

**Do We Drive More Safely When Accidents are More Expensive?
Identifying Moral Hazard from Experience Rating Schemes**

Mark Israel
Kellogg School of Management, Northwestern University
February 2004

Abstract

A major shortcoming of the growing empirical work on asymmetric information is the inability to separately identify moral hazard from adverse selection. Abbring et. al. (2003) point out that dynamic insurance data can help here, by asking whether consumers have fewer claims when they are at a place in the “experience rating” scheme where additional claims are more expensive. However, in the French setting they study, this test boils down to asking whether there is negative state dependence in claims occurrence, and thus requires them to assume away all other forms of state dependence in claims. This paper overcomes this problem by considering U.S data, where claims fall off consumer records after three years, creating an “insurance event” that changes a consumer’s position in the experience rating scheme with no simultaneous claim, and thus allowing identification of moral hazard even with fairly general controls for state dependence and unobserved heterogeneity. In addition, the U.S. data follow consumers for 10 years (vs. 1 in Abbring et. al) and contain a much wider range of claim price effects, both of which increase identification power. The paper’s core finding is a small, but statistically significant, moral hazard effect. This is made more convincing by the fact that the effect grows following a 1997 pricing change that increased the cost of additional claims relative to the first one. Finally, without the controls for state dependence in claims, this effect disappears. This suggests that the lack of evidence for moral hazard in previous work may have resulted from confounding the negative state dependence associated with moral hazard with some underlying source of positive state dependence in claims occurrence.

I. Introduction

After years of a steadily widening gap between the extremely intricate theoretical work on asymmetric information – solidifying the conventional wisdom that both adverse selection and moral hazard are serious problems in insurance markets – and the extremely thin empirical work measuring its real world importance, the last few years have seen a surge in empirical research on this topic (see Chiappori and Salanie, 2003, for an excellent summary). The bulk of this work uses the so-called “conditional correlation approach,” which regresses both an individual’s insurance claims and her coverage choices on all observables on insurance company files, and then asks if the error terms from the two regressions are correlated. If they are, this provides evidence for asymmetric information, as it means that individuals who have more claims than predicted also buy more insurance.

While a very useful approach, one of its major shortcomings is that it can not distinguish between adverse selection and moral hazard. That is, we can not say whether a consumer knew her claims risk was higher than predicted by observables -- so that she was “getting a deal” on insurance -- and thus bought more as a result (adverse selection), or whether she bought more coverage for some unobservable reason and thus took less care in driving (moral hazard). This is unfortunate as the theoretical work has established that the two phenomena suggest distinct policy reactions.

This paper builds on recent suggestions that the dynamics of consumer-firm relationships can provide the identification power required to separate these phenomena. One approach is to derive the properties of optimal dynamic insurance contracts under each form of asymmetric information and then to compare them to actual contracts. However, the form of these contracts is very hard to derive, particularly after one takes the necessary step of incorporating heterogeneous risk aversion. Even more importantly, a simple look inside an insurance company makes the assumption of dynamically optimal contracts hard to maintain. This is largely because insurance companies are under intense regulatory scrutiny, and much of their focus in designing contracts is to achieve simple, easy to explain, stable contract forms that minimize regulatory attention.

Instead, this paper follows the approach developed in Abbring et. al. (2003). The idea is to directly ask the key moral hazard question – do consumers change their driving behavior, and thus risk, in reaction to incentives in the insurance contract? The fundamental difficulty in empirically answering this question is that the terms of that contract are endogenously chosen by each individual, reflecting full information about her own risk. In light of this, the key advance in Abbring et. al. is to point out that experience rating schemes (where prices are based on claims histories) expose consumers to uninsurable

“pricing risk.” Hence, a consumer’s reaction to changes in the level of this risk, resulting from changes in her claims record, enables separate identification of moral hazard.

Abbring et. al. apply this logic to French auto insurance data. In the French “bonus/malus” system, each consumer faces some base premium as a function of her characteristics. Following each period, this is multiplied by 0.95 if the consumer had no claims, or 1.25 for each claim. Since the 25% increase is bigger when multiplying a bigger premium, each additional claim is effectively more expensive. This implies negative state dependence in claims -- following each claim, subsequent claims become more expensive and thus, under moral hazard, should become more rare -- for which the authors test using data from a single contract year. The bulk of their paper is dedicated to establishing that this can be separately identified from any unobserved, individual effects in claims risk, the basic idea being that the individual effect determines the expected number of claims for the full period, while the state dependence determines the timing of those claims.

The current paper relies on a 10 year panel of more than 30,000 consumers from one U.S. automobile insurer in Illinois to build on this work in several important ways. First, because the test in Abbring et. al. simply looks for negative state dependence, it has to assume away all other forms of state dependence in claims. This includes not only other sources of negative state dependence – say, if consumers learn and thus drive more safely following claims – but also things that would show up in the data as positive state dependence. For example, if beyond the fixed, individual effects, a consumer has specific periods of time when his risk is higher – perhaps because he’s loaning the car to his teenage son, because he’s under extra stress at work, etc. – then claims will cluster in those times, creating the appearance of positive state dependence that might confound attempts to detect the moral hazard induced, negative state dependence. In contrast, in the U.S. data used here, each consumer has a base premium, plus an addition for recent claims. Whether that addition is a low, middle, or high percentage of the base is *determined by claims experience in the last 3 years*, with more recent claims leading to a higher increase. Crucially, this 3 year window creates an *insurance event* – 3 years after each claim, it drops off the consumer’s record, changing his position in the experience rating scheme with no simultaneous claim. Using these insurance events, and the associated changes in the price effects of additional claims, we can identify the impact of moral hazard, even after including very general controls for *both* state dependence in claims and unobserved heterogeneity in risk.

Second, in 1997, the study firm changed the amount of the middle and high percentage increases for claims – from 20% and 50% to 40% and 70%, while leaving the low percentage at 10%, and the rest of the pricing structure unchanged. So, in the presence of moral hazard, this should increase the care

taken by drivers facing the higher increases, and thus make the claim reduction effect in these “high claim cost” periods even larger. This basically creates a “differences in differences” test – as those consumers facing the middle and high increases should see a drop in claims after the change, with no change for those facing the low increase.¹

Third, the data used here are a 10 year panel, vs. the 1 contract-year panel of French consumers used by Abbring et. al., so that each consumer is observed in a wider variety of “positions” in the experience rating scheme, increasing identification power. In contrast Abbring et. al. depend on multiple claims in a given year, something which is fairly rare, to achieve identification.

Finally, one criticism of the Abbring et. al. approach is that the induced changes in price increases from claims are quite small. That is, if the first claim costs 25% of the base, the second claim is more expensive simply because it applies another 25% increase on top of the first, meaning that it is more expensive by $0.25 \times 0.25 = 6.25\%$ of the base. The third claim, then, is more expensive than the second by just under 8.0% of the base. If the base premium were, say, \$250, the price increase for the first claim would be \$62.50, going up to \$78.13 for the second, and \$97.65 for the third. The changes in price increases are much larger in the data used here. For instance, following 1997, the price increase on a \$250 base policy ranges from \$25 for those facing the low increase to \$175 for those facing the high increase, a dramatic difference that ought to induce some reaction if moral hazard has any measurable impact.

The main result of the paper is evidence for a small, but statistically significant, moral hazard effect. That is, there is evidence that controlling for observable characteristics, and allowing for fairly general state dependence in claims and unobserved heterogeneity, claims are less frequent when they are more expensive. This is true both when simply using the position in the rating scheme (low, middle, or high increase) on the right hand side, and when including the actual price increase. Making this more convincing, the size of these effects clearly increases after the pricing changes of 1997. Finally, and perhaps most interestingly, without the controls for state dependence in claims (which show evidence of positive state dependence) this effect disappears. This suggests that the lack of evidence for moral hazard in previous work may have resulted from confounding the moral hazard induced, negative state dependence with some underlying source of positive state dependence.

¹ Actually, this is more complicated due to dynamic effects. That is, consumers facing the 10% increase know that with any claims today, future claims become even more expensive under the new system. This could affect their driving behavior today. However, as discussed in some detail below, we find no evidence that consumers react to any such dynamic incentives.

The remainder of the paper proceeds as follows. Section II provides details on the dataset used in analysis. Section III describes the experience rating scheme used by the study firm. Section IV lays out the empirical framework used in testing. Section V contains the results. Section VI concludes.

II. Data

A. Basics

All analysis relies on a panel of 31,215 consumers joining one Illinois auto insurance firm between January 1989 and December 1998. Each consumer is observed beginning with her initial purchase from the firm. The record for each consumer is divided into 6-month periods, the length of an auto insurance policy.

For each period, the dataset includes the total number of claims, broken down by type. The 4 types are: *liability claims*, the “at-fault” claims which result in bodily injury or physical damage to another vehicle; *collision claims*, which are basically single car accidents; *comprehensive claims*, which include “non-accident” sources of damage such as storm damage, theft, etc.; and *other claims*, primarily made up of roadside breakdowns and damage caused by uninsured motorists.² Table 2.1 contains a frequency distribution of claims by consumer period. Because periods with more than 1 claim of any type are very rare, all analysis in the paper fits the probability of 0 vs. 1+ claims.

There is some debate in the literature about which claims to include in the dependent variable in studies of moral hazard. The occurrence of comprehensive claims and other claims are **not** used in the experience rating scheme, suggesting that firms believe they’re outside individual control, so they are excluded from analysis. The more questionable category is collision claims, which may be under individual control. Because consumers may choose not to turn in each collision as a claim, many studies exclude these. But as Abbring et. al. (2003) point out, that decision represents “ex-post moral hazard” (as opposed to the “ex-ante moral hazard” that refers to care taken in driving) and thus might be included in a comprehensive study of moral hazard. However, these claims also create the problem that – because all that is legally mandated is liability coverage -- some consumers are not even covered for such “first party” damage. So, for simplicity, all analysis in this paper relies on liability claims alone.

The dataset also includes all variables used to set the price of this liability coverage. This begins with the “rating class,” determined by age, gender, marital status, vehicle usage, other miscellaneous discounts (good student status, defensive driving courses, etc.), and an adjustment for claims in the 3

² To be clear, note that accidents in which another insured driver is at fault do not show up as claims, since they are covered under that driver’s insurance policy.

years prior to joining the firm.³ This yields 159 categories. In addition, prices depend on the consumer's "rating territory," generally defined by county, although broken down by zip-code in some highly populated counties. Finally, price is multiplied by an experience rating factor, described in detail in Section III.

B. Attrition

One major complication in analysis is that a substantial portion of consumers depart before the end of the sample. This occurs for 2 equally common reasons, as shown in Table 2.2. First, consumers may move to another state, or receive a new policy number for a variety of administrative reasons. In either case, a notation is made explaining why they have fallen off the sample. Such attrition is considered random here and thus presents no problem in estimation – each consumer simply contributes to the likelihood function for those periods when he is on the sample.

But, second, some consumers voluntarily drop off with no such notation, which (because auto insurance is mandatory in Illinois) presumably means they have switched to another firm. This is potentially problematic. If some of these consumers are leaving in reaction to the experience rating scheme, results could be confounded by sample selection. Fortunately, over the sample period, firms in Illinois share claims data, so consumers can not switch firms to flee a claims record, lessening this concern. Generally the way this works is that the new firm will adjust the consumer's initial price to reflect claims in the previous 3 years. But, while most firms have an experience rating scheme very similar to the one described here, new consumers generally start in the "low price increase" category whatever their claims history. So while consumers who switch would have to pay a price reflecting their full history, it is possible that consumers switch firms to reset their position in the experience rating scheme.

As a result, all analysis controls for sample selection. That is, the likelihood function is based on the joint probability that a consumer chooses to remain with the firm through a period and that she has a claim during that period, with the individual effects from the two equations allowed to be correlated. The selection equation includes all variables in the claim probability function. In addition, to help separate any selection effects, the difference between the price at the study firm and the average price for the consumer's rating class in the broader market -- which does not directly impact the probability of a claim

³ In order to be sure that a full 3 year period is observed for all consumers, consumers are only included if they are age 20 or over upon joining the firm.

– is also included in the selection equation, to capture the fact that consumers are more likely to depart if the current firm looks worse relative to the market.⁴

The remaining details of this selection equation and the claim probability function are presented in Sections IV , with results in Section V. The most important result is that, after looking at a wide variety of specifications, no variables related to experience rating were ever found to be significant in the selection equation, lessening concerns that selection effects drive the results.

III. Experience Rating Scheme

A. Basics

Any “chargeable claim” – basically liability and collision claims – that an individual submits increases the price of her insurance for the 3 years (6 policy periods) following the claim. The size of the price increase is found by multiplying the base price – determined by the rating class and territory, discussed above – by an experience factor, determined by claims history. As written in the firm’s pricing manual (prior to 1997):

The charge for the specific accident shall be:

- i. 10% if there were no other chargeable accidents during the three years preceding the date this accident became chargeable
- ii. 20% if there was only one other chargeable accident during this period and its charge was 10%
- iii. 50% in all other instances

In 1997, the percentages in (ii) and (iii) were increased to 40% and 70% respectively, with no other changes.

This system is quite simple. New consumers, and individuals who have had no claims in the last 6 policy periods, are in classification (i), thus facing a 10% price increase if they have a chargeable claim. Following this chargeable claim, they move to classification (ii), where an additional claim costs 20% (40% from 1997 on). If they have such an additional claim, they move into classification (iii) where an additional claim costs 50% (70% from 1997 on). If at any point they go 6 policy periods with no

⁴ While primarily a function of a consumer’s rating class, this price difference does vary independently from the rating class for several reasons. First, it changes from year to year due to changes in the price menus at the study firm and other firms in the market. Second, prices also depend on whether consumers receive a multi-car or multi-line discount – which is bigger at the study firm than most other firms – so consumers receiving these discounts should be less likely to switch. Finally, consumers change rating classes frequently, and while a consumer’s current rating class determines her claims risk, past rating classes determine the probability that the consumer has remained in the sample (see Israel, 2003, for a detailed development of the idea of using past rating classes to control for selection effects).

chargeable claim, they move back to classification (i). This is the “insurance event” that allows identification of moral hazard effects with general controls for state dependence in claims.

Note that only the most recent claim matters for the marginal price of a claim. That is, suppose an individual has a claim in period 1, raising his price by 10% and moving him into classification (ii). Then, he has a claim in period 4, raising his price by 20% and moving him into classification (iii). At the end of period 7, the first claim drops off his record lowering his price by 10% of the base, *but the cost of any additional claims remains 50%*. Put simply, while the current insurance price depends on the price increments of all claims in the past 3 years, the cost of the next claim is determined only by the characteristics of the most recent claim (namely, did it occur in the last 3 years and how much did it cost?)

All of this means that we can allow very general controls for state dependence in claims, while still separately identifying the moral hazard effects of interest. It should be obvious that any measure of claims history that doesn’t rely on the number of claims in the last 3 years can be included. But the details discussed above imply that we could even include number of claims in the last 3 years, since consumers with, say, 1 claim in the last 3 years could be in classification (ii) or (iii) depending on the timing of those claims.

In the results presented below, whether the consumer has experienced a claim in the past 6 months and/or the past year have statistically significant, positive coefficients, even with controls for unobserved heterogeneity. But no other measures of claims history were significant in any analysis, and using other measures of claim history had no substantive effect on the results.

B. Current Price vs. Dynamics

All analysis to this point has focused on the immediate impact of a claim on the price of insurance. But as Abbring et. al. point out, claims have dynamic effects as well. That is, a claim not only changes a consumer’s insurance price for the next 3 years, but also moves her to a new classification, changing the price effect of any subsequent claims in that 3 year period and thus impacting her value function going forward. Abbring et. al. proceed by proving that, in the French system, these current price effects and value function effects go in the same direction, so that both imply negative state dependence.

In the current context, however, the dynamic effects do not necessarily go in the same direction as the current price effect. For example, when in classification (iii), an individual faces the largest price increase, but remains in classification (iii) even following a claim. In contrast, an individual in

classification (i) or (ii) faces a smaller price increase, but moves up to more expensive classification following a claim.

The approach of this paper is to empirically test the importance of these dynamic effects. That is, analysis starts from the point of view that, *because additional claims and thus additional price increases are rare*; the most important effect of a claim is to increase the current price. So, current price effects are used as the first order measure of moral hazard effects. But, we then go on to empirically ask whether consumers react to the dynamic incentives.

At first blush, it may seem difficult to separate the dynamic effects from the current price, since both are determined by the consumer's classification. That is, a consumer in classification (ii) necessarily faces both a 20% price increase today and a shift to a 50% increase for future claims. But a closer look at the role of time provides a way to separately test for dynamic effects. To see this most simply, suppose that a consumer is in classification (iii), but has had no claims for the past 5 periods. Then, he is one claims-free periods away from returning to classification (i). If he has a claim, however, his classification will not change, but that count will be reset, so that he'll again be 6 periods away from moving to classification (i). In contrast, for a consumer in classification (iii) who had a claim last period, a new claim this period only increases the wait to return to classification (i) by one period.

Clearly, then, if dynamic effects matter – that is, if consumers worry about the effect of current claims not only on prices, but also on the cost of subsequent claims – then time since most recent claim will impact current driving behavior. That is, consumers in classification (ii) or (iii) should take greater care as they accrue more claims free periods, as a claim means that they'll give up that accrued state. However, the results presented below show that, while there is evidence that the current price effect matters, there is absolutely no evidence for this timing effect, suggesting that consumers do not react to these more subtle dynamic incentives.⁵

IV. Basic Empirical Specification

For each consumer period , the model consists of two equations – the probability of a liability claim and the probability of completing the period with the study firm – that together form a bivariate probit, including separate individual-level random effects for each equation, which are allowed to be correlated. The basic specification of the claims equation, for a given individual in period t , is given by:

⁵ If these timing effects were important, then we would need to worry a great deal about solving the full dynamic model to arrive at the correct measures of dynamic incentives. But, with no evidence that these basic timing effects matter, it is hard to believe that any more complex dynamics play a major role.

$$\Pr(Claim_t) = \Phi(\alpha_1(exp_t) + \alpha_2(dyn_exp_t) + \alpha_3(clm_hist_t) + \alpha_4(RC_t) + \alpha_5(Terr_t) + \eta) \quad (4.1)$$

where $\Phi(\cdot)$ is the standard normal CDF. $Claim_t$ is a 0/1 indicator for a liability claim in the period. The function $\alpha_1(exp_t)$ captures the impact of the individual's position in the experience rating scheme, and $\alpha_2(dyn_exp_t)$ captures the dynamic effect of claims on future claims costs, determined by the number of periods since the most recent chargeable claim. These are the functions of primary interest, and 2 different specifications are reported in the results section, below.

$\alpha_3(clm_hist_t)$ captures any other state dependence in liability claims. The specific form used in all specifications is:

$$\alpha_3(clm_hist_t) = \alpha_3^1 Tot_Clms_t + \alpha_3^2 Clm_lastyr_t + \alpha_3^3 Clm_lastper_t \quad (4.2)$$

where Tot_Clms_t is the individual's total number of observed liability claims; Clm_lastyr_t is an indicator for whether the consumer had a liability claim in the previous year, and $Clm_lastper_t$ is an indicator for whether the consumer had a liability claim in the previous 6-month policy period.⁶

$\alpha_4(RC_t)$ is a set of indicators for all 159 rating classes (explaining why the model has no constant term), while $\alpha_5(Terr_t)$ is a set of indicators for all rating territories except territory 1. These are included as controls, so the results are not reported, but are available from the author on request. Finally η is an constant, individual level random effect, capturing unobserved heterogeneity in claims risk.

The probability of remaining with the firm through the policy period is specified in the same fashion, with an additional term for the difference between the price at the study firm and the average price from other firms in the market, $pdiff_t$:

$$\Pr(Remain_t) = \Phi(\beta_1(exp_t) + \beta_2(dyn_exp_t) + \beta_3(clm_hist_t) + \beta_4(RC_t) + \beta_5(Terr_t) + \beta_6 pdiff_t + \nu) \quad (4.3)$$

where the functions are all defined in the same fashion as the α functions. The estimates of the $\beta_1(exp_t)$ function are presented with each specification below, to show that the decision to remain or switch appears to be uncorrelated with a consumer's position in the experience rating scheme.

⁶ Recall that claims are only observable for 3 years prior to joining the firm, so Tot_Clms_t only covers this period. Also note that a wide range of specifications – including different time frames, indicator variables vs. counts, etc. – were tried. Only Clm_lastyr and $Clm_lastper$ were ever significant, and the main results were robust to the various specifications.

For simplicity, each random effect distribution is approximated with a discrete distribution with 2 points of support. So,

$$\begin{aligned}\Pr(v = v_1) &\equiv \rho_v \\ \Pr(v = v_0) &= (1 - \rho_v)\end{aligned}\tag{4.4}$$

where v_1 and ρ_v are estimated, while v_0 is constrained to equal $-\rho_v v_1 / (1 - \rho_v)$ to satisfy the 0 expected value restriction. To allow correlation, the probabilities of the η distribution are specified to depend on the value of v as:

$$\begin{aligned}\Pr(\eta = \eta_1 | v = v_1) &= \rho_\eta^1 \\ \Pr(\eta = \eta_0 | v = v_1) &= (1 - \rho_\eta^1) \\ \Pr(\eta = \eta_1 | v = v_0) &= \rho_\eta^0 \\ \Pr(\eta = \eta_0 | v = v_0) &= (1 - \rho_\eta^0)\end{aligned}\tag{4.5}$$

where η_1 and the two probabilities are estimated, with η_0 satisfying the 0 expected value restriction:

$$\eta_0 = \frac{-(\rho_v \rho_\eta^1 + (1 - \rho_v) \rho_\eta^0) \eta_1}{(\rho_v (1 - \rho_\eta^1) + (1 - \rho_v) (1 - \rho_\eta^0))}\tag{4.6}$$

These error terms are independent of all other variables in the model, meaning that these are the relevant distributions, even after conditioning on all x variables.⁷

Estimation relies on maximum likelihood. In each period that a consumer is not randomly or right-censored, she adds a term to the likelihood function. At a minimum, this includes the probability of her observed stay/depart decision. If she remains for the full period, she also contributes the probability of her observed claims. If she leaves during the period, the full set of claims is not observed, so no claims term is included. The product of all these stay/depart and claims probabilities for one individual, given a value of v and η , yields a conditional likelihood contribution for the individual. The individual's overall likelihood contribution is the probability weighted sum of these conditional likelihood contribution.

Estimation proceeds by maximizing the sum of the log likelihood contributions for all 31,215 consumers on the sample.

⁷ This means that η is the unobservable portion of claims risk *after controlling for claims in the last 3 year*, and thus is constructed to be independent of that claims history for simplicity. A more complex specification, in which different error distributions were estimated for consumers with 0 or 1+ claims in the last 3 periods was estimated, with no substantive change in the results.

V. Results

A. Categorical Variable Specification

The first specification simply relies on categorical variables for experience classification. That is, $\alpha_1(exp_t)$ is specified as:

$$\alpha_1(exp_t) = \alpha_1^1 * Cl_ii + \alpha_1^2 * Cl_iii + \alpha_1^3 * Cl_i_97 + \alpha_1^4 * Cl_ii_97 + \alpha_1^5 * Cl_iii_97 \quad (5.1)$$

where Cl_ii is an indicator for experience classification ii, Cl_iii is an indicator for experience classification iii, and the remaining variables are indicators for each classification in 1997 or later, after the pricing change. A symmetric form is used for the $\beta_1(exp_t)$ function. For the dynamic effects, we consider whether the probability of a claim falls in the late periods in classification ii or iii, specifying:

$$\alpha_2(dyn_exp_t) = \alpha_2^1 * per_4 + \alpha_2^2 * per_5 + \alpha_2^3 * per_6 \quad (5.2)$$

where per_n is an indicator for the nth period in either classification ii or iii.

Column 1 of Table 5.1 presents the results (with estimates that are significantly different from 0 at the 5% level shown in bold). Before considering the effects of the experience rating scheme, consider the evidence for state dependence in claims, shown on page 2 of the table. While the total number of claims is not significant, recent claims have a strong positive effect on the current claim probability. Using the sample average claim probability (2.9%) as a base, the estimates imply that a consumer who had a claim in the last year has a 0.6 percentage point higher chance of a claim this period, while a claim in the last period increases this probability by another 2.0 percentage points. This positive state dependence is the largest effect found in any estimates. Since it's hard to see a positive, causal connection between claims yesterday and today, the most likely explanation is that individuals have certain periods when their risk spikes up, leading to clustering in claims occurrence.⁸ Whatever the source, this strong effect calls into question attempts to identify moral hazard by looking for negative state dependence in claims.

With these controls in place, the results show evidence that, as predicted by moral hazard, consumers are less likely to have claims in those periods when they are more expensive. As shown on Page 1 of the Table, when a consumer is in Classification (ii), her claims probability is lower by roughly

⁸ If this pattern was the focus of analysis, we could try to specify a richer error term to capture this. But for the purposes here, capturing it with variables for claims in recent periods should be fine.

0.1 percentage points, a small but statistically significant effect. To provide some perspective, note that for an individual paying the sample average 6-month premium of \$250, moving from classification (i) to (ii) increases the cost of a claim by \$25 in the pre-1997 pricing scheme. Moving to classification (iii), which increases the cost of a claim by \$75 relative to classification i, reduces the claim probability by roughly 0.3 percentage points. These results provide strong evidence for moral hazard, particularly since most alternative theories would suggest *more* claims for consumers in the bad parts of the experience rating scheme.

Even more convincing, the results also show that after 1997, when the cost of a claim increased for classifications (ii) and (iii), but not for classification (i), the probability of a claim decreased in classification (ii) and (iii), with no significant effect in classification (i). In both classifications (ii) and (iii), the additional \$50 in the cost of a claim (for an average policy) reduces the claim probability by roughly another 0.1 percentage points.

Column 2 reports results with no controls for state dependence in claims. In this case, none of the experience rating variables are significant. This provides strong evidence that, without such controls, the negative state dependence in claims resulting from moral hazard is confounded with the underlying positive state dependence in claims. Because this may help to explain why previous work has found no evidence of moral hazard, it is an important avenue for additional research.

The estimates of dynamic effects are in the second section of the table. As explained above, if consumers react to the dynamic incentives – that is, to the impact of a claim on the cost of future claims – the claims probability should fall as they accrue periods in classification (ii) or (iii). The estimates show no evidence of such effects, as the claims probabilities in periods 4 through 6 are not significantly different from the period 1-3 base, and there is no pattern of declining probabilities across these 3 periods.

Finally, the last section of the table, on page 2, shows that none of the experience rating variables are significant in the departure equation, lessening any concerns that selection may be driving the results.

B. Continuous Price Variables

While the categorical specification is simple and its results fairly convincing, the change in the cost of a claim and thus the strength of the reaction, should be higher for consumers paying higher prices. Note that this does not rely on the obvious fact that consumers paying higher prices have higher risk. Instead the prediction is that consumers paying those higher prices will see their risk *decrease by more* when they move into one of the higher cost classifications.

To implement this logic, we specify $\alpha_i(exp_i)$ as:

$$\alpha_1(exp_t) = \alpha_1^1 * pr_incr + \alpha_1^2 * (pr_incr)^2 \quad (5.3)$$

where *pr_incr* is simply the amount by which the consumer's price will increase following a claim. Deciding which price to use is somewhat problematic. The full price increase is based on the full price of the consumer's liability and collision coverage. However, the choice of how much coverage to buy is endogenous. In particular, under adverse selection models, consumers who choose to buy more coverage do so because they know their risk is greater, so higher prices could be associated with unobservably higher risk, confounding attempts to measure moral hazard. To fix this, we can rely on the price of the base liability coverage, which all consumers must buy, so that price differences are based only on the observable rating classes. But this leaves out a source of price differences across consumers. So, we estimate the model in both ways – columns 1 and 2 of Table 5.2 use the base liability price, while columns 3 and 4 use the full price of liability and collision coverage. Fortunately the results are not sensitive to this choice, so the remainder of this discussion will focus on columns 1 and 2.

The dynamic effect on future claims costs also varies with the price of insurance. However, it is not clear what the best simple measure is, as the full dynamic effect depends on the number and timing of all future claims. The measure used here computes the cost of one more claim in any of the next 6 periods, given a claim in the current period, minus the cost of such a future claim with no claim in the current period. Assuming that there are no other changes in the consumer's insurance coverage or rating class, it is easy to compute this claim cost for each period using the experience rating scheme⁹ Note that the change in this claim cost depends in part on how quickly the consumer will return to classification (i) and thus follows the logic developed above. The main complication is deciding how to aggregate these period by period claim costs. For simplicity, we simply use the discounted sum of the costs, using a 6-month discount factor of 0.97. Defining this discounted sum as *incr_cost*, we define $\alpha_2(dyn_exp_t)$ as:

$$\alpha_2(dyn_exp_t) = \alpha_2^1 * incr_cost + \alpha_2^2 * (incr_cost)^2 \quad (5.4)$$

Many other specifications were attempted – using different discount factors, allowing multiple claims, etc. – with no change in substantive results.

Results (presented in Table 5.2, with prices in \$100) confirm the findings from the categorical model. A larger price increase from a claim reduces the claim probability by a small, but statistically

⁹ We assume that consumers do not anticipate the pricing structure change in 1997, and thus use the *current* pricing scheme to compute all future price changes.

significant amount. The positive quadratic term suggest that this effect declines at higher prices, although this term is only significant when using the full policy price. If we evaluate the effect at the sample average claim probability for an individual paying \$150 for the core liability coverage (roughly the sample average), the results suggest that a \$100 increase in the cost of a claim reduces the claim probability by 0.4 percentage points. The magnitude of this effect is consistent with the findings in the categorical specification.

All the other results are consistent as well. The underlying, positive state dependence in claims remains the largest measured effect. Models that don't control for this effect (columns 2 and 4) fail to find significant moral hazard effects. There is no evidence that consumers react to dynamic incentives. And there is no evidence that the experience rating variables impact the decision to depart the firm.

VI. Conclusion

In recent years, empirical work on asymmetric information has begun to close the gap with theoretical work, at least in terms of quantity. However, the fairly blunt tests that are common still lag far behind the extremely nuanced theoretical literature. The rich dynamic data available from insurance company files – containing several years worth of both purchase choices and claims, and following consumers as they change rating classes – provide a mostly untapped opportunity to close this gap.

In this spirit, and following the lead of Abbring et. al. (2003), this paper uses dynamic data to separate moral hazard from other sources of asymmetric information. The key idea is that experience rating schemes – through which insurance prices are a function of a consumer's claims history – imply that the price increase associated with a consumer's next claim varies with the consumer's claims history. And crucially, in the presence of moral hazard, claims should be less common when they are more expensive.

This paper contributes to this approach in several ways. First, while Abbring et. al. have to assume away all other sources of state dependence in claims, the current paper relies on “insurance events” – when claims fall off consumer records after 3 years – to identify the moral hazard effects even with general controls for both state dependence and unobserved heterogeneity in claims risk. Second, the panel used here runs for 10 years (vs. 1 in Abbring et. al.) and the experience rating scheme raises prices much more steeply with claims, both of which increase the variance in claims costs faced by any given consumer, increasing identification power. Finally, over the course of the panel, the price increase associated with multiple claims was raised, while the price increase for the first claim was left unchanged.

So, we can refine the test by asking whether the risk of subsequent claims fell by more than the risk of the first claim following the pricing change.

The results provide evidence for a small, but robust and statistically significant, moral hazard effect, even with general controls for state dependence and unobserved history in claims risk. And this result becomes measurably larger after the pricing structure change. Most interestingly, without the controls for state dependence in claims – which show that the presence of recent claims make current claims *more likely* – the moral hazard effects are not statistically significant. This suggests that the lack of evidence for moral hazard in previous work may have resulted from confounding the moral hazard induced, negative state dependence with some underlying source of positive state dependence.

It is also interesting to contrast this result with the general lack of evidence for *any* asymmetric information, including moral hazard, in previous work on auto insurance using the conditional correlation approach (Chiappori and Salanie, 2003). That work asks whether consumers who purchase more coverage than predicted by observables have more claims. To reconcile the lack of a relationship there with the evidence for moral hazard found here, we must conclude that there are other factors driving coverage choice, perhaps heterogeneous risk aversion, that confound the conditional correlation prediction. In any case, this suggests that we also need more nuanced tests for adverse selection. Dynamic data can help here as well, relying on changes in risk, prices, coverage choices, and claims through time, with the ultimate goal of simultaneously measuring both adverse selection and moral hazard in one model.

References

- Abbring, J.H, P. A. Chiappori, J. Heckman, and J. Pinquet (2003). “Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?” mimeo.
- Abbring, J.H., P. A. Chiappori, and J. Pinquet (2003). “Moral Hazard and Dynamic Insurance Data,” mimeo.
- Chiappori, P-A and B. Salanie (2003). “Testing Contract Theory: A Survey of Some Recent Work,” in *Advances in Economics and Econometrics – Theory and Applications, Eighth World Congress*. M. Dewatripoint, L. Hansen, and P. Turnovsky, ed. Econometric Society Monographs, Cambridge University Press, Cambridge. pp. 115 – 149.

Table 2.1: Claims Frequency

Variable	Count	Frequency	Percent
Total Consumer Periods	N/A	211,128	100.0
Claims in Period	0	194,026	91.9
	1	15,412	7.3
	2 ⁺	1,690	0.8
Liability Claims in Period	0	204,983	97.1
	1	6,034	2.9
	2 ⁺	111	0.1
Collision Claims in Period	0	208,060	98.5
	1	3,003	1.4
	2 ⁺	65	0.1
Comprehensive Claims in Period	0	208,052	98.5
	1	2,851	1.4
	2 ⁺	225	0.1
Other Claims in Period	0	205,176	97.2
	1	5,714	2.7
	2 ⁺	238	0.1

Table 2.2: Attrition

	Count	Percent
Total Policies	31,215	100.0
Survive Through Dec 1998	15,656	50.2
Randomly Censored	7,786	24.9
Voluntary Consumer Departures	7,773	24.9

Table 5.1: Categorical Classification Variables (Page 1/2)

Parameter	Variable	With Claims History	Without Claims History
α_1^1	Classification (ii)	-0.012 (0.003)	-0.007 (0.004)
α_1^2	Classification (iii)	-0.045 (0.022)	-0.034 (0.023)
α_1^3	Classification (i) 1997 or later	0.004 (0.014)	0.007 (0.012)
α_1^4	Classification (ii) 1997 or later	-0.015 (0.007)	0.004 (0.006)
α_1^5	Classification (iii) 1997 or later	-0.013 (0.007)	-0.011 (0.008)
α_2^1	4 th Period in Classification	0.004 (0.013)	0.007 (0.023)
α_2^2	5 th Period in Classification	-0.007 (0.016)	-0.009 (0.013)
α_2^3	6 th Period in Classification	0.005 (0.009)	0.005 (0.008)

Table 5.1: Categorical Classification Variables (Page 2/2)

Parameter	Variable	With Claims History	Without Claims History
α_3^1	Total Claims	-0.004 (0.004)	N/A
α_3^2	Claim Last Year	0.092 (0.009)	N/A
α_3^3	Claim Last Period	0.168 (0.033)	N/A
β_1^1	Classification (ii)	0.013 (0.065)	0.017 (0.060)
β_1^2	Classification (iii)	-0.107 (0.165)	-0.117 (0.151)
β_1^3	Classification (i) 1997 or later	-0.005 (0.043)	-0.033 (0.039)
β_1^4	Classification (ii) 1997 or later	0.054 (0.213)	0.052 (0.103)
β_1^5	Classification (iii) 1997 or later	-0.061 (0.331)	-0.068 (0.311)

Table 5.2: Continuous Price Variables

Parameter	Variables	Basic Liability Price – With Claims History	Basic Liability Price – No Claims History	Full Price – With Claims History	Full Price – No Claims History
α_1^1	Price Increase	-0.062 (0.022)	-0.023 (0.031)	-0.036 (0.015)	-0.014 (0.022)
α_1^2	(Price Increase) ²	0.001 (0.001)	0.002 (0.003)	0.001 (0.0004)	0.002 (0.001)
α_2^1	Increased cost of Next Claim	-0.005 (0.088)	0.005 (0.059)	-0.011 (0.068)	0.003 (0.049)
α_2^2	(Increased cost of Next Claim) ²	-0.008 (0.031)	-0.018 (0.061)	-0.018 (0.023)	-0.011 (0.56)
α_3^1	Total Claims	-0.004 (0.005)	N/A	0.014 (0.015)	N/A
α_3^2	Claim Last Year	0.101 (0.019)	N/A	0.088 (0.029)	N/A
α_3^3	Claim Last Period	0.158 (0.036)	N/A	0.168 (0.046)	N/A
β_1^1	Price Increase	-0.031 (0.043)	-0.035 (0.046)	-0.023 (0.049)	-0.045 (0.066)
β_1^2	(Price Increase) ²	0.004 (0.103)	0.010 (0.086)	0.005 (0.067)	-0.010 (0.096)

All prices and claim costs are measured in \$100