# Ex Post Moral Hazard in Automobile Insurance Markets with Experience Rating

Qing Gong*

October 15, 2017

**Abstract**

Accounting for unreported accidents due to ex post moral hazard is important for studying asymmetric information in insurance markets. In this paper, I study ex post moral hazard in an automobile insurance market with experience rating. I develop a dynamic model in which policyholders with private information about their risk types choose accident prevention efforts (ex ante moral hazard) and make claim filing decisions when accidents happen (ex post moral hazard). I then estimate the model using a detailed policy-level panel dataset from China. I find that policyholders do not report 24% of all accidents, which account for about 5% of total monetary losses. The degree of ex post moral hazard varies by experience rating: policyholders with the best rating hide 40% of all accidents. Finally, I use counterfactual experiments to evaluate the welfare implications. I find that experience rating improves policyholder welfare, mainly by inducing higher preventive efforts and reducing accidents. When ex post moral hazard is restricted and policyholders are forced to report all accidents, the benefit from increased accident prevention efforts barely outweighs policyholders' loss from increased premiums.

*Key words:* Insurance, moral hazard, discrete choice

*JEL codes:* C25, D82, G22

*Department of Economics, University of Pennsylvania, 3718 Locust Walk, 160 McNeil Building, Philadelphia, PA 19104. E-mail: qinggong@sas.upenn.edu. I owe special thanks to Hanming Fang and Petra Todd for their support, guidance and tremendously helpful advice. I have benefited greatly from discussions with Jaap Abbring, Christopher Flinn, Camilo García-Jimeno, Holger Sieg, Jan Tilly, and Weilong Zhang. I would also like to thank comments and suggestions from the UPenn empirical micro group, participants of the UPenn Empirical Micro Lunch Seminar, and participants at the North American Summer Meeting of the Econometric Society in Philadelphia, 2016. All remaining errors are mine.

# 1 Introduction

Moral hazard is a prominent feature of insurance markets, as well as the center of attention of both researchers and insurers. Many automobile insurance providers nowadays use experience rating to alleviate the moral hazard problem. Experience rating penalizes claims with higher future premiums and rewards safe driving with premium discounts. While it helps to curb *ex ante* moral hazard, where policyholders exert a suboptimal amount of preventive effort, it also introduces *ex post* moral hazard into the market, where policyholders choose not to file certain claims upon the occurrence of accidents. It is important to take the unclaimed losses due to ex post moral hazard into account when assessing the degree of asymmetric information and evaluating welfare, which has long been the interest of theoretical and empirical research. Quantifying ex post moral hazard is also crucial for policy evaluations, as the reported accidents are only a fraction of what actually happened.

In this paper, I focus on automobile insurance markets with experience rating, and want to answer the following questions. First, to what extent do ex ante and ex post moral hazard exist? Second, What are the separate welfare implications of the two types of moral hazard? Third, how does the prevailing experience rating rule compares with alternative policies in terms of welfare cost to both policyholders and insurers?

To this end, I develop a dynamic model of individual policyholder decisions with ex ante and ex post moral hazard. In the model, each policyholder chooses the level of preventive effort in each contract year. Upon the occurrence of an accident, she observes the loss realization and decides whether to file a claim. By filing a claim, she receives full reimbursement of the loss, but ends up in a lower experience rating class and faces higher future premiums. By not filing a claim and paying out of pocket for her loss, she keeps a claim-free record and receives discounts in her next-year premium. The model allows for observed heterogeneities in car characteristics and exogenous driving conditions, which affect the marginal cost of accident prevention. More importantly, the model also allows for unobserved heterogeneity in the risk type of policyholders, which is usually hard to control for in reduced-form models. The model implies that policyholders follow a cutoff rule when deciding whether to file a claim. The threshold loss sizes depend crucially on the current experience rating class, along with other observable characteristics of the car and the policyholder. These thresholds quantify the extent to which there is ex post moral hazard in this market.

I estimate the model using detailed policyholder-level panel data on the mandatory liability insurance (MLI) in China. The MLI market is highly regulated, with a single standard contract across insurers and throughout the country. Thus, there is minimal, if any, selection into contracts or insurers, making it an ideal environment to focus on the interaction between ex ante and ex post moral hazard. I find that car ages, driving intensity, and local weather conditions all have significant impacts on the choice of preventive efforts, and that the probability of being the low-risk type highly correlates with whether the car is privately owned. Moreover, the model estimates imply that policyholders do not file claims for 23.83% of all accidents that occurred, or 2.46-5.86% of total loss values. The fraction of accidents not claimed varies greatly across experience rating classes, with policyholders in higher classes hiding as much as 42.27% of all accidents. While reduced-form findings also show hints of heterogeneity and moral hazard, they cannot pin down the amount of unclaimed accidents or losses, which are important measures of ex post moral hazard.

This paper relates to the large body of literature on asymmetric information in insurance markets. Theoretical research dates back to Arrow (1963), Akerlof (1970), and Rothschild and Stiglitz (1976). While these studies make it clear that asymmetric information generally leads to market inefficiency, empirical findings on the presence and welfare impact of asymmetric information are relatively tenuous. Most existing studies focus on either adverse selection or ex ante moral hazard alone. In these studies, automobile insurance markets have been an important testing ground, marked by the availability and quality of micro-level data, and the relatively clear risk environment. In fact, it is the first testing ground for the existence of asymmetric information in the seminal paper, Chiappori and Salanie (2000).

Earlier papers rely on static models and try to find correlations between coverage generosity and claim intensities (Chiappori and Salanie, 2000; Chiappori et al., 2006). While such tests are easy to implement and shed some light on the existence of asymmetric information, they are not robust to the presence of more than one form of asymmetric information, nor multidimensional private information (Chiappori and Salanie, 2012). Later studies acknowledge the interaction between different types of asymmetric information, namely ex ante moral hazard and adverse selection. They typically build dynamic models and utilize panel data to distinguish between the two, borrowing heavily from econometric tools for disentangling state dependence and unobserved heterogeneity in labor economics (Abbring et al., 2003; Ceccarini, 2008).

Although a consensus regarding whether moral hazard and/or adverse selection exist is yet to be reached (Chiappori and Salanie, 2012), the literature almost uniformly treats observed claims as actual losses. This implicitly assumes away ex post moral hazard. In reality, however, the incentive to not file certain claims can be particularly strong under the experience rating, since the current loss in insurance reimbursement may be more than offset by discounts in future premiums. Ignoring the presence of ex post moral hazard may generate bias in tests for asymmetric information and lead to welfare losses (Abbring et al., 2008).

To the best of my knowledge, Abbring et al. (2008) is the first and only paper so far to model ex post moral hazard explicitly. Their model builds on that in Abbring et al. (2003), adding heterogeneous dynamic changes in insurance holders' incentive to conceal some losses but file claims for others. Then the authors implement reduced-form tests on panel micro-level data from a Dutch automobile insurance company. Results on the number of claims and duration between claims provide evidence for the existence of moral hazard. However, the non-stationary model in their paper cannot be estimated; and the only direct evidence of ex post moral hazard comes from a small fraction of filed claims that are later withdrawn by policyholders, which trivially show that there are unclaimed losses. One other paper that studies ex post moral hazard is Robinson and Zheng (2010), which uses aggregate data from Canada and a difference-in-differences approach to evaluating the effect of introducing experience rating to the market. The authors find that ex post moral hazard contributes 25-45% of the reported decline in car accidents in Canada from 1990 to 2005.

In this paper, I contribute to the existing literature in the following ways. First, I explicitly model ex post moral hazard in a dynamic environment, allowing for ex ante moral hazard and unobserved heterogeneity at the same time. In addition, I am able to estimate the model on micro-level data and recover both the distribution of underlying loss sizes and other parameters of primary interest. These provide quantitative measures of ex post moral hazard in terms of hidden accidents. Finally, building on the model estimates, I am able to conduct counterfactual analyses to evaluate the efficacy of current and alternative experience rating policies.

The rest of this paper is organized as follows. Section 2 develops the dynamic model of policyholder decisions. Section 3 describes the data. Section 4 presents the empirical findings, both preliminary reduced-form evidence of ex ante and ex post moral hazard, and estimation results of the structural model. Section 5 does counterfactual experiments to compare the welfare implications of the prevailing experience rating rule with that under alternative scenarios. Section 6 concludes with plans for future work.

# 2 Model

## 2.1 The primitives

In this section, I develop a dynamic programming model where heterogeneous drivers make accident prevention effort (ex ante moral hazard) and claim (ex post moral hazard) decisions.

Time is discrete, each period being a year. The model starts from the year in which the car is newly purchased, i.e. with age $a = 1$. Because the model characterizes the decision process of an individual policyholder $i$, I do not separate age and calendar year subscripts.[1] Instead, I denote time in terms of the age of car, $a = 1, 2, \ldots, A$, where $A = 30$.

All drivers are policyholders of the mandatory liability insurance (MLI). Each contract lasts one year, with premium $q_a$ and a zero deductible. Premium $q_a = Q(b_a)q_0(x)$ depends on the car's experience rating class at the beginning of the contract year, $b_a \in \{0, 1, 2, 3\}$,

[1] For the same reason, I will also suppress the $i$ subscript unless it is necessary to distinguish between individuals.

as well as a baseline premium, $q_0$ that is a function of time-invariant car characteristics $x$, namely car size, identity of owner (individual or organization), and whether the car is for commercial use. The experience rating schedule rewards safe driving and penalizes accidents via the premium discount function, $Q: \ b \rightarrow \{0.7, 0.8, 0.9, 1\}$. Policyholders in higher experience rating classes receive larger discounts, and those in class $b = 0$ pays the full baseline premium.

The policyholder always starts with $b = 0$ when she first purchases MLI. At the end of each year, the class is updated according to $b_{a+1} = B(b_a, n_a)$, where $n_a \in \{0, 1\}$ is an indicator for having filed a claim during the year. Claim filing leads to downgrading of $b$ (and an increase in premium for the next year). Figure 1 illustrates the MLI experience rating schedule.

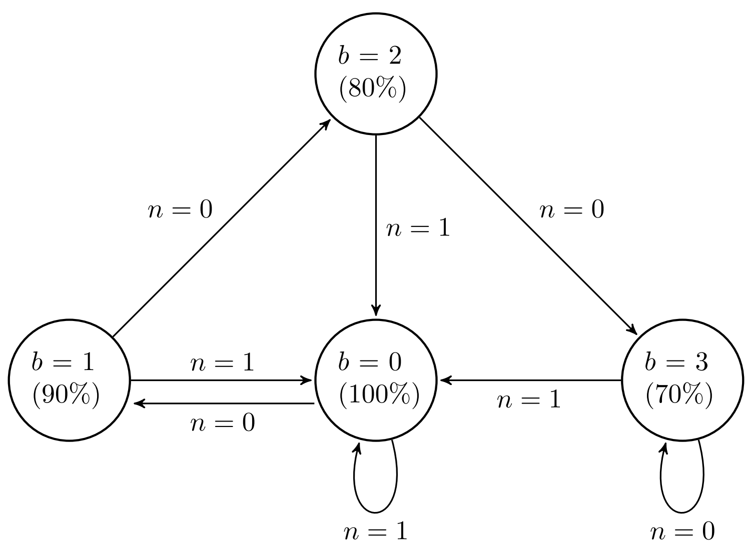


Figure 1: Experience rating schedule

A policyholder with $b = 0$ and paying 100% of the baseline premium will move up to $b = 1$ the next year, paying 90% of the baseline premium, if she doesn't file any claim in the current year ($n = 0$); but will have to stay in $b = 0$ if she has an accident *and* files a claim ($n = 1$). She will continue to move up the ladder one class at at time, until she reaches and stays in class $b = 3$, paying only 70% of the baseline premium, as long as she maintains a claimless record. On the other hand, whenever she files a claim, she will be downgraded all the way to the entry class $b = 0$.

In each period, the policyholder chooses the probability of having an accident, $p_a \equiv \Pr(acc = 1) \in (0, 1)$, by exerting effort

$$e_a = -\gamma(\theta)' X_a \log p_a$$

where $\theta \in \{\theta_N, \theta_G\}$ is $i$'s risk type that is her private information, N standing for the "normal" type and G for the "good" type. The probability that a policyholder is the normal type

$$\Pr(\theta = \theta_N) = \frac{\exp(\beta_1 1\{\text{private}\} + \beta_2 \text{DriverAge})}{1 + \exp(\beta_1 1\{\text{private}\} + \beta_2 \text{DriverAge})} \tag{1}$$

where $1\{\text{private}\}$ is a dummy variable for privately-owned cars, and DriverAge is the initial age of the driver when the car first enters my sample.[2] The intuition for including the private dummy is that drivers of private cars are more likely to belong to the normal type, while drivers of company or government agency cars are more likely to be the good type since some of them drive for a living.

The effort cost of accident prevention also depends on a vector of covariates, $X_a$, which includes car age, $a$, squared car age, $a^2$, an indicator that the car is for commercial use, $1\{\text{commercial}\}$, local annual precipitation $PREC$, and average temperatures in January, $TEMP1$. Note that all covariates evolve deterministically except for the last two weather variables. Long term meteorological studies have found that the precipitation and January temperature time series of the geographical region where my sample comes from are roughly white noises (Li et al., 2005; Sun et al., 2006). Therefore I normalized $PREC$ and $TEMP1$ for each city-level location by first demeaning them, and then dividing by their respective standard deviations. Consequently, when the policyholder calculates her continuation values, both the precipitation and the January temperature have expectations of zero, simplifying the dynamic programming problem.

[2] I do not include the driver's gender as there is limited variation in the data.

Conditional on having an accident, the policyholder draws a loss realization, $l_a \sim F$, and a cost of not filing a claim, $\xi_a \sim G,\ \xi_a \geq 0$. Should the policyholder decide to file a claim, MLI fully reimburses the loss $l_a$. If, however, she chooses not to file a claim but to settle with the other party, she pays for any loss $l_a$ out of pocket, plus any additional cost, $\xi_a$. Regardless of her claim filing decision, however, the policyholder incurs an extra cost, $\omega l_a$ that is proportional in size to the loss $l_a$. This captures non-reimbursable costs of having accidents. The non-reimbursable cost may come in the form of missed work, physical pain, or other non-pecuniary or pecuniary losses.

Policyholders share the same discount factor, $\rho = 0.97$, and derive flow utility from consumption, $u(c_a)$, where

$$c_a = \begin{cases} y_a - q_0 \cdot Q(b_a) & \text{if } acc_a = 0 \\ y_a - q_0 \cdot Q(b_a) - \omega l_a & \text{if } (acc_a = 1, n_a = 1) \\ y_a - q_0 \cdot Q(b_a) - (1+\omega) l_a - \xi_a & \text{if } (acc_a = 1, n_a = 0). \end{cases}$$

Because driver income is not observable in the data, I set $u$ to be a linear utility function, thus getting rid of wealth effects and at the same time accommodating potentially negative consumptions.

Figure 2 summarizes the time line of the model.

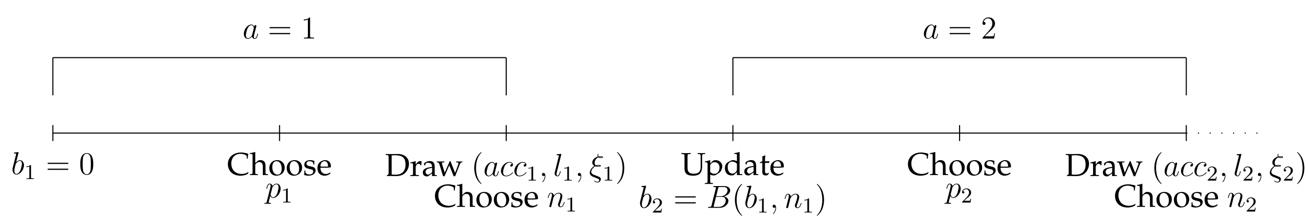


Figure 2: Timing of the model

## 2.2 The policyholder's problem

Denote the policyholder's state space as

$$\Omega_a = \{a, b_a, X_a, y_a\}$$

where all variables except $b_a$ evolves trivially following predetermined rules outside of the model. Thus I only keep $b_a$ in the value functions for notational simplicity.

The policyholder's value function in the final period, $a = A$, is

$$\begin{aligned} V_A(b_A) \quad = \quad & \max_{p_A}(1-p_A)\left[u(y_A - q_0 \cdot Q(b_A))\right] \\ & + p_A \iint \max\{u(y_A - q_0 \cdot Q(b_A) - (1+\omega)l_A - \xi_A), \\ & \qquad u(y_A - q_0 \cdot Q(b_A) - \omega l_A)\}\, \mathrm{d}F(l_A)\mathrm{d}G(\xi_A) \\ & - e(p_A, X_A, \theta) \end{aligned} \tag{2}$$

Her value function in period $a < A$ is

$$\begin{aligned} V_a(b_a) \quad = \quad & \max_{p_a}(1-p_a)\left[u(y_a - q_0 \cdot Q(b_a)) + \rho \mathbb{E} V_{a+1}(B(b_a, 0))\right] \\ & + p_a \iint \max\{u(y_a - q_0 \cdot Q(b_a) - (1+\omega)l_a - \xi_a) + \rho \mathbb{E} V_{a+1}(B(b_a, 0)), \\ & \qquad u(y_a - q_0 \cdot Q(b_a) - \omega l_a) + \rho \mathbb{E} V_{a+1}(B(b_a, 1))\}\, \mathrm{d}F(l_a)\mathrm{d}G(\xi_a) \\ & - e(p_a, X_a, \theta) \end{aligned} \tag{3}$$

where the expectation is taken over future realizations of accidents, losses, as well as weather variables in $X$. As discussed earlier in this section, both weather variables have expectations of zero.

It is easy to see that the policyholder follows a simple cutoff rule of claim filing upon occurrence of an accident, i.e. she files a claim if and only if the sum of $l_a$ and $\xi_a$ exceeds threshold $L_a^*(b_a)$. When $a = A$, $L_A^* = 0$ trivially, thus the optimal accident probability is given by

$$p_A^* = \gamma(\theta)' X_A \cdot \left( u(y_A - q_0 \cdot Q(b_A)) - \int u(y_A - q_0 \cdot Q(b_A) - \omega l) \mathrm{d}F(l) \right)^{-1} \tag{4}$$

More generally for $a < A$, the optimal claim decision threshold is

$$L_a^*(b_a) = \rho \Delta Emax_{a+1}(b_a) \tag{5}$$

where $\Delta Emax_{a+1}(b_a) = \mathbb{E}V_{a+1}(B(b_a, 0)) - \mathbb{E}V_{a+1}(B(b_a, 1))$. Thus, the threshold is the gross loss size that makes the policyholder indifferent between the maximized value after filing the claim and the one after concealing the claim.

Suppressing the state variable, $b_a$, the optimal accident probability is

$$p_a^* = \frac{\gamma(\theta)' X_a}{\rho \Delta Emax_{a+1} \left[1 - \int F(\rho \Delta Emax_{a+1} - \xi) \mathrm{d}G(\xi)\right] + \omega \mathbb{E}[l] + \mathbb{E}[\xi] + \iint_0^{L_a^* - \xi} l \mathrm{d}F(l) \mathrm{d}G(\xi)} \quad (6)$$

Note that the denominator is the expected marginal benefit of reducing accident probability. Denoting this component as $H_a$, then rearranging the above equation gives

$$H_a = \gamma(\theta)' X_a \cdot (p_a^*)^{-1} \quad (7)$$

where the right hand side is exactly the marginal cost of reducing $p$. Hence the optimal accident probability is simply the one that equates the marginal benefit and cost, both of which depends on the state variable, $b_a$, and other covariates.

## 2.3 Likelihood

I estimate the above model using maximum likelihood (ML). The model parameters are

$$\Theta = (F, G, \gamma(\theta_N), \gamma(\theta_G), \omega, \beta)$$

Denote the outcome of policyholder $i$ in period $a$ as $O_{ia}$, where

$$\begin{aligned} O_{ia} &= (n_{ia}, l_{ia}) \quad \text{if } n_{ia} = 1 \\ &= n_{ia} \quad \text{o.w.} \end{aligned}$$

Policyholder $i$'s contribution to the likelihood, given her type $\theta_i$, is then

$$\mathcal{L}_i(O_i, \Theta | \Omega_i, \theta_i) = \prod_{a=0}^{A} \Pr(O_{ia} | \Omega_{ia}, \theta_i) \quad (8)$$

where

$$
\Pr(n_{ia}=1, l_{ia}=l|\Omega_{ia},\theta_i) = p^*_{ia}(\Omega_{ia},\theta_i)\cdot[1-G(L^*_{ia}(\Omega_{ia},\theta_i)-l)]\cdot f(l) \tag{9}
$$

$$
\Pr(n_{ia}=0|\Omega_{ia},\theta_i) = 1+p^*_{ia}(\Omega_{ia},\theta_i)\cdot\Big[\int_0^{L^*_{ia}} F(L^*_{ia}(\Omega_{ia},\theta_i)-\xi)\mathrm{d}G(\xi)-1\Big] \tag{10}
$$

The total log likelihood is

$$
\ln\mathcal{L}(O,\Theta|\Omega)=\sum_{i=1}^{N}\ln\Big(\Pr(\theta_i=\theta_N)\mathcal{L}_i(O_i,\Theta|\Omega_i,\theta_N)+\Pr(\theta_i=\theta_G)\mathcal{L}_i(O_i,\Theta|\Omega_i,\theta_G)\Big) \tag{11}
$$

## 2.4 Identification

The identification of loss distribution, $F$, requires some parametric assumptions, as it essentially has a mixture model component. I choose the Gamma/Gompertz distribution with three parameters, $b_g, s_g, \beta_g > 0$ (Bemmaor and Glady, 2012). This distribution is defined on the positive half of the axis, and is extremely flexible. At the same time, its cdf and pdf have the following nice analytical forms, and its moment generating function well defined.

$$
f(l;b_g,s_g,\beta_g) = b_g s_g \exp(b_g l)\beta_g^{s_g}\Big(\beta_g-1+\exp(b_g l)\Big)^{-(s_g+1)} \tag{12}
$$

$$
F(l;b_g,s_g,\beta_g) = 1-\beta_g^{s_g}\Big(\beta_g-1+\exp(b_g l)\Big)^{-s_g} \tag{13}
$$

The parameters of $F$ can be identified from the observed claim sizes. Although the cutoff rule in the policyholders' claim filing decision leads to truncation on its left tail, the overall shape of $F$ is still preserved. Therefore one could recover the parameters of $F$ by only using the larger observed claims that are free of the selection problem.

As for the distribution of the "concealing cost," $G$, I assume that it is a uniform distribution over $[0, M]$. However, $M$ is not separately identified from other parameters that affect the threshold value for filing a claim. Observing a small claim could either be the result of having a high claiming threshold but at the same time a large $M$; or having a low claiming threshold and small $M$. Because the concealing cost $\xi$ is introduced mainly to avoid degeneracy in the likelihood and thus not of my primary interest, I set value of $M$

to be 1.[3] Having smaller $M$ values could increase the maximized likelihood, but is also more likely to get degeneracy problems.

I identify the parameters for the marginal effort cost of accident prevention, $\gamma(\theta)$, using a standard exclusion restriction argument. Note that the covariates $X$ only affect the probability of having an accident, but not the distribution of losses. The additional cost parameter, $\omega$ is also easily identified, as it only enters the $Emax$ function via the term $\omega\mathbb{E}[l]$, where the expected loss is already pinned down by the loss distribution parameters, $(b_g, s_g, \beta_g)$.

Finally, I take advantage of the panel data structure and identify the type-determining parameters, $(\beta_1, \beta_2)$ from persistence in accidents. To fix ideas, consider a simplified model where $\Pr(\theta = \theta_N)$ is a fixed number $\beta$, and the marginal cost of accident prevention for each type is governed by a single parameter, $\gamma_{0j},\ (j = N, G)$. Because there is no covariate $X$, normalized $\gamma_{0G} = 0$. Consider the extreme case with full persistence in accidents in the sense that cars either always have accidents or never have one. Then the $\beta = \Pr(acc = 1)$, where $\Pr(acc = 1)$ can be recovered from $F$ and model-implied claim filing thresholds; and $\gamma_{0N}$ can be backed out from Equation (6) as well. On the other extreme of the spectrum, suppose there is zero persistence. Then $\beta = 1$ and $\gamma_{0N}$ is again pinned down by Equation (6) given that all other parameters are known.

# 3 Data and Sample Construction

I construct a panel using proprietary data on MLI from a leading property insurance company in China. All automobiles must have MLI starting late 2006. Insured automobiles get a special sticker on the windshields for easy monitoring. The Ministry of Public Security also checks the status of MLI in the annual inspection. Uninsured automobiles, once caught, will be held in police custody until insured, and are subject to additional fines twice the amount of applicable MLI premiums. That said, it is reasonable to believe that all cars have MLI so there is no selection into the market.

Furthermore, the MLI policy is highly regulated, with all insurers offering exactly the

[3]The unit is thousand yuan (deflated to 2006 level).

same contract throughout the country. The policy is the same as the one introduced in Section 2, with baseline premiums only depending on car types, zero deductibles and an experience rating schedule. Notably, the insurers share the claim history information, so a policyholder does not lose her good records or escape from future premium penalties by switching to another insurer. These institutional settings further reduces selection into different insurers, making my sample from one insurer more representative.

The sample randomly selects 5,000 cars from a province in Northeast China and tracks them up to seven years from 2007 to 2013. A sizable fraction of cars do not stay with the insurer for all these years, and may have left because they switched to another insurer, were sold to another person, or left the market for good (e.g. demolished). Given that I do not observe the reason of attrition, and that most cars in my sample are fairly new (the average car age is less than 4 years old), I assume they stay in the market but just switch insurers. This way the policyholder keeps solving the same dynamic programming problem even after she leaves the sample.

For each car-contract year, I observe the set of variables that fully characterize the insurance policy, especially the baseline premium and the current experience rating class. I also observe the accident and loss realizations, i.e. the date and value of all claims filed during the contract year. Merely 0.13% have more than one claim, which is why I restrict the maximum of accidents in a given contract year to 1 in the structural model.

In addition, the data also include a rich set of car and owner characteristics, including type of owner (individual or organization), nature of use (commercial or non-commercial), car age, purchase value, car size, owner age, and city. I supplement this dataset with city-level precipitation and temperature information for the statistical yearbook of the province where my sample comes from.

I drop 17 cars whose experience rating class update violates the aforementioned rules for unknown reasons.[4] In the end, I get an unbalanced panel of 4983 cars and 9987 observations, where each observation is a unique car-contract year combination. Table A1 reports the summary statistics of key variables.

[4]For example, some of these cars filed a claim but were not downgraded to class 0 in the following year, while others had no claim but were not promoted.

# 4 Empirical Results

In this section, I present the empirical results, starting with reduced-form findings that provide preliminary evidence of the prevalence of ex ante and ex post moral hazard in the data. Then I proceed to report estimation results of my structural model, show goodness-of-fit, and discuss the implication of these results.

## 4.1 Reduced-form evidence

To begin with, Table 1 shows the transition probability of experience rating classes. Note that $\Pr(b_{a+1} = 0 \mid b_a)$ is decreasing in $b_a$, indicating that policyholders in better classes tend to file fewer claims.

Table 1: Transition probability across experience rating classes

| | $b_{a+1} = 0$ | 1 | 2 | 3 | $\Pr(b_a)$ |
|---|---|---|---|---|---|
| $b_a = 0$ | 0.128 | 0.872 | 0 | 0 | 0.489 |
| 1 | 0.093 | 0 | 0.907 | 0 | 0.263 |
| 2 | 0.085 | 0 | 0 | 0.915 | 0.119 |
| 3 | 0.072 | 0 | 0 | 0.928 | 0.130 |

This correlation, however, could be the result of heterogeneity in driving ability, as better drivers are more likely to end up in higher classes, and at the same time have fewer accidents, thus fewer claims. Alternatively, it could also be the result of moral hazard, as having an accident and filing a claim when $b_a$ is high means foregoing a larger discount in the next period, let alone ruining a spotless driving record, which gives the policyholder greater incentive to drive carefully or don't file claims even after an accident.

To see this, consider the results in Table 2. Column (1) regresses the binary outcome of filing a claim on the experience rating class dummies and the set of covariates included in the structural model. Comparing with the baseline group, $b = 0$, those in higher classes make significantly fewer claims. The logit estimates in Column (2) are very similar both qualitatively and quantitatively. These results, as those in Table 1, are still a mixture of moral hazard and heterogeneity.

Table 2: Determinants of claim intensity ($Y = 1\{n_a = 1\}$)

| | (1) OLS | (2) Logit | (3) FE | (4) AB |
|---|---|---|---|---|
| 1{b=1} | -0.1189*** | -1.4722*** | -0.3509*** | |
| | [0.0075] | [0.1123] | [0.0123] | |
| 1{b=2} | -0.1796*** | -4.1755*** | -0.5133*** | |
| | [0.0072] | [0.4551] | [0.0153] | |
| 1{b=3} | -0.1996*** | -3.9138*** | -0.6450*** | |
| | [0.0096] | [0.3565] | [0.0181] | |
| $1\{n_{a-1} = 1\}$ | | | | 0.0797** |
| | | | | [0.0348] |
| car age | 0.0240*** | 0.2519*** | 0.1567*** | -0.0201* |
| | [0.0031] | [0.0368] | [0.0089] | [0.0119] |
| car age$^2$ | -0.0017*** | -0.0177*** | -0.0062*** | 0.0008 |
| | [0.0003] | [0.0040] | [0.0008] | [0.0012] |
| 1{private car} | -0.0543*** | -0.7877*** | | |
| | [0.0067] | [0.0907] | | |
| Adj./pseudo $R^2$ | 0.0859 | 0.1624 | 0.3073 | |
| $N$ | 9987 | 9987 | 9987 | 2156 |

Notes: Covariates included but not reported include the value of car, standardized local annual precipitation (demeaned and then divided by the standard deviation), local average temperature in January, driver's age when first sampled, the constant, and the full set of 41 car type dummies and 9 city dummies. "AB" in Column (4) refers to the Arellano-Bond dynamic panel data model. The $R^2$ for Column (2) is the pseudo $R^2$, the others are adjusted $R^2$'s. Robust standard errors are in brackets. $*p < 0.10, **p < 0.05, ***p < 0.01$

Column (3) estimates a fixed effect model to control for the time-invariant heterogeneities. The negative relationship between claim probabilities and experience rating classes persists, and becomes even stronger. This suggests that moral hazard is driving the differences in claim probabilities, given that the individual fixed effects already take care of the unobserved heterogeneities. Column (4) takes a slightly different perspective by looking at the effect of having claims in the previous period on the current period claim probabilities. The moral hazard story would predict that having a previous claim reduces the incentive to avoid future claims, because the previous claim brings the policyholder down to class $b = 0$, thus lowering the stake of filing future claims. Indeed, Arellano-Bond dynamic panel estimates show that having a claim in period $a - 1$ significantly increases the probability of having a claim in the current period by almost 8%. This is highly consistent with the moral hazard story.

So far the results suggest that moral hazard exists in the MLI market under the experience rating rule, but could not distinguish between ex ante and ex post moral hazard effects. To disentangle the two, I exploit information on the size of filed claims.

Figure 3 shows the empirical cdf of filed claim sizes by experience rating class, $b_a$. The five marks represent the minimum, the 25th percentile, the median, the 75th percentile, and the largest value within 1.5 inter-quartile range of the 75th percentile, respectively.[5]

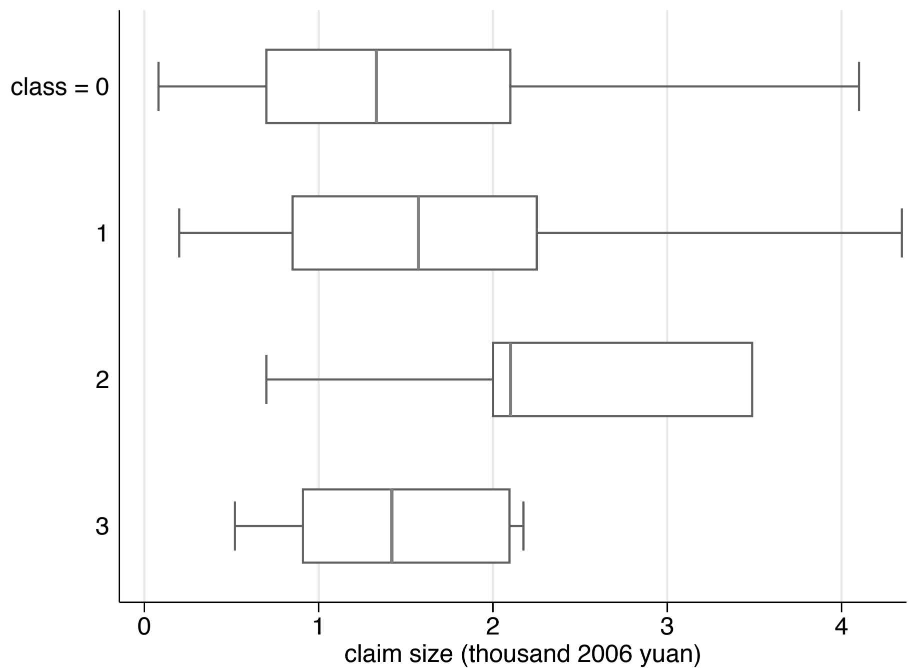


Figure 3: Distribution of claim sizes by experience rating class

Note that the minimum claim size increases with $b_a$, while the rest of the distribution is largely the same. This is consistent with the prediction that policyholders in higher classes have larger threshold values for filing a claim, because they have stronger incen-

[5] Very large extreme values are not shown in the plot.

tives to avoid claims. Hence, the different truncation points in the distribution of observed claims is strongly suggestive of the existence ex post moral hazard. It also shows that it is innocuous to assume that the losses of different policyholders follow roughly the same distribution.

## 4.2 Structural parameter estimates

I estimate the structural model in Section 2 using maximum likelihood. Before presenting the estimation results, I use Monte Carlo simulation to show the efficacy of my estimation strategy and as supplemental evidence for identification. I choose a set of parameters and simulate an unbalanced panel dataset with 3000 policyholders and a maximum length of 7 years.[6] Table A2 compares the "true" parameters and ML estimates. The differences are usually very small, in both absolute and relative terms. The biases in the majority of parameters are less than 5% of the true value. Considering the smaller sample size in the Monte Carlo simulation than in the real data, these results are reasonably reassuring of the estimation results to come.

Table 3 presents the formal ML estimates using the real data. The Gamma/Gompertz distribution parameters are estimated using the 856 observed claim values. The shape parameter $s_g$ is relatively large, indicating that the distribution is highly positively skewed. On the other hand, $\beta_g$ is relatively small, which translates into tails that are not particularly fat.

The point estimate of the square-root of $\omega$ is about 0.5545. This means that the additional, non-reimbursable cost of having an accident is about 30.75% of the loss $l$.

Parameters estimates governing marginal cost of accident prevention for the good type are fairly intuitive. Note first from Equation (6) that higher (lower) marginal cost translates into higher (lower) accident probabilities. Older, hence more experienced cars have lower marginal cost and consequently lower probability of accidents, which suggests the existence of learning. This is qualitatively consistent with the reduced-form results in Table 2 under the Arellano-Bond dynamic panel model. At the same time, the positive estimate

[6] I also introduce considerable attrition in the simulated sample to mimic the short panel structure in the real data.

Table 3: Model estimation results

| | (1) Est. coeff. | (2) s.e. | (3) implied |
|---|---|---|---|
| *Gamma/Gompertz parameters* | | | |
| $b_g$ (scale parameter) | 0.0030 | 0.0419 | |
| $s_g$ (shape parameter 1, "skewness") | 1.2226 | 0.0368 | |
| $\beta_g$ (shape parameter 2, "tail") | 0.0020 | 0.0315 | |
| *Additional cost of accident* ($\sqrt{\frac{\textit{additional cost}}{\textit{loss}}}$) | | | |
| $\omega$ | 0.5545 | 0.0010 | 0.3075 |
| *Marginal cost of accident prevention (good type)* | | | |
| constant | -0.0583 | 0.0027 | |
| car age | -0.0371 | 0.0000 | |
| car age$^2$ | 0.0009 | 0.0008 | |
| 1{commercial} | -0.0174 | 0.0041 | |
| January temperature | -0.0604 | 0.0172 | |
| total precipitation | 0.0369 | 0.0089 | |
| *Marginal cost of accident prevention (normal type,* $\sqrt{\gamma_N - \gamma_G}$*)* | | | |
| constant | 0.4451 | 0.0028 | 0.1398 |
| car age | 0.1673 | 0.0001 | -0.0091 |
| car age$^2$ | 0.0000 | 0.0000 | 0.0009 |
| 1{commercial} | 0.2917 | 0.0025 | 0.0677 |
| January temperature | 0.0981 | 0.0016 | -0.0508 |
| total precipitation | 0.0962 | 0.0020 | 0.0462 |
| *Probability of being normal type* | | | |
| 1{private} | 0.9998 | 0.1196 | |
| initial driver age | 0.0785 | 0.0120 | |
| $N$ = 9987, LR index = 0.7757 | | | |

Notes: $b > 0$ is the scale parameter of Gamma/Gompertz distribution, and $s, \beta > 0$ are the shape parameters. The Gamma/Gompertz parameters are estimated on a subsample of 856 observations.1{commercial} is a dummy variable indicating whether the car is for commercial use, for example taxis and rental cars. January temperature and total precipitation are all measured at the city level, where the temperature variable is standardized monthly average temperatures (demeaned and divided by the standard deviation), and the precipitation variable is standardized annual total precipitation including both rainfall and snowfall. 1{private} is a dummy variable for the car owner being an individual rather than an organization. Initial driver age is the age of the driver when his/her car first entered the sample. Standard errors are computed using numerical Hessian matrices from the maximum likelihood estimation.

in squared car age indicates the downward trend reverses at around $a = 20$, presumably driven by the deteriorating condition of the car, although the estimate is not statistically distinguishable from zero. Cars for commercial use have lower marginal costs, which is seemingly counterintuitive. However, this is unique to the good-type cars only, and is probably because good drivers of commercial cars (e.g. taxis) are especially careful or capable of avoiding accidents. The two weather variables have the expected effects: warmer Januaries usually means less ice on the road and faster melting of snows, thus reducing the probability of accidents; and more precipitation, whether in the form of rain or snow, results in more dangerous road conditions and increases accident occurrence.

The normal type, by definition, have weakly higher marginal cost of accident prevention. The estimates in Column (1) are the square-root of the difference between normal-type and good-type marginal cost parameters, and Column (3) reports the implied normal-type parameter values, $\gamma_N$. To begin with, the constant term is considerably larger than that of the good type. All other covariates have strictly larger partial effects, except for the squared car age: older age cars are still less likely to have accidents, but the effect is much smaller, only one fourth of that of good type; being a normal type car for commercial use increases the accident probability so dramatically that the sign of $\gamma_{1\{\text{commercial}\}}$ switches from negative for the good type to positive for the normal type. This is consistent with the intuition that for "normal" commercial cars, the higher driving intensity makes them more likely to have accidents. Finally, the two weather variables have statistically different effects on the normal type, although the implied marginal cost parameters remain similar to those of the good type, both qualitatively and quantitatively.

Finally, estimates of the type-determining parameters show that drivers of private cars are significantly more likely to be the normal type. This is very intuitive as those driving for companies or government agencies should be more professional. At the same time, older age of the driver is also associated with higher probabilities of being the normal type. Figure 4 illustrates these results by plotting the fraction of normal type drivers against car ownership and the age of driver. While the majority is always the normal type, those driving cars owned by organizations are always more likely, *ceteris paribus*, to be the good type.

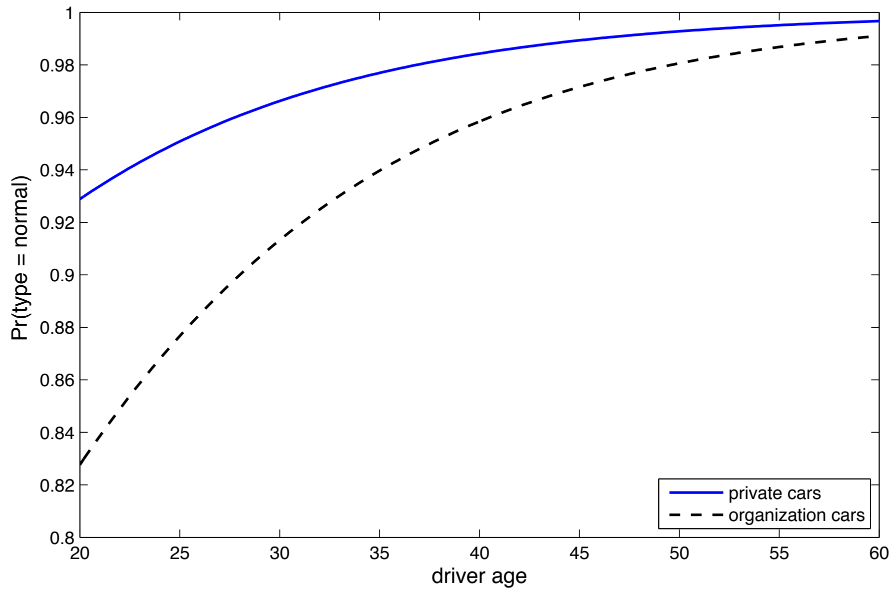


Figure 4: Model implied type distribution by ownership and car age

## 4.3 Goodness-of-fit

As a quick indicator for goodness-of-fit, the likelihood ratio index of the above ML results is $0.7757$, which is fairly close to the right end of the $[0, 1]$ spectrum. For more goodness-of-fit analyses, I use the model estimates to simulate 500 samples with the same size as the data, and compare the key outcomes.

The left panel of Figure 5 plots the model-implied parametric cdf of underlying *losses* and the empirical cdf of the observed claims in the real data. These two distributions have approximately the same shape, but the deviation is still fairly apparent toward the left end. The right panel replaces the model-implied loss distribution with the empirical cdf of *claims* in the simulated data, taking into account the optimal claim filing decisions of policyholders. As a result, this distribution is much closer to the real data. This first shows the presence of ex post moral hazard in the form of hidden claims, and also reassures the choice of Gamma/Gompertz as the underlying loss distribution is reasonably appropriate. Figure A1 in the appendix also compares the entire claim distributions from the model versus those in the data by experience rating class, $b$, and find them very close.[7]

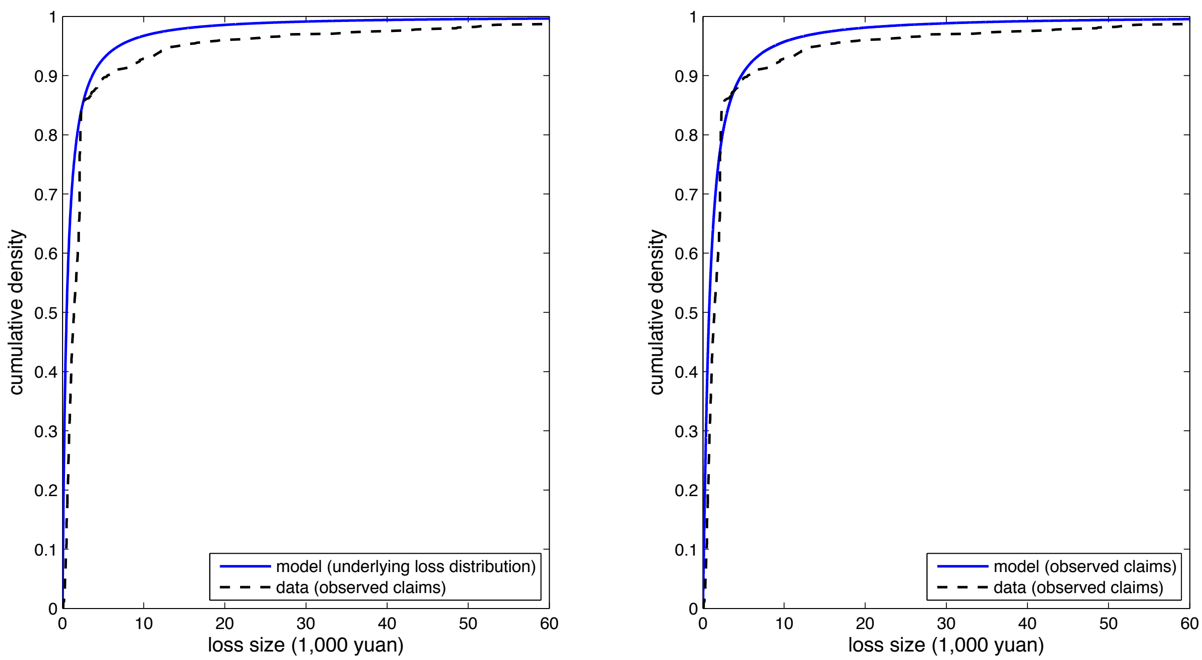


Figure 5: Goodness-of-fit: distribution of actual losses and observed claims

Table 4 reports in the left panel transitions across experience rating classes in the simulated data. The model captures the prominent pattern of transitions in the data, namely those in higher classes have better chances of upgrading to an even higher class. The model is somewhat off in the levels, though, by usually slightly underestimating the fraction that files a claim. However, the model does a good job in predicting the distribution of classes, with simulated $\Pr(b)$ extremely close to those in the data.

[7]Except perhaps for $b = 3$, where few claims are filed in the data.

Table 4: Goodness-of-fit: transition matrix of experience rating class

| | Model | | | | | Data | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $b_{a+1}=0$ | 1 | 2 | 3 | $\Pr(b_a)$ | 0 | 1 | 2 | 3 | $\Pr(b_a)$ |
| $b_a=0$ | 0.080 | 0.920 | 0 | 0 | 47.9% | 0.128 | 0.872 | 0 | 0 | 48.9% |
| 1 | 0.069 | 0 | 0.931 | 0 | 26.3% | 0.093 | 0 | 0.907 | 0 | 26.3% |
| 2 | 0.054 | 0 | 0 | 0.946 | 12.0% | 0.085 | 0 | 0 | 0.915 | 11.9% |
| 3 | 0.094 | 0 | 0 | 0.906 | 13.8% | 0.072 | 0 | 0 | 0.928 | 13.0% |

Notes: The two $\Pr(b_a)$ columns are marginal probabilities of being in class $b_a$. All other numbers are conditional probabilities $\Pr(b_{a+1} \mid b_a)$

Table 5 compares the claim intensity for all cars and by subgroups. Overall, the model-implied claim intensity is slightly higher than that in the data, but the deviation is not significant relative to the standard error of the data moment.[8] As for the comparison between non-commercial and commercial cars, the model correctly captures the significantly higher claim intensity of the latter. This is also consistent with the large estimated effect of the commercial use dummy on the marginal cost of accident prevention for most cars. Finally, claim intensities in the simulated data increases with the age of the car, which is qualitatively similar to the trend in the data, although the slope is slightly higher. Note that the contrast between the upward trend here and the negative effect of car ages on accident probabilities in Table 3. The difference is again driven by the selective filing of claims. On the one hand, older cars are less likely to have accidents; but on the other hand, being closer to the final period, they also have lower thresholds for filing claims, conditional on having an accident. The upward trend in claim intensity shows that the latter of the two opposing effects overpowers the former.

[8]This results does not contradicts those in the previous table, which only uses cars that stay in sample for at least two years, whereas results here uses all observations.

Table 5: Goodness-of-fit: claim intensity

| | Model | Data |
|---|---|---|
| | *Overall* | |
| All | 0.0903 | 0.0857 |
| | (0.0001) | (0.0028) |
| | *By car use* | |
| non-commercial | 0.0840 | 0.0832 |
| | (0.0001) | (0.0031) |
| commercial | 0.1151 | 0.0954 |
| | (0.0003) | (0.0065) |
| | *By car age* | |
| $[1, 3]$ | 0.0876 | 0.0854 |
| | (0.0002) | (0.0038) |
| $[4, 6]$ | 0.0905 | 0.0846 |
| | (0.0002) | (0.0050) |
| $[7, 9]$ | 0.0980 | 0.0885 |
| | (0.0004) | (0.0088) |
| $10+$ | 0.1043 | 0.0909 |
| | (0.0007) | (0.0139) |

Note: Standard errors are in parentheses.

The above results show that the model is able to capture the key features of the data, especially those for the claim probabilities and claim size distributions, which are outcomes of my primary interest.

## 4.4 Discussion

Given that the model fits the data relatively well, I explore the welfare implications of the model estimates. Table 6 compares the model-implied probability of claims with that of underlying accidents as a measure of ex post moral hazard. Column (1) is the same set of predicted claim intensities as in Table 5. Column (2), as a contrast, is the probability of accidents, whether claimed or not. 11.86% of the whole sample have accidents, while the claim intensity is only 9.03%. 2.83% of the sample have had an accident but chose not to file a claim. This means that 23.83% of all accidents that happened are not reported, which is close to the lower bound of estimates by Robinson and Zheng (2010). Granted, the hidden accidents are generally smaller than the ones claimed, so Column (5) shows the fraction of total hidden losses. Policyholders choose to pay out of pocket 2.46% of total reimbursable losses incurred in accidents. This figure is deceivingly small because the distribution is highly skewed, with very large losses happening with non-zero probabilities. If one excludes the top 3% extremely values in the distribution, then the hidden losses account for 5.86% of total loss values.

Table 6: Model implied hidden accidents

| | (1) Pr(claim) | (2) Pr(accident) | (3) Difference ( (2)-(1) ) | (4) % hidden ( (3)/(2) ) | (5) % loss hidden | (6) % loss hidden (trimmed) |
|---|---|---|---|---|---|---|
| | | | *Overall* | | | |
| All | 0.0903 | 0.1186 | 0.0283 | 23.83% | 2.46% | 5.86% |
| | | | *By experience rating class* | | | |
| b=0 | 0.1072 | 0.1173 | 0.0100 | 8.56% | 0.56% | 1.31% |
| b=1 | 0.0701 | 0.1050 | 0.0349 | 33.27% | 3.63% | 8.71% |
| b=2 | 0.0580 | 0.1005 | 0.0425 | 42.27% | 5.18% | 12.75% |
| b=3 | 0.0985 | 0.1649 | 0.0664 | 40.28% | 4.21% | 10.08% |

Note: Standard errors are in parentheses. "% loss hidden" is the total loss of hidden accidents divided by the total loss of all accidents. "% loss hidden (trimmed)" uses the trimmed sample after dropping the largest 3% losses.

More interestingly, the degree of moral hazard varies across subsamples with different incentives to avoid claims. Among the four experience rating classes, the fraction of accidents hidden rises from 8.56% for $b = 0$, to as high as over 40% for $b = 2$ and $b = 3$. The value of hidden losses also increases from around 1% to 5-10%. This is very intuitive given the experience rating rules: filing a claim when in higher classes not only leads to larger foregone premium discounts in the next period, but also shifts the entire trajectory of future premiums upward by ruining a claim-free history. Hence cars in higher classes have stronger incentives to avoid claims than their lower-classes counterparts. They achieve this both by exerting more preventive effort, and by hiding a larger fraction of claims conditional on occurrence of an accident.

Table 6 once again shows the importance of taking ex post moral hazard into account when understanding the behavior of policyholders. Overlooking the hidden accidents will introduce sizable bias, especially in policy evaluations. It also demonstrates the necessity of a structural model in the disentangling of ex ante moral hazard, ex post moral hazard, as well as policyholder heterogeneity.

## 5 Counterfactual Analyses

Now I take the model estimates and conduct counterfactual analyses to compare the efficiency, in terms of welfare cost, of the prevailing experience rating rule with that of alternative ones.

Under the current rules, policyholders have two instruments for self-insurance, ex ante preventive effort and ex post hiding of accidents. Taking the environment as given, the freedom of ex post moral hazard is unambiguously welfare-improving for the policyholders, who voluntarily choose to pay out of pocket for some accidents. But it is not as clear how the current rule adopted in China compares with alternatives, some actually used by insurers in other markets, others purely hypothetical.

I am especially interested in the following alternative scenarios:

No experience rating, with premium fixed at the current baseline level and zero deductible. This is the world with no incentive for ex post moral hazard, and at the same time no longer curbs ex ante moral hazard. It is also the worst scenario from the society's perspective.

No experience rating, with premium fixed at the current baseline level but non-zero deductibles. Recall that an important feature of the baseline model is that policyholders follow cutoff rules when deciding when to file a claim. The resulting outcome that smaller losses are not claimed resembles what would happen when there is a non-zero deductible. Hence it is interesting to compare the two scenarios. The result clearly depends on the choice of deductibles, which could be the same for all cars or depend on observable car and owner characteristics. As a first step, I set the deductible to be the same for all policyholders, and fix it at $D$ such that the cdf evaluated at $D$ is the same as the fraction of hidden accidents in the benchmark case.

With experience rating and mandatory filing of claims for all accidents. This is essentially shutting down ex post moral hazard, and necessarily reduces the welfare of policyholders, although it still encourages accident prevention. Such an experiment helps to decompose the effect of experience rating on ex ante and ex post moral hazard, although it may be hard to implement in the market.

I take the model parameters and simulate 500 samples of the same size as the real data for each of the three scenarios above. Table 7 compares the key outcomes from the counterfactual experiments with those under the benchmark case. Column (1) shows the accident probability, claim probability, fraction of accidents hidden, and fraction of loss values hidden under the benchmark, which are the same set of results discussed in the previous section. Column (2) shows these outcomes when there is no experience rating, with a fixed premium and zero deductible. First, note that now there is no incentive whatsoever not to file a claim conditional on having an accident. Hence the claim probability is the same as the accident probability. Second, the absence of experience rating sets ex ante moral hazard free and ends up with an overall accident probability as high as 0.1976, or almost twice the benchmark value.

Column (3) introduces a non-zero deductible on top of the scenario in Column (2). Because there is still no harness on ex ante moral hazard, preventive effort is still fairly limited. Thus the accident probability is only slightly lower than that in Column (2), but still very high. Nonetheless, a considerable number of accidents are not claimed, reducing the claim probability to a modest 0.1467. Note that although the hidden accidents are similar in number with the benchmark, the total value of hidden losses are significantly smaller. This is because the claim-filing threshold loss value in the benchmark is a function of both the experience rating class and the set of covariates. So some policyholders, especially those in higher classes, may prefer to hide larger losses. On the contrary, the deductible here is the same for everyone, which, under a highly skewed loss distribution, mechanically leads to much smaller values of total loss hidden.

Column (4) shows the scenario under experience rating, but with ex post moral hazard shut down. An immediate implication is again the equality between the accident probability and the claim probability. Despite having one less tool to self-insure against premium increases, the policyholders are subject to the same premium updating rules, thus still having incentives to reduce ex ante moral hazard by exerting more preventive effort. As a result, the accident probability is only 0.1350, which is slightly higher than but more comparable to that under the benchmark case.

Table 7: Counterfactual experiments

| | (1)<br>Benchmark | (2)<br>No EXR, no D | (3)<br>No EXR, with D | (4)<br>EXR, no EPMH |
|---|---|---|---|---|
| Pr(accident) | 0.1186 | 0.1976 | 0.1923 | 0.1350 |
| Pr(claim) | 0.0903 | 0.1976 | 0.1467 | 0.1350 |
| % accident hidden | 23.83% | 0% | 23.74% | 0% |
| % loss hidden | 2.46% | 0% | 0.84% | 0% |
| *Comparing with the benchmark case ...* | | | | |
| Δ(total pecuniary loss) | - | +66.94% | +62.33% | +13.87% |
| Δ(loss paid OOP) | - | - | -44.67% | - |
| Δ(premium) | - | +12.20% | +12.20% | +0.62% |
| Δ(policyholder loss) | - | -15.61% | -15.70% | -0.84% |
| Δ(insurer gross profit) | - | -10.21% | -7.89% | -5.50% |

Note: All comparison-with-the-benchmark results show the percent increase/decrease on top of the benchmark case value. Total pecuniary loss is the sum of all reimbursable losses incurred in accidents, whether claimed or not. Loss paid OOP is the total reimbursable losses that the policyholder could have claimed but paid out of pocket instead. Premium is the total premium paid by policyholders. Policyholder loss is the sum of premium payments, non-pecuniary losses, and loss paid OOP. Insurer gross profit is the premium revenues net of reimbursements.

Finally, I compare the welfare effects of these alternative scenarios with the benchmark. In terms of total pecuniary losses, Columns (2) and (3) are both more than 60 percent higher than the benchmark, which is not surprising given their high accident probabilities; and Column (4) is moderately higher than the benchmark by 13.87 percent. As for out-of-pocket payments by policyholders, only Column (3) has a non-zero number of hidden accidents. But as discussed above, the constant deductible results in smaller values of hidden losses, saving the policyholders 44.67 percent of their our-of-pocket payments under the benchmark. The scenarios also differ in the amount of premiums. In the absence of experience rating in Columns (2) and (3), policyholders no longer receive premium discounts, paying 12.20 percent more than they do under the bench mark. In Column (4), the policyholders still have enough incentives to prevent accidents and stay in good experience rating classes. So they only pay slightly more than the benchmark value, despite no longer having the ex post moral hazard tool to avoid downgrading and increase in premiums.

The last two rows of Table 7 compares the composite welfare change for the policyholder and the insurer, respectively. The policyholders' loss is defined as the sum of premium payments, non-pecuniary losses, and pecuniary losses that they choose to pay out of pocket. Comparing with the benchmark, policyholders are significantly worse-off when there is no experience rating, and considerably worse-off when there is experience rating with mandatory claiming of all accidents. The insurer, on the other hand, are significantly worse-off in all three scenarios. This is because even though they receive more premium revenue, the higher claim probabilities and larger claim sizes also require more reimbursement payments.

Granted, the above are fairly straightforward counterfactual experiments that help to quantify the welfare implications of ex ante and ex post moral hazard. There are certainly other interesting scenarios to explore. For example, it may be interesting to study the welfare effects of having experience rating, but letting future premium depend on the size of claims. The current experience rating rule completely ignores the claim sizes. It would be interesting to see what would be the effect of claim-size-dependent experience rating, which is already adopted in some markets. The extreme case is one where future premiums is a continuous function of past claim sizes. This is essentially introducing coinsurance into the market.

Another interesting scenario to study is one with experience rating that is the same as the current rule, except for accident forgiveness. Some automobile insurance companies in the U.S. highlight their accident forgiveness policies.[9] However, some argue that the terms are not as appealing as they seem.(Smith, 2011) One might expect that the benefit and cost vary among policyholders with different characteristics. So it is helpful to see whether and when accident forgiveness is welfare-improving. I will leave these counterfactual experiments for future work.

# 6 Concluding Remarks

In this paper, I develop a dynamic model of individual policyholders' accident prevention and claim filing decisions, allowing for ex ante and ex post moral hazard, as well

[9]For example, GEICO® and Progressive® both offer policies that forgive the first accident.

as unobserved heterogeneity in risk types. I estimate the model using panel data on a random sample of MLI policyholders in China, and find it fits the data fairly well. The model estimates imply that policyholders do not file claims for 23.83% of all accidents that occurred, or 2.46-5.86% of total loss values. The fraction of accidents not claimed varies greatly across experience rating classes, with policyholders in higher classes hiding as much as 40% of all accidents. These results are consistent with hints of ex ante and ex post moral hazard in reduced-form findings.

The model and empirical results have several limitations. The flow utility is linear in consumption, which is partially because policyholder income is not observed in the data and partially for computational simplicity.[10] Including risk aversion in the model could lead to a better fit of the data and potentially more interesting implications. The model also restricts the unobserved heterogeneity to that in risk types, while previous studies (e.g. Finkelstein and McGarry (2006)) have found evidence for multidimensional private information in long-term care insurance markets, with one dimension being the risk type, and the other being preference for insurance. Finally, the sample size results in only a handful of observed claims from the highest experience rating classes, making some estimates relatively noisy.

In addition to amending these limitations, the next step of the project is to better exploit the non-stationarity in the problem, and to develop a general equilibrium model. By abstracting away the timing of claims *within* a contract year, the current model becomes stationary and is greatly simplified. But at the same time, it loses track of the rich dynamic decision process within the contract year, which, as Abbring et al. (2008) pointed out, could provide valuable information on ex ante and ex post moral hazard. Given the Abbring et al. (2008) model cannot be estimated, I plan to develop a non-stationary model and follow the empirical strategy of Gilleskie (1998), which studies the health care consumption and absenteeism decision of workers and allows for the optimal timing of doctor visits.

Finally, although the supply side of Chinese MLI market is highly regulated, it would

[10] CRRA utility functions are free of wealth effects, but cannot accommodate potentially negative consumption values.

still be interesting to develop a general equilibrium model with the insurers also making profit-maximizing decisions when choosing what policies to offer. This is important for studying other markets, especially for welfare analyses.

# References

Abbring, J., P.-A. Chiappori, and T. Zavadil (2008). Better Safe than Sorry? Ex Ante and Ex Post Moral Hazard in Dynamic Insurance Data. *Discussion Paper 2008-77, CentER, Tilburg*.

Abbring, J. H., P.-A. Chiappori, and J. Pinquet (2003). Moral Hazard and Dynamic Insurance Data. *Journal of the European Economic Association 1*(4), 767–820.

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 488–500.

Arrow, K. J. (1963). Uncertainty and the Welfare Economics of Medical Care. *The American Economic Review*, 941–973.

Bemmaor, A. C. and N. Glady (2012). Modeling Purchasing Behavior with Sudden "Death": A Flexible Customer Lifetime Model. *Management Science 58*(5), 1012–1021.

Ceccarini, O. (2008). *Does Experience Rating Matter in Reducing Accident Probabilities? A Test for Moral Hazard*. Ph. D. thesis, University of Pennsylvania.

Chiappori, P.-A., B. Jullien, B. Salanié, and F. Salanie (2006). Asymmetric Information in Insurance: General Testable Implications. *The RAND Journal of Economics 37*(4), 783–798.

Chiappori, P.-A. and B. Salanie (2000). Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy 108*(1), 56–78.

Chiappori, P.-A. and B. Salanie (2012). Asymmetric Information in Insurance Markets: Empirical Assessments. *Handbook of Insurance, Springer Verlag (ed. Georges Dionne)*.

Finkelstein, A. and K. McGarry (2006). Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market. *American Economic Review 96*(4), 938–958.

Gilleskie, D. B. (1998). A Dynamic Stochastic Model of Medical Care Use and Work Absence. *Econometrica*, 1–45.

Li, X., W. Li, and Z. Zhao (2005). Long-Term Variation Pattern and Future Trend of Precipitation in the Valleys of Songhuajiang and Liaohe Rivers. *Journal of Applied Meteorological Science 16*(5), 593–599.

Robinson, C. and B. Zheng (2010). Moral Hazard, Insurance Claims, and Repeated Insurance Contracts. *Canadian Journal of Economics 43*(3), 967–993.

Rothschild, M. and J. Stiglitz (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, 629–649.

Smith, M. (2011, April). Accident Forgiveness Car Insurance: What's the Catch? Web.

Sun, F., J. Yuan, and S. Lu (2006). The Change and Test of Climate in Northeast China over the Last 100 Years. *Climatic and Environmental Research 11*(1), 101–108.

# Appendix: Supplemental Tables and Figures

Table A1: Sample Descriptive Statistics

| $N$=9987 | Mean | s.e. | Min | Max |
|---|---|---|---|---|
| 1{have filed claim} | 0.086 | 0.280 | 0 | 1 |
| Total amount of claim (thousand 2006 yuan) | 9.105 | 100.972 | 0.1 | 2150 |
| Age or car | 3.837 | 2.731 | 1 | 18 |
| 1{private} | 0.536 | 0.499 | 0 | 1 |
| 1{commercial} | 0.203 | 0.402 | 0 | 1 |
| 1{freight car} | 0.258 | 0.437 | 0 | 1 |
| Baseline premium (thousand 2006 yuan) | 1.006 | 0.693 | 0.1 | 4.7 |
| Purchase value of car (thousand 2006 yuan) | 89.051 | 115.958 | 2 | 2000 |
| Avg. temperature in January (°C) | -15.687 | 1.783 | -21 | -12.4 |
| Total annual precipitation (mm) | 586.258 | 181.463 | 253.2 | 1389.5 |
| Standardized avg. temperature in January | 0.001 | 0.998 | -3.2 | 1.6 |
| Standardized total annual precipitation | 0.000 | 0.999 | -2.4 | 3.8 |
| Initial age of car owner when first sampled | 38.558 | 6.337 | 22 | 56 |
| 1{owner is male} | 0.897 | 0.303 | 0 | 1 |

Table A2: Model estimation results

| | (1) True para. | (2) Est. coeff. | (3) s.e. | (4) $\lvert\text{bias}\rvert$ | (5) $\lvert\frac{\text{bias}}{\text{true}}\rvert$ | (6) $\lvert\frac{\text{bias}}{\text{s.e.}}\rvert$ |
|---|---|---|---|---|---|---|
| *Gamma/Gompertz parameters* | | | | | | |
| $b_g$ (scale parameter) | 0.5 | 0.4781 | 0.1455 | 0.0219 | 0.0438 | 0.1505 |
| $s_g$ (shape parameter 1) | 6 | 6.7181 | 1.9631 | 0.7181 | 0.1197 | 0.3658 |
| $\beta_g$ (shape parameter 2) | 10 | 10.4123 | 1.5531 | 0.4123 | 0.0412 | 0.2655 |
| *Additional cost of accident* ($\sqrt{\frac{\text{additional cost}}{\text{loss}}}$) | | | | | | |
| $\omega$ | 1 | 0.9909 | 0.1406 | 0.0091 | 0.0091 | 0.0646 |
| *Marginal cost of accident prevention (good type)* | | | | | | |
| constant | 0.1 | 0.0920 | 0.0141 | 0.0080 | 0.0804 | 0.5700 |
| car age | 0.1 | 0.0968 | 0.0234 | 0.0032 | 0.0317 | 0.1356 |
| car age$^2$ | -0.01 | -0.0100 | 0.0087 | 0.0000 | 0.0006 | 0.0007 |
| 1{commercial} | 1 | 1.0560 | 0.163 | 0.0560 | 0.0560 | 0.3434 |
| January temperature | -0.05 | -0.0520 | 0.0172 | 0.0020 | 0.0393 | 0.1143 |
| total precipitation | 0.05 | 0.0527 | 0.0268 | 0.0027 | 0.0532 | 0.0992 |
| *Marginal cost of accident prevention (normal type,* $\sqrt{\gamma_N - \gamma_G}$*)* | | | | | | |
| constant | 0.5 | 0.5051 | 0.0994 | 0.0051 | 0.0102 | 0.0515 |
| car age | 0.2 | 0.1972 | 0.0222 | 0.0028 | 0.0141 | 0.1271 |
| car age$^2$ | 0 | 0.0006 | 0.0019 | 0.0006 | - | 0.3380 |
| 1{commercial} | 0.3 | 0.3126 | 0.0182 | 0.0126 | 0.0420 | 0.6920 |
| January temperature | 0.1 | 0.1039 | 0.0284 | 0.0039 | 0.0391 | 0.1377 |
| total precipitation | 0.1 | 0.1048 | 0.0172 | 0.0048 | 0.0476 | 0.2765 |
| *Probability of being normal type* | | | | | | |
| 1{private} | 1.6 | 1.5900 | 0.2037 | 0.0100 | 0.0062 | 0.0489 |
| initial driver age | 0.5 | 0.4005 | 0.2008 | 0.0995 | 0.1990 | 0.4955 |

Notes: $b > 0$ is the scale parameter of Gamma/Gompertz distribution, and $s, \beta > 0$ are the shape parameters. The Gamma/Gompertz parameters are estimated on a subsample of 772 observations. 1{commercial} is a dummy variable indicating whether the car is for commercial use, for example taxis and rental cars. January temperature and total precipitation are all measured at the city level, where the temperature variable is standardized monthly average temperatures (demeaned and divided by the standard deviation), and the precipitation variable is standardized annual total precipitation including both rainfall and snowfall. 1{private} is a dummy variable for the car owner being an individual rather than an organization. Initial driver age is the age of the driver when his/her car first entered the sample. Standard errors are computed using numerical Hessian matrices from the maximum likelihood estimation.

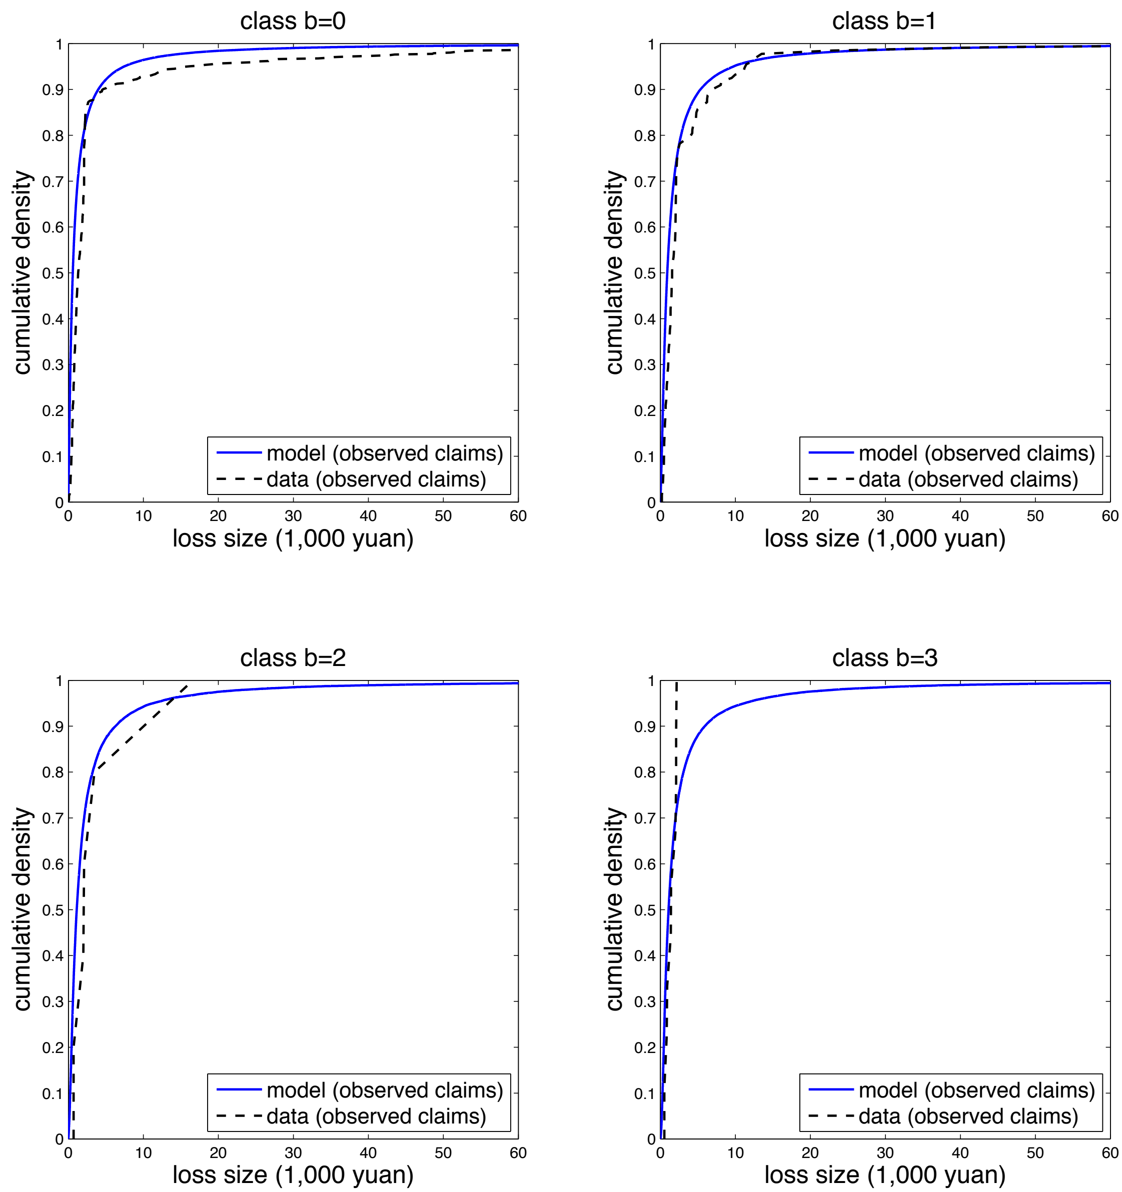


Figure A1: Goodness-of-fit: distribution of actual losses and observed claims by class