Ex Post Moral Hazard and Bayesian Learning in Insurance

Michael Ludkovski[1] and Virginia R. Young

Department of Mathematics
University of Michigan
Ann Arbor, Michigan 48109

E-mail: mludkov@umich.edu, vryoung@umich.edu

[1]corresponding author

EX POST MORAL HAZARD AND BAYESIAN LEARNING IN INSURANCE

Abstract. We study a dynamic insurance market with asymmetric information and *ex post* moral hazard. In our model, the insurance buyer's risk type is unknown to the insurer; moreover, the buyer has the option of not reporting losses. The insurer sets premia according to the buyer's risk rating that is computed via Bayesian estimation based on buyer's history of reported claims. Accordingly, the buyer has strategic incentive to withhold information about his losses. We construct an insurance market information equilibrium model and show that a variety of reporting strategies are possible. The results are illustrated with explicit computations in a two-period risk-neutral case study.

Date: May 20, 2008.

Key words and phrases. Keywords: moral hazard, Bayesian learning, reporting strategies, experience rating, multi-period insurance.

1. Introduction

Buyers of casualty and property insurance possess varying levels of risk. Risk type is used by the insurer to set premia and is the main parameter in pricing the resulting insurance contract. Unfortunately, the intrinsic riskiness of the buyer is typically unknown from the point of view of the insurer. This leads *ex ante* to problems of moral hazard and adverse selection. Namely, higher-risk buyers will attempt to enter into contracts designed for low-risk buyers, and once they obtain insurance, all buyers have little incentive to act prudently. However, many insurance contracts (notably automobile insurance) have a recurring nature. Thus, the issue of information asymmetry is partially mitigated by implementing an *experience-rating* or *bonus-malus system* (see Lemaire (1995) for a comprehensive review of such insurance schemes), through which the insurer gives incentives to the buyer to act in the "best" behavior. Through such incentives, a self-serving buyer may reveal his[1] risk type or exercise an optimal amount of preventive efforts.

A second level of *ex post* information asymmetry arises in connection with reporting losses. After an insurable event occurs, the buyer has the option of *not reporting* the loss in the hopes of signaling that he is of a lower-risk type and, hence, deserves a lower future premium. If the gain from lowering his perceived riskiness (and the corresponding future premia) outweighs the cost for settling the loss out-of-pocket, the buyer will not report the loss. This might happen, for instance, if the risk loading on the insurance is high enough and leads to *ex post moral hazard*. The presence of this non-reporting option has serious implications since it dramatically alters the information received by the insurer. Instead of acting to reveal his risk type, the buyer strategically manipulates reports. In response, the insurer now needs a learning mechanism to infer the risk type of the buyer based on submitted claims. Moreover, a rational insurer recognizes that non-reporting might occur and sets the experience-rating update and premium to account for the optimal reporting strategy of the buyer. The converse of non-reporting is insurance fraud whereby the buyer may manufacture false claims. The information asymmetry of insurance fraud is usually resolved by claim verification and monitoring; by contrast, verifying non-reported losses is usually impractical or against the law.

There is anecdotal empirical evidence for such strategic behavior by both buyers and insurers, especially in the private passenger automobile and homeowner's insurance industries. For instance, after minor car accidents of the "fender-bender" type, it is commonplace that insurance agents advise their clients to pay for repairs themselves and not file a claim, so as to maintain their high "rating." Conversely, it is often observed that the actual malus

[1] We make the standing assumption that the insurer is a *she* and the buyer is a *he*.

penalties for reporting auto claims are relatively severe and would be excessive in a world with perfect reporting (that is, they compensate for the fact that only major accidents are reported).

In this paper, we study the implications of this double information-asymmetry insurance problem in a multi-period model with the insurer learning through Bayesian updating. To focus on the *ex post* moral hazard aspect of the problem, we rule out adverse selection by postulating mandatory insurance (no opt-out clause) and avoid *ex ante* moral hazard by making the claim probabilities completely exogeneous. Mandatory insurance also allows us to sidestep the issue of participation constraint as all buyers are required to nominally subscribe to the insurance market, even if they will then choose to partially self-insure. Instead, we focus on determining the optimal reporting strategies of the buyers of insurance, while assuming that insurers apply the classical Bayesian learning rule to update the premia.

Analysis of games with incomplete information dates back to the pioneering work of Harsanyi (1967, 1968a,b) on Bayesian game theory. In the context of insurance with asymmetric information, recall the seminal paper of Rothschild and Stiglitz (1976) who consider an insurance model with exogenous accident probabilities, constant claim sizes, and publicly known risk preferences of the buyer. In this context, in a one-period model, Rothschild and Stiglitz (1976) find a separating equilibrium in which the insurance contract is designed to *reveal* information, with higher risks buying full coverage at their actuarially fair premium and lower risks buying partial coverage at their actuarially fair premium. Rothschild and Stiglitz (1976) focus on adverse selection, so that the only choice of the buyer is to select among a group of possible insurance contracts.

An extensive literature (Dionne 2002) has also considered moral hazard, whereupon the buyer controls some variable related to his insurance losses, in contrast to adverse selection where the unknown in the model is the static risk parameter. In a typical example, the buyer can exert (unobservable) effort that lowers his accident probability; the aim is, then, to design an insurance contract that forces the buyer to exert a first-best effort level even under the information asymmetry.

The extension of the model of Rothschild and Stiglitz (1976) to multiple periods has been considered in Townsend (1982), Dionne (1983), Dionne and Lasserre (1985), and Cooper and Hayes (1987). The presence of multiple periods leads to questions of contract commitment and renegotiation. With *ex post* moral hazard, a typical buyer cannot commit to more than one period since his reports are intrinsically non-verifiable. Accordingly, with multiple periods, one can focus either on contracts that vary coverage (looking for separating equilibria via the insurer offering a menu of contracts) or on contracts that vary premia. Varying the premia corresponds to an experience rating or bonus-malus mechanism (Hey 1985, Dionne and Vanasse 1992) which is often employed in ratemaking for automobile insurance. The bonus-malus system penalizes occurence of claims thereby providing incentives for prudent behavior.

The second aspect of an experience rating system consists of the insurer attempting to *learn* the type of the buyer over time. The classical example of Bayesian learning has been explicitly considered by Hosios and Peters (1989) and Watt and Vasquez (1997). Hosios and Peters (1989) consider the Rothschild and Stiglitz (1976) framework with varying coverage. In their model, low risks always act myopically, report all claims, and purchase the one-period contract; high risks will choose a randomized reporting strategy. This setup leads to a stochastic game between the monopoly insurer who offers varying coverage and the high risk buyer who strategically reports claims. In contrast to our work, the model of Hosios and Peters (1989) does not require studying the joint interaction of heterogeneous buyers and, hence, an equilibrium strategy always exists. From a different point of view, Watt and Vasquez (1997) show that with multiple periods, a Bayesian bonus-malus contract dominates the multi-period extension of Rothschild and Stiglitz (1976); therefore, the former is a more plausible contract structure (in other words, it is beneficial to vary premia rather than coverage). However, Watt and Vasquez (1997) assume that no underreporting is possible and, thus, treat only the case of adverse selection.

Our work is most allied with a recent paper by Robinson and Zheng (2006) who consider the effect of experience-rating on precautionary behavior by insurance buyers. Robinson and Zheng (2006) study a two-period model with compulsory insurance and high- and low-risk types, where the buyer faces both *ex ante* moral hazard (through an unobservable effort variable) and *ex post* moral hazard (through a possibility of not reporting a claim). In addition, Robinson and Zheng (2006) also assume that the buyers do not know their own type; moreover, the experience rating is solely based on the occurrence of a claim during period one (thus, there is no "prior" risk type). The authors show that even if the effort is unobservable by the insurer, the buyer will engage in precaution in order to maintain a good rating and avoid future costs, mitigating *ex ante* moral hazard, while creating *ex post* moral hazard.

Starting with the same framework, we choose instead to focus solely on *ex post* moral hazard. Accordingly, we drop the option of the buyer to influence accidence probabilities and make the latter completely exogenous. Moreover, in contrast to Robinson and Zheng (2006), we assume that the buyers are aware of their risk type and, therefore, *strategically* manipulate their reporting. In order to focus on underreporting, we rule out the other mentioned type of *ex post* moral hazard, namely insurance fraud; see Moreno et al. (2006) for a study of effect of a bonus-malus contracts on fraud. As in Hosios and Peters (1989), we assume that the insurer uses a Bayesian experience rating that depends *both* on the latest claim report and on the prior of the buyer's risk type. Thus, the insurer has initial ratings assigned to each buyer, which are public knowledge, and the buyer can manipulate his experience rating by reporting or not reporting losses, modeled as a mixed game strategy (in contrast to only pure strategies considered in Robinson and Zheng (2006)). The explicit model of the learning mechanism of the insurer allows us to compute the optimal reporting strategies in closed-form.

In our setting, existence of equilibrium is non-trivial since it requires a compatability between the optimal behavior for buyers of all possible risk types (in terms of resulting randomized reporting strategies) and the global experience-rating mechanism employed by the insurer across all (unknown to her) risk types of buyers and their strategies. One of the main contributions of our model is showing that this compatability is non-trivial and can sometimes fail. Hence it is possible that for some buyers the insurer has no consistent mechanism to update premia and ratings. Moreover, the model features abrupt phase transitions between different reporting strategies as a function of perceived buyer risk.

Finally, we mention some recent empirical literature that attempts to test the significance of various information asymmetries in actual insurance markets around the world. Of particular note are the recent studies of Abbring et al. (2003), Chiappori et al. (2006) and Abbring et al. (2007) who consider the automobile insurance markets in France and the Netherlands, respectively. The Netherlands system, in particular, incorporates a complicated bonus-malus class structure that can be viewed as a proxy for a Bayesian updating mechanism. The cited papers find conflicting evidence of *ex ante* and *ex post* moral hazard and discuss the econometric difficulties of disentangling the multiple information asymmetries. We point out that in those studies, a key emphasis is made on the claim *severity*, a feature that is absent in our model where all claims are of unit size.

The structure of this paper is as follows. In Section 2, we describe the multi-period insurance game that we study; then, in Section 3, we present the mathematical details of our model. In Sections 4 and 5, we study a two-period example in detail and show that an equilibrium might not exist even in this simple case. We also study the types of equilibria that can occur and show that in the first period of the two-period example, low risks either always report or never report their losses, while high risks report their losses with a positive probability. In other words, high risks might try to make themselves look more like low risks by hiding some of their losses, but they never completely mimic low risks by never reporting their losses. In Section 6, we discuss how the equilibria found in the two-period model might change in a more general setting, for example, under a non-linear premium rule or with more than two periods. Section 7 concludes the paper.

2. Insurance Market with *ex post* Moral Hazard

In this section, we provide an economic description of the dynamic insurance market model we consider. The precise mathematical construction is then presented in Section 3. Our notation follows Rothschild and Stiglitz (1976) and Hosios and Peters (1989).

2.1. **Information Asymmetries.** To focus on the stated information asymmetry, we assume that the buyer has no control over his own risk type (that is, no *ex ante* moral hazard). Moreover, the buyer acts rationally in that he chooses a *reporting strategy* to minimize his total expected costs over a given finite time period. For simplicity, all losses are of fixed size 1, so that the issue of partial reporting is non-existent. We also assume that the buyer can hide losses but cannot manufacture them, nor can the insurer discover losses that are not reported. In other words, a reported claim is verifiable; a non-reported one is not.

The main level of information asymmetry is that the buyer is aware of his risk type $X \in E$; in contrast, the insurer never knows X precisely (unless the buyer by his actions reveals X) but starts with a prior distribution $q(0)$ at time 0 over possible values of X. Thus, from the insurer's point of view, X is a random variable, whose distribution is updated in response to incoming insurance information. Buyers differ only in their probability of a loss; a buyer of risk type $X = x$ has probability of loss equal to p_x.

Since the buyer knows his type (that is, his probability of loss), he can strategically manipulate the insurer's beliefs in the manner most advantageous to him. We denote by $q(k)$ the insurer's beliefs at time k, with the space of values of $q(k)$ denoted by

$$D = \{(q^1, q^2, \dots, q^d) \in \mathbb{R}^d \colon \sum_{x \in X} q^x = 1,\ 0 \le q^x \le 1\}, \quad d = |E|.$$

Here, q^x is the probability that a risk type is of type $x \in X$ according to the beliefs of the insurer. We write $q^x(k)$ if we want to emphasize that the belief is held at time k, that is, at the beginning of period $(k+1)$. We assume that at time zero, the buyer knows his own initial rating $q(0)$.

For a given rating $q = q(k-1)$ at the beginning of period k, denote by $d_X(q,k)$ the probability that a realized loss will be reported to the insurer at the end of period k, that is, at time k. Therefore, from the insurer's point of view, the probability that a claim is filed at time k equals $\sum_{x \in E} q^x(k-1)\, p_x\, d_x(q,k)$.

2.2. **Insurance Market.** The insurance market operates as follows. The description below is based on a given fixed collective reporting strategy (across all risk types x), $\vec{d} \equiv \{\vec{d}(q,k)\} = (d_x(q,k))_{x \in E}$, for $q \in D$ and $k \in \{1, \dots, T\}$, and a fixed insurer ratemaking policy. Since verifying non-reported claims is impossible, the buyer cannot credibly commit to a multiple-period contract. Consequently, we only consider one-period contracts with the different periods linked by reputational considerations (insurer beliefs). The insurance contract consists of an update mechanism G that implements the experience-learning scheme and a (contingent) premium schedule $f(q(k-1), \vec{d}(\cdot, k))$ for the kth period that determines the premium payable at time $k-1$ based on the current rating *and* the heterogeneous reporting strategy that will be implemented at time k. The insurer anticipates partial non-reporting and will lower upcoming premia to reflect lower insurance costs, thus making "lying" more profitable for the buyer.

At the beginning of each period $k = 1, 2, \dots, T$, the premium $f(q(k-1), \vec{d}(q(k-1), k))$ is first collected. Second, the buyer either does or does not experience an insurance event (recall that all losses are of size 1). The occurrence of a loss is modeled as an independent Bernoulli trial with probability p_X of occurring, the probability being a function of the true risk type X. If a loss occurs during period k, then the buyer has the option of (a) reporting the 'claim' and receiving reimbursement at time k, or (b) settling the loss himself and reporting 'no claim' at time k. If no loss occurs, the buyer reports 'no claim.' From the insurer's point of view, the latter outcome is identical to case (b) above. The reporting strategy of the buyer is represented by a randomization policy with parameter $d_X(q(k-1), k)$: conditional on the type X of the buyer and an event occurring in the kth period, the buyer performs an independent Bernoulli trial (hidden from the insurer), and with probability $d_X(q(k-1), k)$ reports the event (case (a)) to the insurer at time k. Thus, with probability $1 - d_X(q(k-1), k)$ the event is not reported and case (b) takes place.

Let $Y_k \in \{0, 1\}$ represent the indicator of a reported claim at the end of the kth period (time k) and denote by $\mathcal{F}_k = \sigma\{Y_s \colon s = 1, 2, \dots, k\}$ the claim information received by the insurer by the end of period k. For typographical convenience, we will write $\vec{d}(k)$ for $\vec{d}(\cdot, k)$. Based on the buyer's announcement of 'claim' or 'no claim' (as well as the collective reporting strategy $\vec{d}(k)$) the insurer now updates the buyer's rating using a Markovian update mechanism $G(q(k-1), Y_k; \vec{d}(k))$. More directly, we will write $q^+(q(k-1), \vec{d}(k)) = G(q(k-1), 1; \vec{d}(k))$ and $q^-(q(k-1), \vec{d}(k)) = G(q(k-1), 0; \vec{d}(k))$ to indicate the updated ratings in the event of a 'claim' ('no-claim') report. We reiterate that premia are paid at the beginning of each period and claims are settled at the end; thus at time k (the end of the kth period), the premium $f(q(k), \vec{d}(q(k), k+1))$ for period $(k+1)$ and any losses for period k (either self-absorbed or reimbursed by the insurer) are settled simultaneously.

After all these adjustments, the market moves to the next round. After the last round, the Tth, the game stops (that is, no more insurance is issued). The overall mechanism is illustrated in Figure 1.

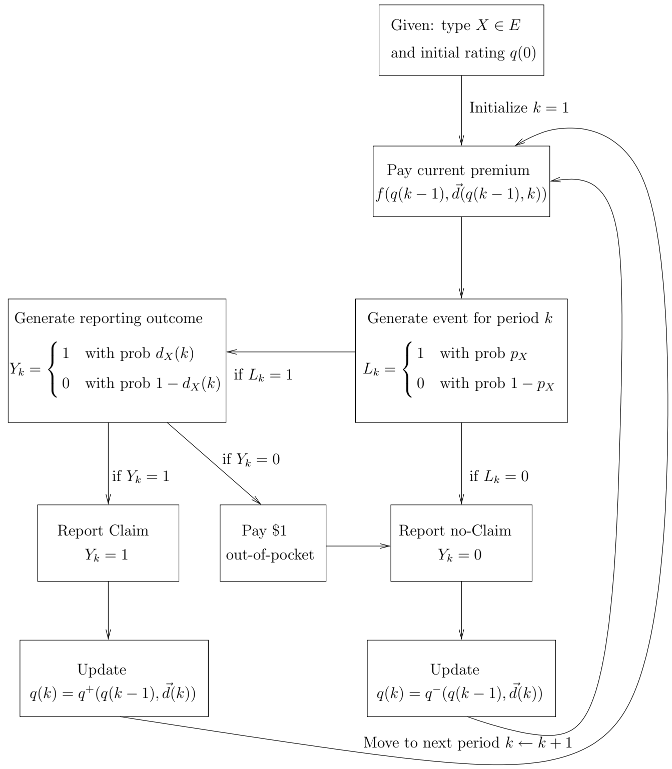

FIGURE 1. Flow chart illustrating the insurance market.

2.3. **Insurance Market Equilibrium.** In order to create an equilibrium, two items are now necessary:

(1) The seller must *ex ante* commit to a verifiable and public update mechanism G regarding ratings.
(2) The buyers must be able to collectively "announce" their entire reporting strategy $(\vec{d}(\cdot, k)_{k=1,\dots,T})$ at time 0, such that it is credible. That is, for any buyer of type $X = x$, the reporting frequency $d_x(\cdot, k)$ is optimal at period k assuming all other buyers implement $\vec{d}$. In other words, deviation from the collective $\vec{d}$ must be sub-optimal for any given buyer.

Our first assumption is easily enforced on the grounds that an insurer is regulated and, hence, faces mandatory rules regarding ratings. Alternatively, there is a competitive insurance market with public ratings that weeds out any unreliable insurers. This situation is generally true in practice where experience rating systems are publicly announced and always honored for settling claims. They also tend to be "sticky" over time, so that changes to the update mechanism are made rarely.

The requirement of collective compatability is clearly necessary to prevent deviations from buyers (as verifying their reporting strategy is obviously impossible). The aspect of the buyers collectively "announcing" their reporting strategy is to resolve an issue of potentially several co-existent equilibria. The insurer must know the prevailing equilibrium; otherwise, she does not have a consistent way of updating her beliefs as she does not know the precise information set with which she is working. As we will show below, in some cases it is possible to prove that at most one equilibrium exists and, therefore, no public announcement is necessary by a rational buyer. With a unique equilibrium, the insurer can directly infer the only collectively compatible strategy $\vec{d}^*$.

Because premia are mandatory and paid at the beginning of each period, the reporting strategy is required to be *ex post* optimal. Thus, even though the period-k premium $f(q(k-1), \vec{d}(k))$ is affected by the reporting strategy $\vec{d}(k)$ employed at time k, from the point of view of the buyer who effectively makes his reporting decision at time k, this premium is a sunk cost. Indeed, at the beginning of a period, the buyer may commit to a reporting strategy at the end of that period *only if* it is *ex post* optimal once premium is paid. This fact is most starkly manifested during the last period. At time $T-1$, a rational buyer may be better off self-insuring and, therefore, could claim that his premium should be zero. However, insurance is mandatory, so once premium is paid, the buyer clearly has no incentive to self-insure at time T and will report all claims. Hence, in the last period the only *ex post* optimal strategy is full reporting.

One key feature of our model postulates that the updating of the buyer's rating is based on a Bayesian mechanism (Hosios and Peters 1989, Watt and Vasquez 1997); that is, given beliefs $q(k-1)$ in the beginning of the kth period, the announced reporting strategy $\vec{d}(k)$ for period k, and the the actual report Y_k at time k, the new insurer beliefs $q(k)$ are computed using a relative-likelihood approach based on conditional probabilities of respective events. Namely, $q(k) \triangleq \mathbb{P}^{q(k-1)}(X|Y_k, \vec{d}(k))$. Thanks to this non-linear mechanism and to the second level of information asymmetry that gives the buyer control over this updating (via the reporting strategy), the resulting insurance market exhibits a rich and interesting structure. Other update mechanisms are also possible (for example, a non-Markovian mechanism, a moving-average mechanism such as $q(k) = \beta q(k-1) + (1-\beta)Y_k$, etc.).

We choose classical Bayesian updating because (a) the resulting $q(k)$ is *the* vector of conditional probabilities for the buyer type given the reporting strategy $\vec{d}$; and (b) $q(k)$ is non-linear in the prior rating $q(k-1)$, which results in non-trivial reporting strategies even under linear preferences, as we show later. Property (a) justifies Bayesian updating from the point of view of learning; it is well known that conditional probabilities minimize the square-error loss, so that an insurer that wishes to best estimate the buyer's risk type will use this method (see Mikosch (2004) for references). Features of Bayesian updating can be detected in some of real-world experience rating systems. For instance, in the Netherlands car insurance market, Abbring et al. (2007) report that the updates are more dramatic for middle rating classes than for the extreme classes, as is also obtained in a Bayesian setting.

Remark 1. We argue that a straight bonus-malus system is not consistent with our asymmetric information model. Indeed, because the update in bonus-malus only depends on the previous experience rating and the latest claim information, it fails to take into account the different strategic behaviors of heterogeneous buyers. A bonus-malus model assumes that the reporting strategy of a particular buyer is independent of the strategies of the other buyers; thus, it ignores the fact that the observed probability of claims is a function of the reporting strategies across buyers. In other words, a consistent update mechanism G *must* be a function of all $d_x(k)$, $x \in E$, since it should weigh the latest claim observation by its relative likelihood for each buyer type.

2.4. **Markov Strategies.** In general, the reporting strategy $\vec{d}(k)$ may be a function of all the past information (in particular, the report history), rather than just a function of the latest rating $q(k-1)$. Our Markovian assumption can be justified by assuming that the report history is not public knowledge (with verification costly or impossible), so that the only credible strategy announcements for reports in the kth period must be based on current rating $q(k-1)$, which *is* public knowledge. This parallels the real-life setting of confidential credit histories that are summarized with a single public credit score for any commercial credit enquiry.

Also, the premium rule f may also be a function of the report history, that is, a minimal requirement is that the premium paid at time k be $\mathcal{F}_k$-adapted. Again, such a possibility can be ruled out that if there is a competitive insurance market with different insurers present at each period. If the report histories are not public knowledge, premia must be based only on the latest rating and implicitly on the anticipated reporting strategy. For example, in private passenger automobile insurance in the United States, insurers and other economic agents rely on a demerit or points system, which is a proxy for risk rating.

3. Probabilistic Model

In this section, we provide the mathematical details of our model.

3.1. **Bayesian Learning.** For simplicity, we consider the two-state case, whence $X \in E \triangleq \{L, H\}$ so that risks are either of low or high type. Let $0 < p_L < p_H \leq 1$[2] be the time-homogeneous probabilities of an insurance event for the respective risk types.

[2] If $p_L = p_H$ then the two risk types are the same, thereby making the problem degenerate.

With a slight abuse of notation relative to the previous section, in this section and the next, let $q(k) \equiv q^H(k) = \mathbb{P}(X = H|\mathcal{F}_k)$ be the belief of the insurer concerning whether the buyer is of high type based on the information available at time k. If the time k is irrelevant or understood, then we also write q for $q(k)$. Also, let $(d_L(k), d_H(k))$ be the reporting strategy of the buyers in the kth period conditional on $q(k-1)$. Recall that $Y_k \in \{0,1\}$ indicates whether a loss is reported at time k. Then, the probability of a reported claim at time k equals

$$\mathbb{P}(Y_k = 1) = p_H \cdot d_H(k) \cdot q(k-1) + p_L \cdot d_L(k) \cdot (1 - q(k-1));$$

therefore, by Bayes' formula the updated beliefs at the end of the kth period $q^+(k)$ and $q^-(k)$ conditional on whether a claim was reported are $G(q, d_L(k), d_H(k)) =$

$$\begin{cases} q^+(d_L, d_H) \triangleq \mathbb{P}^q(X = H|Y_k = 1) = \dfrac{p_H \cdot d_H \cdot q}{p_H \cdot d_H \cdot q + p_L \cdot d_L \cdot (1-q)}, \\ q^-(d_L, d_H) \triangleq \mathbb{P}^q(X = H|Y_k = 0) = \dfrac{(1 - p_H \cdot d_H) \cdot q}{(1 - p_H \cdot d_H) \cdot q + (1 - p_L \cdot d_L) \cdot (1-q)}. \end{cases} \tag{1}$$

Again, we stress that the update mechanism, $G : (q, d_L, d_H) \in D \times [0,1]^2 \to (q^-, q^+) \in D^2$, depends on the announced strategy $(d_L(k), d_H(k))$ since it is based on the relative likelihood of the corresponding risk types filing a claim. The Bayesian update G in (1) is non-linear and asymmetric in the low- and high-risk type parameters. Also, note that if $d_L = 0$ (resp. $d_H = 0$), then a claim is a clear signal that the buyer is of high type, $q^+ = 1$ (resp. low type, $q^+ = 0$); this is the only situation in which the buyer anambiguously reveals his type.

3.2. **Objective of the Buyer.** Let T be the finite horizon of the insurance game, and let $v \leq 1$ be the discount factor for future payments. We assume that the buyer possesses time-additive risk-averse (or risk-neutral) preferences with respect to insurance costs. Let U be the convex loss function of the buyer, with $U(0) = 0$. Let $L_k \in \{0,1\}$ be the actual realized loss in period k (as opposed to Y_k reported to the insurer; note that $Y_k \leq L_k$). Then, at time k (end of the kth period), conditional on his true risk type X, the buyer's risk-adjusted costs are:

$$C_k = \begin{cases} U\left(f(q^+(k), \vec{d}(q^+(k), k+1))\right) & \text{with prob. } p_X d_X(k), \\ U\left(1 + f(q^-(k), \vec{d}(q^-(k), k+1)\right) & \text{with prob. } p_X(1 - d_X(k)), \\ U\left(f(q^-(k), \vec{d}(q^-(k), k+1)\right) & \text{with prob. } 1 - p_X. \end{cases} \tag{2}$$

As argued before, since the market terminates at time T, it is *ex post* optimal to report all the period-T losses; thus, we assume henceforth that $L_T = Y_T$, which implies that the time-T costs are $U(0) = 0$. Recall that the buyer selects his reporting policy just before the end of the k-period, at time $k-$. The losses are then settled at time k, and immediately the rating is updated to $q(k)$. Thus, for a given reporting strategy $\vec{d}$, at time $k-$ the total discounted expected costs from time k through time T for a buyer of type X are

$$V_X(k, q(k-1))(\vec{d}) \triangleq \sum_{n=k}^{T-1} v^{n-k} \mathbb{E}^{q(k-1)} \Big[U \Big(L_n - Y_n + f(q(n), \vec{d}(q(n), n+1) \Big) \Big], \tag{3}$$

with the premium at time T understood to be 0. Note that the premium for period k, namely $f(q(k-1), \vec{d}(q(k-1), k))$, is not part of $V(k, q(k-1))$ since by time $k-$ it is a sunk cost and does not enter into the optimization problem that we introduce below.

By using the time-additivity of utility, (3) can be recursively written as $V_X(T, q) = 0$ and

(4)
$$\begin{aligned} V_X(k, q(k-1))(\vec{d}) = p_X \Big[d_X(k) \Big(U(f(q^+(k), \vec{d}(q^+(k), k+1))) + v V_X(k+1, q^+(k))(\vec{d}) \Big) \\ + (1 - d_X(k)) \Big(U(1 + f(q^-(k), \vec{d}(q^-(k), k+1))) + v V_X(k+1, q^-(k))(\vec{d}) \Big) \Big] \\ + (1 - p_X) \Big[U(f(q^-(k), \vec{d}(q^+(k), k+1))) + v V_X(k+1, q^-(k))(\vec{d}) \Big]. \end{aligned}$$

The first term represents the value in the case of a report, the second term represents the cost of the self-imposed loss that occurs when an event is not reported, and the last term represents the expected value in the case of no-loss. The values $(q^-(k), q^+(k)) = G(q(k-1), d_L(k), d_H(k))$ are obtained from (1).

A rational buyer chooses his reporting strategy to minimize expected future costs. However, if the buyer is of type X, his control is just $\{d_X(q,k)\}_{k=1,\ldots,T}$, while V_X depends on the entire collective $\vec{d}$. Thus, a consistency condition is needed, namely that

$$\begin{cases} d_L^*(q,k) \in \arg\min_{d_L} V_L(k,q)(d_L, d_H^*(k), \vec{d}^*(k+1), \ldots, \vec{d}^*(T)), \\ d_H^*(q,k) \in \arg\min_{d_H} V_H(k,q)(d_L^*(k), d_H, \vec{d}^*(k+1), \ldots, \vec{d}^*(T)) \end{cases} \tag{5}$$

simultaneously, so that a given buyer has no rational reason to deviate from the announced collective strategy $\vec{d}^*(q,k)$ regardless of his true risk type. In other words, for a fixed $d_H^*(q,k)$ (resp. $d_L^*(q,k)$), if the buyer is low risk (resp. high risk), then it is optimal for him to employ the reporting strategy corresponding to $d_L^*(q,k)$ (resp. $d_H^*(q,k)$) at time k, and $d_L^*(\cdot, k+1), \ldots, d_L^*(\cdot, T)$ (resp. $d_H^*(\cdot, k+1), \ldots, d_H^*(\cdot, T)$) thereafter.

The insurance market will exhibit a stable equilibrium if the above collective-compatability condition (5) is satisfied for exactly *one* pair $(d_L^*(k), d_H^*(k))$. In such a case we say that an *equilibrium* exists, and we define the *value function* of a buyer of type X at time $k-$ to be

$$W_X(k,q) \triangleq V_X(k,q)(\vec{d}^*(k), \ldots, \vec{d}^*(T)) \tag{6}$$

for a given rating $q = q(k-1)$, $k = 1, \ldots, T$.

4. Case Study

We now study a simple example in which there are two periods, the buyers are risk neutral, and the premium equals the expected discounted payout with a proportional risk load. Specifically $T = 2$, the loss function U is given by $U(x) = x$, and the time-k premium f is given by

$$f(q(k), \vec{d}(q(k), k+1)) = (1+\theta)v\left[p_H d_H(k+1)q(k) + p_L d_L(k+1)(1-q(k))\right], \tag{7}$$

for some proportional risk load $\theta \geq 0$ and for $k = 0, 1$. For simplicity, we assume that the discount v used by the buyer and the insurer are identical. Any difference between them could be absorbed in the factor of $(1+\theta)$ by allowing θ to take values greater than -1 instead of 0.

Since there are only two periods, in the second (last) period the buyer has no incentive to non-report, so that trivially $(d_L^*(q,2), d_H^*(q,2)) = (1,1)$. In the first period, the premium paid is a function of the insurer's prior $q(0)$ and the reporting strategy at time 1; however, at time 1, the buyer chooses his strategy to minimize future expected costs, so the premium paid at time 0 is irrelevant at that time. Note that the premium paid at time 1 equals $f(q(1), \vec{d}(q(1), 2)) \equiv f(q) = (1+\theta)v(p_H q + p_L(1-q)) = (1+\theta)vp_L + \alpha^2 q$, in which $\alpha \triangleq \sqrt{(1+\theta)v(p_H - p_L)}$. In other words, the insurer charges a base premium $(1+\theta)vp_L$ plus a surcharge proportional to the updated risk rating $q = \mathbb{P}^{q(0)}(X = H|\mathcal{F}_1)$ of the agent, the proportionality constant being α^2.

Henceforth, we focus on the expected costs of the low and high risks at time 1. For typographical convenience, we omit time-dependence and q-dependence of the strategies at time 1 and write $d_H = d_H(\cdot, 1)$ and $d_L = d_L(\cdot, 1)$. From (4), the expected costs at time 1 for a risk of type $X = L, H$ equals

$$\begin{aligned} V_X(q)(d_L, d_H) &= p_X(1-d_X) + (1-p_X d_X)f(q^-) + p_X d_X f(q^+) \\ &= \alpha^2 q^- + \alpha^2 p_X\, d_X\,(q^+ - q^-) - p_X\, d_X + \bar{C}_X, \end{aligned} \tag{8}$$

in which $q = q(0)$, the updated beliefs $(q^-, q^+) = G(q, d_L, d_H)$ are defined in (1), and the constant $\bar{C}_X = p_X + (1+\theta)vp_L$ is independent of (d_L, d_H).

4.1. **Uniqueness of Equilibrium.**

Our first result shows that at most one equilibrium (d_L^*, d_H^*) satisfying (5) exists in this case.

Lemma 1. *The function V_L is concave in the reporting strategy d_L; the function V_H is convex in the reporting strategy d_H.*

Proof. We directly compute

(9)
$$\frac{\partial V_L}{\partial d_L} = -p_L + \alpha^2 p_L q^2 \left\{ \frac{(p_H d_H)^2}{[p_H d_H q + p_L d_L (1-q)]^2} - \frac{(1-p_H d_H)^2}{[(1-p_H d_H)q + (1-p_L d_L)(1-q)]^2} \right\},$$

which is a strictly decreasing function of d_L. Indeed, the first term in the brackets is of the form $C_1/(C_2+d_L)^2$, and the second term is of the form $-C_3/(C_4-d_L)^2$, for some constants $C_i \geq 0$.

Similarly,

(10)
$$\begin{aligned}\frac{\partial V_H}{\partial d_H} &= -p_H + \alpha^2 p_H (1-q)^2 \left\{ \frac{(1-p_L d_L)^2}{[(1-p_H d_H)q + (1-p_L d_L)(1-q)]^2} - \frac{(p_L d_L)^2}{[p_H d_H q + p_L d_L (1-q)]^2} \right\} \\ &= \frac{C_1}{(C_2 - d_H)^2} - \frac{C_3}{(d_H + C_4)^2} + C_5,\end{aligned}$$

for some $C_i \geq 0$; thus, $\frac{\partial V_H}{\partial d_H}$ is a strictly increasing function of d_H. □

The above lemma immediately implies the following key theorem.

Theorem 1. *The optimal reporting strategy for low risks satisfies $d_L^* \in \{0, 1\}$. More precisely, by comparing V_L at $d_L = 0$ with the cost at $d_L = 1$, we obtain that low risks report all claims, namely $d_L^* = 1$, if and only if*

(11)
$$1 \geq \frac{\alpha^2 q(1-q)}{(1-p_H d_H)q + (1-q)} \left\{ \frac{(1-p_L)(1-p_H d_H)}{(1-p_H d_H)q + (1-p_L)(1-q)} + \frac{p_H d_H - p_L(1-p_H d_H)}{p_H d_H q + p_L(1-q)} \right\}.$$

Thus, for the low-risk types, pure reporting strategies (either always report or never report) are the only optimal strategies. On the other hand, since V_H is convex, typically d_H^* will be strictly between zero and one, so high-risk types use a randomized strategy.

Corollary 1. *In equilibrium, the high-risk buyer reports some of his losses, that is, $d_H^* > 0$.*

Proof. The possibility that $d_L^* = 0$ and $d_H^* = 0$ is disallowed since in that case, no claims are ever reported and the insurer has no mechanism to update beliefs. In particular, q^+ is undefined, $q^- = q$, and (11) holds because the right-hand side is negative when $d_H = 0$, which contradicts $d_L^* = 0$. So, we can assume $d_L^* > 0$, and by evaluating (10), we obtain

$$\begin{aligned}\left.\frac{\partial V_H}{\partial d_H}\right|_{d_H=0} &= -p_H + \alpha^2 p_H (1-q)^2 \left\{ \frac{(1-p_L d_L)^2}{[q + (1-p_L d_L)(1-q)]^2} - \frac{(p_L d_L)^2}{[p_L d_L (1-q)]^2} \right\} \\ &= -p_H + \alpha^2 p_H \left\{ \frac{[(1-p_L d_L)(1-q)]^2}{[q + (1-p_L d_L)(1-q)]^2} - 1 \right\} < 0.\end{aligned}$$

Thus, as a function of d_H, V_H is initially decreasing and achieves its minimum at a point strictly greater than 0. □

This corollary relies on the assumption that $p_L > 0$. See Section 5.2 for study of equilibria in the degenerate case for which $p_L = 0$.

The next two lemmas study the case $d_L^* = 0$ and $0 < d_H^* < 1$, whereby low risks self-insure during period 1 and high risks employ a mixed reporting strategy.

Lemma 2. *The optimal strategy is*

$$d_L^* = 0, \quad \textit{and} \quad 0 < d_H^* = \widetilde{d_H} \triangleq \frac{1-(1-q)\alpha}{p_H q} < 1,$$

if and only if $\alpha > 1$ *and the initial rating satisfies* $\tilde{q} < q < \frac{\alpha-1}{\alpha-p_H}$, *in which*

$$\tilde{q} \triangleq 1 - \frac{\alpha-1}{\alpha-p_L} \cdot \frac{2\alpha - p_L - \alpha \cdot p_L}{2\alpha(\alpha-1)+p_L}. \tag{12}$$

If $\tilde{q} > \frac{\alpha-1}{\alpha-p_H}$, *then there are no prior q's such that* $d_L^* = 0$ *and* $0 < d_H^* < 1$ *is optimal.*

Proof. Assuming that optimal d_H^* is strictly between 0 and 1 implies that the first-order condition for V_H must hold. Therefore

$$\begin{aligned} 0 = \frac{\partial V_H}{\partial d_H}\Big|_{d_H = d_H^*, d_L = 0} &= -p_H + \alpha^2 p_H (1-q)^2 \frac{1}{[(1-p_H d_H^*)q + (1-q)]^2} \\ \Leftrightarrow \qquad 1 &= \frac{\alpha^2(1-q)^2}{(1-p_H d_H^* q)^2} \\ \Leftrightarrow \qquad d_H^* = \widetilde{d_H} &= \frac{1-(1-q)\alpha}{p_H q}. \end{aligned}$$

To satisfy the constraint $0 < \widetilde{d_H} < 1$, one must have $\alpha > 1$ and $\frac{\alpha-1}{\alpha} < q < \frac{\alpha-1}{\alpha-p_H}$. At the same time, to satisfy $d_L^* = 0$ at $d_H^* = \widetilde{d_H}$, we must have

$$\begin{aligned} & 1 < \frac{\alpha^2 q(1-q)}{(1-p_H\widetilde{d_H})q + (1-q)} \left\{ \frac{(1-p_L)(1-p_H\widetilde{d_H})}{(1-p_H\widetilde{d_H})q + (1-p_L)(1-q)} + \frac{p_H\widetilde{d_H} - p_L(1-p_H\widetilde{d_H})}{p_H\widetilde{d_H}q + p_L(1-q)} \right\} \\ & \Leftrightarrow 1 < \alpha \left[\frac{(1-p_L)(\alpha-1)}{\alpha-p_L} + \frac{1-(\alpha-p_L+\alpha p_L)(1-q)}{1-(\alpha-p_L)(1-q)} \right] \\ & \Leftrightarrow \tilde{q} < q. \end{aligned}$$

It can be shown that $\tilde{q} > \frac{\alpha-1}{\alpha}$ always, so that the lower bound for both conditions to hold is $\tilde{q} < q$ and the lemma follows. □

Note that for $\alpha - 1$ small enough, $\tilde{q} > \frac{\alpha-1}{\alpha-p_H}$; therefore, the two conditions above cannot be simultaneously satisfied, so that the scenario $d_L^* = 0$ and $0 < d_H^* < 1$ is ruled out.

Lemma 3. *If* $d_L^* = 0$ *and* $0 < d_H^* < 1$, *then* $W_L(1, q) = p_L - \alpha + (1+\theta)v p_H$.

Proof. From (8) with $X = L$, one can check that $V_L(1,q)(d_L = 0, d_H = \widetilde{d_H})$ simplifies to the above expression. □

4.2. **Perfect Reporting.** We term the case $(d_L^*, d_H^*) = (1,1)$ the first-best outcome since it represents the situation for which the buyer has no incentive to withhold information and, therefore, avoids *ex post* moral hazard. When $d_X^* < 1$, the implicit deductibles are inefficient and are an externality of the information asymmetry; that is, the buyer pays the losses himself because the insurer cannot efficiently monitor his type. Thus, the presence of an equilibrium with $d_X^* < 1$ for either $X = L$ or H is a sign of something "breaking down" in the system. Accordingly, it is important to understand the parameter values that lead to full reporting and avoid such externalities.

Lemma 4. *We have that $(d_L^*, d_H^*) = (1,1)$ for any parameter values when q is sufficiently small (close to zero) or sufficiently large (close to one). Moreover, the interval of q's for which $(d_L^*, d_H^*) = (1,1)$:*

- *Increases in width as the risk loading θ decreases.*
- *Increases in width as the high-type riskiness p_H decreases;*
- *Increases in width as the low-type riskiness p_L increases;*

Proof. In order for $(d_L^*, d_H^*) = (1,1)$, it is necessary and sufficient that (11) hold with $d_H = 1$ and that (10) be non-positive at $(d_L, d_H) = (1,1)$. By substituting $(d_L, d_H) = (1,1)$ into (11) and (10), these two conditions become, respectively,

$$1 \geq I_L, \quad \text{and } 1 \geq I_H, \tag{13}$$

in which

$$I_L = \frac{\alpha^2 q(1-q)}{1-p_H q}\left[\frac{(1-p_L)(1-p_H)}{1-p_L-(p_H-p_L)q} + \frac{p_H - p_L(1-p_H)}{p_L+(p_H-p_L)q}\right], \tag{14}$$

and

$$I_H = \alpha^2(1-q)^2\left[\frac{(1-p_L)^2}{[1-p_L-(p_H-p_L)q]^2} - \frac{p_L^2}{[p_L+(p_H-p_L)q]^2}\right]. \tag{15}$$

It is an easy check that as $q \to 0+$ or $q \to 1-$, I_L and I_H converge to zero; therefore, because I_L and I_H are continuous functions of q, the inequalities in (13) are satisfied for $q \sim 0$ and $q \sim 1$.

Moreover, since $\alpha^2 = (1+\theta)v(p_H - p_L)$ is a linear multiplier in I_L and I_H, smaller values of θ decrease I_L and I_H, thereby, making the inequalities in (13) more likely. Similarly, one can show that the derivatives of I_L and I_H are positive with respect to p_H and negative with respect to p_L leading to the comparative statics stated in the lemma. □

The economic intuition behind these results is that when q is very small or very large, the gain from non-reporting is minimal since it is very hard to change the beliefs of the insurer (that is, q^-, q^+ are very close to q); therefore, full reporting is best. This is consistent with empirical observations that show that very low-risk and very high-risk buyers typically report all their accidents since they "have nothing to lose." For instance, in the Netherlands auto insurance study by Abbring et al. (2007), the premia of the safest drivers (experience rating 20) and of the worst drivers (experience rating 1) do not increase after an accident; thus, there is essentially no gain from under-reporting.

Similarly, if the risk loading θ is relatively small, the benefit from non-reporting (through lowering one's future premium) is not enough to outweigh the cost that comes from concealing information about one's type, so full reporting is optimal. Also, if the difference $p_H - p_L$ shrinks, then there is less opportunity to distinguish high and low types and, therefore, less gain from manipulating $q^{\pm}$ via non-reporting.

Our second result, which uses the above proof, shows that the buyer only has incentive to settle out-of-pocket (self-insure) when the risk loading is sufficiently large.

Corollary 2. *$(d_L^*, d_H^*) = (1, 1)$ for all $q \in [0, 1]$ whenever $\alpha^2 = (1 + \theta)v(p_H - p_L) \leq 1$.*

Proof. I_L in (14) is a decreasing function of p_L and is an increasing function of α. Moreover, when $p_L = 0$ and $\alpha = 1$, I_L equals $\frac{1-q}{1-p_H q}\left[\frac{q(1-p_H)}{1-p_H q} + 1\right]$, which is less than or equal to 1 because $\frac{q(1-p_H)}{1-p_H q} \leq \frac{1-p_H q}{1-q} - 1 = \frac{q(1-p_H)}{1-q}$.

Similarly, I_H in (15) is increasing in p_H, and by taking $p_H = 1$ and $\alpha = 1$, we find that

$$I_H = 1 - \frac{p_L^2(1-q)^2}{(q + p_L(1-q))^2} \leq 1.$$

Therefore, both equalities in (13) are satisfied whenever $\alpha \leq 1$, so $d_L^* = d_H^* = 1$ in that case. □

The sufficient condition in Corollary 2 is consistent with the intuition that non-reporting is a form of market break-down and should only occur if the market is sufficiently inefficient (that is, insurance is mandatory but carries an excessive risk loading).

The set of q's for which both inequalities in (13) are satisfied can consist of *two or three* disjoint intervals. Figure 2 illustrates this phehomenon. We see that the second inequality in (13), namely $1 \geq I_H$, dominates for small q and the first inequality, namely $1 \geq I_L$, dominates for large q. That is, there is $\underline{q}$ such that $d_H^* = 1$ and $d_L^* = 1$ for $q \in [0, \underline{q}]$, and just above $q = \underline{q}$ we switch to $d_L^* = 1$ and $d_H^* < 1$. Similarly, there is a $\bar{q}$ such that for $q \in [\bar{q}, 1]$, we again have $d_L^* = 1$ and $d_H^* = 1$, and just below $q = \bar{q}$ we switch to $d_L^* = 0$ and $d_H^* = 1$.

Additionally, for some values of α, there may be a *third* interval of q's in which both inequalities in (13) are satisfied. This is illustrated when $\alpha = 1.215$ for q near 0.33 in Figure

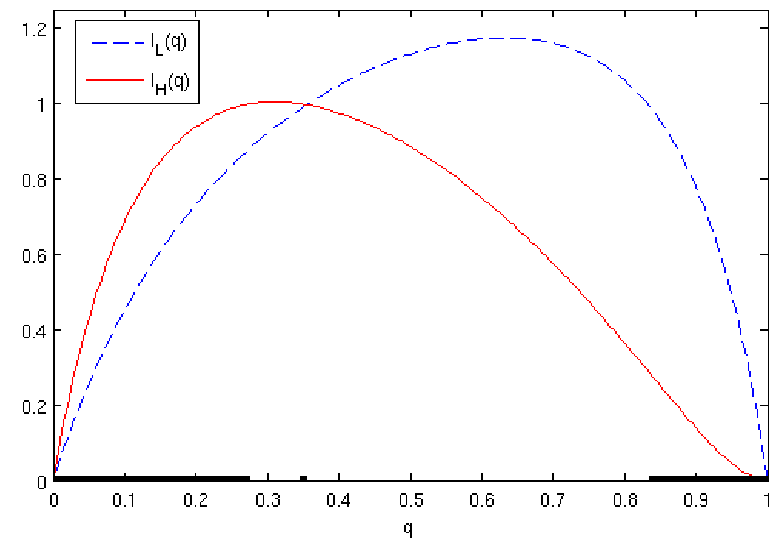

FIGURE 2. I_L and I_H in (14) and (15), respectively, as functions of the initial belief q for $p_L = 0.2$, $p_H = 0.8$, and $\alpha = 1.215$. The three intervals of q where $\max(I_L, I_H) < 1$ are highlighted in black on the x-axis.

2 and is also illustrated in the second panel of Figure 3 for $\alpha = 1.215$. Note that the first inequality in (13) is cubic in q and the second is quartic in q, so that although analytic expressions exist for $\underline{q}$ and $\bar{q}$ above, they are quite complicated.

5. Classification of Optimal Strategies in the Case Study

A variety of reporting strategies are possible in the simple two-period model from the previous section.

(1) Full reporting: $d_L^* = 1$ and $d_H^* = 1$, that is, all buyers report all accidents. This equilibrium holds if the cost from self-imposed deductibles outweighs any gain from information manipulation. Since information sensitivity is small for q's close to 0 or 1, in particular, this outcome happens for buyers whose risk type is essentially known *a priori* as demonstrated in Lemma 4.

(2) Obscuring high risk: $d_L^* = 1$ and $d_H^* < 1$, that is, low-risk buyers report all losses, and some of the losses generated by high risks are self-absorbed. The effect of such a policy is to reduce the information gain from observing a claim as the rate of reporting losses $p_X d_X$ is similar for both risks.

(3) Obscured signaling: $d_L^* = 0$ and $0 < d_H^* < 1$, that is, low-risk buyers absorb all losses, and some of the insurable events of the high-risk buyers are also not reported. The effect is that *a reported claim is a clear signal that the buyer is a high risk.* However, due to a large proportion of out-of-pocket settlements, losses are filed very

infrequently. Such a policy may arise if the expected out-of-pocket costs are small for low-risk types relative to the gain of moving q^-.

(4) Clear signaling: $d_L^* = 0$ and $d_H^* = 1$, that is, low risks report no losses, while high risks reports all losses. Again, a reported loss is a clear signal that the buyer is a high risk. It may occur for medium values of q at which the cost for a high risk is sufficiently high to make Case 3 above sub-optimal.

(5) No equilibrium (No-Eqm): There is no pair (d_L^*, d_H^*) that satisfies the consistency condition (5). Thus, in this case Bayesian learning essentially fails since there is no reporting strategy that is rationally optimal. In our numerical experiments, we observe that No-Eqm arises as a transition phase between Cases 2 and 3 above. In the transition from Case 2 to Case 3, both d_L^* and d_H^* change simultaneously, with the former changing discontinuously, which is impossible.

Due to the discrete nature of d_L^*, as model parameters vary, the equilibrium strategy exhibits instantaneous *phase transitions* whence an infinitesimal change in a parameter causes a shift from one type of equilibrium to another. In particular, negligible changes in the prior q can lead to radically different equilibria. In the next subsection, we investigate these transitions more closely.

In the next lemma, we show that even though high risks might attempt to mimic low risks by settling some losses out-of-pocket, as in Cases 2 and 3, they always appear more risky to the insurer. This is intuitive since one expects that reporting a loss is a signal of riskiness, rather than its opposite, so that $q^+ > q$.

Lemma 5. *The reporting frequency observed by the insurer $p_X d_X$ maintains the risk order, that is, $p_L d_L^* < p_H d_H^*$.*

Proof. We have two cases to consider: (1) $d_L^* = 1$, and (2) $d_L^* = 0$. If $d_L^* = 1$, then by substituting $d_L = 1$ and $d_H = p_L/p_H$ into (10), we obtain

$$\frac{\partial V_H}{\partial d_H} = -p_H < 0,$$

which implies that $d_H^* > p_L/p_H$ by convexity of V_H. It follows that $p_H d_H^* > p_L = p_L d_L^*$.

We know from Corollary 1 that $d_H^* > 0$, so that if $d_L^* = 0$, then we automatically have $p_H d_H^* > p_L d_L^*$. □

Note that in the case for which $d_L^* = 0$, it is possible to have $p_H d_H^* < p_L$. Indeed, by substituting $d_H = \widetilde{d_H}$ and $d_L = 0$ into (10), we obtain $\frac{\partial V_H}{\partial d_H} = 0$. Thus, $p_H d_H^* = p_H \widetilde{d_H} = \alpha - \frac{\alpha - 1}{q}$. It is possible that $\alpha - \frac{\alpha-1}{q} < p_L$; for example, take $\alpha = 1.5$, $q = 0.5$, and $p_L = 0.8$.

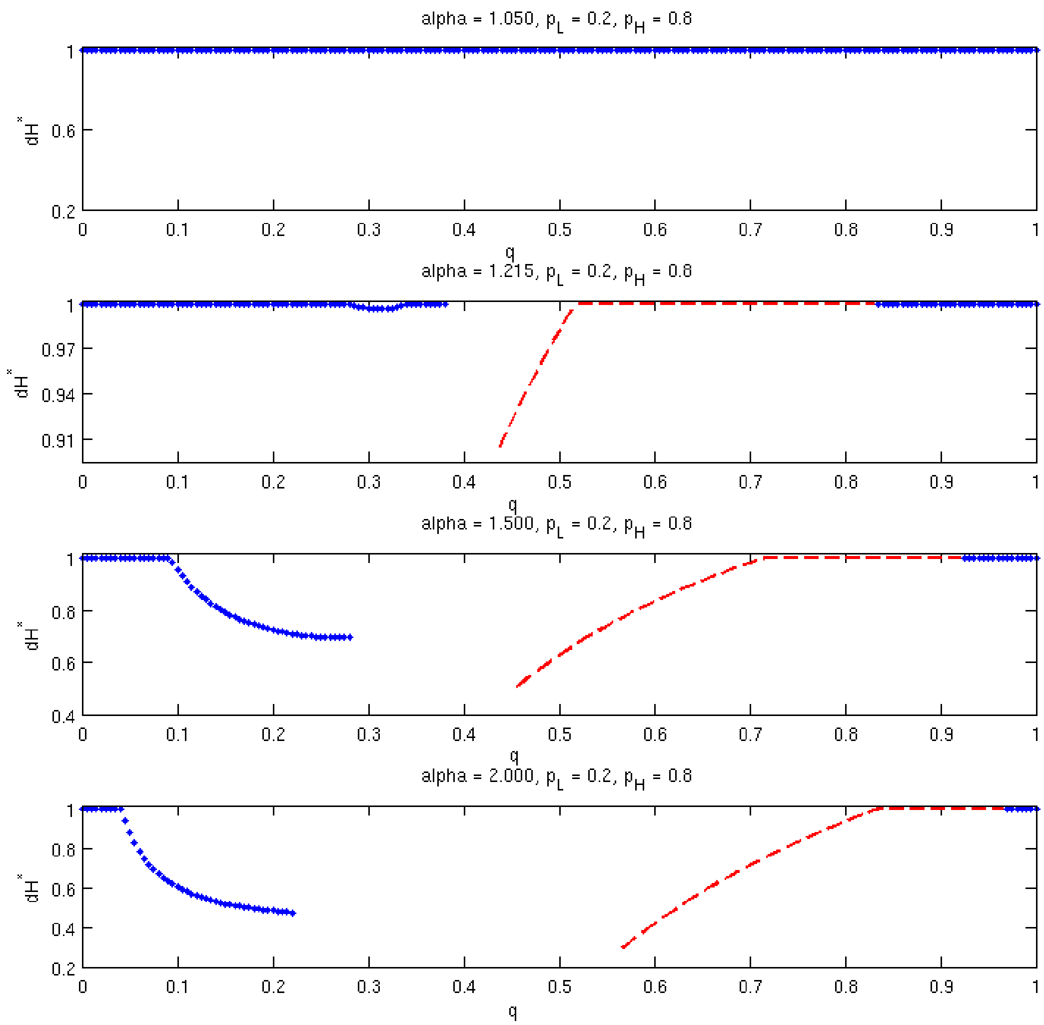

FIGURE 3. Optimal buyer strategies as a function of the prior $q(0)$ for different values of α. We plot d_H^*, with a dashed red line indicating where $d_L^* = 0$ and blue diamonds indicating where $d_L^* = 1$.

5.1. **Classification of Phase Transitions according to α.** In our numerical work, we find four possible type of phase transitions, three of which appear in Figure 3 (the phase labels refer to the classification in the beginning of Section 5):

(1) $\alpha \leq 1$: Corollary 2 implies that the equilibrium is always Case 1, full reporting for all risk types. This also happens for $\alpha - 1$ being small, see top panel of Figure 3;
(2) $\alpha \geq 1$ small (for example, $\alpha = 1.15$): $1 \to 4 \to 1$ (high risk always reports, no randomized strategies);

(3) $\alpha \geq 1$ moderate (for example, $\alpha = 1.215$): $1 \to 2 \to 1 \to 5 \to 3 \to 4 \to 1$ (second panel);
(4) $\alpha \geq 1$ large (for example, $\alpha = 2$): $1 \to 2 \to 5 \to 3 \to 4 \to 1$ (third and fourth panels;)

5.2. **Limiting Cases.** In this section, we study a variety of limiting cases with respect to the parameters of the model.

$\mathbf{p_L = 0}$. In this case, low-risk types never have accidents; accordingly, d_L^* is indeterminate. Indeed, when $p_L = 0$, $V_L = v(1+\theta)p_H(1-p_H d_H)q/(1-p_H d_H q)$, which is independent of d_L. Although there is no unique value of d_L that minimizes V_L when $d_H = d_H^*$, we still say that an equilibrium exists with $d_H^* \in [0,1]$ equal to the minimizer of V_H because the (intended) actions of the low risks are irrelevant when $p_L = 0$.

It is easy to show that $d_H = \widetilde{d_H} = \frac{1-\alpha(1-q)}{q\,p_H}$ minimizes V_H when $p_L = 0$. By examining the possible values of $\widetilde{d_H}$, we obtain the following:

(1) If $q \leq \frac{\alpha-1}{\alpha}$, then $\widetilde{d_H} \leq 0$, which implies that $d_H^* = 0$;
(2) If $\frac{\alpha-1}{\alpha} < q < \frac{\alpha-1}{\alpha-p_H}$, then $0 < \widetilde{d_H} < 1$, which implies that $d_H^* = \widetilde{d_H}$;
(3) If $q \geq \frac{\alpha-1}{\alpha-p_H}$, then $\widetilde{d_H} \geq 1$, which implies that $d_H^* = 1$.

No discontinuous phase transitions are present because V_H is continuous in q. See the illustration in the top panel of Figure 4.

$\mathbf{p_L = p_H}$. In this case, $\frac{\partial V_L}{\partial d_L} = -p_L < 0$, and $\frac{\partial V_H}{\partial d_H} = -p_H < 0$. Thus, $d_L^* = 1 = d_H^*$. This first-best outcome makes sense because there is no reason not to report a loss due to the fact that there is no difference between the risk classes.

$\mathbf{p_H = 1}$. In this case, high-risk types are guaranteed to have claims. This accentuates *ex post* moral hazard because high risks have a strong incentive to hide their claims and "masquerade" as low risks. We find the following:

(1) $d_L^* = 0$ and $d_H^* = \widetilde{d_H} = \frac{1-\alpha(1-q)}{q} \in (0,1)$ if and only if $\alpha > 1$ and $\tilde{q} < q < 1$, in which $\tilde{q}$ is given in (12).
(2) $d_L^* = 1$ and $d_H^* = \widetilde{d_H} = \frac{1-\alpha(1-q)}{q} \in (0,1)$ if and only if (11) holds at $d_H = \widetilde{d_H}$ and $\frac{\partial V_H}{\partial d_H} > 0$ at $d_L = 1 = d_H$. The latter condition holds if and only if $\alpha > 1$ and $q \in \left(p_L \frac{\alpha\sqrt{\alpha^2-1}-\alpha^2+(1-p_L)}{\alpha^2(1-2p_L)-(1-p_L)^2},\ 1\right)$.

   As α approaches 1, the lower bound on the interval for q converges to 1. The lower bound decreases with respect to α, so that as the risk loading θ increases, it is more likely that $d_H^* < 1$ in this case. Also, if $\alpha^2(1-2p_L) = (1-p_L)^2$, then the interval for q becomes $\left(\frac{1-2p_L}{2(1-p_L)}, 1\right)$.

To find the values of q such that (11) holds at $d_H = \widetilde{d_H}$ requires solving a cubic inequality in q, so although there is an analytical solution for such an inequality, it is not simple.

(3) $d_L^* = 1 = d_H^*$ if and only if $\alpha \le 1$ or both $\alpha > 1$ and $q \in \left[0,\ p_L \frac{\alpha\sqrt{\alpha^2-1}-\alpha^2+(1-p_L)}{\alpha^2(1-2p_L)-(1-p_L)^2}\right] \cup \{1\}$.

(4) No-Eqm is also possible; see the bottom panel of Figure 4.

Note that when $p_H = 1$, an equilibrium with $d_L^* = 0$ and $d_H^* = 1$ is not possible. Indeed to have it, one needs

$$\begin{cases} 1 < \dfrac{\alpha^2 q}{q + p_L(1-q)}, & \text{that is, } d_L = 0 \text{ is optimal when } d_H = 1; \\ 1 \le \widetilde{d_H} = \dfrac{1-\alpha(1-q)}{q}, & \text{that is, } d_H = 1 \text{ is optimal when } d_L = 0; \end{cases}$$

which are impossible to satisfy together for any $0 \le q \le 1$.

$\mathbf{q = 0}$. $d_L^* = 1$ because (11) holds when $q = 0$. Also, $\frac{\partial V_H}{\partial d_H} = -p_H < 0$ when $q = 0$ and $d_L = 1$; thus, $d_H^* = 1$. This result makes sense because when $q = 0$, the updated Bayesian belief is still $q^+ = 0$, so there is no penalty for reporting a loss.

$\mathbf{q = 1}$. As in the previous case, $d_L^* = 1$ because (11) holds when $q = 1$. Also, $\frac{\partial V_H}{\partial d_H} = -p_H < 0$ when $q = 1$ and $d_L = 1$; thus, $d_H^* = 1$. This result makes sense because when $q = 1$, the updated value $q^- = 1$, so there is no possible gain from *not* reporting a loss.

θ **or** α **large.** Our final result shows that as the risk loading θ increases, which makes α increase, the game breaks down and no equilibrium is possible.

Lemma 6. *For any fixed $q \in (0,1)$, there is an α large enough, such that there is no equilibrium for the two-period model.*

Proof. Fix $q \in (0,1)$. If an equilibrium exists, then either $d_L^* = 0$ or $d_L^* = 1$. Suppose $d_L^* = 0$; then, when $d_L = 0$, $\frac{\partial V_H}{\partial d_H} = -p_H + p_H \frac{\alpha^2(1-q)^2}{(1-p_H d_H q)^2}$, which is positive for all $d_H \in [0,1]$ for $\alpha > 1/(1-q)$, from which it follows that $d_H^* = 0$. This result contradicts Corollary 1, so we cannot have $d_L^* = 0$.

Next, suppose $d_L^* = 1$; by looking at (10), we see that if α is large enough, then $\partial V_H / \partial d_H$ is determined by the sign of the expression in the curly brackets. It is easy to check that that expression is negative for $d_H < p_L/p_H$ and positive for $d_H > p_L/p_H$. Therefore, for α large, V_H achieves its minimum at $d_H = p_L/p_H + \epsilon$, in which $\epsilon \to 0+$ as $\alpha \to \infty$. Specifically, α and ϵ are related as follows:

$$\alpha^2 = \frac{1}{(1-q)^2} \left[\left(\frac{1-p_L}{1-p_L-\epsilon p_H q} \right)^2 - \left(\frac{p_L}{p_L+\epsilon p_H q} \right)^2 \right]^{-1}.$$

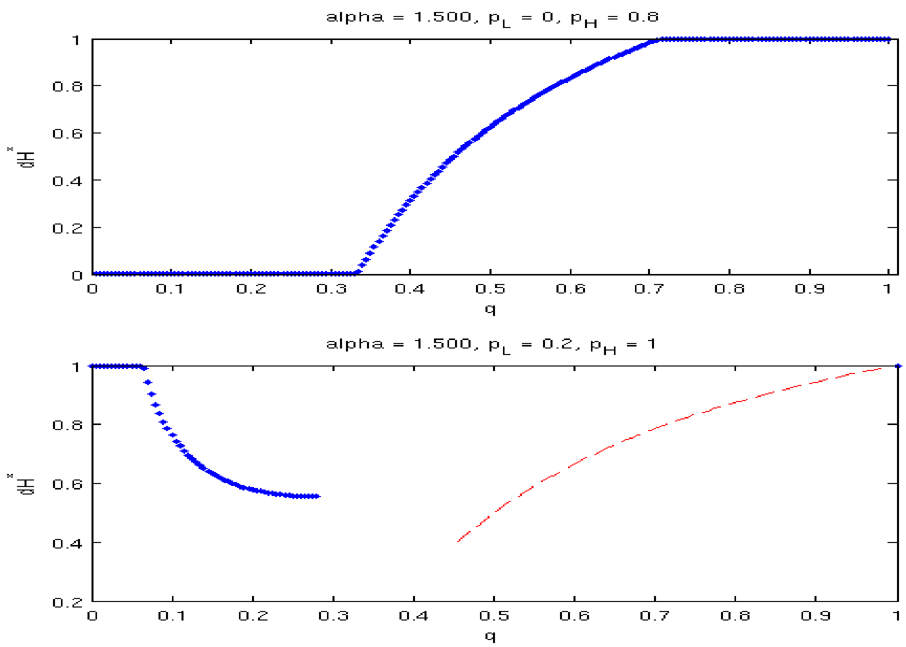

FIGURE 4. Optimal (d_L^*, d_H^*) for limiting cases: $p_L = 0$ and $p_H = 0.8$ for the top panel; $p_L = 0.2$ and $p_H = 1$ for the bottom panel. The format of the figure is same as that of Figure 3.

By substituting $d_H = p_L/p_H + \epsilon$ into (11), we obtain $1 \geq \frac{\alpha^2 q(1-q)}{1-p_L q} \cdot (1 + O(\epsilon))$, which is a contradiction when α is large. Thus, $d_L^* = 1$ is also impossible. □

5.3. **Value Functions.** By solving for d_L^* and d_H^* and substituting these values into V_L and V_H, we obtain the value functions $W_L(q)$ and $W_H(q)$, respectively. In the case of No-Eqm, W_L and W_H are undefined, but one can set $W_L = W_H = +\infty$ there, so that individuals avoid having an updated q lie in the region of No-Eqm. This is not important in a two-period model but becomes relevant once more periods are considered.

We find that W_L is a non-decreasing continuous function of q; so that for a low-risk type, the higher the original rating, the lower the expected cost to the buyer. Also, recall from Lemma 3 that W_L is constant on the set $\{q\colon d_L^* = 0, d_H^* < 1\}$. Surprisingly, we find that W_H is not monotone, nor continuous with respect to q. In particular, W_H experiences a negative jump at $\hat{q}$ where we switch from $d_L^* = 0, d_H^* = 1$ to $d_L^* = 1, d_H^* = 1$ (around $q \sim 0.95$ in Figure 5).

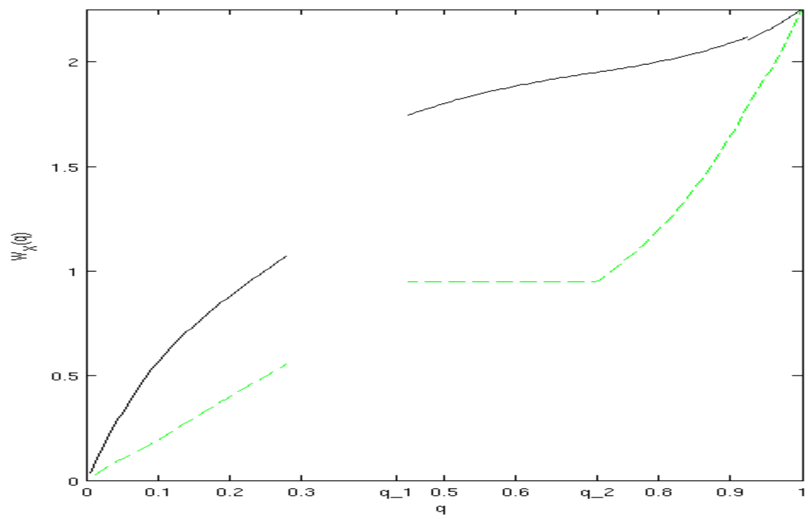

FIGURE 5. Optimal value functions $W_L(1, q)$ and $W_H(1, q)$ as functions of the prior q for same parameters as top-left panel of Figure 3; we plot W_H with a black solid line and W_L with a dashed green line. Inside the No-Eqm region in the middle the value functions are undefined. We also indicate on the x-axis the region $[q_1, q_2] = [\tilde{q}, \frac{\alpha-1}{\alpha-p_H}] = [0.4344, 0.7143]$ of Lemma 3. Note the discontinuity of W_H at $q = 0.924$.

6. EXTENSIONS OF THE CASE STUDY

The analysis of the case study presented in the previous two sections provides insight into the more general settings of our model. In this section, we discuss how the above results carry over in general.

6.1. **General premium rules.** First, we consider general premium rules rather than the linear function used in Sections 4 and 5. These rules might arise due to non-linear pricing mechanisms employed by the insurer. For instance, it is common practice that the premium paid is most elastic for buyers with uncertain risk profiles (medium q's), while it is relatively inelastic for buyers that belong to a given type with high probability; see, for example, Table 1 in Abbring et al. (2007). Such pricing rules can be represented by an "arctangent"-shaped premium function $f(q)$. Alternatively, insurance regulations might state that high-risk buyers should be penalized for being "careless" (in terms of *ex ante* moral hazard), so that $f(q)$ would be convex in that case.

A non-linear premium rule can be viewed as an extension of the linear case, with the risk-loading α being a function of the current rating q. Consider first a piecewise linear $f(q)$; in that case, since the main problem (5) is optimized locally for each q, one may treat each linear interval of $f(q^{\pm})$ separately and then combine the results. Therefore, the resulting reporting strategies will be a combination of the outcomes observed in Sections 4 and 5. Of course, for some ratings q, depending on the value of $\vec{d}$, $q^{\pm}$ will be on different segments of f, in which case a more complicated equilibrium may emerge. In particular, V_L might no longer be concave. As a result, it is possible that $0 < d_L < 1$ is optimal; also multiple solutions of (5) may exist.

6.2. **More than two periods.** Non-linear pricing rules also arise in the case of a multi-period model with more than two periods. In that case, the next-period value function $V_X(k+1,\cdot)$, as well as the next period premium $f(q(k),\vec{d}(k+1))$, enter into the optimization problem. Even if f itself is linear, the non-linear features in $\vec{d}(k+1)$ (which arise as soon as non-reporting is optimal for some q's), and in V_X make the "effective" next-period premium non-linear in current rating. This can be seen in Figure 5, where the value functions are highly non-linear, including points of discontinuity and intervals of constancy.

With more than two periods, the buyer has a stronger incentive to manipulate his rating early on, since he anticipates that it will affect his expected costs for several periods into the future. For instance, with $T = 3$ and a small $\alpha = 1.1$, the solution structure is that $d^*_X(2) = 1$ for $X = L$ and H, but $d^*_X(1)$ belongs to a variety of equilibrium phases with the overall picture similar to the third and fourth panels of Figure 3.

A multi-period problem also requires considering the precise interpretation of what happens in the case of No-Eqm or multiple solutions to (5). An orthodox solution is to set V to be undefined in such cases; this would rule out solutions for nearly all q after several periods, since the non-defined regions expand as more periods are added because at each step all possible updates $q^{\pm}$ must be inside the defined region.

One may also consider the infinite-horizon problem, whence (4) is replaced by the time-stationary version

$$\begin{aligned}H_X(q)(\vec{d}) = p_X\Big[d_X\Big(U(f(q^+,\vec{d})) + vH_X(q^+)(\vec{d})\Big) + (1-d_X)\Big(U(1+f(q^-,\vec{d})) + vH_X(q^-)(\vec{d})\Big)\Big] \\ + (1-p_X)\Big[U(f(q^-,\vec{d})) + vH_X(q^-)(\vec{d})\Big].\end{aligned}$$

Heuristically, one expects that $W(k,q) \to H(q)$ as $k \to \infty$. However, given the complications arising already with two or three periods, this does not appear to be a feasible way of computing H. Also, on an infinite horizon, the learning problem of the insurer must be treated differently. Indeed, given a sufficiently large claim history and an equilibrium solution to the insurance reporting problem, the insurer can determine the risk type of the buyer to an arbitrary degree of precision (since she can infer his reporting strategy). In other words,

$q(k) \to 1_X$ as $k \to \infty$ almost surely. Therefore, with an infinite horizon (and hence infinite past history), the insurer will, in fact, know the true type of the buyer with certainty.

6.3. **Other extensions.** Finally, many other cases can be considered. The buyers can belong to multiple risk types, in which case the insurer's beliefs are not one-dimensional as in case study, but live in the general space D. In that case, in the two-period setting, in the first period, we still expect the lowest risk type either to always report or to never report losses, with the same ordering of the reporting frequency $p_X d_X$ that we obtained in Lemma 5.

The claim sizes Y_i may also be varying; under the assumption that the buyer must either report the full amount or nothing at all (so that underreporting is not possible), this becomes a variation of our model in which the claim probabilities p_X are modified by the claim size distribution $f_X(y)$ and reporting strategy $\vec{d}(y)$ depends on the claim size. In that case, the update rule (1) becomes

(16)
$$\begin{cases} q^+(d_L, d_H)(y) = \mathbb{P}^q(X = H|Y = y) = \dfrac{p_H f_H(y) \cdot d_H(y) \cdot q}{p_H f_H(y) \cdot d_H(y) \cdot q + p_L f_L(y) \cdot d_L(y) \cdot (1-q)}, \\ \quad q^-(d_L, d_H) = \mathbb{P}^q(X = H|Y = 0) \\ \qquad\qquad\qquad = \dfrac{(1 - p_H \int_0^\infty d_H(y) \cdot f_H(dy)) \cdot q}{1 - \int_0^\infty (p_H d_H(y)\, q\, f_H(dy) + p_L d_L(y)\,(1-q)\, f_L(dy))}. \end{cases}$$

The resulting problem is now infinite-dimensional in the buyer's strategy $d_X(y), X \in E, y \in \mathbb{R}_+$, and we will study this problem in future work. Intuitively, higher claims are more costly to self-absorb; therefore, we expect a self-imposed deductible $Ded(q)$ to emerge such that $d_X(q, y) = 1$ for all $y > Ded(q)$ and $d_X(q, y) < 1$ for $y \leq Ded(q)$.

7. Summary and Conclusions

In this paper, we studied an asymmetric information insurance problem with *ex post* moral hazard. Since the buyer has the option of not reporting losses, and since his premia are based on his experience rating, the buyer will strategically manipulate the claim information received by the insurer. Equilibrium reporting strategies depend on the joint strategic behavior of all risk types, and upon the assumption that the insurer updates the risk ratings via Bayesian conditioning and anticipates the optimal reporting strategies when setting the premia.

We find that even in a simple case study with two periods, linear risk premia, and risk-neutral (but captive) buyers, non-trivial reporting strategies might be optimal. In particular, the behavior of different risk types can be quite varied, ranging from total non-reporting, full reporting, and mixed strategies. We also find that the equilibrium strategies are highly sensitive to the buyer's current risk rating and exhibit abrupt phase shifts as the risk rating varies. Such features are expected to be even more pronounced in more general model settings and showcase the complexity of the underlying insurance markets. Our work provides probabilistic underpinnings to anecdotal evidence of *ex post* moral hazard in real-life private automobile and homeowner's insurance markets.

References

Abbring, J. H., P. A. Chiappori, and J. Pinquet (2003), Moral hazard and dynamic insurance data, *Journal of the European Economic Association*, 1 (4): 767–820.

Abbring, J. H., P. A. Chiappori, and T. Zavadil (2007), Better Safe than Sorry? Ex Ante and Ex Post Moral Hazard in Dynamic Insurance Data. Working Paper, Department of Economics, Columbia University.

Chiappori, P. A. (2001), Econometric models of insurance under asymmetric information, In G. Dionne (ed.), *Handbook of Insurance*, Kluwer Academic Publishers, Boston, Chapter 12, pp. 363–392.

Chiappori, P. A., B. Jullien, B. Salanié, and F. Salanié (2006), Asymmetric information in insurance: general testable implications, *RAND Journal of Economics*, 37 (4): 783–798.

Cooper, R. and B. Hayes (1987), Multi-period insurance contracts, *International Journal of Industrial Organisation*, 5: 211–231.

Dionne, G. (1983), Adverse selection and repeated insurance contracts, *Geneva Papers on Risk and Insurance*, 8: 316–333.

Dionne, G. (ed.) (2002), *Handbook of Insurance*, Huebner International Series on Risk, Insurance and Economic Security, Kluwer Academic Publishers, Boston.

Dionne, G. and P. Lasserre (1985), Adverse selection, repeated insurance contracts and announcement strategy, *Review of Economic Studies*, 52 (4): 719–724.

Dionne, G. and C. Vanasse (1992), Automobile insurance ratemaking in the presence of asymmetrical information, *Journal of Applied Econometrics*, 7 (2): 149–165.

Harsanyi, J. C. (1967), Games with incomplete information played by "Bayesian" players. I. The basic model, *Management Science*, 14: 159–182.

Harsanyi, J. C. (1968), Games with incomplete information played by "Bayesian" players. II. Bayesian equilibrium points, *Management Science*, 14: 320–334.

Harsanyi, J. C. (1968), Games with incomplete information played by "Bayesian" players. III. The basic probability distribution of the game, *Management Science*, 14: 486–502.

Hey, J. (1985), No claims bonus?, *Geneva Papers on Risk and Insurance*, 10: 209–228.

Hosios, A. and M. Peters, (1989), Repeated insurance contracts with adverse selection and limited commitment, *Quarterly Journal of Economics*, 104: 229–253.

Lemaire, J., (1995), *Bonus-Malus Systems in Automobile Insurance*, Kluwer Academic Publishers, Boston.

Moreno, I., F. J. Vázquez, and R. Watt (2006), Can bonus-malus alleviate insurance fraud?, *Journal of Risk and Insurance*, 73 (1): 123–151.

Mikosch, T. (2004), *Non-life insurance mathematics. An introduction with stochastic processes*, Universitext, Springer-Verlag, Berlin.

Rothschild, M. and J. E. Stiglitz (1976), Equilibrium in competitive insurance markets: an essay on the economics of imperfect information, *Quarterly Journal of Economics*, 90: 629–649.

Robinson, C. and B. Zheng (2006), Moral hazard and repeated insurance contracts, Working Paper, Department of Economics, University of Western Ontario.

Townsend, R.M. (1982), Optimal multiperiod contracts and the gain from enduring relationships under private information, *Journal of Political Economy*, 90 (6): 1166–1186.

Watt, R. and F.J. Vasquez (1997), Full insurance, Bayesian updated premiums, and adverse selection, *Geneva Papers on Risk and Insurance Theory*, 22: 135–150.