Bayesian Learning with Forgetting: An Empirical Analysis of Automobile Insurance Policyholders*

Chao Ma†

ma.271@buckeyemail.osu.edu

First version: October 2021. This version: February 2025

Abstract

Analyzing data from an automobile insurance company, we find that policyholders' past accident experience reduces their accident probabilities in the current period. Our explanation is that accidents make policyholders update their beliefs about their innate risk types upward and hence exert more caution. We rule out alternative explanations, including changes in financial incentives to prevent future accidents, underreporting of minor accidents, more limited funds after an accident, etc. Moreover, we find that the effect of a past accident is smaller in magnitude if the accident occurred longer ago, suggesting a forgetting effect in the learning process. Structural estimations show that the monthly forgetting rate is 0.1385. Additionally, we find that when a policyholder's belief of being high-risk is low, a new accident will be perceived as an occasional occurrence of a rare event instead of conveying rich information about her/his risk; at an intermediate-level belief, a new accident will reinforce the belief by a substantial amount and hence have a prominent impact on her/his caution level; if the policyholder already believes almost for sure that she/he is high-risk, a new accident will add little information and hence have little effect on her/his caution level.

Keywords: learning; forgetting; information asymmetry; automobile insurance; car accidents; adverse selection; moral hazard

JEL codes: D81; D82; D83; D9; L8; G22; G4; G52; M31; R41

---

* I thank Hengchen Dai, Teck-Hua Ho, Cheng Hsiao, Sunny Huang, Kohei Kawaguchi, Jianan Li, Mengyun Lin, John List, Francesca Molinari, Andrew Pua, Rui Tang, Xi Qu, Yi Xin, Menghan Xu, Jubo Yan, Li Zhao, Songfa Zhong, and seminar participants at China Industrial Organization Seminar, Royal Economic Society Annual Conference, Asia-Pacific IO Conference, Asian Meeting of Econometric Society, International Symposium on Advance in Panel Data and Time Series Econometrics in Honor of Cheng Hsiao's 80th Anniversary, Workshop on Recent Advances in Panel data, Microeconometrics and Applied Microeconomics, Hong Kong University of Science and Technology, Nanyang Technological University, Shanghai Jiaotong University, and Xiamen University for their helpful comments and suggestions. All errors are my own.

† WISE, Gregory Chow Institute, and School of Economics, Xiamen University, China. Email: ma.271@buckeyemail.osu.edu. Webpage: http://chaoma2014.weebly.com.

1. Introduction

Information asymmetry usually refers to the fact that agents' innate types (e.g., endowment, ability, risk) and actions (e.g., effort) are not observed by others and can only be partially inferred from some observable outcomes. A typical underlying assumption is that agents have complete information about their own innate types. However, in reality, even agents themselves may not accurately know their own ex-ante innate types and hence need to infer from observable ex-post outcomes. For example, fund managers may not accurately know their own abilities and need to infer from realized fund returns; employees may not accurately know their own abilities and need to infer from realized performance; policyholders of health insurance may not accurately know their probabilities and potential severities of sickness and need to infer from claim events. Given agents' incomplete information about their own innate types, information asymmetry may not be that asymmetric.

In this study, we examine the automobile insurance market. Adverse selection is one of the major concerns in insurance markets. Agents with better information than insurers about their innate risk types may be more likely to purchase insurance or choose high-coverage plans, which will result in incorrect pricing and losses to insurance companies. Multiple studies have examined the presence of adverse selection in automobile insurance markets (e.g., Puelz and Snow, 1994; Dionne and Doherty, 1994; Chiappori and Salanie, 2000; Cohen, 2005; and Dionne et al., 2001) and have drawn different conclusions.

However, although policyholders observe more characteristics about themselves than do insurers, they cannot accurately translate their characteristics into quantitative accident probabilities because they are not as expert as risk analysts in insurance companies. Insurance companies maintain large databases and possess information on the population of policyholders (including policyholders' characteristics and claim records), which enables insurance companies to estimate very sophisticated econometric models. In contrast, a policyholder possesses the driving experience of only herself/himself. Moreover, car accidents are infrequent events in each driver's daily life. Drivers can acquire sufficient driving skills after years of driving experience but may still not accurately know their accident probabilities, and hence they continue learning and updating their beliefs over time. Each accident they experience can substantially alter their beliefs about their innate risk types.

We study a unique dataset from an automobile insurance company in China and find that policyholders' past accident experience can reduce their accident probabilities in the current period. Our explanation is that accidents make policyholders update their beliefs about their innate risk types in an upward direction and hence exert more caution or effort to reduce risk. We also find that the magnitude of the negative effect of a past accident on the accident probability in the current period is smaller if the accident occurred longer time ago; this empirical pattern suggests a forgetting effect in the learning process. We also rule out alternative explanations, including changes in financial incentives to prevent future accidents, underreporting of minor accidents, more limited funds after an accident, etc. (see Section 7).

To explain the reduced-form empirical pattern, we build theoretical models of Bayesian learning with forgetting. Additionally, we conduct structural estimation by indirect inference and gauge the key parameters in the theoretical model to match the reduced-form empirical pattern. The estimated monthly forgetting rate (depreciation rate of past information in policyholders' belief formation) is 0.1385.

Moreover, we find a nonmonotonic pattern that the learning effect of a new accident is stronger when a policyholder has two accidents within one year than when a policyholder does not encounter an accident for years or when a policyholder has three or more accidents within one year. The explanation is that when a policyholder's belief of belonging to the high-risk type is very low, a new accident is more likely to be perceived as an occasional occurrence of a rare event instead of conveying rich information of belonging to the high-risk type; if two accidents occur within one year, the belief of high-risk type will reach a certain level such that a new accident will reinforce the belief by a substantial amount and hence have a more prominent impact on the policyholder's caution level; if three or more accidents occur within one year, the policyholder will believe almost for sure that she/he is in the high-risk type, and hence the next accident will provide almost no additional information about the policyholder's true risk type.

In addition, we find that the negative occurrence dependence is stronger for female drivers than for male drivers because female drivers are likely to be more risk-averse and hence are more responsive to the update of their beliefs caused by an accident. The negative occurrence dependence is weaker for drivers of institution-owned vehicles than for drivers of individual-

owned vehicles because, for the former drivers, the losses uncovered by insurance will be paid by their employers instead of by the drivers themselves.

The results in this study yield three important managerial implications for insurance companies. First, policyholders' incomplete information about their own innate risk types can mitigate adverse selection. In the theoretical framework developed by Einav et al. (2010) and Einav and Finkelstein (2011), adverse selection emerges if individuals who have the highest willingness to pay for insurance are those who have the highest risk, whereas advantageous selection may emerge if the individuals who are willing to pay the most for insurance are those who are the most risk-averse. Given the fact that individuals know their risk preferences but are not fully aware of their innate risk types, the real-world insurance markets may incline to exhibit advantageous selection.

A second implication is that policyholders' learning about their own types helps reduce their ex-post risk and hence mitigates moral hazard. While an accident partially reveals a policyholder's high-risk type to the insurer, the accident also partially reveals the policyholder's risk type to herself/himself and makes her/him exert more caution or effort to reduce future risk. Therefore, the insurer should increase the policyholder's premium by only a limited amount. The current "bonus-malus" pricing schemes only consider information revealed by an accident to insurers but fail to account information revealed to policyholders and their responses.

A third implication is that insurers' current pricing schedules are inconsistent with policyholders' risk dynamics. After an accident, because of the forgetting effect, a policyholder's risk should reach its lowest level in the next contract year and then increase thereafter; however, the policyholder's premium will reach the highest level in the next contract year and then decrease thereafter.

Although many economic models assumed Bayesian rules for agents' belief formation and whether people's belief formation follows Bayesian rules has been intensively examined in experiments,[1] empirical research on belief updating (learning about probabilities or innate types)

[1] For theoretical models assuming Bayesian rules for agents' belief formation, see Welch (1992), Dekel et al. (2004), Acemoglu et al. (2011), Golub and Jackson (2012), and Babus and Kondor (2018). For experimental studies examining whether people's belief formation follows Bayesian rules, see Camerer and Ho (1999), Cox et al. (2001), Charness and Levin (2005), Chandrasekhar et al. (2020), and Grimm and Mengel, (2020).

is limited, especially on belief updating based on infrequent events. Several empirical studies examined belief updating based on frequent events, such as consumers' learning about their cellphone usage demand based on the monthly bills (Grubb and Osborne, 2015; Narayanan et al., 2007; Miravete, 2003), physicians' learning about drugs' effectiveness based on patient feedbacks (Chan et al., 2013), and investors' learning about their ability from their investment experiences (Anagol et al., 2021). Gallagher (2014) studied residents' learning of the probability of large regional floods, which are a type of natural event common to every resident; the learning by residents will not change the probability of the event. In contrast, we study car accidents, which are a type of individual-specific infrequent event; the learning by policyholders can change their caution level and hence alter the probability of the event.

Although multiple theoretical studies have already modeled forgetting in agents' learning process (Camerer and Ho, 1999; Malmendier and Nagel, 2011), as laboratory experiments can only be conducted within a short period, it is difficult to detect forgetting in the Bayesian learning process, especially forgetting about information learned from infrequent events. In contrast, our study uses long-term field data and finds the forgetting phenomenon in individuals' Bayesian learning based on infrequent car accidents.[2] [3] Note that Mehta et al. (2004) empirically examined

[2] Based on aggregate-level data over 28 years, Gallagher (2014) found that the community-level take-up rates of flood insurance (policies divided by population) spike the year after a flood and then gradually decline. However, as discussed by the author, the reason for the decline of the community-level take-up rate could be a mixture of forgetting by current residents and migration. New residents immigrating from other locations after the flood year of the current location did not experience the flood and hence are less likely to take up flood insurance. In addition, among residents who were residing in the location in the flood year, over the years after the flood year, elder residents gradually pass away, and new grown-ups have vague memories of flood experience in their childhood and hence are less likely to take up insurance when they reach the age facing the take-up decision. In the structural estimation of Gallagher (2014), the hypothesis of no forgetting was not rejected. In contrast, our study is based on individual-level data. Policyholders learn their own accident probabilities rather than the probability of a public event in an area. The learning starts in policyholders' adulthood because drivers must be more than 18 years old in China. Therefore, the decreasing in magnitude of negative dependence of accidents should be caused by forgetting. Our structural estimation rejects the hypothesis of no forgetting at a high significance level.

[3] Forgetting in the Bayesian learning process examined in our study is different from the recency effect in the behavioral economic literature (e.g., Bordalo et al., 2020; Jin et al., 2023). In Bordalo et al.'s (2020) model, when people make decisions in the current situation, they recall past experiences. How much attention people pay to a certain past experience depends on both how recent (recency effects) and how similar (similarity effects) it is to the current situation. Therefore, in their model, people could pay too much attention to the most recent experience even if it is not very relevant to the current situation. In contrast, in our model, drivers only recall past car accident experiences and do not recall other adverse events (e.g., floods, earthquakes, sicknesses, financial crises) in their Bayesian updating process. The reason is that

consumers' forgetting about their prior brand evaluations while making their purchase decisions in frequently purchased product categories.

This study is not the first to examine the occurrence dependence of car accidents. Abbring et al. (2003b) studied automobile insurance data in France from 1987-1989 and did not find negative occurrence dependence. Dionne et al. (2013) studied automobile insurance data in France from 1995-1997 and found negative occurrence dependence. They attributed the negative occurrence dependence to changes in financial incentives after an accident under the multiplicative "bonus-malus" schemes in France. In contrast, under the additive "bonus-malus" schemes in China, an accident will not cause changes in financial incentives (see detailed discussions in Section 7.1). Therefore, we can attribute the negative occurrence dependence to learning.

Dionne et al. (2013) and Cohen (2005) identified learning through the phenomenon that policyholders tend to increase insurance coverage after accidents. In contrast, our study identified learning through the negative occurrence dependence of accidents. These two identification channels differ in two aspects. First, increasing insurance coverage after learning tends to raise future accident probabilities, whereas the negative dependence of accidents due to learning found in our study indicates that policyholders also exert more caution or effort to reduce risk after learning. Therefore, our channel implies that insurers should increase a policyholder's premium by only a limited amount after an accident, whereas the previous channel in the literature did not generate this managerial implication.

Second, we provide evidence not only on learning but also on forgetting based on the negative occurrence dependence of accidents that is decreasing in magnitude over time. It is difficult to identify forgetting based on changes in insurance plans over time. Given that a policyholder switches to a plan with higher coverage immediately after an accident (causing a discrete update of belief), due to switching costs, gradually forgetting over time does not necessarily lead the policyholder to switch back to a plan with a discrete reduction in coverage.

Another strand of the learning literature concerns learning certain skills instead of learning about probabilities or innate types, including learning by exporting (Clerides et al., 1998), learning

drivers' prior and post beliefs in the Bayesian updating process are both only about their innate risk of encountering car accidents and hence other adverse events are irrelevant.

by trading in financial markets (Seru et al., 2010), and learning by doing conducted by entrepreneurs (Lafontaine and Shaw, 2016), firms (Levitt et al., 2013), and employers and employees (Nagypál, 2007). We attribute the negative dependence of car accidents mainly to policyholders' learning about accident probabilities instead of their learning of driving skills. Learning driving skills is a continuous process and experienced drivers may not continue to acquire more driving skills over time because they have already acquired sufficient skills. In contrast, belief updates caused by accidents are discrete jumps; experienced drivers can still learn accident probabilities and update their beliefs after an accident because accidents are infrequent events in a driver's daily life (see the detailed discussion in Section 7.2).

A third strand of the learning literature concerns learning on certain outcomes, mechanisms, or schedules instead of learning on underlying probabilities. Haselhuhn et al. (2012) found that renters at video stores are more likely to return their videos on time if they have recently been fined for returning a video late. Agarwal et al. (2008) studied credit card behavior and found that paying a fee in the previous month reduces fee payments in the current month. Ketcham et al. (2012) studied enrollees' learning to better utilize Medicare Part D.

This study also contributes to the literature on the automobile insurance industry and other automobile-related industries (e.g., taxies and ride-sharing). Studying the automobile insurance industry is of great importance not only because the monetary volume of this industry is tremendous but also because driving safety is a crucial concern of the entire society. Most previous studies in the literature employed rational models with incomplete information to explain policyholders' behavior.[4] A few recent studies started to examine drivers' behavioral patterns and employed behavioral models to explain them. Ma (2021) found the sunk cost fallacy held by policyholders that exacerbates their moral hazard. Gao et al. (2019) examined policyholders' reference-dependence. Shum and Xin (2020) studied drivers' time-varying risk preferences. Ho et al. (2017) studied the effect of the color yellow on taxi accident rates. Duong et al. (2023) studied

[4] For example, Abbring et al. (2003a); Abbring et al. (2003b); Cohen and Dehejia (2004); Cohen (2005); Hoyt et al. (2006); Dionne et al. (2013); Weisburd (2015); Jeziorski et al. (2017); and Soleymanian et al. (2019).

taxi drivers' reference-dependent preferences.[5] Our study also employs a model with a behavioral component (forgetting) to explain policyholders' accident patterns.

The remainder of this paper is organized as follows. In Section 2, we describe the data. In Section 3, we estimate static and dynamic panel models and obtain the reduced-form empirical results of the negative dependence of car accidents. In Section 4, we build a theoretical model of Bayesian learning with forgetting and conduct simulations to explain the reduced-form results obtained in Section 3. In Section 5, we conduct heterogeneity analyses. In Section 6, we estimate the structural parameters (e.g., the forgetting rate) in the theoretical model by indirect inference to match the reduced-form empirical pattern. In Section 7, we rule out alternative explanations and address some concerns. Section 8 discusses the managerial implications of the results of this study. Finally, we conclude in Section 9.

2. Background and Data

In China, the term of an automobile insurance policy contract is one year, and the contract can start on any day of a calendar year. Each automobile insurance company has access to the entire past accident history of a policyholder, including her/his history from a previous period covered by a different insurer. An automobile insurance contract covers property damages and bodily injuries (including deaths).[6] The contract can have two layers of coverage: mandatory and optional. Prices are intensively regulated by the China Banking and Insurance Regulatory Commission. It provides formulas for calculating suggested prices based on vehicle characteristics and drivers' demographic characteristics and driving histories. Insurance companies cannot deviate from the suggested prices generated by the formulas for mandatory coverage and can only deviate within a narrow floating range (usually −15% to +15%) for optional coverage.

The pricing formulas have a "bonus-malus" scheme; the premium is the product of two factors: the "base premium" and the "bonus-malus coefficient" at the time of renewal. The "base premium" is determined based on factors such as vehicle characteristics and policyholder demographic characteristics. The "bonus-malus coefficient" is additively determined based on a

[5] Behavioral models have also been employed to explain policyholders' behavior in health insurance markets, such as salience theory (see Dalton et al., 2019; Chang et al., 2018).
[6] Health insurance contracts usually have an exclusion clause that they do not cover the medical costs of bodily injuries caused by car accidents.

policyholder's accident history: for example, the coefficient for a policyholder with one, two, or three accidents in the previous year is 1.25, 1.50, and 1.75, respectively. Thus, the increment in the "bonus-malus coefficient" caused by each accident within a contract year is constant (0.25).

After an accident occurs, the increased premium of the next contract is not the marginal cost of the next accident because the policyholder needs to pay at least the increased premium for the next contract regardless of whether the policyholder encounters a next accident. The marginal cost of the next accident is related to the increment in the next premium caused by the next accident rather than the level of the next premium itself. Therefore, under the additive "bonus-malus" scheme in China, although an accident will increase the premium of the next contract, it will not increase the marginal cost of a policyholder's next accident and hence an accident would not cause an increase in a policyholder's financial incentive to avoid future accidents.

Note that some countries such as France have a multiplicative "bonus-malus coefficient" instead of an additive one as in China. The coefficient for a policyholder with one, two, or three accidents in the previous year is 1.25, 1.25^2, and 1.25^3, respectively, for example; correspondingly, the increment of the coefficient is 0.25, 0.3125, and 0.390625, respectively ($1.25-1$, $1.25^2-1.25$, and $1.25^3-1.25^2$). Therefore, not only the level of premium but also the increment of premium caused by an accident (the marginal cost of an accident) for a policyholder is increasing in the number of accidents that have occurred to the policyholder; an accident will cause **an increase** in a policyholder's financial incentive to avoid future accidents, which can generate the negative occurrence dependence of accidents. Dionne et al. (2013) studied automobile insurance data in France and attributed the negative occurrence dependence to changes in financial incentives after an accident under the multiplicative "bonus-malus" schemes in France (see detailed discussions in Section 7.1)

Our data come from an automobile insurance company in China and cover policies sold during 2010-2017 in nine major cities. The policy information in the data includes the starting and ending dates, the characteristics of the automobile (e.g., number of seats, manufacturer, model, and manufacture year), the characteristics of the policyholder (e.g., gender, age, and individual owner vs. company or government owner), and the characteristics of the policy contract (e.g., premium and coverage). The settlement information includes the accident date, the type of accident,

the report date, the settlement date, and the settlement proceeds paid by the insurance company.[7]

We select policyholders with two or more contract years in the data, resulting in approximately 147,000 drivers and 368,000 policy contracts.[8] [9] Table 1 reports the descriptive statistics.

[Insert Table 1 here]

For each policy, the dataset provides the policyholder's previous accident history, which is classified into the following categories: "no accident in the previous three or more years," "no accident in the previous two years," "no accident in the previous year," "one accident in the previous year," "two accidents in the previous year," "three or more accidents in the previous year," and "new driver." Table A.2 in Appendix A displays the distribution of the accident history categories for all the policy contracts and the distribution for policyholders' first contracts in the data.

3. Main empirical evidence

We first estimate static panel models in Section 3.1. We run regressions of the accident indicator on policyholders' accident history categories. Because the key regressors (accident history categories) are determined by the lags of the dependent variable, the coefficients suffer from downward bias. Therefore, we estimate a dynamic panel model in Section 3.2.

3.1. Static panel model

We estimate the following static panel model:

$$y_{it} = H \cdot AccdntHstr_{it} + B \cdot X_{it} + a_i + \varepsilon_{i,t}. \quad (3.1)$$

[7] In China, drivers are required to call the police and the insurance company within 48 hours after an accident occurs. The accident date and the report date are accurately recorded as different variables in the insurance company's database. In the data, most report dates are within two days after the accident dates (see Figure 6 in Ma, 2021).

[8] We also delete policies with early termination. Approximately only 3% of contracts were terminated before the end of the contract term. The main reasons for early termination include vehicle ownership changes and the vehicle being scrapped. See footnote 14 in Ma (2021). Our results are robust regardless of whether we include policies with early termination or not.

[9] The distribution of contract years per driver in the data is displayed in Table A.1 in Appendix A.

If policyholder i encountered an accident in period t, $y_{it} = 1$; otherwise, $y_{it} = 0$. $AccdntHstr_{it}$ is a set of dummy variables indicating the accident history category of policyholder i for the current contract year. Accident history is classified into the following categories: "no accident in the previous three or more years," "no accident in the previous two years," "no accident in the previous year," "one accident in the previous year," "two accidents in the previous year," "three or more accidents in the previous year," and "new driver."[10] a_i represents the individual fixed effect. $\varepsilon_{i,t}$ is the residual. X_{it} is a set of control variables.

When we control for the individual fixed effects (a_i), X_{it} has only time-variant control variables, including the premium, the coverage, a measure of driving experience,[11] and the period fixed effects. When we do not control for the individual fixed effects (a_i), X_{it} also includes the seat number and age of the vehicle, the gender and age of the policyholder, the vehicle owner type (individual vs. company or government), the vehicle model fixed effect, and the city fixed effect.

[Insert Table 2 here]

Table 2 reports the results of the contract-year-level regressions. The different results in columns 1 and 2 of Table 2 reflect the importance of disentangling true state dependence from individual heterogeneity, which is a well-known issue in the applied microeconomic literature (see Heckman, 1991). In column 1 of Table 2, when we do not control for the individual fixed effects (a_i), the worse the accident history is, the higher the accident probability for the current contract cycle. Without controlling for individual fixed effects, the accident history partially reveals a policyholder's unobserved innate risk type (a_i). Risky drivers are more likely to encounter accidents both in the past and in the current period (heterogeneity). The result in column 1 is consistent with insurers' pricing strategy that policyholders with a worse accident history are charged higher premiums.

[10] Traffic violations are not counted as accidents.

[11] We do not directly observe policyholders' time length of driving experience; thus, we use the number of periods since the start of their first contract with this insurance company as a proxy. Their true periods of driving experience should equal the sum of periods of driving before their first contract with this insurance company and periods since the start of their first contract with this insurance company. The unobserved former part is controlled for by the individual fixed effects. The proxy is equal to the true length of driving experience for the policyholders of whom the accident history for their first policies with this insurance company is categorized as "new driver."

In contrast, in column 2 of Table 2, when we control for the individual fixed effects (a_i) and hence the variation of accident history is within a driver, a worse accident history leads to a lower accident probability for the current contract cycle. This result indicates that policyholders learn about their own innate risk types from the accidents they previously experienced and correspondingly exert more caution or effort to reduce their accident probabilities in the current contract cycle (state dependence).[12]

3.2. Dynamic panel model

Because the key regressors (accident history categories) in equation (3.1) are determined by the lags of the dependent variable, the coefficients suffer from a downward bias. Thus, we estimate a dynamic panel model as follows:

$$y_{it} = \sum_{\tau=1}^{m} \xi_\tau y_{i,t-\tau} + B \cdot X_{it} + a_i + \varepsilon_{i,t}. \tag{3.2}$$

In equation (3.2), ξ_τ captures the effect of the accident in the τth period before the current period on the accident probability in the current period. Because the lagged dependent variable $y_{i,t-\tau}$ appears on the right-hand side of equation (3.2), the within estimator for the fixed-effect model has a downward bias. Therefore, we follow the standard procedure (see Anderson and Hsiao, 1981; Arellano and Bond, 1991) to estimate dynamic panel models.

Taking the first-order difference for y_{it} based on equation (3.2), we have

$$\Delta y_{it} = y_{it} - y_{i,t-1} = \sum_{\tau=1}^{m} \xi_\tau \Delta y_{i,t-\tau} + B \cdot \Delta X_{it} + \Delta \varepsilon_{i,t}, \tag{3.3}$$

where $\Delta y_{i,t-\tau} = y_{i,t-\tau} - y_{i,t-\tau-1}$, $\Delta X_{it} = X_{it} - X_{i,t-1}$, and $\Delta \varepsilon_{it} = \varepsilon_{it} - \varepsilon_{i,t-1}$. In $\sum_{\tau=1}^{m} \xi_\tau \Delta y_{i,t-\tau}$, $\Delta y_{i,t-1} = y_{i,t-1} - y_{i,t-2}$; thus, $\Delta y_{i,t-1}$ is correlated with $\varepsilon_{i,t-1}$ and hence $\Delta \varepsilon_{it}$, i.e., $\Delta y_{i,t-1}$ in equation (3.3) is endogenous. Following Anderson and Hsiao (1981) and Arellano and Bond

[12] Given that we include only policyholders with two or more years of records in the data, selective attrition could be a potential concern. Tables A.3 and A.4 in Appendix A report the descriptive statistics separately for policyholders with only one year of records and those with two or more years of records. The two groups are similar in characteristics and risk. We also use policyholders with only one year of records to estimate the static panel model without individual fixed effects similar to column 1 of Table 2 and obtain similar results. Even if there were selective attrition, our results would still be meaningful to insurance companies because their profits and risk do not depend on drivers who have already dropped out.

(1991), we use $\Delta y_{i,t-m-1}, \Delta y_{i,t-m-2}, \ldots, \Delta y_{i,t-m-l}$ as instruments for $\Delta y_{i,t-1}$, while using $\Delta y_{i,t-2}$, $\Delta y_{i,t-3}$,…, $\Delta y_{i,t-m}$, and ΔX_{it} as instruments for themselves.

[Insert Table 3 here]

Table 3 reports the quarter-level regression results for both the within estimator and the 2SLS estimator of ξ_τ, $\tau = 1,2,\ldots,m$. We set $m = 7$, i.e., we estimate the effects of accidents that occurred within two years before the end of the current period on the accident probability in the current period. We use $\Delta y_{i,t-8}$ and $\Delta y_{i,t-9}$ as the instrumental variables for the endogenous regressor $\Delta y_{i,t-1}$.[13] The 2SLS estimator shows that the $\xi_\tau s$ are all significantly negative. The results indicate that policyholders may learn their innate risk types through their past accidents and exert more caution or effort to prevent accidents in the future. Furthermore, the 2SLS estimator shows that the ξ_τs decrease in magnitude as τ increases. The results indicate a forgetting effect: the effect of a past accident on reducing the accident probability in the current period is smaller if the past accident occurred longer time ago.

When seven lags of y_{it} are included as regressors ($m = 7$) in equation (3.2), the resulting ε_{it}s do not have a statistically significant serial correlation, suggesting that no additional lags of y_{it} are omitted. Therefore, drivers' effective memory duration for an accident is 7 quarters. Events beyond 7 quarters ago are completely forgotten in terms of effectiveness. A historical event only affects drivers' actions within two years.

We also rule out alternative explanations, including changes in financial incentives to prevent future accidents, underreporting of minor accidents, more limited funds after an accident, etc. (see Section 7).

The within estimator of ξ_τs exhibits a similar pattern as that of the 2SLS estimator, but the magnitude of the within estimator is larger than that of the 2SLS estimator. The reason is that the within estimators for dynamic panel models have downward biases.

[Insert Figure 1 here]

[13] The estimation results are robust to changes in the instrumental variables (see Table A.5 in Appendix A).

Figure 1 reports the 2SLS estimator of ξ_τs for the month-level regression for equation (3.3). The ξ_τs are negative and decreasing in magnitude with τ.

4. Theoretical models

In Section 4, we build theoretical models and conduct simulations to explain the reduced-form empirical pattern obtained in Section 3. We gradually add the components of Bayesian learning and forgetting into the theoretical models and discuss the roles played by the key parameters in the theoretical models.

4.1. Bayesian learning

Assume that there are two innate types of policyholders: high-risk type with θ=θ_H and low-risk type with θ=θ_L ($\theta_H > \theta_L$). The probability that an accident occurs to policyholder i in period t given her/his innate type θ is

$$Prob(y = 1|\theta, c_t) = \theta e^{-\eta c_t}, \tag{4.1}$$

where c_t represents the level of effort or caution the policyholder exerts to prevent accidents in a month. The parameter η captures the effectiveness of caution in reducing risk.[14] [15] [16] Assume that the policyholder will incur a loss L if an accident occurs.[17]

A policyholder does not know her/his innate risk type θ but holds a belief about the distribution of θ in the current period, denoted as $Prob\ (\theta|Y_{t-1})$, given the policyholder's accident history $Y_{t-1} = \{y_{t-1}, y_{t-1}, \dots, y_1\}$. If the policyholder is a new driver, her/his accident history is an empty set, i.e., $Y_0 = \{\}$.

Assume that the policyholder is risk-averse and her/his Bernoulli utility function $u(x) = -e^{-\gamma x}$. x includes losses in accidents and costs (hassles) of exerting effort to avoid accidents.[18] We assume that $0 < \gamma < \eta$ to ensure that the optimization solution exists. In each period, given

[14] The innate risk type can be interpreted as an exogenous type that a driver cannot change. It can depend on not only the driver's characteristics but also the environment or background in which the driver resides. If the environment has busy traffic and poor road conditions or other drivers in the location have a high tendency of negligence on average, the driver can be in the high exogenous risk type. The innate risk type can also include a driver's innate ability to avoid accidents caused by other drivers' negligence.

[15] We do not distinguish accidents at fault and accidents not at fault. We are able to distinguish for part of the observations in the data but not for all. However, the accident probability in a driver's belief should cover any potential accident instead of only accidents at fault because a driver can suffer from losses caused by accidents not at fault (e.g., bodily injuries, death, and monetary losses not fully covered by the at-fault driver's insurance). A not-at-fault accident can make a driver update her/his belief about the background risk (included in her/his exogenous risk type; see the previous footnote) and exert more caution in driving.

[16] Besides innate risk type and driving caution, driving skills or experience is a third factor affecting the accident probability. In the reduced-form empirical analyses, we control for the fixed effects of periods of driving experience. We also divide the sample into first-year drivers and experienced drivers and run separate regressions; the results are similar between the two groups. Given that our study focuses on innate risk type and driving caution, driving experience is abstracted from the theoretical model.

[17] For simplicity, I assumed that experiencing an accident alters the accident probability in policyholders' beliefs but not the probability distribution of the loss given an accident in policyholders' beliefs; effort only affects the accident probability and losses given an accident are fixed. As drivers are maximizing their expected utilities, affecting the accident probability and affecting the loss given an accident are equivalent in terms of affecting the expected utility. The model can be extended to variant loss sizes but all the properties derived from the model will not change. This simplification is standard in the literature (e.g., Shum and Xin (2020)).

[18] This utility functional form has been widely used in the empirical literature on risk and insurance (e.g., Shum and Xin, 2020; Cohen and Einav, 2007; Barseghyan et al., 2013). Although constant relative risk aversion (CRRA) utility possesses better theoretical properties, it requires individual-level wealth as an input in empirics, which is usually unobserved in insurance data. As mentioned by Cohen and Einav (2007), "CRRA requires us to make additional assumptions about the relevant (and unobserved) wealth level of each individual, making it less attractive." Barseghyan et al. (2013) observed policyholders' home values and assumed that wealth is proportional to home value. They found that the main results of structural estimation based on CRRA utility are similar to those based on constant absolute risk aversion (CARA) utility.

her/his belief on θ (or given her/his accident history), the policyholder's expected utility in belief as a function of the caution level is as follows:

$$V(c_t|Y_{t-1}) = Prob\ (\theta_H|Y_{t-1})[\theta_H e^{-\eta c_t} u(-c_t - L) + (1 - \theta_H e^{-\eta c_t}) u(-c_t)] + Prob\ (\theta_L|Y_{t-1})[\theta_L e^{-\eta c_t} u(-c_t - L) + (1 - \theta_L e^{-\eta c_t}) u(-c_t)] \tag{4.2}$$

The policyholder optimally chooses her/his effort level c_t by maximizing the expected utility in belief:

$$\max_{c_t} V(c_t|Y_{t-1})$$

From the first-order condition w.r.t. c_t, we have

$$e^{-\eta c_t} = \frac{\gamma}{(\eta - \gamma)(e^{\gamma L} - 1)E[\theta|Y_{t-1}]}\ , \tag{4.3}$$

where $E[\theta|Y_{t-1}] = \theta_H Prob\ (\theta_H|Y_{t-1}) + \theta_L Prob\ (\theta_L|Y_{t-1})$. Plugging equation (4.3) into equation (4.1), we obtain the probability of encountering an accident in period t given the policyholder's innate risk type θ and accident history Y_{t-1}:

$$Prob(y = 1|\theta, Y_{t-1}) = \frac{\gamma\theta}{(\eta - \gamma)(e^{\gamma L} - 1)E[\theta|Y_{t-1}]}\ . \tag{4.4}$$

The higher the policyholder's innate risk type θ is, the larger the accident probability will be. The higher the expected θ in the policyholder's belief is, the smaller the accident probability will be because the policyholder will exert more caution.

Assume that, for a new driver with the accident history $Y_0 = \{\}$, the prior belief of the distribution of θ is as follows:

$$Prob\ (\theta = \theta_H|Y_0) = p_H$$

$$Prob\ (\theta = \theta_L|Y_0) = p_L$$

Assume that the policyholder updates her/his belief using the Bayesian rule. The belief at the beginning of the second period given the accident history in the first period ($\{Prob\ (\theta_H|y_1), Prob\ (\theta_L|y_1)\}$) is as follows:

$$Prob\ (\theta_H|y_1 = 1) = \frac{Prob(y_1 = 1|\theta_H)Prob\ (\theta_H|Y_0)}{Prob(y = 1|\theta_H)Prob\ (\theta_H|Y_0) + Prob(y = 1|\theta_L)Prob\ (\theta_L|Y_0)}$$
$$= \frac{\theta_H e^{-\eta c_1} p_H}{\theta_H e^{-\eta c_1} p_H + \theta_L e^{-\eta c_1} p_L} = \frac{\theta_H p_H}{\theta_H p_H + \theta_L p_L} \tag{4.5}$$

$$Prob\ (\theta_L|y_1 = 1) = \frac{Prob((y_1 = 1|\theta_L)Prob\ (\theta_L|Y_0)}{Prob(y = 1|\theta_H)Prob\ (\theta_H|Y_0) + Prob(y = 1|\theta_L)Prob\ (\theta_L|Y_0)}$$
$$= \frac{\theta_L e^{-\eta c_1} p_L}{\theta_H e^{-\eta c_1} p_H + \theta_L e^{-\eta c_1} p_L} = \frac{\theta_L p_L}{\theta_H p_H + \theta_L p_L} \tag{4.6}$$

$$Prob\ (\theta_H|y_1 = 0) = \frac{Prob(y_1 = 0|\theta_H)Prob\ (\theta_H|Y_0)}{Prob(y = 0|\theta_H)Prob\ (\theta_H|Y_0) + Prob(y = 0|\theta_L)Prob\ (\theta_L|Y_0)}$$
$$= \frac{(1 - \theta_H e^{-\eta c_1}) p_H}{(1 - \theta_H e^{-\eta c_1}) p_H + (1 - \theta_L e^{-\eta c_1}) p_L} \tag{4.7}$$

$$Prob\ (\theta_L|y_1 = 0) = \frac{Prob(y_1 = 0|\theta_L)Prob\ (\theta_L|Y_0)}{Prob(y = 0|\theta_H)Prob\ (\theta_H|Y_0) + Prob(y = 0|\theta_L)Prob\ (\theta_L|Y_0)}$$
$$= \frac{(1 - \theta_L e^{-\eta c_1}) p_L}{(1 - \theta_H e^{-\eta c_1}) p_H + (1 - \theta_L e^{-\eta c_1}) p_L} \tag{4.8}$$

More generally, the belief at the beginning of the tth period given the accident history in the first $t - 1$ periods ($\{Prob\ (\theta_H|Y_{t-1}), Prob\ (\theta_L|Y_{t-1})\}$) is as follows:

$$Prob\ (\theta_H|y_{t-1}, Y_{t-2}) = \frac{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2})}{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2}) + Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})}$$
$$= \frac{p_H \prod_{\tau=1}^{t-1} {\theta_H}^{y_{t-\tau}}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}{p_H \prod_{\tau=1}^{t-1} {\theta_H}^{y_{t-\tau}}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}} + p_L \prod_{\tau=1}^{t-1} {\theta_L}^{y_{t-\tau}}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}$$
$$= \frac{p_H {\theta_H}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}{p_H {\theta_H}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}} + p_L {\theta_L}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}$$

(4.9)

$$Prob\ (\theta_L|y_{t-1}, Y_{t-2}) = \frac{Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})}{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2}) + Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})}$$
$$= \frac{p_L \prod_{\tau=1}^{t-1} {\theta_L}^{y_{t-\tau}}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}{p_H \prod_{\tau=1}^{t-1} {\theta_H}^{y_{t-\tau}}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}} + p_L \prod_{\tau=1}^{t-1} {\theta_L}^{y_{t-\tau}}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}$$
$$= \frac{p_L {\theta_L}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}{p_H {\theta_H}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}} + p_L {\theta_L}^{\sum_{\tau=1}^{t-1} y_{t-\tau}} \prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{1-y_{t-\tau}}}$$

(4.10)

Based on equations (4.9) and (4.10), each accident in the past should have a similar impact on the accident probability in the current period instead of having a monotonically decreasing impact as τ increases. Although the impact of $y_{t-\tau} = 1$ relative to $y_{t-\tau} = 0$ is related to $c_{t-\tau}$, $c_{t-\tau}$ should zigzag rather than evolve roughly monotonically with τ because if the random variable $y_{t-\tau-1}$ is realized to be 1, $c_{t-\tau}$ will shift substantially in response.

We set values for the parameters in the model and simulate the monthly accident path y_t. Given that our empirical sample for the regression of equation (3.3) has 53,934 policyholders, we randomly generate 53,934 policyholders; each policyholder's θ has a 50% probability to equal θ_H and a 50% probability to equal θ_L. Given that policyholders in the empirical sample have a four-year record on average, we simulate the monthly accident path y_t for four years for each policyholder. Then, we use the simulated y_t to estimate a dynamic panel model similar to Section 3 and obtain the estimates for ξ_τs, as displayed in Figure 2. While these ξ_τs are generally negative, they are not decreasing in magnitude as τ increases, which is different from the empirical results in Section 3 (Figure 1).

[Insert Figure 2 here]

We calculate the average monthly accident probability for each policyholder in the empirical sample who has more than two years of records in the data. The means of policyholders above and below the median are 0.0470794 and 0.002535, respectively. Accordingly, to generate Figure 2, we set that $\theta_H = 0.0470794$ and $\theta_L = 0.002535$ and that the true distribution of innate risk type θ in the population is that $Prob(\theta = \theta_H) = 0.5$ and $Prob(\theta = \theta_L) = 0.5$. We assume that each policyholder's initial prior belief about her/his θ is that $p_H = 0.1$ and $p_L = 0.9$. The average payment by the insurance company for an accident in the sample is RMB7103.73. Accordingly, we set $L = \frac{7103.73}{1000}\frac{0.15}{1-0.15} = 1.2536$ because the out-of-pocket cost of an accident for policyholders in China is 10-20% of the total damage caused by the accident.[19] We set $\gamma = 5.2$ and $\eta = 6$ to make the simulated $e^{-\eta c_t}$ fluctuate within a reasonable range (from a number reasonably lower than 1 to a number reasonably higher than 1).

4.2. Forgetting

To explain the downward trend of the magnitudes of ξ_τs in the empirical results, we add a forgetting component to the theoretical learning model. Instead of the Bayesian updating rule stated in equations (4.9) and (4.10), policyholders update their beliefs following a rule with a forgetting rate $1-\delta$ (or a remembering rate δ, $0 < \delta \leq 1$) as follows:

$$Prob\ (\theta_H|y_{t-1}, Y_{t-2}) = \frac{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2})^\delta}{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2})^\delta + Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})^\delta}$$

$$= \frac{{p_H}^{\delta^{t-1}}\theta_H^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}}}{{p_H}^{\delta^{t-1}}\theta_H^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}} + {p_L}^{\delta^{t-1}}\theta_L^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}}} \quad (4.11)$$

$$Prob\ (\theta_L|y_{t-1}, Y_{t-2}) = \frac{Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})^\delta}{Prob(y_{t-1}|\theta_H)Prob\ (\theta_H|Y_{t-2})^\delta + Prob(y_{t-1}|\theta_L)Prob\ (\theta_L|Y_{t-2})^\delta}$$

$$= \frac{{p_L}^{\delta^{t-1}}\theta_L^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}}}{{p_H}^{\delta^{t-1}}\theta_H^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_H e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}} + {p_L}^{\delta^{t-1}}\theta_L^{\sum_{\tau=1}^{t-1} y_{t-\tau}\delta^{\tau-1}}\prod_{\tau=1}^{t-1}(1-\theta_L e^{-\eta c_{t-\tau}})^{(1-y_{t-\tau})\delta^{\tau-1}}}$$

[19] See Shum and Xin (2020).

(4.12)

The presence of δ in equations (4.11) and (4.12) makes accidents in the past have lower weights in influencing a policyholder's current belief.[20]

[Insert Figure 3 here]

We set a value for the remembering rate δ while keeping the other parameters the same as in Figure 2; we simulate the monthly accident path y_t. Then, we use the simulated y_t to estimate a dynamic panel model similar to that in Section 3 and obtain the estimates for ξ_τs, as displayed in Figure 3. The blue, red, orange, and purple lines display the ξ_τs for the cases of $\delta = 0.85$, $\delta = 0.70$, $\delta = 0.55$, and $\delta = 0.40$, respectively. The ξ_τs are negative and decreasing in magnitude as τ increases, similar to the empirical ξ_τs. Given that the past information depreciates exponentially (see equations (4.11) and (4.12)), under a higher forgetting rate (lower δ), the line of ξ_τs is more concave, i.e., the ξ_τs are decreasing in magnitude at the beginning at a higher speed as τ increases.

5. Heterogeneity Analyses

In Section 5, we conduct heterogeneity analyses for policyholders with different initial accident history categories (the accident history category for each policyholder's first contract year in the data), different genders, and different vehicle owner types (individual vs. company or government).

5.1 Initial accident history

In Subsection 5.1, we first theoretically derive the relationship between the magnitude of the effect of an accident on the accident probability in the next period and the policyholder's prior belief before the accident. Then, we test the relationship empirically.

In period $t-1$, given a policyholder's belief based on the accident history during the first $t-2$ periods ($Prob\ (\theta_H|Y_{t-2})$ and $Prob\ (\theta_L|Y_{t-2})$), the accident indicator y_{t-1} will make the policyholder update her/his belief following the Bayesian rule, as follows:

[20] Camerer and Ho (1999), Malmendier and Nagel (2011), and Gallagher (2014) also modeled the forgetting effect with a weighting parameter that discounts past information.

$$Prob\ (\theta_H|y_{t-1}=1,Y_{t-2})=\frac{\theta_H Prob\ (\theta_H|Y_{t-2})}{\theta_H Prob\ (\theta_H|Y_{t-2})+\theta_L Prob\ (\theta_L|Y_{t-2})}$$

$$Prob\ (\theta_L|y_{t-1}=1,Y_{t-2})=\frac{\theta_L Prob\ (\theta_L|Y_{t-2})}{\theta_H Prob\ (\theta_H|Y_{t-2})+\theta_L Prob\ (\theta_L|Y_{t-2})}$$

$$Prob\ (\theta_H|y_{t-1}=0,Y_{t-2})=\frac{(1-\theta_H e^{-\eta c_{t-1}})Prob\ (\theta_H|Y_{t-2})}{(1-\theta_H e^{-\eta c_{t-1}})Prob\ (\theta_H|Y_{t-2})+(1-\theta_L e^{-\eta c_{t-1}})Prob\ (\theta_L|Y_{t-2})}$$

$$Prob\ (\theta_L|y_{t-1}=0,Y_{t-2})=\frac{(1-\theta_L e^{-\eta c_{t-1}})Prob\ (\theta_L|Y_{t-2})}{(1-\theta_H e^{-\eta c_{t-1}})Prob\ (\theta_H|Y_{t-2})+(1-\theta_L e^{-\eta c_{t-1}})Prob\ (\theta_L|Y_{t-2})}$$

The updated expectations of innate risk type given y_{t-1} are as follows:

$$E[\theta|y_{t-1}=1,Y_{t-2}]=\frac{{\theta_H}^2 Prob\ (\theta_H|Y_{t-2})+{\theta_L}^2 Prob\ (\theta_L|Y_{t-2})}{\theta_H Prob\ (\theta_H|Y_{t-2})+\theta_L Prob\ (\theta_L|Y_{t-2})}$$

$$E[\theta|y_{t-1}=0,Y_{t-2}]=\frac{\theta_H(1-\theta_H e^{-\eta c_{t-1}})Prob\ (\theta_H|Y_{t-2})+\theta_L(1-\theta_L e^{-\eta c_{t-1}})Prob\ (\theta_L|Y_{t-2})}{(1-\theta_H e^{-\eta c_{t-1}})Prob\ (\theta_H|Y_{t-2})+(1-\theta_L e^{-\eta c_{t-1}})Prob\ (\theta_L|Y_{t-2})}$$

Given that the policyholder's caution conditional on her/his belief is determined by equation (4.3) and that the resulting accident probability is determined by equation (4.4), the effect of an accident in period $t-1$ relative to no accident in period $t-1$ on the accident probability in period t through altering the belief at the beginning of period t is:

$$\begin{aligned}&Prob(y_t=1|\theta,y_{t-1}=1,Y_{t-2})-Prob(y_t=1|\theta,y_{t-1}=0,Y_{t-2})\\&\quad=\frac{\gamma\theta}{(\eta-\gamma)(e^{\gamma L}-1)}\left[\frac{1}{E[\theta|y_{t-1}=1,Y_{t-2}]}-\frac{1}{E[\theta|y_{t-1}=0,Y_{t-2}]}\right]\end{aligned}$$

It is easy to show the following:

$$Prob(y_t=1|\theta,y_{t-1}=1,Y_{t-2})-Prob(y_t=1|\theta,y_{t-1}=0,Y_{t-2})<0$$

Furthermore, we have the following proposition:

Proposition 1: $\exists$ a $\tilde{p}\in(0,1)$ such that

$$\frac{d[Prob(y_t=1|\theta,y_{t-1}=1,Y_{t-2})-Prob(y_t=1|\theta,y_{t-1}=0,Y_{t-2})]}{dProb\ (\theta_H|Y_{t-2})}\begin{cases}<0 & if\ Prob\ (\theta_H|Y_{t-2})\in[0,\tilde{p})\\ \geq 0 & if\ Prob\ (\theta_H|Y_{t-2})\in[\tilde{p},1]\end{cases}$$

The proof is provided in Appendix B.

$\frac{d[Prob(y_t = 1|\theta, y_{t-1} = 1, Y_{t-2}) - Prob(y_t = 1|\theta, y_{t-1} = 0, Y_{t-2})]}{dProb\,(\theta_H|Y_{t-2})}$ measures how the prior belief at the beginning of the last period can alter the impact of an accident in the last period relative to no accident in the last period on the accident probability in the current period. When the prior belief of the high-risk type at the beginning of the last period is very low, a new accident is more likely to be perceived as an occasional occurrence of a rare event instead of conveying rich information of belonging to the high-risk type. As the prior belief of high-risk type at the beginning of the last period increases, a new accident will reinforce the belief more and hence have a more prominent impact on the policyholder's caution level. In contrast, when the prior belief of high-risk type at the beginning of the last period is sufficiently high (the policyholder is already almost sure that she/he is in the high-risk type), as the prior belief increases, a new accident will provide less information about the policyholder's true risk type and hence have a less prominent impact on the policyholder's caution level.

To test whether the theoretical prediction is consistent with the empirical pattern, we run the month-level regression of the dynamic panel model in equation (3.3) separately for policyholders with different initial accident history categories (the accident history for each policyholder's first contract cycle in the sample). Figure 4 displays the 2SLS estimates of ξ_τs for each of the following groups: "no accident for three or more years," "no accident for two years," "no accident for one year," "one accident in the last year," "two accidents in the last year," and "three or more accidents in the last year."

[Insert Figure 4 here]

From the group with the best driving records at the beginning of the sample period through the group with the worst driving records, the belief of being high-risk type at the beginning should be monotonically increasing. Therefore, that belief afterward should also be generally increasing from the former group through the latter group. However, as shown in Figure 4, the magnitudes of ξ_τs (the learning effect) are roughly increasing from the category of "no accident for three or more years" to the category of "two accidents in the last year" and then decreasing from the category of "two accidents in the last year" to the category of "three or more accidents in the last year." This nonmonotonic empirical pattern is consistent with the theoretical prediction.

The empirical turning point of the nonmonotonicity is "two accidents in the last year." First, when a policyholder does not encounter an accident for years, the policyholder's belief of belonging to the high-risk type is very low such that a new accident is more likely to be perceived as an occasional occurrence of a rare event instead of conveying rich information of belonging to the high-risk type.

Second, if two accidents occur within one year, the belief of high-risk type will reach a certain level such that a new accident will reinforce the belief by a substantial amount and hence have a more prominent impact on the policyholder's caution level.

Third, if three or more accidents occur within one year, the policyholder will believe almost for sure that she/he is in the high-risk type; correspondingly, the next accident will provide almost no additional information about the policyholder's true risk type and hence have a less prominent impact on the policyholder's caution level.[21]

5.2. Gender

We run the regressions separately for female and male drivers. The estimates of ξ_τs for each group are displayed in Figure 5. Female drivers have a stronger learning effect than do male drivers. One possible explanation is that female drivers are more risk-averse (greater γ in equation (4.4)) than

[21] One possible concern is why the line for "2 accidents in the previous year" is far away from the other lines in Figure 4. The reason is that the event probabilities change highly nonlinearly in percent from "no accident for 3 years" to "3 accidents in one year." Suppose that the monthly accident probability for a driver is 0.0248 (the sample average); as shown in Figure A.14 in Appendix A, the probabilities of one, two, or three accidents in a year are 0.2258, 0.0316, or 0.0027, respectively. Correspondingly, "two accidents in a year" is an event with a probability smaller than 5%, and "three accidents in a year" is an event with a probability smaller than 1%. An occurrence of a small-probability event will substantially alter a driver's belief, making it significantly deviate from the average. Structural estimation allowing heterogeneous prior beliefs (Table A.6 in Appendix A) shows that the prior beliefs of being high-risk at the beginning of their first contract in the sample are 0.3016 or 0.6003, respectively, for policyholders in the accident history categories of "2 accidents in the previous year" or "3 or more accidents in the previous year" for their first contract, highly deviated from the sample average (0.1019, see Table 4 for structural estimation assuming homogenous prior beliefs) and significantly different from each other. In contrast, the prior belief of the remaining policyholders (0.1005) is close to the sample average.

male drivers. Consequently, female drivers' caution level and accident probability are more sensitive to the update of their beliefs on their innate risk types.[22] [23]

[Insert Figure 5 here]

5.3 Individual-owned vs. company-/government-owned vehicles

We run the month-level regression of the dynamic panel model in equation (3.3) separately for individual-owned and company-/government-owned vehicles. The estimates of ξ_τs for each group are displayed in Figure 6. Divers of individual-owned vehicles have a stronger learning effect than do drivers of company-/government-owned vehicles. One possible explanation is that the losses incurred to a company-/government-owned vehicle that are not covered by insurance companies are borne partially by the company/government and hence the driver bears only a small out-of-pocket loss (smaller L in equation (4.4)). Consequently, the driver's caution level and accident probability are less sensitive to the update of her/his belief on her/his innate risk type.

[Insert Figure 6 here]

6. Structural estimation

In this section, we structurally estimate the key parameters in the theoretical model of Bayesian learning with forgetting by indirect inference.[24]

---

[22] Cohen and Einav (2007) found that females are more risk-averse than males in the context of automobile insurance coverage choices. The literature also documented that females tend to be more risk-averse than males in multiple other contexts (see Castillo et al., 2020; Huang and Bao, 2020; Charness and Gneezy, 2012; Eckel and Grossman, 2008; Dohmen et al., 2011; Gneezy and Rustichini, 2004; Gneezy et al., 2003).

[23] In equation (4.4), the exponential effect of γ in $e^{\gamma L}$ dominates the linear effects of γ in $\eta - \gamma$ and $\gamma\theta$. Therefore, a larger γ leads to a higher sensitivity to changes in $E[\theta|Y_{t-1}]$.

[24] Indirect inference was widely used in estimating or calibrating structural models in industrial organization (Ackerberg and Gowrisankaran, 2006; Goettler, and Gordon, 2011; Collard-Wexler, 2013; Greenwood and Hanson, 2014), international trade (Ruhl and Willis, 2017), labor economics (Topa, 2001), household finance (van der Klaauw and Wolpin, 2008; Low, and Pistaferri, 2015), corporate finance (Hennessy and Whited, 2005; Coles et al., 2012), asset pricing (Phillips et al., 2011), macroeconomics (Cooper and Haltiwanger, 2006; Lentz and Mortensen, 2008), law economics (Voena, 2015), and political economics (Martin and Yurukoglu, 2017). The technical details about this method can be found in Gourieroux et al. (1993) and Gourieroux and Monfort (1996).

6.1. Estimation procedure and results

Policyholders' initial prior belief about their innate risk types, which is represented by $\{\theta_H, \theta_L, p_H, 1-p_H\}$ in the theoretical models, is unobservable. Given that the means of average monthly accident probabilities for policyholders in the actual data above and below the median, respectively, are 0.0470794 and 0.002535, we set $\theta_H = 0.0470794$ and $\theta_L = 0.002535$ and leave p_H to be estimated. We set $L = \frac{7103.73}{1000}\frac{0.15}{1-0.15} = 1.2536$ based on the logic discussed in Section 4.1. The structural estimation results in this study are robust to changes in the assumed initial prior belief. [25] [26]

Next, we estimate γ, η, δ, and p_H. We randomly generate 53,934 drivers (the same as the number in the empirical regression sample); each driver's θ has a 50% probability to equal θ_H and a 50% probability to equal θ_L. Given a set of values for the parameters $\{\gamma, \eta, \delta, p_H\}$, we simulate the monthly accident path y_t for five years for each policyholder. Then, we use the simulated y_t to estimate a dynamic panel model similar to that described by equation (3.2) in Section 3 and obtain the estimates for ξ_τs, denoted as vector $\tilde{\Xi}$. Denote the empirical ξ_τs of the monthly regression in Figure 1 in Section 3 as column vector $\hat{\Xi}$. Denote the average accident probability of the simulated data as $\tilde{\mathrm{g}}(\gamma, \eta, \delta, p_H)$ and the average accident probability of the empirical sample as $\hat{\mathrm{g}}$. Let the column vector

$$\mathbb{G} = \begin{bmatrix} \Xi \\ \mathrm{g} \end{bmatrix}$$

The parameters $\{\gamma, \eta, \delta, p_H\}$ are estimated by solving the following distance minimization problem:

[25] In many structural estimations in the literature, initial distributions are unobserved and hence have to be assumed, but the estimation results are robust to changes in assumed initial distributions. For example, Hendel and Nevo (2006) estimated a dynamic demand model for storable goods. They assumed the distribution of consumers' initial inventories and found that the structural estimation results were robust to changes in the assumed distribution.

[26] As robustness checks, we further allow policyholders in different accident history categories in the first period of the data to have different initial prior beliefs about their innate risk types. A similar estimate for the remembering rate δ is obtained (In Appendix A, see Table A.6 for estimation results and Figure A.3 for the fit).

$$\min_{\{\gamma,\eta,\delta,p_H\}} \left(\tilde{\mathbb{G}} - \hat{\mathbb{G}}\right)'\left(\tilde{\mathbb{G}} - \hat{\mathbb{G}}\right)$$

The estimated values for $\{\gamma, \eta, \delta, p_H\}$ are reported in Table 4. The standard errors are estimated by bootstrapping. The simulated $\tilde{\Xi}$ given the estimated $\{\gamma, \eta, \delta, p_H\}$ is plotted as the dash blue line in Figure 7, which has a similar pattern as that of its empirical counterpart $\hat{\Xi}$ (the solid blue line). The remembering rate δ is 0.8615, significantly smaller than 1 at a 5% level; thus, we can reject the hypothesis that no forgetting occurs in the learning process. Correspondingly, the monthly forgetting rate (monthly depreciation rate of past information, $1 - \delta$) is 0.1385.

[Insert Table 4 here]

[Insert Figure 7 here]

As robustness checks, we allow that the coefficient of risk aversion γ to be estimated is different across genders. A similar estimate for the remembering rate δ is obtained (In Appendix A, see Table A.7 for estimation results and Figure A.4 for the fit).

In the structural model, we do not explicitly model policyholders' decisions on insurance coverage because our focus is how learning on accident probabilities alters drivers' caution levels conditional on coverage; abstracting from coverage decisions can make the structural model substantially parsimonious. Essentially, we use simulated ξ_τs from the structural model without the component of coverage decisions to match ξ_τs from the real-data reduced-form regression in which the effect of coverage on accident probabilities has already been controlled for.

The framework is extendable to modeling coverage decisions and allowing the degree to which belief updating affects caution exerting to differ across different coverage plans. The former component can be estimated based on the coverage choice data. To incorporate the latter component, we can allow ξ_τs in the real-data reduced-form regression to differ under different coverage plans; then, we match these ξ_τs with their counterparts simulated from the structural models. [27]

[27] Changes in beliefs caused by an accident can make policyholders change their insurance coverage for the next contract, which will in turn alter their accident probabilities in the future. We control for policyholders' insurance coverage in the real-data reduced-form regression. The structural parameters are estimated by matching the simulated ξ_τs from the structural model without the component of coverage decisions to the ξ_τs from the real-data reduced-form regression in which the effect of coverage on accident probabilities has

6.2. Identification

Because we cannot observe drivers' effort levels, the model cannot be nonparametrically identified.[28] However, after we assume the functional form of the Bernoulli utility function and the functional form of how effort or caution helps reduce accident probabilities (equation 4.1), the model can be parametrically identified.

Based on equation (4.4), the coefficient of risk aversion γ, the parameter for the effectiveness of caution η, and the initial prior belief p_H jointly determine the accident probability and the responsiveness to new information. Therefore, those three parameters can be identified by the average accident probability and the levels of two of ξ_τs. The remembering rate δ can be identified by the concavity of the line of ξ_τs, i.e., the pattern of how ξ_τs decay in magnitude over time.

7. Concerns

7.1. Learning or changes in financial incentives?

Dionne et al. (2013) studied automobile insurance data in France during 1995-1997 and also found negative occurrence dependence.[29] However, they attributed the negative dependence to changes in financial incentives caused by accidents instead of learning because of the multiplicativity of the French "bonus-malus" schemes.

Under the multiplicative French "bonus-malus" schemes, the premium is the product of two factors: the "base premium" and the "bonus-malus coefficient" at the time of renewal. The "base premium" is determined based on factors such as vehicle characteristics and policyholder demographic characteristics. The "bonus-malus coefficient" is multiplicatively determined based

already been controlled for. We also conduct estimations using policyholders who did not change their coverages and obtain similar results.

[28] Note that Wang and Yang (2024) have derived the formal identification conditions for structural learning models in which agents' discrete choices are observable. In contrast, in our case, drivers' choices on the caution level are continuous and unobservable. We only observe the outcomes (accidents) partially governed by the choices.

[29] Abbring et al. (2003b) also studied automobile insurance data in France during 1987-1989 and did not find negative occurrence dependence.

on a policyholder's accident history: each year without an accident decreases the "bonus-malus coefficient" by a fixed factor (0.95); each accident increases the coefficient by a fixed factor (1.25). For example, the coefficient of a policyholder with one, two, or three accidents is 1.25, 1.25^2, and 1.25^3, respectively; correspondingly, the increment of the coefficient is 0.25, 0.3125, and 0.390625, respectively ($1.25 - 1$, $1.25^2 - 1.25$, and $1.25^3 - 1.25^2$). Therefore, not only the level of premium but also the increment of premium caused by an accident (the marginal cost of an accident) for a policyholder is increasing in the number of accidents that have occurred to the policyholder; an accident will cause **an increase** in a policyholder's financial incentive to avoid future accidents, which can generate the negative occurrence dependence of accidents.

The Chinese "bonus-malus" scheme differs from the French scheme mainly in that the "bonus-malus coefficient" is additive instead of multiplicative. Simply speaking, the increment in the "bonus-malus coefficient" caused by each accident within a contract year is constant (0.25) instead of increasing exponentially as in the French scheme. Therefore, in general, although an accident will increase the premium of the next contract, it will not increase the marginal cost of a policyholder's next accident and hence an accident would not cause **an increase** in a policyholder's financial incentive to avoid future accidents.

After an accident occurs, the increased premium of the next contract is not the marginal cost of the next accident because the policyholder needs to pay at least the increased premium for the next contract regardless of whether the policyholder encounters a next accident. The marginal cost of the next accident is related to the increment in the next premium caused by the next accident rather than the level of the next premium itself.

To illustrate, suppose the following case: a policyholder's base premium in period t is $BasePremium_t$ and her/his "bonus-malus coefficient" in period t is 1;[30] if there is no accident in period t, her/his premium in period $t+1$ will remain $BasePremium_t$; the first accident in period t will increase her/his premium in period $t+1$ to $1.25BasePremium_t$; the second accident in period t will increase her/his premium in period $t+1$ to $1.5BasePremium_t$. The two accidents have the same marginal costs because they both increase the premium in the next period by

[30] In China, the "base premium" is determined based on the model and age of the vehicle, the owner type (individual-owned vs. company-/government-owned vehicle), and the gender, age, years of driving experience, residing city, and traffic violation history of the driver.

$0.25BasePremium_t$. Although a policyholder with two accidents in period t will pay a higher premium in period $t+1$ than will a policyholder with only one accident in period t, the first accident in period t has already made the former policyholder need to pay at least $1.25BasePremium_t$ in period $t+1$, which is not caused by the second accident in period t. The financial incentive to avoid the next accident for a policyholder with one accident in period t is the same as that for a policyholder with no accident in period t.[31]

Under the Chinese "bonus-malus" scheme, after an accident and the resulting premium increase, a policyholder's premium will gradually decrease over years if the policyholder can continuously maintain excellent driving records with no accident. Correspondingly, one may concern whether the decreasing in the magnitude of the negative occurrence dependence of accidents over time is caused by the remedy in the "bonus-malus" scheme. In fact, what affects drivers' caution decisions should be the change in the future premium that can be caused by a potential accident in the current contract year, instead of the level of the current premium, which is a sunk cost because drivers pay that level of premium for the current contract year regardless of whether an accident occurs.

Consider the following two situations (as shown in Tables A.9 and A.10). In situation 1, given that one accident occurs in year t, if no accident occurs in year $t+1$, the policyholder will go to the category of "no accident for one year" and the corresponding "bonus-malus coefficient" for year $t+2$ is 0.85; alternatively, if one accident occurs in year $t+1$, the policyholder will go to the category of "one accident last year" and the corresponding "bonus-malus coefficient" for year $t+2$ is 1. Therefore, the increment in the "bonus-malus coefficient" caused by the follow-up accident is 0.15 $(=1-0.85)$ if the follow-up accident occurs in year $t+1$.

[31] The "bonus-malus" scheme in China is set by the China Banking and Insurance Regulatory Commission. The complete scheme is shown in the first two columns of panel A of Table A.8 in Appendix A. In general, the marginal cost of each accident is not increasing in the ordinal of the accident in a policyholder's lifetime. For example, for a policyholder in the category of "no accident last year," the increments in the "bonus-malus coefficient" caused by each accident in the current contract year are 0.3, 0.25, 0.25, 0.25, 0.25, 0, 0, 0…Regarding accidents across different years, as shown in Table A.9, if a policyholder encounters one accident in years t, t+1, and t+2, respectively, the increment in the "bonus-malus coefficient" caused by the second accident is 0.15, and the increment caused by the third accident is also 0.15. Therefore, in general, the increment in the "bonus-malus coefficient" caused by each accident in different contract years is not increasing in the ordinal of the accident in the policyholder's lifetime.

In situation 2, given that one accident occurs in year t, if no accident occurs in year $t+1$ and year $t+2$, the policyholder will go to the category of "no accident for two years" and the corresponding "bonus-malus coefficient" for year $t+3$ is 0.7; alternatively, if no accident occurs in year $t+1$ and one accident occurs in year $t+2$, the policyholder will go to the category of "one accident last year" and the corresponding "bonus-malus coefficient" for year $t+3$ is 1. Therefore, the increment in the "bonus-malus coefficient" caused by the follow-up accident is 0.3 $(=1-0.7)$ if the follow-up accident occurs in year $t+2$.

Therefore, the increment in the "bonus-malus coefficient" caused by an accident could be increasing in how many years it occurs after the previous accident. In other words, the follow-up accident will make a larger difference (have a larger marginal cost) if it occurs more years after the previous accident. This change in financial incentives over time should tend to drive the magnitudes of ξ_τ s to increase as τ increases. However, empirically, we observe that the magnitudes of ξ_τs decrease as τ increases, which is consistent with forgetting.

Another piece of evidence suggesting that our empirical results are not driven by changes in financial incentives comes from drivers of company-/government-owned vehicles. While drivers of company-/government-owned vehicles care about other potential losses caused by accidents (e.g., bodily injuries), they should not care much about potential premium increases caused by accidents because the premiums are paid by their employers. However, we also observe a negative dependence of accidents for those drivers (see Section 5.3).

7.2. Bayesian learning about accident probabilities or learning driving skills?

One concern is that the negative dependence of accidents may be driven by driving-skill accumulation over time instead of learning about accident probabilities from an accident. First, in the quarter-level regression reported in Table 3, we have already controlled for quarters of driving experience.

Second, the speed of driving skill accumulation is fast for first-year drivers and is not significant for experienced drivers because they have already acquired sufficient skills.[32] Therefore, if the negative dependence of accidents is substantially driven by learning driving skills, we should observe larger ξ_τs in magnitude for first-year drivers than for experienced drivers. We run the month-level regressions of the dynamic panel model in equation (3.3) separately for first-year drivers and experienced drivers. As shown in Figure 8, the ξ_τs do not differ substantially across the two groups: the two lines cross multiple times; and the difference between the two groups is much smaller than the difference across genders (see Figure 5) and owner types (see Figure 6).[33]

[Insert Figure 8 here]

Although experienced drivers do not continue to prominently learn driving skills, they can still learn their accident probabilities through accidents because accidents are infrequent events in one's daily life.

Third, drivers should not forget their driving skills because they mostly drive every day.[34] In contrast, drivers can forget their accident probabilities over time because accidents are rare events in one's daily life and drivers may not pay attention to estimating their accident probabilities every day, which is consistent with the empirical pattern that the magnitudes of ξ_τs decrease as τ increases.

7.3. Underreporting minor accidents

[32] We run the month-level regression of the accident indicator on policyholders' months of driving experience and control variables, using policyholders who started as new drivers in the sample. The results are displayed in Figure A.1 in Appendix A. The fixed effects of months of driving experience are decreasing in the first ten months and fluctuate thereafter. This result is consistent with the pricing scheme based on years of driving experience. The years-of-driving-based coefficient for the premium is 1.05, 1.02, and 1 for drivers with less than one year, more than or equal to one year but less than three years, and more than or equal to three years of driving experience, respectively.

[33] Unlike in developed countries, where first-year drivers are usually 18 years old, first-year drivers in China can be any age. The reason is that the rate of private car ownership in China was close to zero before 2000 and then started to grow rapidly but is currently still lower than the rates in developed countries. Many individual drivers gain access to a car at an age much older than 18.

[34] If drivers forgot their driving skills, we should have observed an upward trend after a certain point in Figure A.1 in Appendix A.

One concern is that the learning effect in our empirical results could be driven by the possibility that a policyholder with a previous accident is more likely to underreport minor accidents in the future. However, the reality is the opposite: if underreporting exists, as a policyholder experiences more accidents, she/he will become less likely to underreport minor accidents. In Table A.11 in Appendix A, the first two columns show the distributions of payments by the insurance company for the first and second accidents, respectively, to policyholders with two accidents in the data; the last three columns show the distributions of payments by the insurance company for the first, second, and third accidents, respectively, to policyholders with three accidents in the data. The proportion of reported accidents with minor damages is increasing in the ordinal of the accident experienced by a policyholder. This pattern is consistent with the findings in Gong (2017). Therefore, the existence of underreporting can only cause the learning effect to be underestimated rather than overestimated.

The empirical pattern that a policyholder who experiences more accidents will become less likely to underreport minor accidents is consistent with the fact that the marginal cost of each accident (the financial incentive to avoid each accident) can be decreasing in the ordinal of the accident. For example, as shown in Table A.8, for a policyholder in the category of "no accident for two years," if no accident occurs in the current year, the policyholder will go to the category of "no accident for three years" and the corresponding "bonus-malus coefficient" for the next year is 0.6; alternatively, if one accident occurs in the current year, the policyholder will go to the category of "one accident last year" and the corresponding "bonus-malus coefficient" for the next year is 1; if two accidents occur in the current year, the policyholder will go to the category of "two accidents last year" and the corresponding "bonus-malus coefficient" for the next year is 1.25. Therefore, the marginal cost of the first accident in the current year (1-0.6=0.4) is larger than that of the second accident (1.25-1=0.25).

Another concern is that the forgetting effect in the empirical results could be driven by the fact that a policyholder may become less likely to underreport minor accidents as she/he experiences more accidents. We rerun the 2SLS regression in Table 3, counting only accidents with settlement proceeds above a certain threshold. The rationale is that accidents with substantial damages will not be underreported regardless of how many accidents a policyholder has already

experienced. The results (see Table A.12 in Appendix A) are similar to the last column of Table 3: ξ_τ is significantly negative and is decreasing in magnitude over time.

7.4. More limited funds after an accident

One concern is that policyholders may have more limited funds after an accident and hence become more cautious while driving, which may drive the empirical results. Therefore, we only examine the effect of past accidents with small damages (settlement proceeds below a certain threshold) on the accident probabilities in the current period. Accidents with small damages should not incur large out-of-pocket losses for policyholders and hence should not make them have more limited funds. The regression results are reported in Table A.13 in Appendix A and exhibit similar results to the last column in Table 3.

Another piece of evidence suggesting that our results on learning are not fully driven by more limited funds after an accident can be found in Figure 4. In principle, policyholders with three or more accidents last year, on average, should have more limited funds than policyholders with two accidents last year and hence should be more responsive to a new accident in their caution levels. However, as shown in Figure 4, the negative effect of an accident on the future accident probabilities is much stronger for policyholders with two accidents last year than for policyholders with three or more accidents last year, which is consistent with learning (see discussions in Section 5.1) instead of more limited funds.[35]

[35] A related possible concern is whether the empirical results are driven by the wealth effect of increased premiums. On the one hand, if the wealth effect is due to changes in the current wealth or net worth in a driver's personal account, an accident will not reduce the driver's current wealth because it causes increases only in premiums of the next annual policy cycle. However, we observe that the learning effect is immediate after the accident instead of starting in the next policy cycle. On the other hand, if the wealth effect is due to changes in the permanent wealth, first, experiencing three accidents in the last year and experiencing only one accident in the last year will lead to different increases in premiums of the next year and hence different reductions in a driver's permanent wealth. However, we do not observe that the learning effects in these two cases are significantly different from each other (see Figure 4). Second, the reduction in permanent wealth is permanent and does not decrease over time. However, we observe that the learning effect is decreasing over time (the forgetting effect). In fact, generally, an increase in premiums caused by an accident is less than 1% of a driver's annual income; the increase is even smaller compared to a driver's permanent wealth. Therefore, the wealth effect caused by premium changes should not essentially drive the empirical results.

7.5. Learning or changes in risk aversion?

One concern is whether the negative dependence of accidents found in this study is driven by increases in policyholders' risk aversion caused by the psychological impact of an accident. To the best of our knowledge, we are unaware of any previous empirical study with administrative data that separately identifies learning and increases in risk aversion.[36] Our study finds some evidence suggesting that the effect of learning exists.

First, the nonmonotonic empirical pattern in Figure 4 from the group with the best driving records at the beginning of the sample period through the group with the worst driving records supports the learning story (theoretically proven in Section 5.1): when a policyholder's belief of belonging to the high-risk type is very low, a new accident is more likely to be perceived as an occasional occurrence of a rare event instead of conveying rich information of belonging to the high-risk type; if two accidents occur within one year, the belief of high-risk type will reach a certain level such that a new accident will reinforce the belief by a substantial amount and hence have a more prominent impact on the policyholder's caution level; if three or more accidents occur within one year, the policyholder will believe almost for sure that she/he is in the high-risk type, and hence the next accident will provide almost no additional information about the policyholder's true risk type.

In contrast, if we use the risk aversion story to explain the nonmonotonic pattern, we need to require that the increment in risk aversion caused by an accident following two consecutive accidents in the previous year should be dramatically larger than the increment caused by an accident following one or three consecutive accidents in the previous year, which is an odd requirement. There are no findings in the economics or psychology literature that support this odd nonmonotonic evolution in risk aversion (see Cameron and Shah, 2015 on the impact of the number of natural disasters; Brown et al., 2019 on the impact of the number of local homicides per 10,000 inhabitants).

[36] Some empirical studies employed survey questions separately relative to risk aversion and information updating. They found that the occurrence of adverse events both changes risk aversion and increases people's subjective probabilities of adverse events (see Guiso et al., 2018 on the impact of financial crisis; Cameron and Shah, 2015 on the impact of natural disasters).

Second, Shum and Xin (2020) found that the increases in risk aversion caused by "near-miss" accidents vanish in approximately 2-3 weeks. In contrast, we find that the negative effect of an actual accident on future actual accident probabilities lasts for two years, although with a decaying magnitude. This difference suggests that our results may be mainly driven by updates on belief, which could be a long-lasting effect, although subject to forgetting.

One may argue that changes in risk aversion caused by actual accidents could last longer than those caused by "near-miss" accidents because actual accidents cause more dramatic psychological impacts on drivers. However, we do not find that the effects of accidents with different levels of impacts last different lengths of time. One example is that based on Figure 4, an accident following two consecutive accidents in the previous year should generate a larger impact than an accident following one or three consecutive accidents in the previous year, but their effects decay to similar levels in two years, which is more likely to be generated by depreciation of information (forgetting) than recovering from psychological impacts.

Another example is that accidents of larger sizes (measured by settlement proceeds) generate more substantial impacts and hence should have longer lasting effects. However, as shown in Figure A.2 in Appendix A, the effects of accidents with settlement proceeds below and above the median decay to similar levels in two years.

7.6. Mean reversion?

One question is whether our findings are simply phenomena of mean reversion. The answer is no. If we pick all the drivers with accidents in the first year and find that their accident rate in the second year is not much higher than that of other drivers instead of being close to 100%, this is a phenomenon of mean reversion. In contrast, however, we include both drivers with accidents and drivers without accidents in the regressions; and our finding is that the expected accident rate in the second year of drivers with accidents in the first year is lower than that of drivers without accidents in the first year, which is different from mean reversion.

7.7. Point accumulation

One possible concern is whether the empirical results could be driven by point accumulation. First, point accumulation and driver license revocation are caused by traffic violations rather than accidents. An accident will lead to point accumulation only if the accident causes bodily injuries or deaths and the driver escapes. An accident will lead to driver license revocation only if the accident committed a crime.

Second, points do not decay continuously over time. Instead, they are reset to zero every 12-month cycle. Therefore, even if point accumulation were caused by accidents, accidents more than 12 months ago should not have any effect on the current period accident probability through this channel. However, as shown in Table 3, ξ_τs beyond four quarters ago are still significantly negative with a considerable magnitude.

Third, within a 12-month cycle, a driver can choose to take a course to reduce the points. If a driver's points reach 12 points, the license will be withheld (not revoked) and the driver will be forced to take the course immediately. After finishing the course, the driver can get her/his license back.

8. Managerial implications

The results in this study provide multiple managerial implications for automobile insurance companies. First, policyholders' incomplete information about their own innate risk types can mitigate adverse selection. In the theoretical framework developed by Einav et al. (2010) and Einav and Finkelstein (2011), adverse selection emerges if individuals who have the highest willingness to pay for insurance are those who have the highest risk, whereas advantageous selection may emerge if the individuals who are willing to pay the most for insurance are those who are the most risk-averse. Given the fact that Individuals know their risk preferences but are not fully aware of their innate risk types, the real-world insurance markets may incline to exhibit advantageous selection.

Second, policyholders' learning about their own types helps reduce their ex-post risk and hence mitigates moral hazard. On the one hand, an accident partially reveals a policyholder's innate risk type to insurers; thus, insurers should increase the policyholder's premium. On the other hand, an accident also partially reveals the policyholder's innate risk type to herself/himself and

makes her/him exert more caution or effort to reduce future risk; in this sense, insurers should increase the policyholder's premium by only a limited amount. The current claim-based pricing schemes only consider information revealed by an accident to insurers but fail to account information revealed to policyholders and their responses.

Third, the Chinese "bonus-malus" scheme counts only the accidents in the previous contract year when determining the "bonus-malus" coefficient for the current premium, which allows a policyholder to substantially remedy her/his poor record through a year with no accident. For example, if a policyholder encounters one accident in the current contract year, the policyholder's premium will reach the highest level in the next contract year and then decrease thereafter. However, because of the forgetting effect, the policyholder's risk should reach its lowest level in the next contract year and then increase thereafter. Therefore, insurers' pricing schedules are inconsistent with policyholders' risk dynamics. Insurance companies should take into account both policyholders' learning and forgetting when they build their claim-based pricing system.

Fourth, when a policyholder purchases an insurance policy, in addition to offering a characteristics-based premium, insurance companies should provide the policyholder with an estimated accident probability and expected losses conditional on the policyholder's characteristics. In this way, policyholders learn their innate risk types not only through their own driving experience but also through the large database maintained by insurance companies, which may induce more driving caution and fewer accidents.

One concern is whether providing this information to policyholders will make them more informed and hence exacerbate adverse selection. As this intervention will make policyholders know more information that insurers have already known rather than make policyholders know more information that insurers do not know, the intervention will not exacerbate adverse selection.

Fifth, to mitigate the forgetting effect, as reminders, insurance companies could periodically send a policyholder emails or text messages with the estimated accident probability, the number of accidents experienced, and the increment in the estimated accident probability caused by each accident the policyholder has experienced. [37]

[37] Field experiments have found that text message reminders are effective in multiple scenarios, including reducing drivers' traffic violations (Chen et al., 2017; Lu et al., 2016), improving borrowers' repayments

Future research could conduct field experiments to investigate the effectiveness of providing upfront information on estimated accident probabilities and sending periodic reminders in reducing accident occurrence.

9. Conclusion

We study a unique dataset from an automobile insurance company in China and find that policyholders' past accident experience can reduce their accident probabilities in the current period. Our explanation is that accidents make policyholders update their beliefs about their innate risk types in an upward direction and hence exert more caution or effort to reduce risk. Additionally, we find that the magnitude of the negative effect of a past accident on the accident probability in the current period is smaller if the accident occurred longer time ago; this empirical pattern suggests a forgetting effect in the learning process that results in an approximately 7-quarter effective memory duration for an accident. We conduct a structural estimation and the estimated monthly forgetting rate (monthly depreciation rate of past information, $1-\delta$) is 0.1385.

We also rule out alternative explanations, including changes in financial incentives to prevent future accidents, underreporting of minor accidents, more limited funds after an accident, etc. In addition, we find that the negative occurrence dependence is stronger for female drivers than male drivers because female drivers are likely to be more risk-averse and hence are more responsive to the update of their beliefs caused by an accident.

Moreover, we find a nonmonotonic pattern: when a policyholder's belief of belonging to the high-risk type is very low, a new accident is more likely to be perceived as an occasional occurrence of a rare event instead of conveying rich information about belonging to the high-risk type; if two accidents occur within one year, the belief of high-risk type will reach a certain level such that a new accident will reinforce the belief by a substantial amount and hence have a more prominent impact on the policyholder's caution level; if three or more accidents occur within one year, the policyholder will believe almost for sure that she/he is in the high-risk type, and hence

for online P2P lending (Du et al., 2020; Huang and Bao, 2020), microloans (Cadena and Schoar, 2011), and credit cards (Bursztyn et al., 2015), enhancing compliance with tax (Hallsworth et al., 2014) and TV license fees (Fellner et al., 2013), boosting COVID-19 vaccinations (Dai et al., 2021), and increasing commitment attainment for banks' clients with commitment savings accounts (Karlan et al., 2016).

the next accident will provide almost no additional information about the policyholder's true risk type.

This study generates three managerial implications for automobile insurance companies. First, policyholders' incomplete information about their own innate risk types can mitigate adverse selection. Given the fact that Individuals know their risk preferences but are not fully aware of their innate risk types, the real-world insurance markets may incline to exhibit advantageous selection.

Second, policyholders' learning about their own types helps reduce their ex-post risk and hence mitigates moral hazard. While an accident partially reveals a policyholder's high-risk type to the insurer, the accident also partially reveals the policyholder's risk type to herself/himself and makes her/him exert more caution or effort to reduce future risk. Therefore, the insurer should increase the policyholder's premium by only a limited amount. The current "bonus-malus" pricing schemes only consider information revealed by an accident to insurers but fail to account information revealed to policyholders and their responses.

Third, in the Chinese "bonus-malus" scheme, if a policyholder encounters one accident in the current contract year, the policyholder's premium will reach the highest level in the next contract year and then decrease thereafter. However, because of learning with forgetting, the policyholder's risk should reach its lowest level in the next contract year and then increase thereafter. Therefore, insurers' pricing schedules are inconsistent with policyholders' risk dynamics. Insurance companies should take into account both policyholders' learning and forgetting when they build their claim-based pricing system.

References

Abbring, Jaap H.; Pierre-André Chiappori; James J. Heckman and Jean Pinquet. 2003a. "Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?" *Journal of the European Economic Association*, 1(2/3), 512-21.

Abbring, Jaap H.; Pierre-André Chiappori and Jean Pinquet. 2003b. "Moral Hazard and Dynamic Insurance Data." *Journal of the European Economic Association*, 1(4), 767-820.

Acemoglu, Daron; Munther A. Dahleh; Ilan Lobel and Asuman Ozdaglar. 2011. "Bayesian Learning in Social Networks." *The Review of Economic Studies*, 78(4), 1201-36.

Ackerberg, Daniel A. and Gautam Gowrisankaran. 2006. "Quantifying Equilibrium Network Externalities in the Ach Banking Industry." *The RAND Journal of Economics*, 37(3), 738-61.

Agarwal, Sumit; John C Driscoll; Xavier Gabaix and David Laibson. 2008. "Learning in the Credit Card Market." NBER Working Paper.

Anagol, Santosh; Vimal Balasubramaniam and Tarun Ramadorai. 2021. "Learning from Noise: Evidence from India's Ipo Lotteries." *Journal of Financial Economics*, 140(3), 965-86.

Anderson, T. W. and Cheng Hsiao. 1981. "Estimation of Dynamic Models with Error Components." *Journal of the American Statistical Association*, 76(375), 598-606.

Arellano, Manuel and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies*, 58(2), 277-97.

Babus, Ana and Péter Kondor. 2018. "Trading and Information Diffusion in over-the-Counter Markets." *Econometrica*, 86(5), 1727-69.

Barseghyan, Levon; Francesca Molinari; Ted O'Donoghue and Joshua C. Teitelbaum. 2013. "The Nature of Risk Preferences: Evidence from Insurance Choices." *American Economic Review*, 103(6), 2499-529.

Bordalo, Pedro; Nicola Gennaioli and Andrei Shleifer. 2020. "Memory, Attention, and Choice*." *The Quarterly Journal of Economics*, 135(3), 1399-442.

Brown, Ryan; Verónica Montalva; Duncan Thomas and Andrea Velásquez. 2019. "Impact of Violent Crime on Risk Aversion: Evidence from the Mexican Drug War." *The Review of Economics and Statistics*, 101(5), 892-904.

Bursztyn, Leonardo; Stefano Fiorin; Daniel Gottlieb and Martin Kanz. 2015. *Moral Incentives: Experimental Evidence from Repayments of an Islamic Credit Card*. The World Bank.

Cadena, Ximena and Antoinette Schoar. 2011. "Remembering to Pay? Reminders Vs. Financial Incentives for Loan Payments." NBER Working Paper.

Camerer, Colin and Teck Hua Ho. 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica*, 67(4), 827-74.

Cameron, Lisa and Manisha Shah. 2015. "Risk-Taking Behavior in the Wake of Natural Disasters." *The Journal of Human Resources*, 50(2), 484-515.

Castillo, Marco; Greg Leo and Ragan Petrie. 2020. "Room Composition Effects on Risk Taking by Gender." *Experimental Economics*, 23(3), 895-911.

Chan, Tat; Chakravarthi Narasimhan and Ying Xie. 2013. "Treatment Effectiveness and Side Effects: A Model of Physician Learning." *Management Science*, 59(6), 1309-25.

Chandrasekhar, Arun G.; Horacio Larreguy and Juan Pablo Xandri. 2020. "Testing Models of Social Learning on Networks: Evidence from Two Experiments." *Econometrica*, 88(1), 1-32.
Chang, Tom Y.; Wei Huang and Yongxiang Wang. 2018. "Something in the Air: Pollution and the Demand for Health Insurance." *The Review of Economic Studies*, 85(3), 1609-34.
Charness, Gary and Uri Gneezy. 2012. "Strong Evidence for Gender Differences in Risk Taking." *Journal of Economic Behavior & Organization*, 83(1), 50-58.
Charness, Gary and Dan Levin. 2005. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect." *The American Economic Review*, 95(4), 1300-09.
Chen, Yan; Fangwen Lu and Jinan Zhang. 2017. "Social Comparisons, Status and Driving Behavior." *Journal of Public Economics*, 155, 11-20.
Chiappori, Pierre-Andre and Bernard Salanie. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1), 56-78.
Clerides, Sofronis K.; Saul Lach and James R. Tybout. 1998. "Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico, and Morocco." *The Quarterly Journal of Economics*, 113(3), 903-47.
Cohen, Alma. 2005. "Asymmetric Information and Learning: Evidence from the Automobile Insurance Market." *The Review of Economics and Statistics*, 87(2), 197-207.
Cohen, Alma and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review*, 97(3), 745-88.
Cohen, Alma and Rajeev Dehejia. 2004. "The Effect of Automobile Insurance and Accident Liability Laws on Traffic Fatalities." *The Journal of Law & Economics*, 47(2), 357-93.
Coles, Jeffrey L.; Michael L. Lemmon and J. Felix Meschke. 2012. "Structural Models and Endogeneity in Corporate Finance: The Link between Managerial Ownership and Corporate Performance." *Journal of Financial Economics*, 103(1), 149-68.
Collard-Wexler, Allan. 2013. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica*, 81(3), 1003-37.
Cooper, Russell W. and John C. Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *The Review of Economic Studies*, 73(3), 611-33.
Cox, James C.; Jason Shachat and Mark Walker. 2001. "An Experiment to Evaluate Bayesian Learning of Nash Equilibrium Play." *Games and Economic Behavior*, 34(1), 11-33.
Dai, Hengchen; Silvia Saccardo; Maria A. Han; Lily Roh; Naveen Raja; Sitaram Vangala; Hardikkumar Modi; Shital Pandya; Michael Sloyan and Daniel M. Croymans. 2021. "Behavioural Nudges Increase Covid-19 Vaccinations." *Nature*, 597(7876), 404-09.
Dalton, Christina M; Gautam Gowrisankaran and Robert J Town. 2019. "Salience, Myopia, and Complex Dynamic Incentives: Evidence from Medicare Part D." *The Review of Economic Studies*, 87(2), 822-69.
Dekel, Eddie; Drew Fudenberg and David K. Levine. 2004. "Learning to Play Bayesian Games." *Games and Economic Behavior*, 46(2), 282-303.
Dionne, Georges and Neil A. Doherty. 1994. "Adverse Selection, Commitment, and Renegotiation: Extension to and Evidence from Insurance Markets." *Journal of Political Economy*, 102(2), 209-35.
Dionne, Georges; Christian Gouriéroux and Charles Vanasse. 2001. "Testing for Evidence of Adverse Selection in the Automobile Insurance Market: A Comment." *Journal of Political Economy*, 109(2), 444-53.

Dionne, Georges; Pierre-Carl Michaud and Maki Dahchour. 2013. "Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France." *Journal of the European Economic Association*, 11(4), 897-917.
Dohmen, Thomas; Armin Falk; David Huffman; Uwe Sunde; Jürgen Schupp and Gert G. Wagner. 2011. "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences." *Journal of the European Economic Association*, 9(3), 522-50.
Du, Ninghua; Lingfang Li; Tian Lu and Xianghua Lu. 2019. "Prosocial Compliance in P2p Lending: A Natural Field Experiment." *Management Science*, 66(1), 315-33.
Duong, Hai Long; Junhong Chu and Dai Yao. 2023. "Taxi Drivers' Response to Cancellations and No-Shows: New Evidence for Reference-Dependent Preferences." *Management Science*, 69(1), 179-99.
Eckel, Catherine C. and Philip J. Grossman. 2008. "Chapter 113 Men, Women and Risk Aversion: Experimental Evidence," C. R. Plott and V. L. Smith, *Handbook of Experimental Economics Results.* Elsevier, 1061-73.
Einav, Liran and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives*, 25(1), 115-38.
Einav, Liran; Amy Finkelstein and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *The Quarterly Journal of Economics*, 125(3), 877-921.
Fellner, Gerlinde; Rupert Sausgruber and Christian Traxler. 2013. "Testing Enforcement Strategies in the Field: Threat, Moral Appeal and Social Information." *Journal of the European Economic Association*, 11(3), 634-60.
Gallagher, Justin. 2014. "Learning About an Infrequent Event: Evidence from Flood Insurance Take-up in the United States." *American Economic Journal: Applied Economics*, 6(3), 206-33.
Gao, Feng; Jaimie Lien and Jie Zheng. 2019. "Exaggerating to Break-Even: Reference-Dependent Moral Hazard in Auto Insurance Claims." *Woring Paper*.
Gneezy, Uri; Muriel Niederle and Aldo Rustichini. 2003. "Performance in Competitive Environments: Gender Differences." *The Quarterly Journal of Economics*, 118(3), 1049-74.
Gneezy, Uri and Aldo Rustichini. 2004. "Gender and Competition at a Young Age." *American Economic Review*, 94(2), 377-81.
Goettler, Ronald L. and Brett R. Gordon. 2011. "Does Amd Spur Intel to Innovate More?" *Journal of Political Economy*, 119(6), 1141-200.
Golub, Benjamin and Matthew O. Jackson. 2012. "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *The Quarterly Journal of Economics*, 127(3), 1287-338.
Gong, Qing. 2017. "Ex Post Moral Hazard in Automobile Insurance Markets with Experience Rating." Working Paper.
Gourieroux, Christian; and Alain Monfort. 1996. *Simulation Based Econometric Methods* (Oxford University Press, Oxford, U.K.).
Gourieroux, C.; A. Monfort and E. Renault. 1993. "Indirect Inference." *Journal of Applied Econometrics*, 8(S1), S85-S118.
Greenwood, Robin and Samuel G. Hanson. 2014. "Waves in Ship Prices and Investment *." *The Quarterly Journal of Economics*, 130(1), 55-109.
Grimm, Veronika and Friederike Mengel. 2018. "Experiments on Belief Formation in Networks." *Journal of the European Economic Association*, 18(1), 49-82.
Grubb, Michael D. and Matthew Osborne. 2015. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *American Economic Review*, 105(1), 234–71.

Guiso, Luigi; Paola Sapienza and Luigi Zingales. 2018. "Time Varying Risk Aversion." *Journal of Financial Economics*, 128(3), 403-21.
Hallsworth, Michael; John A. List; Robert D. Metcalfe and Ivo Vlaev. 2017. "The Behavioralist as Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance." *Journal of Public Economics*, 148, 14-31.
Haselhuhn, Michael P.; Devin G. Pope; Maurice E. Schweitzer and Peter Fishman. 2012. "The Impact of Personal Experience on Behavior: Evidence from Video-Rental Fines." *Management Science*, 58(1), 52-61.
Heckman, James J. 1991. "Identifying the Hand of Past: Distinguishing State Dependence from Heterogeneity." *The American Economic Review*, 81(2), 75-79.
Hendel, Igal and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica*, 74(6), 1637-73.
HENNESSY, CHRISTOPHER A. and TONI M. WHITED. 2005. "Debt Dynamics." *The Journal of Finance*, 60(3), 1129-65.
Ho, Teck-Hua; Juin Kuan Chong and Xiaoyu Xia. 2017. "Yellow Taxis Have Fewer Accidents Than Blue Taxis Because Yellow Is More Visible Than Blue." *Proceedings of the National Academy of Sciences*, 114, 3074-78.
Hoyt, Robert E.; David B. Mustard and Lawrence S. Powell. 2006. "The Effectiveness of State Legislation in Mitigating Moral Hazard: Evidence from Automobile Insurance." *The Journal of Law and Economics*, 49(2), 427-50.
Huang, Difang and Zhengyang Bao. 2020. "Gender Differences in Reaction to Enforcement Mechanisms: A Large-Scale Natural Field Experiment." Working Paper.
Jakiela, Pamela and Owen Ozier. 2019. "The Impact of Violence on Individual Risk Preferences: Evidence from a Natural Experiment." *The Review of Economics and Statistics*, 101(3), 547-59.
Jeziorski, Przemyslaw; Elena Krasnokutskaya and Olivia Ceccarini. 2017. "Adverse Selection and Moral Hazard in the Dynamic Model of Auto Insurance." Working Paper.
Jin, Lawrence; Rui Tang; Han Ye; Junjian Yi and Songfa Zhong. 2023. "Path Dependency in Physician Decisions." Forthcoming at *The Review of Economic Studies*.
Karlan, Dean; Margaret McConnell; Sendhil Mullainathan and Jonathan Zinman. 2016. "Getting to the Top of Mind: How Reminders Increase Saving." *Management Science*, 62(12), 3393-411.
Ketcham, Jonathan D.; Claudio Lucarelli; Eugenio J. Miravete and M. Christopher Roebuck. 2012. "Sinking, Swimming, or Learning to Swim in Medicare Part D." *American Economic Review*, 102(6), 2639-73.
Lafontaine, Francine and Kathryn Shaw. 2016. "Serial Entrepreneurship: Learning by Doing?" *Journal of Labor Economics*, 34(S2), S217-S54.
Lentz, Rasmus and Dale T. Mortensen. 2008. "An Empirical Model of Growth through Product Innovation." *Econometrica*, 76(6), 1317-73.
Levitt, Steven D.; John A. List and Chad Syverson. 2013. "Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant." *Journal of Political Economy*, 121(4), 643-81.
Low, Hamish and Luigi Pistaferri. 2015. "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off." *American Economic Review*, 105(10), 2986-3029.
Lu, Fangwen; Jinan Zhang and Jeffrey M. Perloff. 2016. "General and Specific Information in Deterring Traffic Violations: Evidence from a Randomized Experiment." *Journal of Economic Behavior & Organization*, 123, 97-107.

Ma, Chao. 2021. "Be Cautious in the Last Month: The Sunk Cost Fallacy Held by Car Insurance Policyholders." *International Economic Review*, 62(3), 1199-236.
Malmendier, Ulrike and Stefan Nagel. 2011a. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *The Quarterly Journal of Economics*, 126(1), 373-416.
Martin, Gregory J. and Ali Yurukoglu. 2017. "Bias in Cable News: Persuasion and Polarization." *American Economic Review*, 107(9), 2565-99.
Mehta, Nitin; Surendra Rajiv and Kannan Srinivasan. 2004. "Role of Forgetting in Memory-Based Choice Decisions: A Structural Model." *Quantitative Marketing and Economics*, 2(2), 107-40.
Miravete, Eugenio J. 2003. "Choosing the Wrong Calling Plan? Ignorance and Learning." *American Economic Review*, 93(1), 297-310.
Nagypál, Éva. 2007. "Learning by Doing Vs. Learning About Match Quality: Can We Tell Them Apart?" *The Review of Economic Studies*, 74(2), 537-66.
Narayanan, Sridhar; Pradeep K. Chintagunta and Eugenio J. Miravete. 2007. "The Role of Self Selection, Usage Uncertainty and Learning in the Demand for Local Telephone Service." *Quantitative Marketing and Economics*, 5(1), 1-34.
Phillips, Peter C. B.; Yangru Wu and Jun Yu. 2011. "Explosive Behavior in the 1990s Nasdaq: When Did Exuberance Escalate Asset Values?" *International Economic Review*, 52(1), 201-26.
Puelz, Robert and Arthur Snow. 1994. "Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market." *Journal of Political Economy*, 102(2), 236-57.
Ruhl, Kim J. and Jonathan L. Willis. 2017. "New Exporter Dynamics." *International Economic Review*, 58(3), 703-26.
Seru, Amit; Tyler Shumway and Noah Stoffman. 2009. "Learning by Trading." *The Review of Financial Studies*, 23(2), 705-39.
Shum, Matthew and Yi Xin. 2020. “Time-Varying Risk Aversion? Evidence from Near-Miss Accidents.” *Review of Economics and Statistics*, forthcoming
Soleymanian, Miremad; Charles B. Weinberg and Ting Zhu. 2019. "Sensor Data and Behavioral Tracking: Does Usage-Based Auto Insurance Benefit Drivers?" *Marketing Science*, 38(1), 21-43.
Topa, Giorgio. 2001. "Social Interactions, Local Spillovers and Unemployment." *The Review of Economic Studies*, 68(2), 261-95.
van der Klaauw, Wilbert and Kenneth I. Wolpin. 2008. "Social Security and the Retirement and Savings Behavior of Low-Income Households." *Journal of Econometrics*, 145(1), 21-42.
Voena, Alessandra. 2015. "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review*, 105(8), 2295-332.
Wang, Zhide and Nathan Yang. 2024. "Identification of Structural Learning Models." *Working Paper*.
Weisburd, Sarit. 2015. "Identifying Moral Hazard in Car Insurance Contracts." *The Review of Economics and Statistics*, 97(2), 301-13.
Welch, Ivo. 1992. "Sequential Sales, Learning, and Cascades." *The Journal of Finance*, 47(2), 695-732.

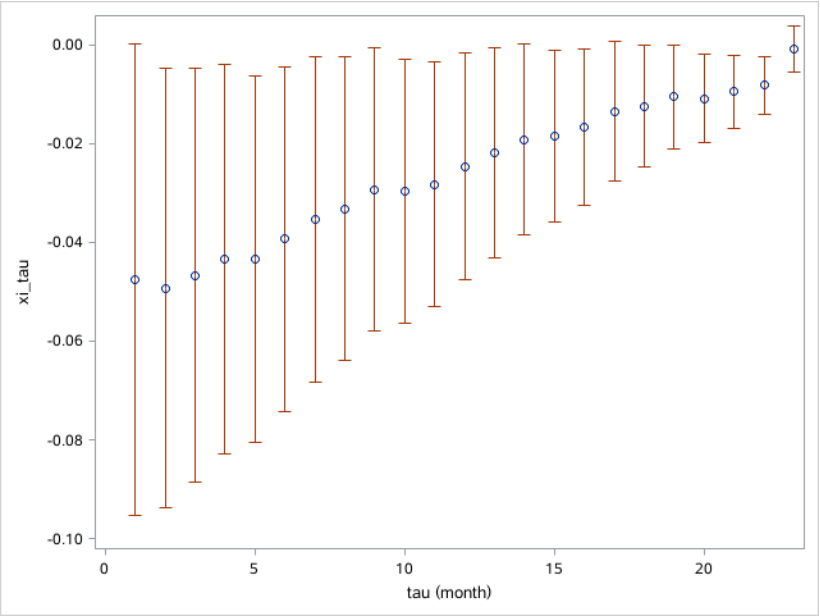


Figure 1. Month-level regression. This figure reports the 2SLS estimator of ξ_τs for the month-level regression of the dynamic panel model displayed in equation (3.3). Vertical bars represent the 95% confidence intervals.

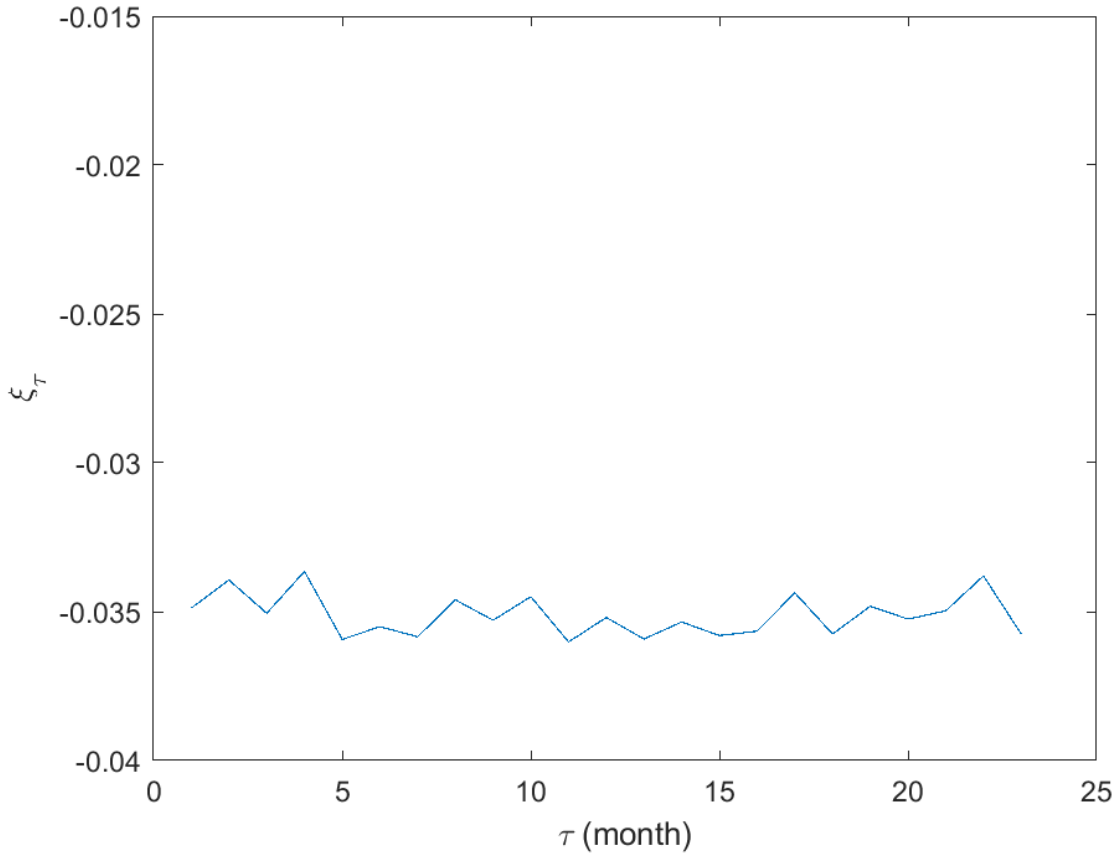


Figure 2. Simulated accident dependence structure for the model of Bayesian learning without forgetting. We set $\theta_H = 0.0470794$ and $\theta_L = 0.002535$ and assume that the true distribution of innate risk type θ in the population is that $Prob(\theta = \theta_H) = 0.5$ and $Prob(\theta = \theta_L) = 0.5$, given the fact that the means of each policyholder's monthly average accident probabilities above and below the median in the sample are 0.0470794 and 0.002535, respectively. We set $L = \frac{7103.73}{1000}\frac{0.15}{1-0.15} = 1.2536$ because the average payment by the insurance company for each accident in the sample is RMB7103.73 and the out-of-pocket cost of an accident for policyholders in China is 10-20% of the total damage caused by the accident. We assume that each policyholder's initial prior belief about her/his θ is that $p_H = 0.1$ and $p_L = 0.9$. We set $\gamma = 5.2$ and $\eta = 6$ to make the simulated $e^{-\eta c_t}$ fluctuate within a reasonable range (from a number reasonably lower than 1 to a number reasonably higher than 1). We randomly generate 53,934 drivers (the same as the number in the empirical regression sample); each policyholder's θ has a 50% probability to equal θ_H and a 50% probability to equal θ_L. We simulate the monthly accident path y_t for four years for each policyholder (equal to the average length of record in the empirical sample). Then, we use the simulated y_t to estimate a dynamic panel model similar to Section 3.2 and obtain the estimates for ξ_τs, as displayed in this figure. Because the theoretical model has only learning and does not have forgetting, these ξ_τs are generally negative but are not decreasing in magnitude as τ increases, which is different from the empirical results in Section 3.2 (Figure 1).

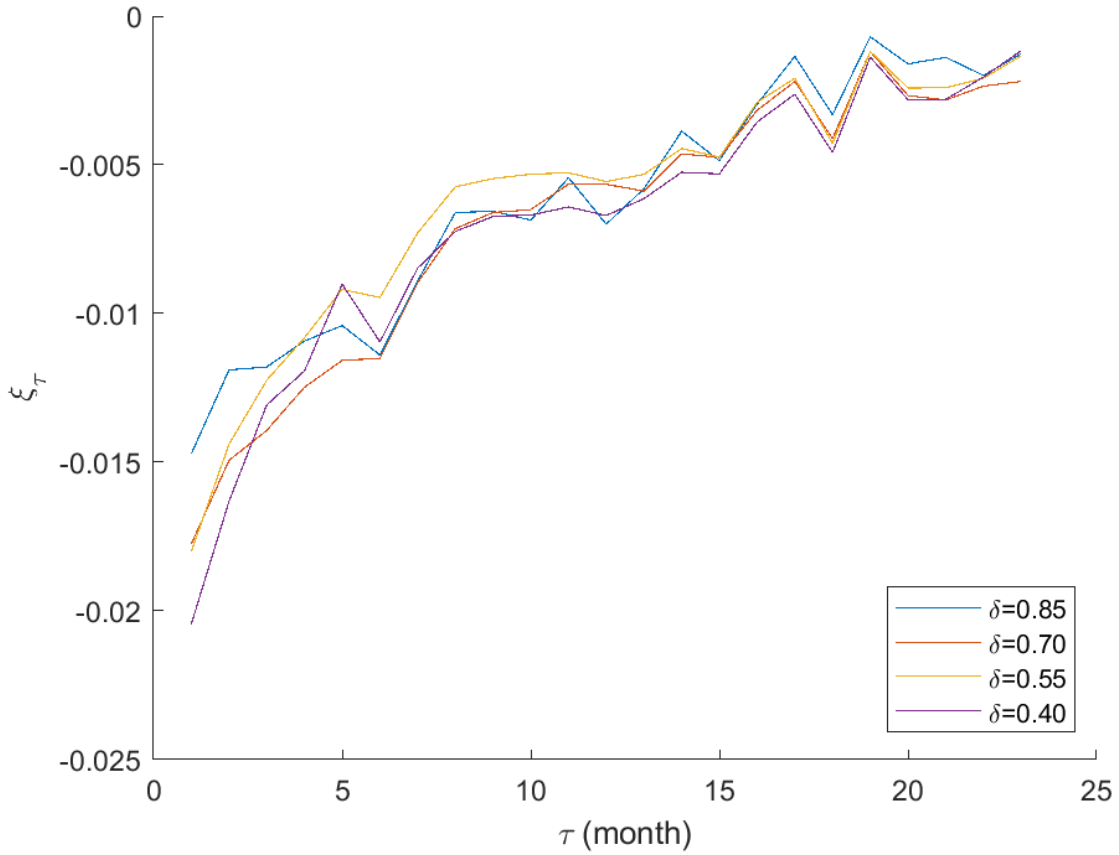


Figure 3. Simulated accident dependence structure for the model of Bayesian learning with forgetting. The remembering rate δ is set to 0.85, 0.70, 0.55, and 0.40, respectively; correspondingly, the forgetting rate (depreciation rate of past information, $1-\delta$) is 0.15, 0.30, 0.45, and 0.60, respectively. The other parameters are the same as those for Figure 2. We simulate the monthly accident path y_t for each policyholder. Then, we use the simulated y_t to estimate a dynamic panel model similar to Section 3 and obtain the estimates for ξ_τs, as displayed in the figure. The ξ_τs are negative and are decreasing in magnitude as τ increases, similar to the empirical ξ_τs in Section 3.2 (Figure 1). Under a higher forgetting rate (lower δ), the line of ξ_τs is more concave, i.e., the ξ_τs are decreasing in magnitude at the beginning at a higher speed as τ increases.

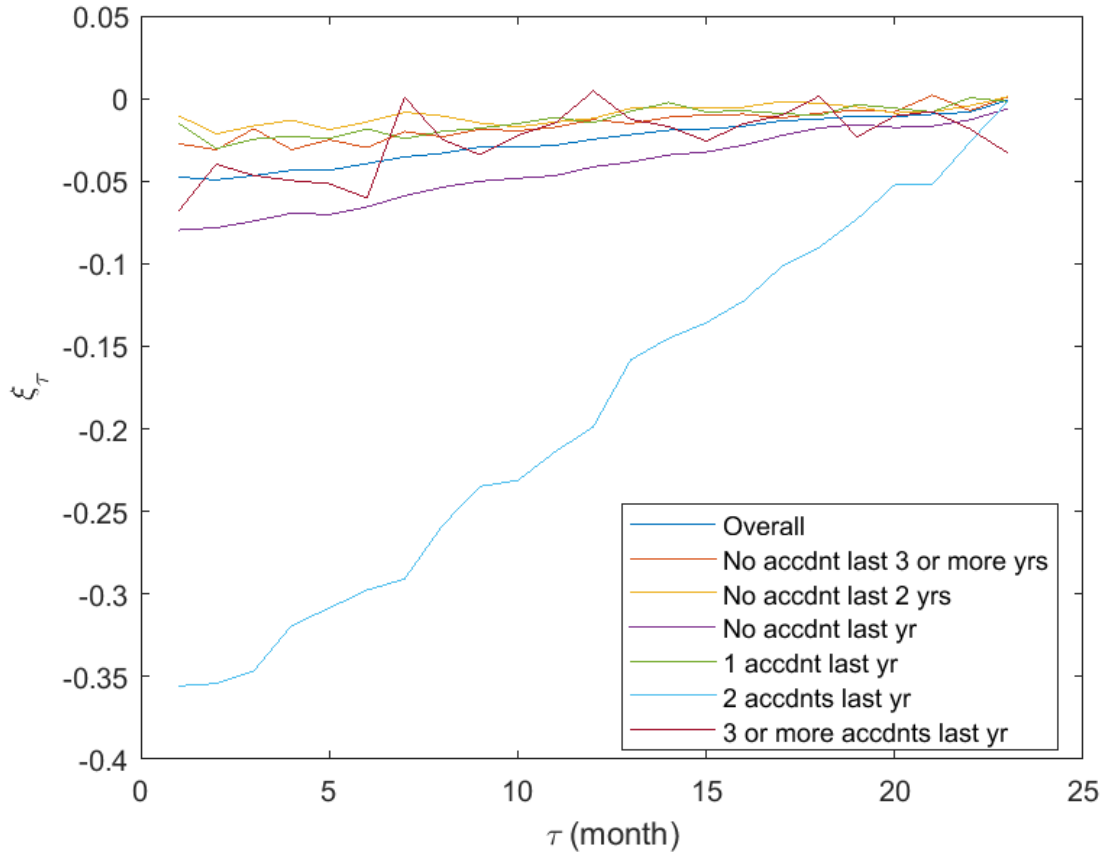


Figure 4. Heterogeneity analyses: initial accident history category. We run the month-level regression of the dynamic panel model in equation (3.3) separately for policyholders with different initial accident history categories (the accident history category for each policyholder's first contract cycle in the sample). This figure displays the 2SLS estimates of ξ_τs for each of the following groups: "no accident for three or more years," "no accident for two years," "no accident for one year," "one accident in the last year," "two accidents in the last year," and "three or more accidents in the last year." The average of ξ_τs for the category of "two accidents in the last year" is significantly different from that of other categories at a level of 1%. The overall results are not mainly driven by this category because only 1.74% of drivers in the sample belong to this category (see Table A.2 in Appendix A). There are other categories with ξ_τs in larger magnitude than the overall ξ_τs. We also exclude the category of "two accidents in the last year" and obtain similar overall results on learning with forgetting.

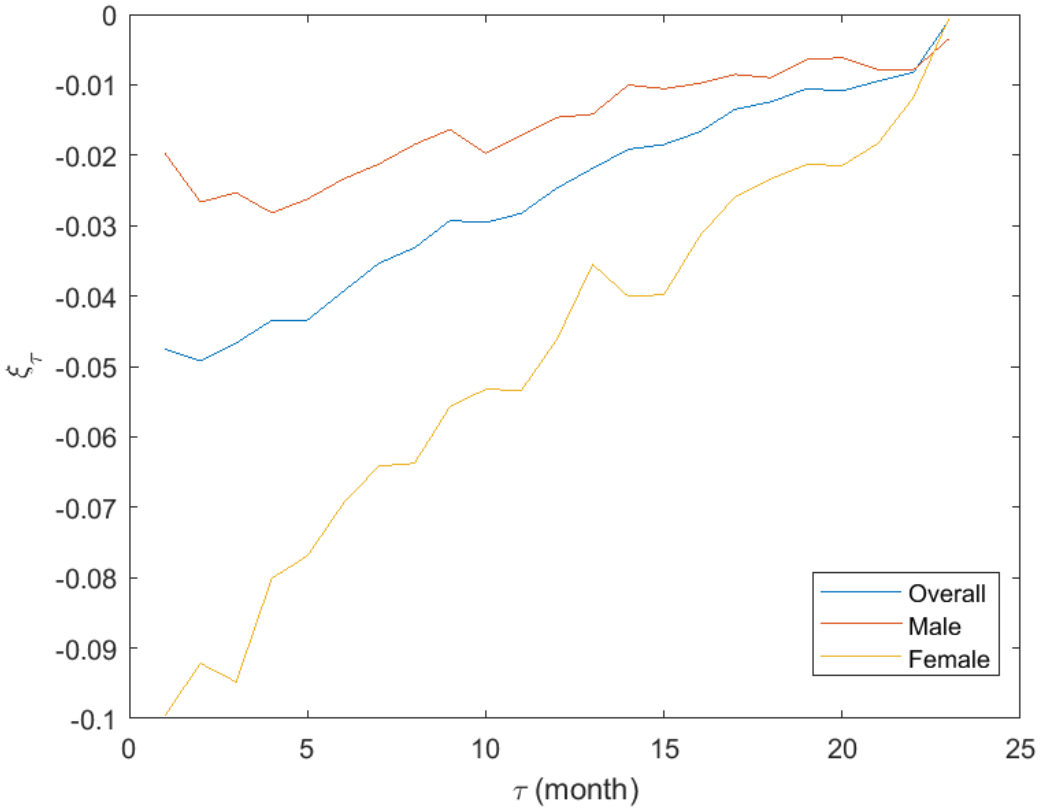


Figure 5. Heterogeneity analyses: gender. We run the month-level regression of the dynamic panel model in equation (3.3) separately for female and male drivers. The figure displays the estimates of ξ_τs for each gender group. The average of ξ_τs for females is significantly different from that for males at a level of 5%. Because female drivers are more risk-averse than male drivers, female drivers' caution level and accident probability are more sensitive to the update of their beliefs on their innate risk types.

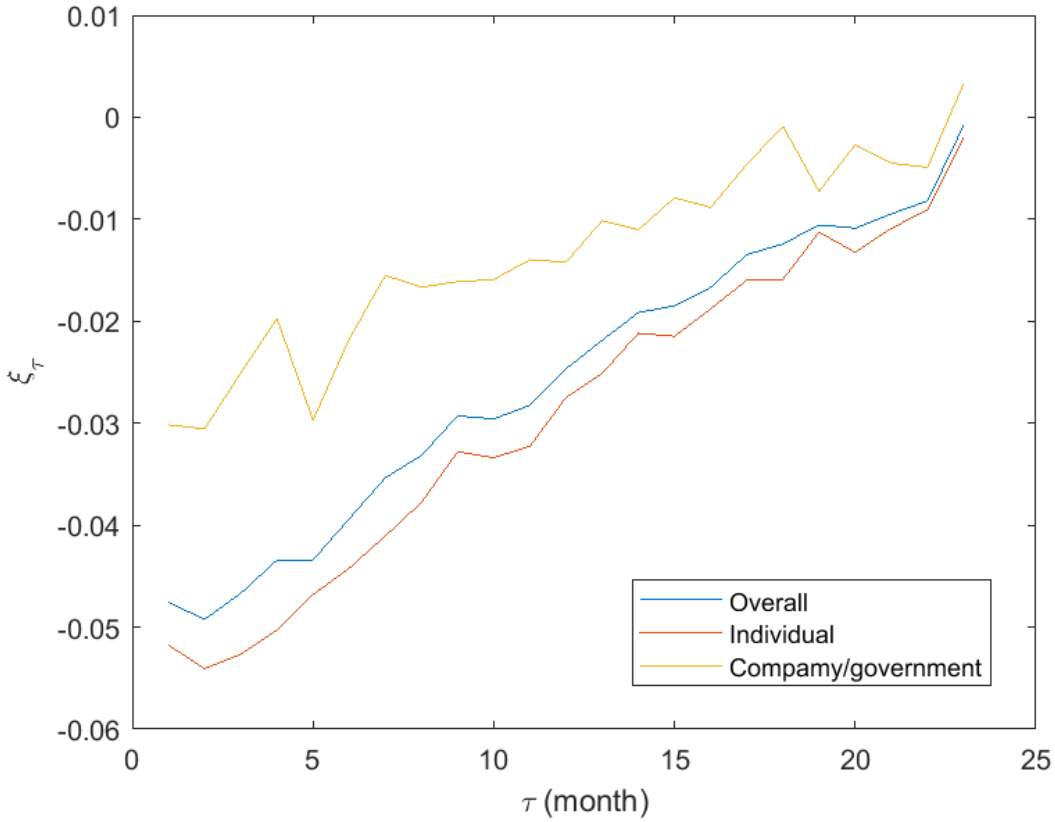


Figure 6. Heterogeneity analyses: individual-owned vs. company-/government-owned vehicles. We run the month-level regression of the dynamic panel model in equation (3.3) separately for individual-owned and company-/government-owned vehicles. The figure displays the estimates of ξ_τs for each group. Because drivers of company-/government-owned vehicles bear smaller out-of-pocket losses than drivers of individual-owned vehicles, the caution level and accident probability of drivers of company-/government-owned vehicles are less sensitive to the updates of their beliefs on their innate risk types.

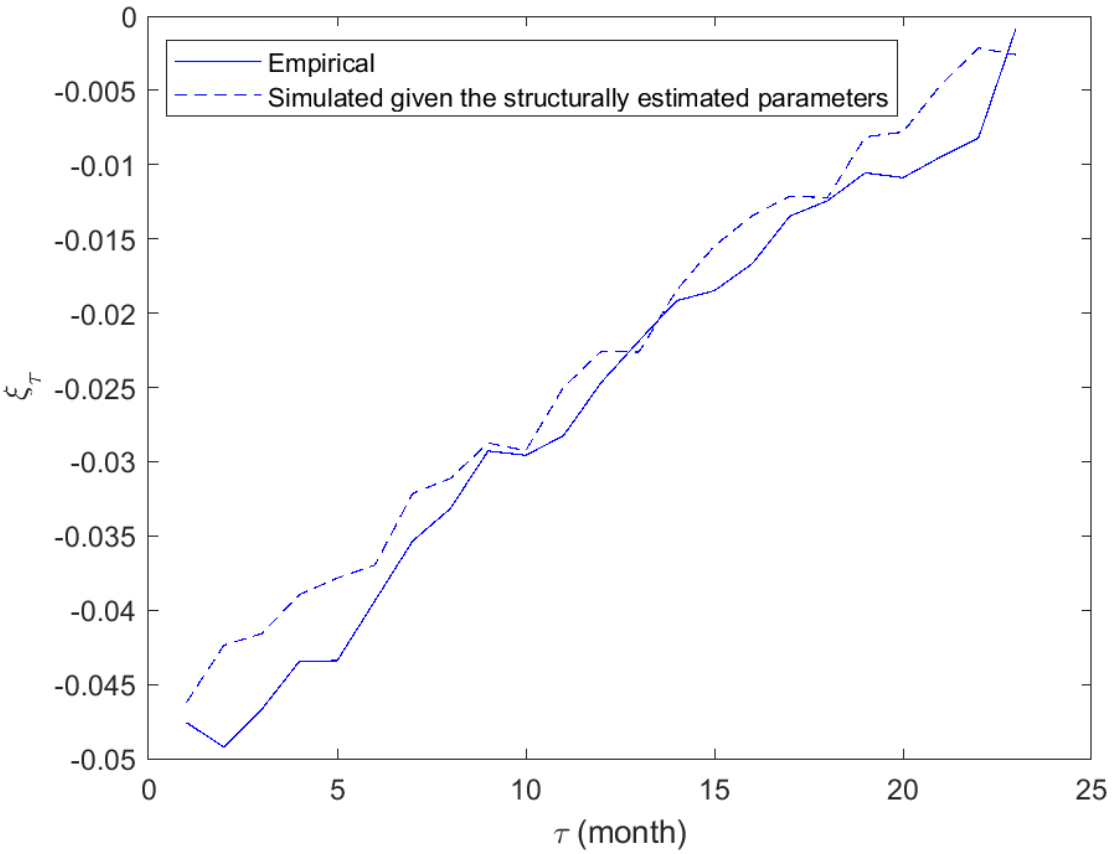


Figure 7. Fit of the structural estimation. The simulated $\tilde{\Xi}$ given the estimated $\{\gamma, \eta, \delta, p_H\}$ is plotted as the dash blue line, which has a similar pattern as that of its empirical counterpart $\hat{\Xi}$ (the solid blue line).

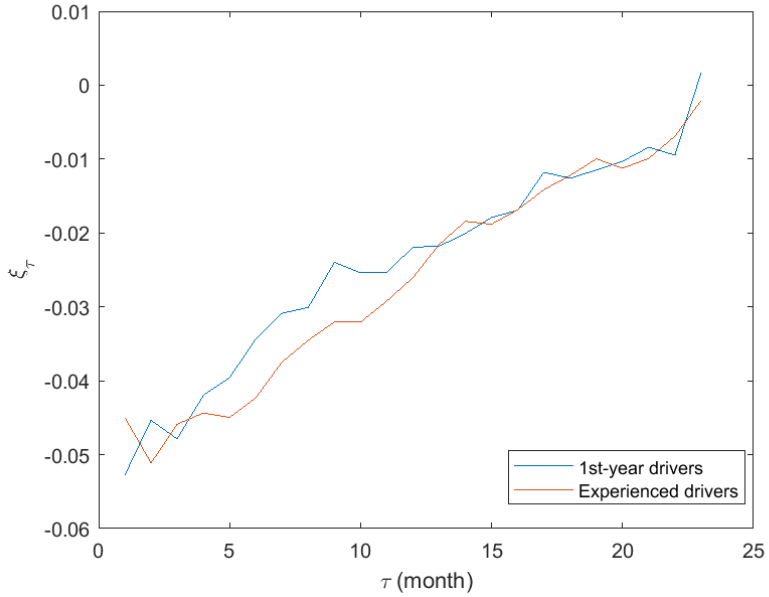


Figure 8. Heterogeneity analyses: first-year drivers vs. experienced drivers. We divide policyholders in the sample into two groups: first-year drivers and drivers with more than one year of experience. We run the month-level regression of the dynamic panel model in equation (3.3) separately for each group. The figure displays the estimates of ξ_τs for each group.

Table 1. Descriptive statistics

Variable	N	Mean	Std Dev
Policyholder age (years)	302627	40.61	9.41
Female	302627	0.20	0.40
Seat number	302627	5.00	1.50
Vehicle age (years)	296726	4.57	3.20
Company- or government-owned	302627	0.12	0.32
Premium (RMB)	302627	3234.11	2793.68
Coverage (RMB)	302627	607292.05	434771.05
Number of accidents per policy	302627	0.29	0.66
Settlement payment per policy (RMB)	302627	1823.26	14626.96
Settlement payment per accident (RMB)	62964	6248.00	18278.87
Settlement payment per policy for policies with accidents (RMB)	49566	11131.96	34679.44

The statistics are based on policyholders with **two** or more years of records in the data.

Table 2. Static panel model - annual analyses

	Column 1	Column 2	Column 3
No accident for 3 or more years	-0.0352**	0.224***	0.220***
	(0.0167)	(0.0357)	(0.0285)
No accident for 2 years	-0.0216	0.152***	0.148***
	(0.0167)	(0.0356)	(0.0284)
No accident for 1 year	-0.00669	0.0835**	0.0824***
	(0.0167)	(0.0355)	(0.0284)
New driver	0.00366	0.0704**	0.0697**
	(0.0167)	(0.0356)	(0.0285)
1 accident last year	0.0500***	-0.140***	-0.137***
	(0.0169)	(0.0357)	(0.0286)
2 accidents last year	0.142***	-0.114***	-0.123***
	(0.0181)	(0.0365)	(0.0293)
3 or more accidents last year	0.173***	0.00800	0.00115
	(0.0308)	(0.0468)	(0.0394)
Coverage (RMB 1 M)	0.0254***	0.233***	0.224***
	(0.00386)	(0.00586)	(0.00507)
Premium (RMB 1 K)	0.0302***	0.00451***	0.00340***
	(0.000829)	(0.00109)	(0.000939)
Policyholder age	✓		
Seat number	✓		
Vehicle age	✓		
City fixed effects	✓		
Gender	✓		
Vehicle owner type (individual vs. company or government)	✓		
Vehicle model fixed effects	✓		
Driving experience	✓	✓	✓
Calendar year fixed effects	✓	✓	✓
Individual fixed effects		✓	✓
Observations	296,655	296,655	367,906
R-squared	0.100	0.538	0.533

This table reports the contract-year-level regressions of the static panel model in equation (3.1). The dependent variable is a 0-1 indicator of whether an accident occurs in policyholder i's contract year t.

In column 1, we do not control for the individual fixed effects (a_i); the results exhibit that the worse the accident history is, the higher the accident probability for the current contract cycle. Without controlling for individual fixed effects, the accident history partially reveals a policyholder's unobserved innate risk type (a_i). Risky drivers are more likely to encounter accidents both in the past and in the current period (heterogeneity).

In column 2, we control for the individual fixed effects (a_i); the results exhibit that a worse accident history leads to a lower accident probability for the current contract cycle. These results indicate that policyholders learn about their own innate risk types from the accidents they previously experienced and correspondingly exert more caution or effort to reduce their accident probabilities in the current contract cycle (state dependence).

In column 3, we add drivers with missing time-invariant characteristics (e.g., seat number), who cannot be included in column 1. The results are similar to column 2.

The different results of columns 1 and 2 of Table 2 reflect the importance of disentangling true state dependence from individual heterogeneity, which is a well-known issue in the applied microeconomic literature (see Heckman, 1991). The standard errors in parentheses are clustered by individual. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level. The sample includes policyholders with two or more years of records in the data. The results are robust if the sample includes only policyholders with three or more years of records in the data. Policyholders with missing accident histories in the sample are the omitted reference group.

Table 3. Dynamic panel model – quarterly analyses

	Within estimator		2SLS		
	Dependent variable: y_{it}		Dependent variable: Δy_{it} Endogenous regressor: $\Delta y_{i,t-1}$ IV: $\Delta y_{i,t-8}$, $\Delta y_{i,t-9}$		
Parameter	Regressor	Estimates	Regressor	1st stage	2nd stage
ξ_{t-1}	$y_{i,t-1}$	-0.1904***	$\Delta y_{i,t-1}$	--	**-0.0438*** **
		(0.0028)			**(0.0165)**
ξ_{t-2}	$y_{i,t-2}$	-0.1845***	$\Delta y_{i,t-2}$	-0.8448***	**-0.0426*** **
		(0.0027)		(0.0030)	**(0.0138)**
ξ_{t-3}	$y_{i,t-3}$	-0.1671***	$\Delta y_{i,t-3}$	-0.7232***	**-0.0352*** **
		(0.0026)		(0.0037)	**(0.0116)**
ξ_{t-4}	$y_{i,t-4}$	-0.1378***	$\Delta y_{i,t-4}$	-0.6097***	**-0.0343*** **
		(0.0025)		(0.0041)	**(0.0096)**
ξ_{t-5}	$y_{i,t-5}$	-0.1067***	$\Delta y_{i,t-5}$	-0.5003***	**-0.0269*** **
		(0.0023)		(0.0042)	**(0.0077)**
ξ_{t-6}	$y_{i,t-6}$	-0.0911***	$\Delta y_{i,t-6}$	-0.3912***	**-0.0222*** **
		(0.0023)		(0.0040)	**(0.0058)**
ξ_{t-7}	$y_{i,t-7}$	-0.0686***	$\Delta y_{i,t-7}$	-0.2974***	**-0.0146*** **
		(0.0021)		(0.0037)	**(0.0043)**
	$coverage_{it}$	0.0762***	$\Delta coverage_{it}$	-0.0406***	0.0873***
	(RMB 1 M)	(0.0033)	(RMB 1 M)	(0.0052)	(0.0063)
	$premium_{it}$	-8.56e-04	$\Delta premium_{it}$	1.81e-03	-1.81e-03
	(RMB 1 K)	(7.33e-04)	(RMB 1 K)	(1.33e-03)	(1.46e-03)
			$\Delta y_{i,t-8}$	-0.1994***	
				(0.0033)	
			$\Delta y_{i,t-9}$	-0.0916***	
				(0.0024)	
Individual fixed effects		✓		✓	✓
Calendar year-quarter fixed effects		✓		✓	✓
Fixed effects for quarters since the first contract		✓		✓	✓
N		383,416		221,687	221,687
R-squared		0.2651		0.4337	

y_t is a 0-1 indicator of whether policyholder i encountered an accident in quarter t. The sample includes policyholders with **three** or more years of records in the data. The first regression reported in this table is the within estimator of equation (3.2). The second and third regressions reported in this table are the first and second stages of the 2SLS estimation for the dynamic panel model displayed in equation (3.3). In the first stage of the 2SLS estimation, the F statistic for the joint null hypothesis that the coefficients of $\Delta y_{i,t-8}$ and $\Delta y_{i,t-9}$ are both zero is 1854.55, which is much higher than the threshold (10) to judge that the instrumental variables are relevant. In both the within estimator and the 2SLS estimator, the individual fixed effects on the accident probability are implicitly controlled for because they are differenced out. The standard errors in parentheses are clustered by individual. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level. Because the regressors in column 1 have seven lags, observations of a policyholder's first seven quarters in the data are not included in the regression.

Table 4. Structural estimation

Structural parameter	Description	Estimate
γ	Coefficient of risk aversion	5.0654***
		(0.8281)
η	Effectiveness of caution or effort	6.2025***
		(0.9574)
δ	Remembering rate	0.8615***
		(0.0633)
p_H	Initial prior belief on being the high-risk type	0.1019**
		(0.0544)

The structural parameters in the theoretical model of Bayesian learning with forgetting are estimated by indirect inference. The standard errors in parentheses are estimated by bootstrapping. We randomly select an 80% subsample and conduct the structural estimation. We repeat the process 100 times and calculate the standard deviation of the estimates that we obtain in the 100 times. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level. The remembering rate δ is significantly smaller than 1 at a 5% level; thus, we can reject the hypothesis that no forgetting occurs in the learning process. Accordingly, the monthly forgetting rate (depreciation rate of past information in policyholders' belief formation, $1-\delta$) is 0.1385.

Internet Appendix

(For Online Publication)

Appendix A. Figures and Tables

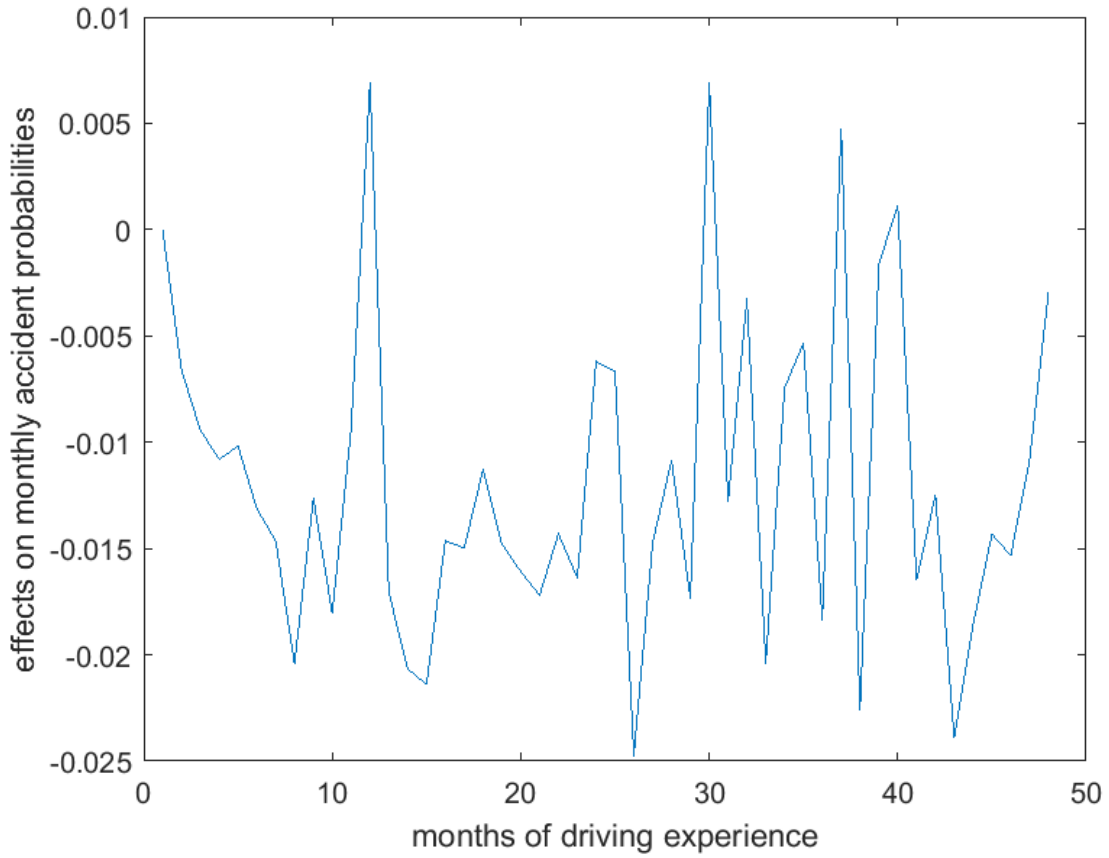


Figure A.1. Effect of driving experience on accident probabilities. Using policyholders who started as new drivers and have four contract years in the data, we run the month-level regression of the accident indicator on dummy variables of policyholders' months of driving experience. We control for the premium, coverage, the age fixed effects, and the calendar-year-month fixed effects. The figure displays the estimates for the coefficients of dummy variables of policyholders' months of driving experience. The fixed effects of months of driving experience are decreasing in the first ten months and fluctuate thereafter.

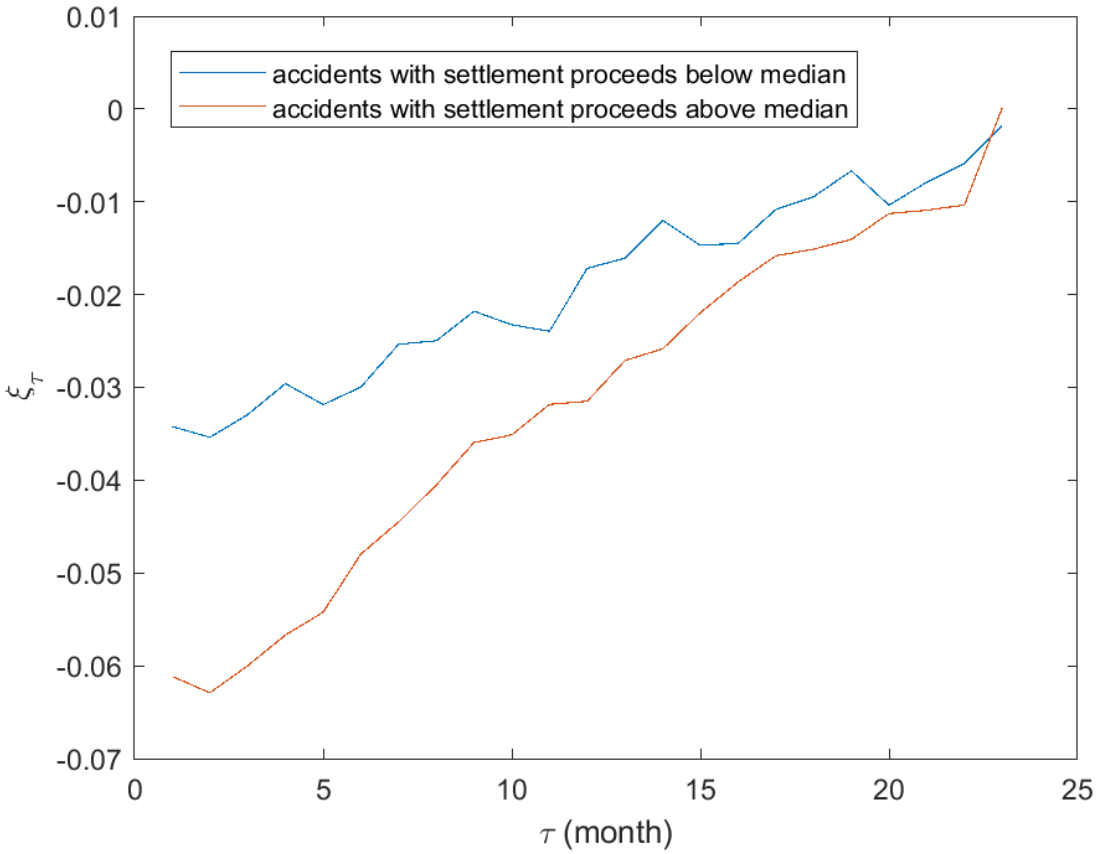


Figure A.2. Heterogeneity analyses: accidents of different sizes. We run the month-level regression of the dynamic panel model in equation (3.3), allowing ξ_τs to be different across the accidents with settlement proceeds below the median and those above the median. The blue line displays ξ_τs for accidents below the median, while the red line displays ξ_τs for accidents above the median.

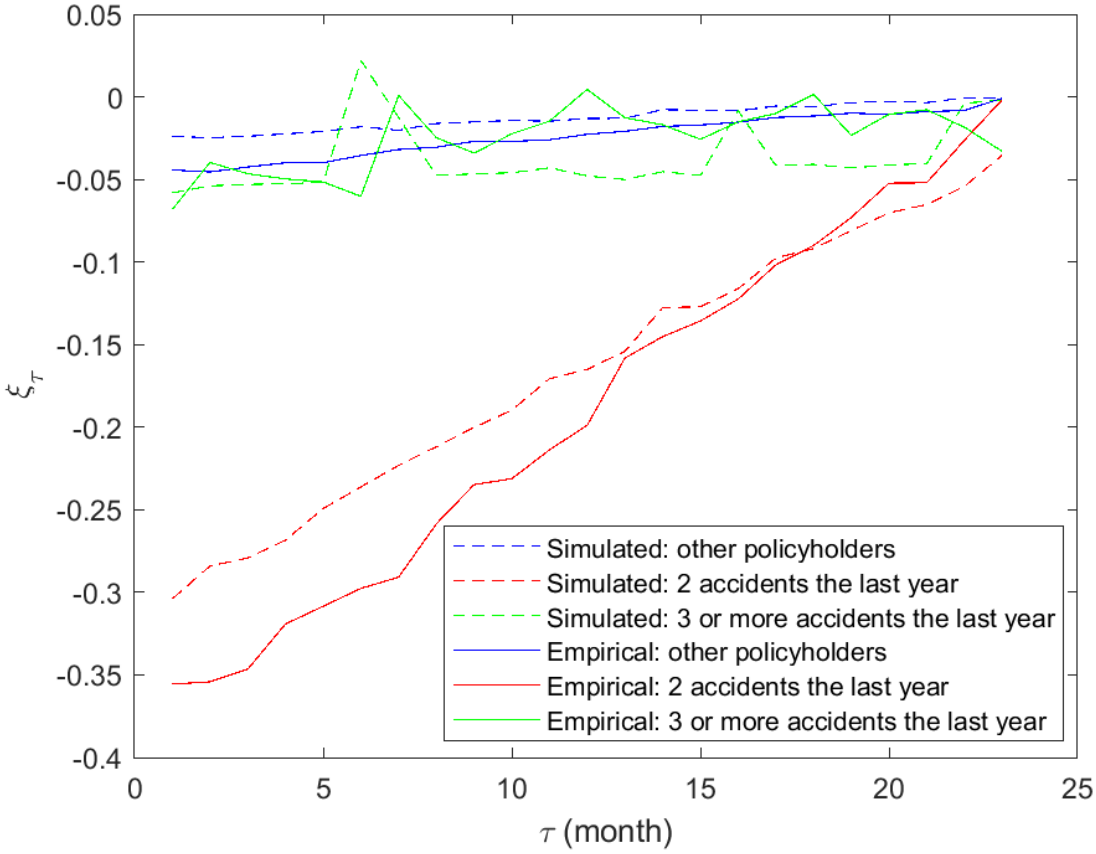


Figure A.3. Fit of the structural estimation allowing heterogeneous initial prior beliefs. The simulated $\tilde{\Xi}$ given the estimated $\{\gamma, \eta, \delta, p_H\}$ is plotted as the dash lines, while its empirical counterpart $\hat{\Xi}$ is plotted as the solid lines. The red, green, and blue colors, respectively, represent policyholders in the following three categories based on their initial accident history for their first contract year in the sample period: "2 accidents in the last year," "3 or more accidents in the last year," and others.

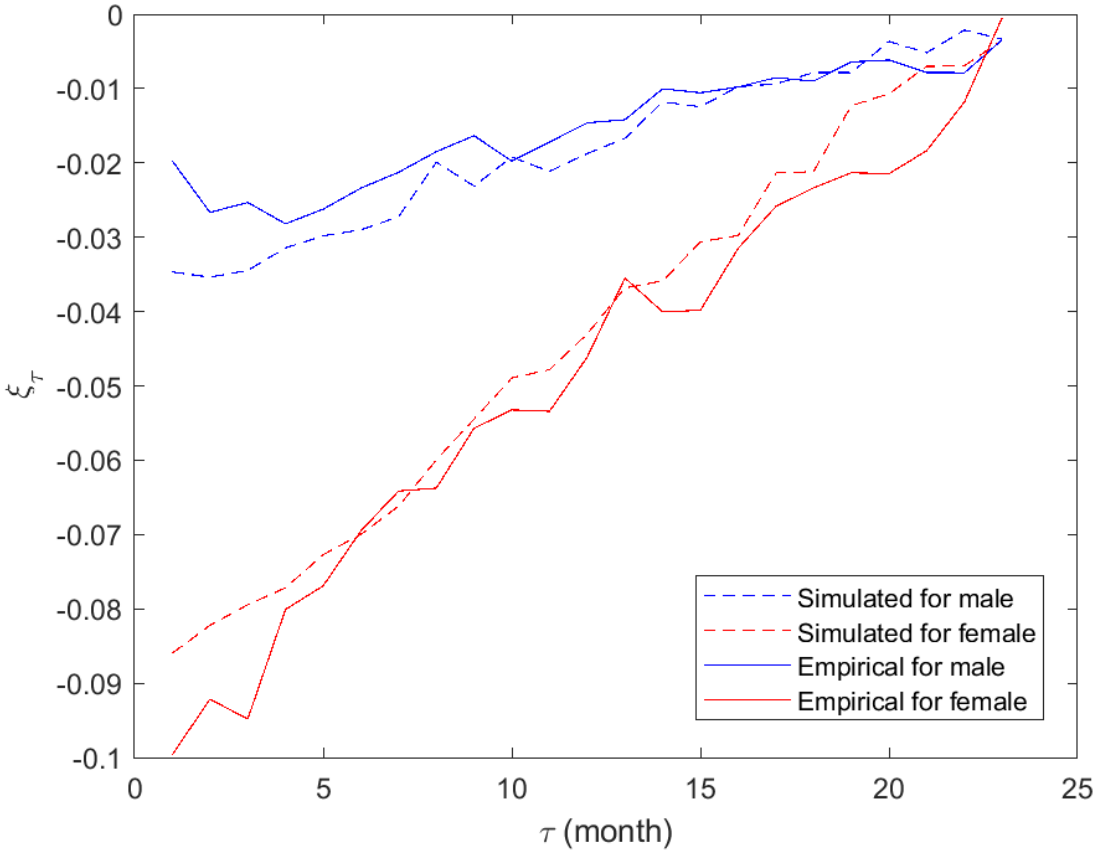


Figure A.4. Fit of the structural estimation – heterogeneous risk aversion across genders. The simulated $\tilde{\Xi}$ given the estimated $\{\gamma, \eta, \delta, p_H\}$ is plotted as the dash lines, while its empirical counterpart $\hat{\Xi}$ is plotted as the solid lines. The blue and red colors, respectively, represent male and female policyholders.

Table A.1. Distribution of contract years per driver in the data

Years of records per policyholder in the sample	Percentage of policyholders	Cumulative percentage of policyholders
2	62.21%	62.21%
3	24.69%	86.9%
4	8.33%	95.23%
5	3.11%	98.34%
6	1.2%	99.55%
7	0.23%	99.78%
8	0.14%	99.92%
9	0.06%	99.97%
10	0.02%	100%

Table A.2. Distribution of accident history categories

Grouping based on accident history category of all policy contract	**Percentage in the sample**
No accident for 3 or more years	32.65%
No accident for 2 years	16.73%
No accident in the previous year	22.73%
New driver	16.07%
1 accident in the previous year	9.93%
2 accidents in the previous year	1.77%
3 or more accidents in the previous year	0.13%

Grouping based on accident history category of policyholders' first contract in the sample	**Percentage in the sample**
No accident for 3 or more years	22.39%
No accident for 2 years	13.51%
No accident in the previous year	22.48%
New driver	30.99%
1 accident in the previous year	8.64%
2 accidents in the previous year	1.74%
3 or more accidents in the previous year	0.26%

Table A.3. Descriptive statistics for subgroups

	Policyholders with 1 year of records			Policyholders with 2 or more years of record		
Variable	N	Mean	Std Dev	N	Mean	Std Dev
Policyholder age (years)	328356	39.21	9.57	302627	40.61	9.41
Female	328356	0.20	0.40	302627	0.20	0.40
Seat number	328355	4.50	1.85	302627	5.00	1.50
Vehicle age (years)	299882	3.85	3.63	296726	4.57	3.20
Company- or government-owned	328356	0.08	0.28	302627	0.12	0.32
Premium (RMB)	328356	2763.85	2815.43	302627	3234.11	2793.68
Coverage (RMB)	328356	514532.45	444357.70	302627	607292.05	434771.05
Number of accidents per policy	328356	0.26	0.65	302627	0.29	0.66
Settlement payment per policy (RMB)	328356	1738.14	15828.37	302627	1823.26	14626.96
Settlement payment per accident (RMB)	84891	6723.07	20277.15	62964	6248.00	18278.87
Settlement payment per policy for policies with accidents (RMB)	47198	12092.22	40222.02	49566	11131.96	34679.44

Table A.4. Distribution of accident history categories for subgroups

Accident history category	Percentage for policyholders with only 1 year of records in the data	Percentage for the first contract of policyholders with 2 or more years of records in the data
No accident for 3 or more years	22.16%	22.39%
No accident for 2 years	10.98%	13.51%
No accident in the previous year	17.59%	22.48%
New driver	40.79%	30.99%
1 accident in the previous year	7.05%	8.64%
2 accidents in the previous year	1.32%	1.74%
3 or more accidents in the previous year	0.10%	0.26%

Table A.5. Dynamic panel model: robustness checks

Parameter	Regressor	Column 1	Column 2
ξ_{t-1}	$\Delta y_{i,t-1}$	-0.0641***	-0.0687***
		(0.0239)	(0.0160)
ξ_{t-2}	$\Delta y_{i,t-2}$	-0.0586***	-0.0633***
		(0.0199)	(0.0135)
ξ_{t-3}	$\Delta y_{i,t-3}$	-0.0448***	-0.0524***
		(0.0166)	(0.0114)
ξ_{t-4}	$\Delta y_{i,t-4}$	-0.0384***	-0.0489***
		(0.0134)	(0.0095)
ξ_{t-5}	$\Delta y_{i,t-5}$	-0.0282***	-0.0372***
		(0.0103)	(0.0078)
ξ_{t-6}	$\Delta y_{i,t-6}$	-0.0235***	-0.0258***
		(0.0075)	(0.0061)
ξ_{t-7}	$\Delta y_{i,t-7}$	-0.0145***	-0.0181***
		(0.0049)	(0.0045)
	$\Delta coverage_{it}$	0.0874***	0.0865***
	(RMB 1 M)	(0.0062)	(0.0062)
	$\Delta premium_{it}$	-1.97e-03	-1.91e-03
	(RMB 1 K)	(1.43e-03)	(1.44e-03)
Individual fixed effects		✓	✓
Calendar year-quarter fixed effects		✓	✓
Fixed effects for quarters since the first contract		✓	✓
N		275,535	167,895

Column 1 is a 2SLS estimation similar to the 2SLS estimation in Table 3 but the instrumental variable for the endogenous regressor $\Delta y_{i,t-1}$ is $\Delta y_{i,t-8}$. Column 2 is a 2SLS estimation similar to the 2SLS estimation in Table 3 but the instrumental variables for the endogenous regressor $\Delta y_{i,t-1}$ are $\Delta y_{i,t-8}$, $\Delta y_{i,t-9}$, and $\Delta y_{i,t-10}$. The results are robust to changes in instrumental variables.

Table A.6. Structural estimation – allowing heterogeneous initial prior beliefs

Structural parameter	Description	Estimate
γ	Coefficient of risk aversion	5.3723***
		(0.9896)
η	Effectiveness of caution or effort	7.2701***
		(3.3761)
δ	Remembering rate	0.8749***
		(0.0750)
p_H^2	Initial prior belief for policyholders with an initial accident history of "2 accidents in the last year"	0.3016**
		(0.1448)
p_H^3	Initial prior belief for policyholders with an initial accident history of "3 or more accidents in the last year"	0.6003**
		(0.2809)
p_H^{other}	Initial prior belief for policyholders with other initial accident histories	0.1005**
		(0.0532)

As robustness checks, we allow policyholders in different initial accident history categories in the first period of the data to have different initial prior beliefs about their innate risk types, which are to be estimated. We divide policyholders into three categories based on their initial accident history for their first contract year in the sample period: "2 accidents in the last year," "3 or more accidents in the last year," and others. Given an initial accident history category k, we assume that the true distribution of innate risk type θ for category k in the population is $Prob^k(\theta = \theta_H)$, the solution of the equation $Prob^k(\theta = \theta_H)\theta_H + \left(1 - Prob^k(\theta = \theta_H)\right)\theta_L = A^k$, where A^k is the average monthly accident probability of policyholders in this category calculated from the actual data. We use simulated ξ_τs from the structural model for those categories, respectively, to match ξ_τs obtained from the reduced-form regressions for those categories. A similar estimate for the remembering rate δ is obtained. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level.

Table A.7. Structural estimation – heterogeneous risk aversion across genders

Structural parameter	Description	Estimate
γ	Coefficient of risk aversion for male drivers	5.1890***
		(0.9643)
$\Delta\gamma_female$	Difference in the coefficient of risk aversion between female and male drivers	0.7098**
		(0.4140)
η	Effectiveness of caution or effort	6.0299***
		(2.2141)
δ	Remembering rate	0.8419***
		(0.0878)
p_H	Initial prior belief on being the high-risk type	0.0991**
		(0.0560)

As robustness checks, we allow that the coefficient of risk aversion γ to be estimated is different across genders. We use simulated ξ_τs from the structural model for males and females, respectively, to match ξ_τs obtained from the reduced-form regressions for males and females. A similar estimate for the remembering rate δ is obtained. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level.

Table A.8. The increment in the "bonus-malus coefficient" caused by each accident within a contract year

Panel A

Accident history for determining premium of year t	coefficient	No accdnt in year t		1 accdnt in year t		
		Resulting category for $t+1$	Coefficient	Resulting category for $t+1$	Coefficient	**Increment**
No accdnt 3 or more yrs	0.6	No accdnt 3 or more yrs	0.6	1 accdnt last yr	1	**0.4**
No accdnt 2 yrs	0.7	No accdnt 3 or more yrs	0.6	1 accdnt last yr	1	**0.4**
No accdnt 1 yr	0.85	No accdnt 2 yrs	0.7	1 accdnt last yr	1	**0.3**
1 accdnt last yr or new driver	1	No accdnt 1 yr	0.85	1 accdnt last yr	1	**0.15**
2 accdnts last yr	1.25	No accdnt 1 yr	0.85	1 accdnt last yr	1	**0.15**
3 accdnts last yr	1.5	No accdnt 1 yr	0.85	1 accdnt last yr	1	**0.15**
4 accdnts last yr	1.75	No accdnt 1 yr	0.85	1 accdnt last yr	1	**0.15**
5 or more accdnts last yr	2	No accdnt 1 yr	0.85	1 accdnt last yr	1	**0.15**

Panel B

Accident history for determining premium of year t	2 accdnts in year t			3 accdnts in year t		
	Resulting category for $t+1$	Coefficient	**Increment**	Resulting category for $t+1$	Coefficient	**Increment**
No accdnt 3 or more yrs	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
No accdnt 2 yrs	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
No accdnt 1 yr	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
1 accdnt last yr or new driver	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
2 accdnts last yr	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
3 accdnts last yr	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
4 accdnts last yr	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**
5 or more accdnts last yr	2 accdnts last yr	1.25	**0.25**	3 accdnts last yr	1.5	**0.25**

Panel C

Accident history for determining premium of year t	4 accdnts in year t			5 accdnts in year t		
	Resulting category for $t+1$	Coefficient	**Increment**	Resulting category for $t+1$	Coefficient	**Increment**
No accdnt 3 or more yrs	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
No accdnt 2 yrs	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
No accdnt 1 yr	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
1 accdnt last yr or new driver	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
2 accdnts last yr	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
3 accdnts last yr	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**

4 accdnts last yr	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**
5 or more accdnts last yr	4 accdnts last yr	1.75	**0.25**	5 or more accdnts last yr	2	**0.25**

Panel D

Accident history for determining premium of year t	6 accdnts in year t			7 accdnts in year t		
	Resulting category for $t+1$	Coefficient	**Increment**	Resulting category for $t+1$	Coefficient	**Increment**
No accdnt 3 or more yrs	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
No accdnt 2 yrs	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
No accdnt 1 yr	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
1 accdnt last yr or new driver	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
2 accdnts last yr	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
3 accdnts last yr	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
4 accdnts last yr	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**
5 or more accdnts last yr	5 or more accdnts last yr	2	**0**	5 or more accdnts last yr	2	**0**

Policyholders are classified into different accident history categories based on previous driving records (see column 1 in panel A). The "bonus-malus coefficient" for each category is displayed in column 2 of panel A. In general, the increment in the "bonus-malus coefficient" caused by each accident within a contract year is not increasing in the ordinal of the accident within the contract year.

For example, for a policyholder in the category of "no accident for one year" for determining the premium of year t, the coefficient for year $t+1$ will become 0.7 if no accident occurs in year t and 1.0 if one accident occurs in year t. Therefore, the increment caused by the first accident in year t is 0.3 $(= 1.0 - 0.7)$. If a second accident occurs in year t, the coefficient for year $t+1$ will become 1.25. Therefore, the increment caused by the second accident in year t is 0.25 $(= 1.25 - 1.0)$. In summary, for this policyholder, the increments in the "bonus-malus coefficient" for year $t+1$ caused by each accident in year t are 0.3, 0.25, 0.25, 0.25, 0.25, 0, 0, 0…

The increasing marginal effect occurs only in one situation: for policyholders with one or more accidents in year $t-1$, the increment in the coefficient for year $t+1$ caused by the first accident in year t is smaller (0.15) than that caused by the second accident in year t (0.25). However, we will obtain similar empirical results if we exclude this situation.

Table A.9. The increment in the "bonus-malus coefficient" caused by each accident across different contract years

1 accdnt in year t	1 accdnt in year $t+1$					1 accdnt in year $t+2$				
Resulting category for determining premium of $t+1$	Resulting category for determining premium of $t+2$	Coef	Category if no accdnt	Coef	Increment caused by the accident	Resulting category for determining premium of $t+3$	Coef	Category if no accdnt	Coef	Increment caused by the accident
1 accdnt last yr	1 accdnt last yr	1	No accdnt 1 yr	0.85	0.15	1 accdnt last yr	1	No accdnt 1 yr	0.85	0.15

If a policyholder encounters one accident in years t, $t+1$, and $t+2$, respectively, the increment in the "bonus-malus coefficient" caused by the second accident is $0.15\ (=1-0.85)$, and the increment caused by the third accident is also $0.15\ (=1-0.85)$. Therefore, in general, the increment in the "bonus-malus coefficient" caused by each accident in different contract years is not increasing in the ordinal of the accident in the policyholder's lifetime.

Table A.10. The effect of the time length since the previous accident

Panel A

1 accdnt in year t	No accdnt in year $t+1$		1 accdnt in year $t+2$				
Resulting category for determining premium of $t+1$	Resulting category for determining premium of $t+2$	Coef	Resulting category for determining premium of $t+3$	Coef	Category if no accdnt	Coef	Increment caused by the accident
1 accdnt last yr	No accdnt 1 yr	0.85	1 accdnt last yr	1	No accdnt 2 yrs	0.7	0.3

Panel B

1 accdnt in year t	No accdnt in years $t+1$ and $t+2$		1 accdnt in year $t+3$				
Resulting category for determining premium of $t+1$	Resulting category for determining premium of $t+3$	Coef	Resulting category for determining premium of $t+4$	Coef	Category if no accdnt	Coef	Increment caused by the accident
1 accdnt last yr	No accdnt 2 yrs	0.7	1 accdnt last yr	1	No accdnt 3 yrs	0.6	0.4

Situation 1. Given that one accident occurs in year t, if no accident occurs in year $t+1$, the policyholder will go to the category of "no accident for one year" and the corresponding "bonus-malus coefficient" for year $t+2$ is 0.85; if one accident occurs in year $t+1$, the policyholder will go to the category of "one accident last year" and the corresponding "bonus-malus coefficient" for year $t+2$ is 1. Therefore, the increment in the "bonus-malus coefficient" caused by the follow-up accident is 0.15 $(=1-0.85)$ if the follow-up accident occurs in year $t+1$.

Situation 2. Given that one accident occurs in year t, if no accident occurs in year $t+1$ and year $t+2$, the policyholder will go to the category of "no accident for two years" and the corresponding "bonus-malus coefficient" for year $t+3$ is 0.7; alternatively, if no accident occurs in year $t+1$ and one accident occurs in year $t+2$, the policyholder will go to the category of "one accident last year" and the corresponding "bonus-malus coefficient" for year $t+3$ is 1. Therefore, the increment in the "bonus-malus coefficient" caused by the follow-up accident is 0.3 $(=1-0.7)$ if the follow-up accident occurs in year $t+2$.

Similarly, the increment caused by the follow-up accident is 0.4 $(=1-0.6)$ if the follow-up accident occurs in year $t+3$. The increment in the "bonus-malus coefficient" caused by an accident could be increasing in how many years it occurs after the previous accident.

Table A.11. Distributions of settlement proceeds for different accidents of a policyholder

	Policyholders with 2 accidents		Policyholders with 3 accidents		
	1st accident	2nd accident	1st accident	2nd accident	3rd accident
Mean (RMB)	6925.87	6084.47	8807.57	7716.40	6506.00
99% percentile (RMB)	84646.51	82000	89491.20	85920	81894.90
95% percentile (RMB)	25335.00	22575.96	31000	28128	24907.10
90% percentile (RMB)	13670.05	11920	18283.60	16128	13350
75% percentile (RMB)	5400.00	4308	7440.50	6300	5000
50% percentile (RMB)	2263.98	1800	3200	2680	2000
25% percentile (RMB)	950.00	500	1400	1050	500
10% percentile (RMB)	0	0	200	0	0
5% percentile (RMB)	0	0	0	0	0
1% percentile (RMB)	0	0	0	0	0

The first two columns show the distributions of payments by the insurance company for the first and second accidents, respectively, to policyholders with two accidents in the data; the last three columns show the distributions of payments by the insurance company for the first, second, and third accidents, respectively, to policyholders with three accidents in the data. The proportion of reported accidents with minor damages is increasing in the ordinal of the accident experienced by a policyholder. This pattern suggests that if underreporting exists, as a policyholder experiences more accidents, she/he will become less likely to underreport minor accidents. Therefore, the existence of underreporting can only cause the learning effect to be underestimated rather than overestimated.

Table A.12. Robustness checks: accidents with settlement proceeds above different thresholds

Parameter	Regressor	Threshold of settlement proceed			
		≥RMB500	≥RMB1000	≥RMB1500	≥RMB2000
ξ_{t-1}	$\Delta y_{i,t-1}$	-0.0420**	-0.0415**	-0.0374*	-0.0361*
		(0.0185)	(0.0196)	(0.0209)	(0.0217)
ξ_{t-2}	$\Delta y_{i,t-2}$	-0.0424***	-0.0419***	-0.0398**	-0.0404**
		(0.0147)	(0.0154)	(0.0161)	(0.0164)
ξ_{t-3}	$\Delta y_{i,t-3}$	-0.0338***	-0.0362***	-0.0359***	-0.0361***
		(0.0121)	(0.0125)	(0.0130)	(0.0131)
ξ_{t-4}	$\Delta y_{i,t-4}$	-0.0365***	-0.0396***	-0.0410***	-0.0449***
		(0.0099)	(0.0103)	(0.0107)	(0.0108)
ξ_{t-5}	$\Delta y_{i,t-5}$	-0.0283***	-0.0278***	-0.0252***	-0.0277***
		(0.0080)	(0.0083)	(0.0086)	(0.0087)
ξ_{t-6}	$\Delta y_{i,t-6}$	-0.0249***	-0.0236***	-0.0214***	-0.0218***
		(0.0060)	(0.0062)	(0.0065)	(0.0066)
ξ_{t-7}	$\Delta y_{i,t-7}$	-0.0162***	-0.0146***	-0.0137***	-0.0138***
		(0.0044)	(0.0046)	(0.0049)	(0.0049)
	$\Delta coverage_{it}$	0.0607***	0.0558***	0.0499***	0.0450***
	(RMB 1 M)	(0.0056)	(0.0054)	(0.0052)	(0.0050)
	$\Delta premium_{it}$	-2.23e-4	-3.87e-4	-9.13e-4	-7.92e-4
	(RMB 1 K)	1.37e-3	(1.35e-3)	(1.30e-3)	(1.27e-3)
Individual fixed effects		✓	✓	✓	✓
Calendar year-quarter fixed effects		✓	✓	✓	✓
Fixed effects for quarters since the first contract		✓	✓	✓	✓
N		221,687	221,687	221,687	221,687

This table reports the 2SLS regressions similar to that in the last column of Table 3. The only difference is that only accidents with settlement proceeds above a certain threshold are counted. That is, $y_t = 1$ if policyholder i encountered an accident in quarter t with settlement proceeds above the threshold; $y_t = 0$ otherwise. The dependent variable is Δy_{it}; the endogenous regressor is $\Delta y_{i,t-1}$; and the IV is $\Delta y_{i,t-8}$ and $\Delta y_{i,t-9}$. The threshold is RMB 500, RMB 1,000, RMB 1,500, and RMB 2,000, respectively, in the four columns. The results exhibit similar patterns to the last column of Table 3. The individual fixed effects on the accident probability are implicitly controlled for because they are differenced out. The standard errors in parentheses are clustered by individual. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level. Because the regressors in column 1 have seven lags, observations of a policyholder's first seven quarters in the data are not included in the regression.

Table A.13. Robustness checks: accidents with settlement proceeds below different thresholds

Parameter	Regressor	Threshold of settlement proceed for past accidents			
		≤RMB2500	≤RMB3000	≤RMB3500	≤RMB4000
ξ_{t-1}	$\Delta y_{i,t-1}$	-0.0513**	-0.0562***	-0.0619***	-0.0663***
		(0.0207)	(0.0200)	(0.0198)	(0.0195)
ξ_{t-2}	$\Delta y_{i,t-2}$	-0.0451***	-0.0494***	-0.0529***	-0.0568***
		(0.0172)	(0.0166)	(0.0164)	(0.0161)
ξ_{t-3}	$\Delta y_{i,t-3}$	-0.0362**	-0.0395***	-0.0433***	-0.0475***
		(0.0144)	(0.0140)	(0.0138)	(0.0136)
ξ_{t-4}	$\Delta y_{i,t-4}$	-0.0279**	-0.0313***	-0.0354***	-0.0384***
		(0.0119)	(0.0115)	(0.0114)	(0.0112)
ξ_{t-5}	$\Delta y_{i,t-5}$	-0.0258***	-0.0269***	-0.0305***	-0.0323***
		(0.0095)	(0.0092)	(0.0091)	(0.0090)
ξ_{t-6}	$\Delta y_{i,t-6}$	-0.0192***	-0.0204***	-0.0228***	-0.0249***
		(0.0073)	(0.0070)	(0.0069)	(0.0068)
ξ_{t-7}	$\Delta y_{i,t-7}$	-0.0158***	-0.0168***	-0.0186***	-0.0200***
		(0.0053)	(0.0052)	(0.0051)	(0.0050)
	$\Delta coverage_{it}$	0.0895***	0.0892***	0.0889***	0.0887***
	(RMB 1 M)	(0.0064)	(0.0063)	(0.0063)	(0.0063)
	$\Delta premium_{it}$	-1.93e-3	-1.94e-3	-1.92e-3	-1.91e-4
	(RMB 1 K)	1.49e-3	(1.49e-3)	(1.49e-3)	(1.48e-3)
Individual fixed effects		✓	✓	✓	✓
Calendar year-quarter fixed effects		✓	✓	✓	✓
Fixed effects for quarters since the first contract		✓	✓	✓	✓
N		221,687	221,687	221,687	221,687

This table reports the 2SLS regressions similar to that in the last column of Table 3. The only difference is that, for past accidents in the regressors, only accidents with settlement proceeds below a certain threshold are counted. The dependent variable is Δy_{it}; the endogenous regressor is $\Delta y_{i,t-1}$; and the IV is $\Delta y_{i,t-8}$ and $\Delta y_{i,t-9}$. The threshold is RMB 2500, RMB 3,000, RMB 3,500, and RMB 4,000, respectively, in the four columns. The results exhibit similar patterns to the last column of Table 3. The individual fixed effects on the accident probability are implicitly controlled for because they are differenced out. The standard errors in parentheses are clustered by individual. * denotes significance at a 10% level. ** denotes significance at a 5% level. *** denotes significance at a 1% level. Because the regressors in column 1 have seven lags, observations of a policyholder's first seven quarters in the data are not included in the regression.

Table A.14. Nonlinearity in event probabilities

Event	Probability	% change relative to the former event
No accident for 3 years	0.4049	
No accident for 2 years	0.5473	35.17%
No accident in 1 year	0.7398	35.17%
1 accident in 1 year	0.2258	-69.48%
2 accidents in 2 year	0.0316	-86.01%
3 accidents in 3 year	0.0027	-91.52%

Appendix B. Proof of Proposition

Proposition 1: ∃ a $\tilde{p} \in (0,1)$ such that

$$\frac{d[Prob(y_t = 1|\theta, y_{t-1} = 1, Y_{t-2}) - Prob(y_t = 1|\theta, y_{t-1} = 0, Y_{t-2})]}{dProb\ (\theta_H|Y_{t-2})} \begin{cases} < 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [0, \tilde{p}) \\ \geq 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [\tilde{p}, 1] \end{cases}$$

Proof:

$$\frac{d[Prob(y_t = 1|\theta, y_{t-1} = 1, Y_{t-2}) - Prob(y_t = 1|\theta, y_{t-1} = 0, Y_{t-2})]}{dProb\ (\theta_H|Y_{t-2})}$$

$$= \frac{\gamma\theta}{(\eta - \gamma)(e^{\gamma L} - 1)}$$

$$\times \frac{(\theta_H - \theta_L)^2 \left[Prob\ (\theta_H|Y_{t-2})^2 {\theta_H}^3 (1 - \theta_H e^{-\eta c_{t-1}}) - \left(1 - Prob\ (\theta_H|Y_{t-2})\right)^2 {\theta_L}^3 (1 - \theta_L e^{-\eta c_{t-1}})\right]}{\left[{\theta_H}^2 Prob\ (\theta_H|Y_{t-2}) + {\theta_L}^2 \left(1 - Prob\ (\theta_H|Y_{t-2})\right)\right]^2 [Prob\ (\theta_H|Y_{t-2})(\theta_H - \theta_L)(e^{-\eta c_{t-1}}(\theta_H + \theta_L) - 1) - \theta_L(1 - \theta_L e^{-\eta c_{t-1}})]^2}$$

(B.1)

Because $\eta > \gamma$, $\frac{\gamma\theta}{(\eta-\gamma)(e^{\gamma L}-1)}$ is positive. The denominator of the third line of equation (B.1) is positive. Denote the numerator of the third line of equation (B.1) as $f\left(Prob\ (\theta_H|Y_{t-2})\right)$. It is easy to see that $f\left(Prob\ (\theta_H|Y_{t-2})\right)$ monotonically increases as $Prob\ (\theta_H|Y_{t-2})$ increases. It is also easy to see that $f(0) < 0$ and $f(1) > 0$. From the intermediate value theorem, ∃ a $\tilde{p} \in (0,1)$ such that $f(\tilde{p}) = 0$. Because $f\left(Prob\ (\theta_H|Y_{t-2})\right)$ monotonically increases as $Prob\ (\theta_H|Y_{t-2})$ increases, we have

$$f\left(Prob\ (\theta_H|Y_{t-2})\right) \begin{cases} < 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [0, \tilde{p}) \\ \geq 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [\tilde{p}, 1] \end{cases}$$

Therefore, we have

$$\frac{d[Prob(y_t = 1|\theta, y_{t-1} = 1, Y_{t-2}) - Prob(y_t = 1|\theta, y_{t-1} = 0, Y_{t-2})]}{dProb\ (\theta_H|Y_{t-2})} \begin{cases} < 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [0, \tilde{p}) \\ \geq 0 & if\ Prob\ (\theta_H|Y_{t-2}) \in [\tilde{p}, 1] \end{cases}$$

It is easy to see that the proposition still holds if we add the forgetting component governed by δ in the Bayesian updating process.