One modification of the EU model that would predict low deductibles has been proposed by Schelsinger and Doherty: it may be that there are “losses” associated with the uncertain event which are not directly insurable. In a second best sense, it may then be preferable to overinsure the insurable event (e.g., monetary damage or cost) to compensate for the inability to insure the other event. The substitution cannot be perfect, of course, but paying for coverage in this indirect way may be better than no coverage at all.

These two different explanations imply two potentially different roles for government. If the EU model were accepted, the discovery of a rational explanation for deductible aversion would obviate the need for government to take corrective action about overinsuring. In contrast, if the explanation were one based on behavioral motivations, welfare economics is probably useless as a basis for offering advice to government. If people have preferences differing from those of the benchmark model, the most “democratic” reaction is to say, “so be it,” and let them waste their

money in the way they choose. It is not always possible to take such a hands-off attitude; if preferences are merely eccentric it is one thing, but if they are harmful (or impoverishing) it is quite another. And on top of this, in a society where government caters to what voters rather than what welfare economists want, the politician maximizing his tenure or political fortunes may well find it desirable to go along with these mistaken preferences, whether it is for first dollar coverage, dread disease insurance, or even poor investment vehicles disguised as life insurance.

These are the thoughts that flow in our dialogs, and motivate continued investigation. Usually it is the extended EU model, rather than the behavioral explanation, that yields refutable hypotheses. Indeed, the circumstantial evidence for the behavioral explanation often is the rejection of the EU hypothesis, and then one turns to the behavioral explanation as the only argument left standing. For example, in principle one could see whether deductible aversion was in fact stronger when there was uninsurable (but not unmeasurable) additional loss. Do people who choose low deductibles experience larger uninsured changes in their wealth (in multiple ways) than people who choose high deductible insurance? The answer to this is not yet known, but could be. If the answer was negative, the behavioral explanation would gain support.

Consumer Warranties.

Another example of a phenomenon that seems to be best explained by a behavioral theory rather than by the EU theory is the purchase of warranties on consumer durables. Here again the cost often exceeds the expected value of the benefit by a substantial amount, and yet many people still buy. The proportion engaging in this behavior here is a minority (in contrast to the majority who choose low deductibles), but is still usually a substantial minority, about a third of buyers. One other important factoid: the demand for warranties is responsive to their prices.

Buy why does this behavior occur? The obvious explanation is that buyers of warranties overestimate the probability of needing repairs (especially when experiencing the hard sell after an already costly purchase). Another possible explanation is that, having made the effort to shop and choose an item, at least some people may attach special value or affection to it, an argument that Howard has been made in a more general framing (Kunreuther and Hsee), but usually with an application to an “old favorite” and difficult to replace item, not a brand-new item still in the box. A warranty also protects the buyer against higher prices for replacements, and may be priced favorably by a manufacturer whose marginal cost of replacement is much below the selling price—a real possibility for many relatively low priced consumer electronics. And, finally, a warranty gives you peace of mind and freedom from regret.

An interesting question is whether the willingness of some buyers to take warranties that would not be a good deal even if they were not overpriced results in lower price of the product itself; could the product be a loss leader (as are razors) with the money to be made on warranties and/or repairs? Perhaps, but this seems like a stretch.

There is some reason to go slow on a behavioral explanation here. In addition to the fact that most people do *not* buy warranties, competitively sold warranties are available for large

purchases like automobiles. It may also be that warranties are a form of price discrimination. A buyer who counts himself as very lucky at getting a good deal on a good product may be more willing to pay for a warranty than a buyer who thinks his particular purchase and purchase price are run of the mill and par for the course. The first buyer having discovered treasure, he may be more willing to invest in protecting it. Finally, the price responsiveness of demand for warranties implies that, even if the preferences are not so rational, consumer behavior in fulfilling those preferences does repeat the standard model.

Regulation of warranty pricing or terms is a possible public policy solution, although at present it is uncommon (though not unknown). An alternative policy intervention would be the provision of information on frequency of repairs, so that buyers could make a better calculation of the expected value of a warranty. Indeed, such information would be useful across a whole range of insurance products, from warranties to long term care insurance, where information on probabilities is either absent or grossly distorted.

The Non-Poor Health Uninsured.

These have been two examples of over purchase of insurance. What about behavioral versus rational explanations for under- or non-purchase? The most familiar example from my perspective is the non-purchase of health insurance. Some of the uninsured population in the US has such low incomes that the explanation for non-purchase is obvious: they can afford neither to pay for care nor to pay for insurance. But a fraction we know is “sizeable” could afford insurance, in the sense that they would have enough left over for other consumption and in the sense that most other people in their circumstances do have and pay for insurance (Bundorf and Pauly; Gruber). Moreover, although they may get some free care if their bill is enormous, they will be at risk for paying themselves up to that point.

So why are the non-poor uninsured that way? The all-purpose explanation is that their tastes for insurance are weaker than those of otherwise similar people, but a “tastes” explanation is distinctly unsatisfying. The consequence of not paying an insurance premium for this population is not to have that much more income available for other types of consumption; there is the prospect of expected out of pocket expenses. That is, the alternative to insurance is not no expense, but is a risky prospect of high expense. (That the uninsured are charged much higher retail or undiscounted prices for hospital inpatient and outpatient care is another reason why insurance is a better buy than non-insurance.) Even with the likelihood of non-payment for catastrophically large expenses, the real chance of severe financial distress should motivate purchase by everyone with income or wealth above some relatively low level. Why doesn't it?

There are plausible explanations from the behavioral side. To begin with, some people may just not be risk averse enough to be willing to pay the administrative loading for health insurance. This loading can be high—if you just look at the first individual insurance policy you are offered on the web—but the great bulk of the uninsured in principle have access to lower prices and tax subsidized employment based group insurance. A substantial fraction of the uninsured turned down group insurance for which they are eligible, and even more could find jobs that carry coverage (even if Starbucks has stopped hiring) if they wanted to get coverage. A more

plausible explanation from a behavioral perspective is that they do not have enough information to see the value of insurance. A study of the non-poor uninsured in California provide some confirming if circumstantial evidence: the non-poor without health insurance are also less likely to have other kinds of insurance (life, homeowners, collision), to be maxed out on their credit cards, and to think that individual insurance is much more expensive for them than it really is. These are people who are just not into risk protection of all sorts. But there is also an economic explanation. Group insurance seems to discriminate against lower risks, in the sense that lower risk employed people are less likely to have coverage than the higher risk employed. (It probably helps the risk pool that a person has at least to be able to work to get group insurance, so there is less of a possibility of severe adverse selection.) A low risk person will face the same explicit premium as a high risk person working in the same firm, but the options available in firms employing mostly lower risk people and the existence of a lower wage offset for coverage provide some semblance of experience rating.

Here the best conclusion is probably a compromise. Insurance pricing for the uninsured is not especially attractive, and they tend to underestimate the value of insurance. Both behavioral factors and rational factors reinforce each other.

The role of public policy here has recently been much debated. Some subsidies to non-poor but not high income uninsured which exceed the subsidy they get for group insurance is a part of every health reform plan. And most plans purport to be able to offer coverage at lower administrative cost through exchanges or pools of some kind. Behavioral issues—motivation, salience, and framing—have interestingly not been a major part of the debate.

One approach that sidesteps these issues in practice but makes them more important in theory is the use of mandates for coverage (for children or adults). The strongest argument for mandates is the spillover costs the uninsured impose on others, costs which recent research shows are much larger than hospital level charity care, but extend to lower quality for insured people in the community as larger fractions of the uninsured make high quality less profitable (whether they use care embodying it or not) (Pagan and Pauly)

Insurance and a Tranquil Life.

A theme in each of these three cases is one that has characterized our recent work: people have other objectives than getting their insurance purchases just right. We formalized this idea in a paper that postulated a decision-making cost to investigating insurance purchasing—collecting information on the best premiums, the best coverage, and the true loss probabilities (Kunreuther and Pauly, 2000). Where a risk really matters—fire insurance on your home, collision or an extended warranty on your fairly new car, insuring your dependents against your premature death—insurance does commonly exist. That is, where your expected utility might really be impacted by a loss, you usually care enough to start to look for insurance and eventually to buy it. How long you persist depends on how much a good price matters to you; we found that low risk young people paid premiums less reflective of their expected expenses for health insurance, a phenomenon consistent with a cost of search hypothesis but also with others.

One important consideration in that work that has both rational and behavioral implications is how people think about insurance pricing. If they believe that insurance premiums generally reflect relatively modest profits for insurers across the board, they can greatly simplify their lives. Given such an assumption, the premium an insurer charges you to insure a given asset against a particular peril tells you what the insurer thinks the loss probability is.

Suppose you are considering buying insurance coverage against the possibility of an event which imposes a loss of \$L. You do not have to worry about how likely the event is, or try to collect some data or ask others about it. Instead, you only need to plug in your estimate of how much of the premium goes for loading, interpret the rest as a measure of expected loss (and expected benefit from insurance), and decide whether or not your risk premium (which you can calculate when you know the loss and its probability) is or is not greater than the loading. This information will also tell you how large a deductible to choose. If a risky prospect has more than one adverse outcome, then you will need more information than is contained in the premium, since you will want to know the frequency distribution of losses of different amounts. But you will still have a lot of the information you need for a good decision, and getting the shape of the loss distribution a little wrong is not going to lead to big mistakes.

This means that the first question a rational person should ask when confronted with an offer of insurance coverage in an unfamiliar setting (say, long term care insurance), is not how much the loss might amount to (since there is almost always some terrible combination of events, however unlikely, that can be associated with a large loss) or even the loss probability or probabilities, but rather what is the loading (including profit) on the offered coverage. Unless you have some inside information that your loss probability is much different from what the insurer would think is average (that is, unless you think you would gain through adverse selection), knowing the loading and knowing the premium gives you almost everything you need, if not for a perfect choice, one that will be pretty good.

Unless there is a conscious attempt at secrecy, finding out the loading embodied in an insurance offer is usually easier than trying to find the data on losses and their frequencies. If the insurance market has multiple sellers, if there is past data on premiums paid in and benefits paid out, and if you expect the future to be much like the past, the needed information is in principle available from the insurers own accounting data. Often this data must be submitted to a state insurance regulator, or, alternatively, a prospective buyer may ask the insurer for this information. If the insurer refuses to furnish it, that alone speaks volumes about its likely value.

Sometimes this strategy may not work. For example, if windstorm insurers have dramatically increased their premiums after an event with large scale damage, information on past loadings may not tell the potential buyer what profit margin (or safety margin) is being built into the new offer, or even what the insurer thinks the next chance of a major storm is. In this case the best strategy may well be a mutual insurance arrangement, in which a large number of insurance buyers determine that their risks are similar but to some extent independent, and then agree to make a contribution toward a common risk pool based on a rough guess about the loss probability (Pauly, Kunreuther, and Vaupel). If they overestimated, the mutual insurer can pay dividends; if they underestimated, the insurer can assess additional contributions (the latter

action being more practically problematic than the former). But the key point is that insurance purchasing can often be based on a small amount of fairly objective data (or assumptions about data) and need not depend on the feelings and facts currently known by insurance purchasers.

Conclusion.

My conclusion from these applications is that the rational and the behavioral need not be in serious conflict if insurance markets are thought to work reasonably well. Given relatively free entry into insurance, “working reasonably well” should be the exception rather than the rule, but spirited discussion about how well is reasonably well will be commonplace. There are some exceptions for relatively small and linked risks, like warranties and rental car insurance. But for those risks that might (if uninsured) make a big difference to a person’s expected utility, we can expect markets to work reasonably well, and usually they do. Of course, a large number of little risks can mount up to a big risk, but insurance firms seem adept at bundling risks to a given asset in a given use (as in the form of all-perils coverage). Things are not going to work well for new, large, unknown losses—but you have probably guessed that.