INSURANCE MARKET EQUILIBRIA WITH CREDIBILITY ADJUSTED PREMIUMS

Mark V. Pauly

February 1983
WP-83-26

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
A-2361 Laxenburg, Austria

INSURANCE MARKET EQUILIBRIA WITH CREDIBILITY ADJUSTED PREMIUMS

Mark V. Pauly

1. INTRODUCTION

Insurance firms base the premiums they charge on advice from actuaries. But despite the sophisticated statistical apparatus that actuaries are trained to manipulate, they can not forecast the profits (or absence of losses) for the insurance firm with perfect accuracy. Indeed, actuaries do not claim such accuracy. Profits often depend on the correctness of guesses that actuaries have made about the expected losses associated with insurance contracts of different types.

The linkage between the actuary's estimate of the expected loss or "pure premiums" and the actual premium the firm would charge is not very clear in the actuarial literature. There is always assumed to be some "loading" or addition to the pure premium. Part of the loading goes to cover administrative costs--selling costs, billing, claims processing, and the like. This part of the loading is, in the more sophisticated treatments, said to be ideally equal to these costs, although the problem of joint costs and their allocation is not explicitly confronted. The crucial question, however, is the determination of the other part of the loading, the excess over the pure premium and expenses. This part is variously alleged to be a "safety loading", to be a fair profit, or to be a way of bringing in that amount of revenue which maximizes some ad hoc firm utility function (Bühlmann, 1970). Discussion in the actuarial literature is very much in the spirit of searching for a cost or disutility justification for a positive margin over costs and expected losses. It almost never defines the optimal premium in terms of the parameters of the firm's demand function as well as its price, precisely because neither expected profit maximization nor competition are assumed.

Economic analyses, in contrast, have typically assumed expected profit maximization, and appealed to the law of large numbers as a justification--with free entry and "large" numbers of insurance contracts, zero expected profit is an equilibrium condition (Rothschild and Stiglitz, 1974). They have also been concerned a great deal with competitive equilibria. However, there has been virtually no explicit attention paid to the question of equilibrium market behavior when the data available to the firm is not sufficient to permit estimation with virtual certainty of the expected loss per contract of a particular identified class of insureds. This is despite the focus in the economic literature on situations in which the firm is imperfectly informed about the loss probability of a set of individuals; the economic literature has not analyzed the process by which insurance firms obtain and use information to sharpen their estimates.

This paper will attempt to use some old and some new notions of the process of premium adjustment, or "credibility" as it is called in the actuarial literature, to specify how insurance firms might estimate loss probabilities. But then it will also describe the characteristics of market equilibrium premia in a world where many insurers are engaged in the same premium estimation process.

One conclusion of the paper will come as no surprise to economists, though it is at variance with the usual actuarial approach. The process by which the equilibrium premium is determined will be shown to depend on demand-side conditions as well as on expected losses (however forecasted). Another, perhaps less obvious, conclusion is that this equilibrium insurance premium can be represented (ignoring administrative costs) as the sum of the expected loss (called the "pure premium") and an additional amount to be added for risk, *even if the insurer is risk neutral.* That is, the optimal and equilibrium pricing strategy for any insurance firm will involve adding such an addition to its best estimate of the pure premium. Perhaps even more surprisingly, it can be shown that the average premium actually charged when all firms follow such a process will be just equal to the pure premium that would be estimated by a person who had all of the information available to all firms available to him. Sometimes the actual premium will even be exactly equal to the fully informed pure premium. A final section shows that insurer desires for more accurate data (as embodied, for example, in industry rating bureaus) may involve excessive expenditures on data management and unnecessary temptations to deviate from the competitive path.

What little we know about the real world suggests that actual premium setting behavior falls somewhere between following actuarial advice to the letter and competitive equilibrium in a large-numbers situation. For example, in automobile rate hearings in New Jersey, an insurance firm official gave the following description of the use of data relating losses to a driver's traffic violation record:

> We have followed at times very closely the experience being developed by the driver record classes, and other management have come in and said, no, there are such strong competitive disadvantages to that kind of operation ... that you end up with a class of business that is so bad in relation to what other companies are willing to write that you just price yourself completely out of the market. So in the past several years we have not followed the actuarial indications in our driver record class relativities. We have kept them in line with our competition and also looked at our experience to see just how that was working out ... (State of New Jersey Department of Insurance, 1981).

The relevance on competitive prices of other firms and actual profit experience suggests that something more than credibility is needed to explain insurer pricing behavior.

2. INSURER MAXIMIZATION, ACTUARIES, AND THE PREIMUM DETERMINATION PROCESS

An insurer should set its price for insurance so as to achieve its objectives. Actuarial theory as such is obviously not capable of saying what firm objectives are or ought to be. And yet actuaries are hired to offer advice on what premiums the insurer should charge. Formally speaking, then, the actuary ought to have some model of the firm's objectives in mind in order to judge what information he should convey. In practice, however, this model has often been implicit in the rules of thumb suggested. Where there has been some explicit concern for objectives, that concern has been satisfied simply by choosing a firm utility function in an ad hoc way (Freifelder, 1975) without reference to the markets in which ownership of insurance stocks is bought and sold or to the markets in which firms compete to sell insurance.

The specific part of actuarial theory with which we shall deal is concerned with determining premiums for insureds whose expected loss or probability of loss is known to be identical, but whose value is *not* known with (virtual) certainty. However, there is some data suggesting what that expected loss is. This data could either be information on the prospective insured's recent loss experience, or the experience of others who are similar in particular ways. In either case, the number of observations on risks known to be identical is assumed to be insufficient to permit appeal to the law of large numbers to determine the premium entirely from the data. The way in which partial or incomplete information should be incorporated into premium determination is called "credibility theory" (see, for example, Langley-Smith, 1962). A brief summary of several versions or interpretations of credibility will be provided below.

We will consider a market with stock insurance firms selling insurance to cover large numbers of independently (but not identically) distributed losses. If the capital market is assumed to function well, the most plausible model of the individual firm is that it should behave so as to maximize expected profits. While it is true that the probability that losses exceed premiums by more than a given fixed dollar amount rises as the number of persons insured at a given premium increases, the risk per share and shareholder approaches a negligible amount if the number of shares and shareholders increases proportionately. In addition, even if the risk per share (say, the variance in per-share earnings due to underwriting losses) is not zero, if stockholders hold diversified portfolios (including small fractions of the stock of many insurance firms), the random risk for any one firm should be treated as if it were negligible.

This model may well not be consistent with how insurance firm managements actually do behave, though I do not believe there is any definitive evidence. In any case, if present reward structures are alleged to induce management to behave in an excessively risk averse way, that only raises the further question of why stockholders choose reward structures that produce results at variance with their own interests. Perhaps the true model is one in which firm managers have a good deal more control over underwriting profits than we will give them here.

There is a final consideration in the determination of actual premiums, which we will not be able to discuss--whether the market can reasonably be modelled as competitive on the demand side. If it cannot, because of imperfections in consumer information, then firms may price to some extent as discriminating monopolists, overcharging those with few alternatives more than those with many. This characterization fits some of the studied facts, but not all--it means that firms should seek more strenuously after the business of those who will be overcharged.

3. VERSIONS OF CREDIBILITY THEORY

Credibility becomes an issue whenever the number of persons of a particular risk type becomes too small to obtain an estimate of expected losses for these persons in period t+1 by merely observing the average loss in period t. (Credibility of a different but related sort also becomes relevant when there is reason to suspect that there has been a structural change between the two periods.) When all losses cannot be estimated perfectly accurately, some alternative way must be found to get an estimate of the expected loss. The small-size sample mean could still be used, but it may be desirable to combine that information with other information or judgment the firm may have.

For reasons of simplicity, I will treat only a one-period model in which all firms announce their premiums simultaneously at the beginning of the period, and have exogenously given sets of data. Considerations of using one period's information about losses or other firms' prices to set premiums for future periods will not be incorporated here.

The objective of the firm is to maximize the expected value of the difference between the premiums it collects and the claims and expenses it pays out. If we ignore the administrative expenses, the expected profit can be written most simply as:

$$E^j(\Pi^j) = \sum_{\Pi} P_i^j X_i^j - \sum_i \phi_i X_i^j \qquad (1)$$

where P_i^j is the firm j's premium per dollar of insurance charged to person i, X_i is the amount of insurance purchased by persion i from firm j, and ϕ_i is person i's probability of loss per dollar of insurance coverage. Here we are assuming that the contract is one which pays X_i dollars in the event "a loss" occurs with probability ϕ_i.

The first term is the firm's total revenue, and the second term is the firm's expected loss on its portfolio of insurance contracts. Since P_i^j and X_i^j are known to firm j, its expected profit depends only on the ϕ_i.

If all firms know ϕ_i for every person, competitive equilibrium obviously requires that $P_i^j = \phi_i$ for all i and j. If this condition is satisfied, not only does (1) hold for the current portfolio of contracts, but it is also easy to see, if customers only purchase from firms with the lowest P_i, that there is no set of prices other than $P_i^j = \phi_i$ which also satisfies (1).

The more realistic problem that credibility theory was intended to solve is the one in which the firms do not know all of the ϕ_i beforehand. How might a firm estimate ϕ_i, and what would the resulting pattern of prices be? A firm contemplating selling a portfolio of contracts X_i does not really know what its expected loss or "pure premium" is. Instead, at best it will have some data on past losses for the i persons, or persons similar to them in particular ways. (It may also have data on current precautions taken to prevent loss, but I will assume here that the ϕ_i are fixed.) A necessary (but not sufficient) condition for understanding market equilibrium is therefore an understanding of how each firm j determines $\hat{\phi}_i^j$, its estimate of potential customer i's probability of loss.

"Credibility" is the label in the insurance literature for this topic. A general definition is that credibility means "the systematic adjustment of insurance premiums as claims experience is obtained." (Hickman, 1975, p.181). But what is the system, and what is the objective of the system? A more operational definition of credibility is that it is "a linear estimate of the true (inherent) expectation derived as a result of compromise between hypothesis and observation." (Hewitt, 1963). For example, suppose a firm j knows that there are N persons with identical characteristics facing the possible loss of $1, and that their total losses have amounted to $S in a recent past time period. Then we could formulate the expected loss $\hat{\phi}_i^j$ per dollar of coverage as:

$$\hat{\phi}_i^j = Z(\frac{S}{N}) + (1-Z)m \quad . \tag{2}$$

Here Z is the "credibility" weight, and is to be thought of as a function of N or S, while m is some prior or hypothesized value of $\hat{\phi}_i^j$. The presumption is usually made that estimating $\hat{\phi}_i^j$ is enough to tell the firm how to set P_i^j. For example, Freifelder (1975) interprets the action that should follow from observation of loss data as the setting of a premium. While this result would occur in Freifelder's model because of the postulated utility function for the firm and because of the assumed absence of an effect of premium variation on the firm's demand, I will argue that the action of setting the premium will, in competitive equilibrium, depend on more than the firm's estimate of expected losses. In particular, I will argue that, *even if the firm is risk neutral*, it will still make a "risk adjustment" of its estimate of expected losses to determine the competitive equilibrium premium it will charge. That is, its premium will still exceed its expected loss. This risk adjustment will depend in a particular way on the characteristics of the insured hazard, the nature of the variation in the ϕ_i across individuals, and the number of insurance firms.

How are Z and m to be determined in traditional credibility theory? Rules of thumb are used to select a number of observations that would be given "full credibility". For example, if the distribution of accidents were thought to be Poisson, then one version of full credibility would require a type of risk to have experienced 1084 losses (SRI, 1976) for full credibility. At this number of losses, we would have 90 percent confidence of making an error of no more than 5 percent. (Note that as the probability of loss falls, more *exposures* will be needed to generate the required number of losses).

If there are fewer observations than this in a cell, but there are more than 1054 losses in the firm's entire portfolio of contracts, the traditional procedure, still used with regard to so-called "secondary" rating factors, is to estimate the pure premium as a weighted average of the cell mean and the overall mean. Suppose, for example, that there are 2 classes, one with 18,000 observations, enough for full credibility, and a ϕ_1 of 0.0944, and one with 2000 observations, 300 losses, and an observed ϕ_2 of 0.15. The overall ϕ is then 0.1. The preliminary pure premium for class 2 is given by $E(\phi_2) = .5(.1) +$ $+ .5(.15) = .125$, where 0.5 is the relative weight recommended by credibility theory. To maintain rate adequacy, both the .125 and .0944 values must then be "scaled up" so that their average is again 0.1, yielding the credibility adjusted pure premiums of $\phi_1 = .0969$ and $\phi_2 = .128$.

There are two types of questions one might raise about this procedure:

(a) What statistical decision theory (if any) would justify this rule-of-thumb procedure, or are there better ways of using the same information?

(b) If all firms follow this rule of thumb and base actual premiums on the pure premium thus calculated, what will market equilibrium be like?

The first question is one that has recently been investigated by actuarial theorists and statisticians (Chaing and Fairley, 1979, Tomberlin, 1981), but an implicit assumption in virtually all that literature (the only expection I have seen is Taylor, 1975) is that the insureds whose experience generated the data will remain with the firm regardless of the premium structure it selects. But it is exactly the implausibility of this assumption that the second question addresses, and so we will begin with it. We use an example based on the SRI report (1976) to illustrate.

Assume that there is a large number of insurance firms all of which have identical numbers of type-1 and type-2 customers in their initial portfolios. Each firm only knows the experience of its own insureds, and it can distinguish type-1 and type-2 persons perfectly. The true probabilities of loss are .15 and .0944 respectively, *each firm* has 2000 type-2 customers and 18,000 type-1 customers.

We can suppose therefore that each firm tries to use data on the loss experience of the sample of 2000 customers to estimate the type-2 population mean (which actually is .15). There will, however, be sampling errors, so that some firms will observe more than 300 losses, and others fewer. In the next round, if firms construct their premiums based on their experience by using credibility rules, the firms that experienced fewer than 300 losses will set P_2 below .128, and P_1 below 0.0969. In contrast, firms with more than 300 losses will set both P_1 and P_2 higher.

The question then is: can this pattern of prices, which are structured to produce rate adequacy for each firm, represent a competitive equilibrium? The answer is obviously not. Type-2 consumers will purchase from the lower priced firms, who are charging them less than 0.128. Even if these firms should also get all type-1 customers and charge them .0969, they will suffer from rate inadequacy, and if they obtain less than all low risks, or charge them less than .0969, their deficits will be even worse. Thus the use of traditional credibility methods is not consistent with competitive equilibrium; those firms that are "lucky" enough to have large market shares will suffer losses.

If customers are not distributed uniformly across firms, the results are even more striking. Suppose firms on average have 10 percent high risks in their portfolios, but there are some with very few high risk (type-2) customers. Then they will charge a premium close to 0.0944 to type-1 customers, and will attract almost all of them. But, because the experience of their few type-2 customers gets a low credibility weight, they will charge type-2 customers a premium that is also close to 0.0944 (though it will be higher than the premium for type-1 customers). Consequently, such firms will get virtually all customers of both types. But since their average premium will be below 0.1, they will suffer losses.

If firms sell to both high and low risk groups, then the final level of profit using traditional credibility approaches depends on the mix of such groups. Empirically, it seems that less than fully credible groups are more likely to be high risk than low risk, so that the prediction of negative profits is still warranted.

This type of credibility bias has been criticized in the literature, and firms now frequently are advised to use methods which yield unbiased estimates of cell means. Some of these methods are versions of least-squares curve fitting (to either a linear or multiplicative specification), and there have also been some empirical Bayes methods that take into account the credibility of large cells (Tomberlin, 1981). While there may still be some slight systematic bias in methods actually used--Fairley et al. (1980) allege that the multiplicative method usually employed still overcharges high risks--the bias appears to be small. Does this mean competitive equilibrium exists with $P_i \approx \phi_i$?

Unfortunately, the answer is negative. To see this, consider an unbiased estimation method that yields a $\hat{\phi}_i^j$ with the following properties:

$$\hat{\phi}_i^j = \phi_i + \varepsilon \tag{3}$$

where ε is an error term with mean zero and variance σ_ε^2.

If credibility is at issue, there will be some firms whose sample will, under ideal procedures, yield a $\hat{\phi}_i^j$ which is less than ϕ_i. But if they base their premium on this estimate, then they will attract all of the customers but will sustain a loss. The larger is σ_ε^2, the more serious a problem this will be. In any case, simply basing actual premium on estimated pure premiums cannot be consistent with competitive equilibrium once firms begin to detect that they always lose. That is, the use of credibility rules is not consistent with competitive equilibrium.

4. LONG-RUN COMPETITIVE EQUILIBRIUM WITH LESS-THAN-FULLY CREDIBLE DATA

The previous section showed that, under a wide variety of credibility procedures, firms will be likely to lose money if they based their actual premiums on pure premiums. The low-bidder always understimates; he is subject to the "winner's curse". One may suppose that firms will recognize this. They will recognize that, should they be the low bidder, they will get most of the business but suffer losses. From this speculation it is possible to think of a firm developing the following strategy. It does not know whether it will be the firm with the sample with the lowest mean, but if it is, it wants to make sure that it will not lose money. The firm can estimate, on average, how far away the lowest mean of N samples of some size will be from the true mean. Without more, it can expect to at least break even by adding this difference to its pure premium. If it is the low estimator because it had the lowest sample, it will get all the business and breaks even. If it is not the low estimator, it gets no business of this type and breaks even. One equilibrium strategy is therefore for firms to add this "risk premium" (the same for all firms) to their pure premium.

If there are sufficiently many firms that expected profits get bid down to zero after this correction, at what price will insurance finally be purchased? The price will be the estimate of the true value made by the firm with the lowest sample mean. (As noted above, however, this estimate exceeds the credibility pure premium.) For expected profits to be zero in the long run, those estimates must have an expected or average value which is equal to the true value. This is so whether "traditional" or "improved" credibility rules are used. Thus, in the long run equilibrium, P_i does indeed equal ϕ_i on average, even though, for all firms, $P_i^j > \hat{\phi}_i^j$.

While this equilibrium price will have an expected value of ϕ_i, the accuracy is also important. Is the equilibrium price for less-than-credible risk classes always equal to ϕ_i, or is their substantial variance around ϕ_i, variance which induces people to purchase non-optimal amounts of insurance? The answer depends on how precisely the lowest sample tracks the true value.

Unfortunately, there appears to be no general treatment of this matter in the bidding literature. Instead, results have been obtained with fairly arbitrary assumptions about the relationship between $\hat{\phi}_i$ of the lowest firm and ϕ_i. The work of Wilson (1977) and Smiley (1979) provides examples of these special cases. In what follows I restate their models in terms of insurance.

Consider a population of R potential insureds of a particular risk class, all of whom will purchase full coverage insurance at any premium below some reservation premium $\underline{P}$. There are N insurance firms, and each firm has obtained a sample of data on past losses of size R/N. Each sample of data can be thought of as random sample (without replacement) of the experience of the population R. We assume that a sample size R/N is not "fully credible".

Call μ_i the mean or average loss in the data firm i has observed. Given that it has observed μ_i in its sample, each firm i could estimate the distribution of possible values of the population mean μ and come up with a best guess estimate of μ, or $\hat{\mu}_i$. This process is really what credibility theory is intended to discuss, and we can represent it in a general way by the use of Bayes' formula, as

$$f_{\mu}^{i}(\mu/\mu_i) = \frac{f_{\mu_i}(\mu_i/\mu)g(\mu)}{\int_{\mu} f_{\mu_i}(\mu_i/\mu)g(\mu)d\mu} \tag{4}$$

where $f_{\mu}^{i}(\mu/\mu_i)$ is the posterior distribution of μ for firm i conditional on its having observed μ_i, and $g(\mu)$ is its prior probability density function, assumed to be the same for all firms. Then its estimate of μ, or $\hat{\mu}_i$, is just the mean of $f_{\mu}^{i}(\mu/\mu_i)$, or $\frac{\int_{\mu} \mu f_{\mu}^{i} d_{\mu}}{\int_{\mu} f_{\mu}^{i} d_{\mu}}$. Of course, we can represent the posterior distribution either as $f_{\mu}^{i}(\mu|\mu_i)$ or $f_{\mu}^{i}(\mu|\hat{\mu}_i)$.

The bidder's estimates $\hat{\mu}_i$, conditional on μ, are assumed to be iid with a cumulative distribution function (c.d.f.) $F_{\mu_i}(\hat{\mu}_i|\mu)$. The form of F_{μ_i} depends on the form of the initial distribution of loss-production events. It is assumed that F_{μ_i} is known to all firms, but aht no firm knows the $\hat{\mu}_i$ of another firm. (Thus all firms know the form of F_{μ_i} but not

μ itself. They perceive μ as an unkown parameter of the distribution of $\hat{\mu}_i$).

Each firm uses its knowledge of μ, μ_i, F_{μ_i}, $g(\mu)$, and the number of firms to formulate a premium bid ρ_i which will maximize the expected profit from selling insurance. The firm's profit is $(\rho_i-\mu)$ if it submits the lowest premium bid and zero otherwise. The rule for transforming μ_i or $\hat{\mu}_i$ into ρ_i is the i-th firm's pricing strategy, and it obviously corresponds to the crediblity problem that actuaries are trying to solve. The bidding rule is represented by the function $\rho_i(\hat{\mu}_i)$. Bidders are assumed to recognize their interdependence but not to collude. Each firm selects a pricing strategy function, $\tilde{\rho}_i(\hat{\mu}_i)$ $(=\tilde{\rho}_i(\mu_i))$, which maximizes its expected profit, given $\tilde{\rho}_j(\hat{\mu}_j)$. This process defines the set of Nash equilibrium strategies $[\tilde{\rho}_i]$. Finally, all bidders use the same equilibirum strategy $\tilde{\rho}$ so the equilibrium will be a symmetric equilibrium. If $\tilde{\rho}$ is differentiable and $\tilde{\rho}_i > 0$, the expected profit of the i-th firm, conditional on ρ_i and $\hat{\mu}_i$, is

$$E(\rho_i,\hat{\mu}_i) = \int_\mu (\rho_i(\hat{\mu}_i)-\mu) f_\mu(\mu|\hat{\mu}_i) F_{\hat{\mu}_i}^{n-1}(\tilde{\rho}^{-1}(\rho_i(\hat{\mu}_i))|\mu)d\mu \quad (5)$$

where $f_\mu^i(\mu|\hat{\mu}_i)$ is the posterior distribution of μ for firm i conditional on $\hat{\mu}_i$ (or μ_i), $\tilde{\rho}^{-1}$ is the inverse of the equilibrium strategy of the (n-1) other firms and $F_{\hat{\mu}_i}^{n-1}(\tilde{\rho}^{-1}(\rho(\hat{\mu}_i))|\mu)$ is the probability that firm i submits the lowest bid, given $\hat{\mu}_i$, μ, ρ_i and $\tilde{\rho}$.

Differentiation with respect to $\rho_i(\hat{\mu}_i) = \tilde{\rho}(\hat{\mu}_i)$ yields the necessary condition for an optimal bid:

$$0 = \int_\mu [\tilde{\rho}'(\hat{\mu}_i) F_{\hat{\mu}_i}^{n-1}(\hat{\mu}_i|\mu) f_{\hat{\mu}_i}(\hat{\mu}_i|\mu) -$$

$$(\tilde{\rho}(\hat{\mu}_i)-\mu)(n-1) F_{\hat{\mu}_i}^{n-2}(\hat{\mu}_i|\mu) f_{\hat{\mu}_i}(\hat{\mu}_i|\mu)] g(\mu) d\mu.$$

The terms in the square brackets capture the two offsetting effects on the firm's expected profit of raising its premium: the profit earned if it is the lowest bidder increases, but the likelihood of its being the lower bidder decreases.

Wilson shows that, if certain regularity assumptions are made about the relationship between the lowest μ_i and μ, the lowest bid which wins all the business converges in probability to μ as N approaches infinity. For finite but large N, this means that the price at which insurance is sold should be quite close to the true pure premium μ even though no firm knew the value of μ to start with.

This result has exceedingly strong and interesting implications. It means that, if the regularity conditions hold, then the price at which business is transacted is independent of the prior distribution $g(\mu)$; in effect, each bidder's knowledge that there is enough information in the system to estimate μ almost exactly induces bidders to follow behavior for which $\rho \approx \mu$ is the outcome.

The assumed regularity conditions do, however, put some limits on the application of Wilson's theorem to the insurance-credibility problem. The critical condition is that the minimum $\hat{\mu}_i$ that can be observed must be a strictly decreasing function of μ. That is, there must be a one-to-one correspondence between the lowest possible μ_i and μ. If the "samples" of firms are very small, and if μ is small and has a distribution like the Poisson, then the smallest possible sample mean will be $\hat{\mu}_i^{min} = 0$. Even if μ is varied over some range, the smallest sample is still likely to have zero losses in it, so that the regularity condition is not satisfied. Clearly, if this happens, then firms observing such a sample will have to base their bid on $g(\mu)$, and so the minimum bid will be somewhere between μ and $g(\mu)$. If the sample sizes are sufficiently large that the probability that the lowest sample mean is zero becomes very small, then the impact of $g(\mu)$ on the winning bid disappears. It would be desirable to establish analytically the relationship between the winning bid, μ, $g(\mu)$, and the sample size; one suspects that the result would be something similar to credibility, but established on a basis quite different from that of credibility theory.

The other desirable excercise would be to solve analytically for the winning bid as a function of n for specific values of R, the distribution of losses, and $g(\mu)$. The only attempt to do this of which I am aware is in Smiley's study of oil bidding (Smiley, 1979). He actually used a model designed by Rothkopf (1969), which is a special case of Wilson's model under the assumption that each bidder knows the distribution of the ratios of the bidders' estimate to the a priori estimate of the true value. Even with this restrictive assumption, and even after assuming convenient forms for the distribution of $F(\hat{\mu}_i|\mu)$, Smiley is able to derive an analytical result using Rothkopf's procedures only by assuming further that $g(\mu)$ is a flat and diffuse prior, so that "the bidder's prior expectations about (μ) do not shift the posterior expected value away from the estimate." Such an assumption seems quite restrictive, and leaves open the question of what will happen if $g(\mu)$ is not so loose.

One suspects that there may be some intermediate value of N which makes the actual price closest, on average, to the correct price. For if N = 1, the monopolist correctly estimates μ, but charges the monopoly price. But if R/N is small the lowest μ_i is almost surely going to be zero, so that the estimate of μ will depend primarily on imprecise prior beliefs. Some value of N between these two extremes may achieve an appropriate compromise between accuracy of estimate and reduction in monopoly distortion.

There are two extensions to this analysis that move in the direction of more realism, but also more complexity. One extension is to assume that the low bidder does not receive all of the insurance business, but rather only a fraction which is larger than the fraction received by the next highest bidder. The second extension is to assume that the total amount of insurance purchased is an inverse function of price; indivudual demand curves are not perfectly inelastic. We consider the first extension first.

To get a zero profit equilibrium when some insurance is sold at a ρ in excess of μ, we could assume that there is some fixed cost c to being in business at all. This avoids the problem that posting a very high price always yields a positive expected profit. Let ρ_i now be the *unit* price of insurance charged by firm i, and let X_i be the total number of units bought. The firm's expected profit is defined as

$$\hat{\Pi}_i \equiv \int_\mu (\rho_i X_i - \mu_i X_i - c) f_i(\mu) d_\mu \quad . \tag{6}$$

One simple way to set things up is to define X_i as a function of ρ_i, ρ_j, and n. Equilibrium then obtains with a *distribution* of ρ_i's such that expected profit is zero at every ρ_i, which in turn implies that the number of firms offering any price rises as the price rises. This gets us into the well-known complexities of specifying models in which there is an equilibrium distribution of prices. All we can say is that, in equilibrium, the premiums will still be of the form $\rho_i > \mu_i$ but now the difference will vary across firms. Whether the difference will be related to the $\hat{\mu}_i$ (and the μ_i), and, if so, how, is a topic that we will not pursue here. It is obvious that the average difference for any class of insureds is going to be related to search behavior by consumers as well as to the level and distribution of the μ_i, so here there will be an additional reason why ρ_i will diverge from μ_i in equilibrium.

If aggregate demand $\sum_i X_i$ is a function of price, then the equilibrium *quantities* will obviously depend on this responsiveness. How the equilibrium price distribution varies with overall demand elasticity needs to be incorporated into the solution of a problem like that is described in the previous paragraph.

5. INDUSTRY RATING BUREAUS, THEIR IMPACT ON COMPETITION, AND THEIR RELATIONSHIP TO CREDIBILITY

Insurance industry rating bureas in the United States typically perform five functions:

1) They compile and average past loss data for prespecified policyholder characteristics. (It is not clear if they could or would produce special tabulations for other non-standard characteristics, or whether, if they did, they would communicate the requests to other insurers).
2) They trend past loss data to furnish estimates of future losses.
3) They compile and average expense data.
4) Based on (2) and (3), and on assumptions about profit margins, they publish suggested premiums.
5) They file these suggested premiums in so-called "prior approval" states, where the premiums are usually approved by the state regulatory body as maximum and minimum premiums unless a firm can support a request for a deviation.

In all of these functions the bureaus are protected from anti-trust action by state law and by the McCarran-Ferguson Act exemption from federal scrutiny.

The question is whether some or all of these activities ought to be so exempted. The collection and compilation of past data is generally thought to be justified on two grounds:

1) "Collection of past cost data by an industry association has generally been considered lawful."

2) "In the case of joint pooling and calculation of past loss data, efficiencies are likely to be great and the anti-competitive potential small... Such collective activity is likely to have a procompetitive effect ... where many firms will not have a sufficiently large policyholder base to make their own actuarially sound computations." (National Commission for the Review of Antitrust Laws and Procedures, 1979).

Trending for future losses has been viewed as an open question. The real criticism of rating bureaus on anti-trust grounds is applied to bureau collection of data on expenses, and the use of that data to project future prices. (These prices might be "suggested" either to all insurance firms or to a state regulatory commission.) It is alleged, for example, that in Illinois where only past loss data may be exchanged, the market has functioned as well as or better than in states where all data is exchanged. (Ibid).

Consider an insurer firm i considering entering a "small cell" market, where it has only a partial sample of data. It must estimate a breakeven premium $\overline{\rho}_i = X_{ij}+e_{ij}$, where X_{ij} is its expected underwriting loss and e_{ij} is its expected expense. The crucial questions would appear to be:

1) Without any new data, how does it estimate X_{ij} and e_{ij}, and how does it respond to uncertainty?
2) What data would "sharpen" its estimate of either? If it knew the industry-wide X_{ij}, would there still be a great deal of variance in its estimate of e_{ij}, and how would that affect its pricing behavior? How would knowing *past* e_{ij}'s be different from other information on past cost data? We need to know how category-specific are the e's.
3) In what way is communication of the e's different from the approved communication of past loss costs that has generally not been found to be illegal?

In what follows I ignore the legality issue and concentrate on the social costs and benefits. The primary point I wish to make is that the social gain from collection and aggregation of loss or cost data is generally much less than the private gain to expected profit maximizing insurance firms. Restriction of their ability to assemble such data, far from reducing the overall efficiency of market operation, may actually improve welfare. Any gains from cutting the likelihood of collusion would be added to these production efficiency improvements. Consequently, the notion that anti-trust action would compromise legitimate efficiency advantages from information exchange is not well founded.

If insurance firms' owners are risk neutral, as would be expected in a world of diversified investment portfolios, then the primary social cost of incorrect premium estimation would be the excess burden associated with the purchase by consumers of the "wrong" amount of insurance. If the premium actually charged is above the risk neutral premium based on the full set of data, then some mutually beneficial transaction will not have occurred; too little insurance will be bought. If premiums are set too low, then an excessive amount of risk will have been transferred ex post, in the sense that less insurance and a lump sum transfer would be preferred both by insurer and by insurance firm if there are loading costs. In any case, it is clear that there can be some distortions in demand. However, these welfare costs are likely to be relatively small if demand is fairly inelastic, or if the quantity of insurance is constrained by legal rules (e.g., compulsory auto liability insurance). In contrast, the costs of producing information can be large, especially in a competitive equilibrium.

To see this, we need to model the process by which information about the loss experiences of firms is generated and aggregated. For simplicity, we will consider here only the two extreme possiblities: (1) *Separate firm equilibrium:* All firms use only their own past data, which is available at zero cost (the model discussed in the previous section), or (2) *Rating bureau equilibrium:* Every firm gets data on the industry experience in return for furnishing its own data and paying a pro-rated share of the cost of maintaining and using the data pool. We will not investigate here the possibility of combining subsets of data larger than a single firm's sample but smaller than the industry set, with some firms participating and some firms not. We only observe that there may well be equilibria with such combinations, and those equilibria may dominate either of the two extreme cases.

To illustrate how inefficiency can occur, we now show that a rating bureau equilibirum may be stable, but represent lower welfare than a separate firm equilibrium. Suppose that the information collected from all n firms in the industry would permit exact estimation of μ, but that collection of this information, analysis of the data, and distribution of the results to the firms has a total opportunity cost of \$C, or a cost per firm of $c = C/n$ for each of n identical firms.

If c is sufficiently small, the rating bureau equilibrium in which each firm provides its data, pays \$c to the rating bureau, and charges a premium $p = \mu+c$ may be a Nash equilibrium. Suppose all other n-1 firms except firm i already belong to the rating bureau; they will charge a price approximately equal to p, and firm i could join at a price of c. To decide whether it should do so or not, it must compare the profits it could achieve if it joined (zero) with the profits it would earn if it priced as best it could using only knowledge of μ_i but paying no rating bureau membership cost. Firm i knows that all other firms will charge $\mu+c$, but it does not know what μ is; it only knows μ_i.

Suppose that a firm observes some μ_i. Given this observation, the firm estimates a distribution of μ given by $f_i(\mu|\mu_i)$, and therefore a distribution $f(\mu+c|\mu_i)$. With this estimated distribution, the firm can then formulate a bid $p_i = E(\mu+c) - \varepsilon$. (We suppose that the firm is sufficiently small that $E(\mu+c)$ is unaffected by its sample.) The firm's expected profit conditional on having observed μ_i would therefore be

$$\Pi(\mu_i) = \int_{\underline{\mu}}^{\mu+c-\varepsilon} (p_i(\mu_i) - \mu) f(\mu|\mu_i) d\mu$$

where $\underline{\mu}$ is the lowest possible μ. We can think of this expression as having two parts: (1) over the range from $\underline{\mu}$ to μ, the firm loses money, while (2) over the range from μ to $\mu+c-\varepsilon$ it makes positive profits. (Above $\mu+c-\varepsilon$ it gets no business.) If it happens that the second part is greater than the first, the firm has positive expected profits for this μ_i. One can think of estimated expected profits for all values of μ_i, given the distribution of μ_i (or $f(\mu_i|\mu)g(\mu)$). If it happens that this overall expected profit is positive, then the firm is better off by not participating in the rating bureau. In such a case, a rating bureau with full participation is not a Nash equilibrium. But if expected profits are negative, as would occur if c is small, then firm i is better off participating than not participating. But since all firms are identical, no firm will gain by not participating, so the rating bureau equilibirum is a Nash equilibrium.

In the rating bureau equilibrium, all firms charge a price $p = \mu+c$. In the separate firm equilibrium described in the previous section, the market equilibrium price has a mean of μ but a variance $\sigma \geq 0$. The tradeoff is obvious: by sacrificing \$C of resources society can eliminate the variation in premiums about its true value. The welfare loss from premium variation has two parts. First, risk averse persons are worse off if the premium is a random variable. But if the premium is small relative to wealth, this change is trivial. Second, one may compare the consumers' surplus from a price that is not always equal to μ to that from a price always set at $\mu+c$. The difference could be either positive or negative, depending on the size of c relative to σ^2. Even if more consumers' surplus is lost under the separate firm equilibrium, this amount must still be compared to the real cost of operating the bureau.

It is obviously possible that the latter exceeds the former, so that there is a welfare loss from collection of data. The reason for the loss is that knowledge of μ, which is what is being bought, does not affect what μ turns out to be. The information is useless in affecting the actual amount of real resource loss, however useful it is in estimating beforehand how much that loss will be. If people are risk averse, there is some utility gain from knowing the value of the expected loss beforehand, but that gain can well be less than the gain (in the sense of loss avoided) to firms from such knowledge. As in the Gaskins

(1976) model of auctioning oil leases, the private incentives to firms do not necessarily correspond to social benefits. If we add the possibility that the existence of rating bureaus may make cartel pricing behavior more likely, then there is still another efficiency cost to allowing rating bureaus. In short, prohibition of rating bureaus, or of similar exchanges of data, will do a little harm. But that harm may be much offset by the benefits in terms of resource savings and removal of temptation to deviate from competitive pricing.

6. CONCLUSION

These results suggest that one should be cautious in trying to infer what firms *do* do from what they *say* (and probably think) they are doing. Firms may actually think that they are adding a rule-of-thumb safety loading to their best actuarial estimate of losses, when in fact they are in the equilibrium of a bidding game. The inability of a firm to defend its actual rating practice is likely to be common.

In those cases in which firms have been asked to defend their rating and classification practices, as in the New Jersey automobile case, they have in fact been quite unable to do so. As the insurance commissioner of the State of New Jersey noted, for many of the rating practices, not a single industry witness could explain or point to data which justified them. All they could say was that they seemed consistent with adequate profits. The theory in this paper does not necessarily predict this behavior, but it would make it easier to understand. The indefensible equilibrium is not necessarily wrong.

REFERENCES

Bühlmann, H. 1970. Mathematical Methods in Risk Theory. New York: Springer Verlag.

Chaing, L., and W. Farley. 1979. Pricing Automobile Insurance under Multivariate Classification. Journal of Risk and Insurance, pp. 75-93.

Fairley, W., T.J. Tomberlin, and H.I. Weisberg. 1981. Pricing Automobile Insurance under a Cross-classification of Risks. Evidence from New Jersey. Journal of Risk and Insurance, pp. 505-520.

Freifelder, L. 1975. Statistical Decision Theory and Credibility Theory Procedures. In P.M. Kahn, ed., Credibility: Theory and Applications. New York: Academic Press, pp. 71-88.

Hewitt, C.C. 1963. Credibility for Severity. Proceedings of the Casualty Actuarial Society.

Hickman, J-C. 1975. Introduction and Historical Overview of Credibility. In P.M. Kahn, ed., Credibility: Theory and Applications. New York: Academic Press, pp. 181-192.

National Commission for the Review of Anti-trust Laws and Procedures. U.S. Department of Justice. 1979. Report to the President and the Attorney General, pp. 225-251.

New Jersey Department of Insurance. 1981. Final Determination, Analysis and Report. In Re: Hearing on Automobile Insurance Classification and Related Methodologies, April 1981.

Rothkopf, M.H. 1969. A Model of Rational Competitive Bidding. Management Science (March), pp. 362-373.

Smiley, A. 1979. Competitive Bidding under Uncertainty: The Case of Offshore Oil. Cambridge: Ballinger Publishing Co.

SRI, Inc. 1976. The Role of Risk Classifications in Property and Casualty Insurance, Final Report.

Taylor, G.C. 1975. Credibility under Conditions of Imperfect Persistency. In P.M. Kahn, ed., Credibility: Theory and Applications. New York: Academic Press, pp. 391-400.

Tomberlin, T.J. 1981. Empirical Bayes Estimation for Poisson Parameters in a Two-way Model. Analysis and Inference, Inc., November 5, 1981.

Wilson, R. 1977. A Bidding Model of Perfect Competition. Review of Economic Studies (October), pp. 511-518.