Optimal Auditing with Scoring:
Theory and Application to Insurance Fraud*

Georges Dionne, Florence Giuliano, and Pierre Picard

Canada Research Chair in Risk Management, HEC Montreal, Montreal (Qc) Canada, georges.dionne@hec.ca
Department of Economics, Paris X University, Paris, France, fgiuliano@riskdynamics.eu
Department of Economics, Ecole Polytechnique, Palaiseau, France, pierre.picard@polytechnique.edu

This version: February 6, 2008

Abstract

This article makes a bridge between the theory of optimal auditing and the scoring methodology in an asymmetric information setting. Our application is meant for insurance claims fraud, but it can be applied to many other activities that use the scoring approach. Fraud signals are classified based on the degree to which they reveal an increasing probability of fraud. We show that the optimal auditing strategy takes the form of a "Red Flags Strategy" which consists in referring claims to a Special Investigative Unit (SIU) when certain fraud indicators are observed. The auditing policy acts as a deterrence device and we explain why it requires the commitment of the insurer and how it should affect the incentives of SIU staffs. The characterization of the optimal auditing strategy is robust to some degree of signal manipulation by defrauders as well as to the imperfect information of defrauders about the audit frequency. The model is calibrated with data from a large European insurance company. We show that it is possible to improve our results by separating different groups of insureds with different moral costs of fraud. Finally, our results indicate how the deterrence effect of the audit scheme can be taken into account and how it affects the optimal auditing strategy.

Keywords: Audit, scoring, insurance fraud, red flags strategy, fraud indicators, suspicion index, moral cost of fraud, deterrence effect, signal manipulation.

JEL numbers: D81, G14, G22, M40.

* *This research has benefited from the financial support of the Fédération Française des Sociétés d'Assurances (FFSA), of the Risk Foundation–AXA Chair on Large Risks in Insurance, and of the Center for Research on E-finance (CREF) at HEC Montréal. We thank Claire Boisvert, Denise Desjardins, and Mathieu Maurice for their efficient assistance, Richard Butler, Richard Derrig, Frederic Nguyen Kim, Jean Pinquet, an associate editor, and two referees for their comments and suggestions.*

1. Introduction

Auditing has been a major topic of interest in the economic and financial literature since the path-breaking articles of Townsend (1979) and Gale and Hellwig (1985). The emphasis has been put on the informational asymmetries between principals (bankers, insurers, regulators, tax inspectors…) and agents (borrowers, insureds, regulated firms, tax payers…), asymmetries which lead to implement costly state-verification strategies. The trade-off between the reduction of monitoring costs and the mitigation of informational asymmetries has thus been the core of the economic analysis of auditing.

On the empirical side, the importance of auditing in corporations, financial institutions or governmental agencies has given rise to serious analysis of the design of optimal auditing procedures (e.g. Should auditing be internal or external? How should auditors be rewarded? How can collusion between auditors and those audited be avoided? How frequent should auditing be? …) and has motivated firms and governments to search for relevant information on ways to cut auditing costs. Nowadays, the search for optimal auditing procedures is a major concern for a number of players: banks and insurance companies seeking better risk assessment of their customers; prudential regulators of the banking and insurance industries investigating risk management activities; financial markets introducing enforcement actions against accounting misconduct; governments pursuing better compliance by taxpayers; and regulatory agencies in the field of environmental law, food safety or working conditions.

On the theoretical side, many extensions of the basic models have been proposed. In particular, Mookherjee and Png (1989) have shown that random auditing dominates deterministic models. Among many other issues, the effect of collusion between auditees and auditors has received special attention (Kofman and Lawarrée, 1993). The consequences of commitment vs. no-commitment assumptions in an auditing procedure have also been examined with the analytical tools of modern incentives theory (Graetz, Reinganum and Wilde, 1986; Melumad and Mookherjee, 1989).

Optimal auditing has been extensively analyzed in the accounting literature. Strategic auditing in financial reporting games is developed by Morton (1993), Newman, Patterson and Smith (2001) and Pae and Yoo (2001), just to name a few. The emphasis is put on the interaction between internal and external monitoring and on the trade-off between types 1 or 2 errors in financial auditing by Shibano (1990), Baiman and Rajan (1994), and Caplan (1999), on the links between accounting systems, audit policies and managers'incentives by Christensen, Demski and Frimor (2002), on the auditor's legal liability by Dye (1993) and Narayanan (1994), and on the conflicts of interest when auditors also provide management advisory services by Simunic (1984) and more recently by Antle et al. (2006) among others.

In such an asymmetric information setting, it is in the interest of principals to use signals on agents' types or actions when deciding whether a costly verification should be performed. Scoring techniques are then most useful since they help to identify suspicious files to be audited as a priority by associating scores. However, almost all theoretical approaches have completely ignored the role of scoring in the design of an optimal auditing strategy[1]. The present paper tries to fill this gap by bringing together auditing and scoring within a unified costly state verification approach and by applying the model to insurance claim fraud. In doing so, we shall build a model of optimal auditing which is much more closely related to auditing procedures actually used by insurers (or other principals) than abstract auditing models. Scoring will signal whether an audit should be performed or not, depending on the signals perceived by the insurer: This will be called a "red flags strategy". In fact, we shall build a model where the optimal auditing strategy actually takes the form of such a red flags strategy.

Insurance fraud provides a fascinating case study for the theory of optimal auditing, and particularly for connecting the scoring methodology with the theory of optimal costly state verification. In recent years, the economic analysis of insurance fraud has developed along two branches. The first branch is mostly theoretical. It aims at analyzing the strategy adopted by insurers faced with claims or application fraud. The most usual setting for this literature is a costly state-verification model in which insureds have private information about their losses and insurers can verify claims by incurring an audit cost. Important assumptions are made relative to the ability of insurers to commit to an auditing policy and to the skill of defrauders in manipulating audit costs, i.e. to make the verification of claims more difficult.[2] The second branch of the literature on insurance fraud is more statistically based: It focuses mainly on the significance of fraud in insurance portfolios, on the practical issue of how insurance fraud can be detected, and on the scope of automated detection mechanisms in lowering the cost of fraudulent claims.[3] The scoring methodology is one of the key ingredients in this statistical approach.

Using scoring techniques to predict fraudulent behavior is also the purpose of active research in other fields and particularly in accounting. A number of studies have investigated the relationship between financial statement fraud and corporate governance features (Dunn, 2004; Dechow et al., 1996; Beasley, 1996) or financial-related warning variables (Dechow et al., 1996; Beneish, 1997; Summers and Sweeney, 1998). More recently, Skousen and Wright (2006) and Dechow et al. (2007) have used data on Accounting and Auditing Enforcement Releases (AAERs) issued by the Security and Exchange Commission to identify variables that are correlated with accounting fraud. In particular, Dechow et al. (2007) develop a logistic model to estimate the probability of misstatements as a function of accruals quality measures, performance measures and market-related measures. The output of this analysis is a fraud score, i.e. the probability that a firm has engaged in an earnings misstatement, which can be used as a screening device to signal the need for more thorough audit investigation.

The present paper shares a common interest with this strand of research in the sense that we develop a scoring approach to fraud detection. However instead of limiting ourselves to a purely statistical approach (i.e. identifying fraud signals and estimating the probability of fraud), we put the insurance fraud detection problem into the context of a cost minimizing insurance company where the behavior of defrauders is explicitly described. We characterize the optimal scoring based audit strategy in a setting where the audit cost, the deterrence effect of fraud investigation, the insurer's commitment to fraud detection and the ability of defrauders to manipulate signals are taken into account.

Our approach reconciles an important result of optimal audit theory with the widespread practice of insurance companies. Indeed, on one hand, theory predicts that an optimal auditing strategy should be random. In other words, claims should be audited with probability less than one. Indeed, filing fraudulent claims is a strictly dominated strategy for opportunistic policyholders if claims are always audited (assuming that audit allows the insurer to detect fraud with certainty). In such a case, slightly decreasing the audit probability would allow insurers to reduce their audit costs without inciting policyholders to fraud. On the other hand, factual evidence suggests that insurers tend to systematically audit certain claims and to directly pay the other ones, depending on the available information on policyholders' and on the perceived red flags. Our model will show that such a red flag strategy is in fact the key ingredient of a random auditing strategy, when individuals cannot perfectly control the fraud signals perceived by the insurers. From the insurers' standpoint, the auditing strategy is deterministic (the optimal audit decision is actually a non-random function of red flags) but it is random for the policyholders.

We will also derive some consequences of our results for the implementation of the optimal auditing strategy. Firstly, it will be shown that such a strategy requires a strong commitment from insurance companies. More explicitly, the optimal strategy *does not* amount to perform an audit if and only if the expected benefits of auditing (i.e. the expected value of possibly detected fraudulent claims) exceed the audit cost. Indeed, claims are verified in order to detect fraud but auditing also acts as a deterrence device. The deterrence objective requires auditing some claims although the investigation cost may exceed the expected benefits that the insurer may get from verification. Secondly, we will show that the optimal proportion of successful audits depends on the policyholders' type (more precisely on the particulars that can be observed by the insurer)[4], which affects the incentive scheme that should be offered to the insurance staffs in charge of verifying claims.

In the empirical part of the paper, we will calibrate our model by using data on automobile insurance (theft and collision) from a large European insurance company and we will derive the optimal auditing strategy. As a final outcome, our analysis yields an easily automated procedure which may be viewed as a prototype for the kind of insurance fraud detection mechanisms that are now more and more used by large insurance companies.

Section 2 presents the theoretical model while Section 3 derives the optimal auditing strategy and its main implications. Section 4 extends the model to the cases where defrauders can manipulate fraud signals or have imperfect information about the insurer's auditing strategy. The application of the model to the portfolio of an insurer begins with Section 5 where the regression analysis is presented. Sections 6 and 7 outline the model calibration and its results. Section 8 concludes. Proofs, robustness tests, and variables description are in the appendices available at http://neumann.hec.ca/gestiondesrisques/Opt-Auditing.pdf.

2. The model

We consider a population of policyholders who differ from one another in the moral cost of filing a fraudulent claim. For the sake of notational simplicity, all individuals own the same initial wealth and they all face the possibility of a monetary loss L with probability π with $0 < \pi < 1$. We simply describe the event leading to this loss as an "accident".

All individuals are expected utility maximizers and they display risk aversion with respect to their wealth. Let u be the state dependent utility of an individual drawn from this population. u depends on final wealth W but it also depends on the moral cost incurred in case of insurance fraud: $u = u(W, \omega)$ in case of fraud and $u = u(W, 0)$ otherwise, where ω is a non-negative parameter which measures the moral cost of fraud to the policyholder. We assume that $u_1 > 0, u_{11} < 0$ and $u_2 < 0$ where u_i is for the first derivative and u_{ii} is for the second derivative and that ω is distributed over R_+ among the population of policyholders. In other words, individuals who choose to defraud incur more or less high moral costs. Some of them are purely opportunistic (their moral cost is very low) whereas others have a higher sense of honesty (their moral cost is thus higher). Note that moral cost is private information held by the insured.[5]

For notational simplicity we assume that all the individuals have taken out the same insurance contract. This contract specifies a level of coverage t in case of an accident and a premium that should be paid to the insurer. If there is no fraud, we have $W = W_0 - L + t$ in case of an accident and $W = W_0$ if no accident occurs where W_0 is the initial wealth net of the insurance premium. Each individual in the population is characterized by a vector of observable exogenous variables θ, with $\theta \in \Theta \subset R^m$. The moral cost of fraud may be statistically linked to some of these variables.[6] Let $H(\omega | \theta)$ be the conditional cumulated distribution of ω over the population of type-θ individuals, with a density $h(\omega | \theta)$.

Our model describes insurance fraud in a very crude way. A defrauder simply files a claim to receive the indemnity payment *t* although he has not suffered any accident. If a policyholder is detected to have defrauded, he will receive no insurance payment and must in addition pay a fine *B* to the government.[7] Let Q^f be the probability that a fraudulent claim is detected; this probability is the outcome of the insurer's antifraud policy and it depends on the observable variables θ as we shall see in Section 3.

When an individual has not suffered any loss, his utility is written as $u(W_0,0)$ if he does not defraud. If he files a fraudulent claim (i.e. if he claims to have suffered an accident although this is not true), his final wealth is $W = W_0 + t$ if he is not detected and $W = W_0 - B$ if he is detected. An individual with moral cost ω decides to defraud if he expects greater utility from defrauding than staying honest, which is written as:

$$\left(1-Q^f\right)u\left(W_0+t,\omega\right)+Q^f u\left(W_0-B,\omega\right)\geq u\left(W_0,0\right).$$

This inequality holds if $\omega \leq \phi\left(Q^f\right)$, where function $\phi(p):[0,1]\to R^+$ is implicitly defined by:

$$(1-p)u\left(W_0+t,\phi\right)+pu\left(W_0-B,\phi\right)=u\left(W_0,0\right) \text{ if } 0\leq p\leq p_0$$

$$\text{and } \phi(p)=0 \text{ if } p_0\leq p\leq 1 \text{ where } p_0 \text{ is given by}$$

$$\left(1-p_0\right)u\left(W_0+t,0\right)+p_0u\left(W_0-B,0\right)=u\left(W_0,0\right)$$

with $0 < p_0 < 1$ if $B > 0$ and $p_0 = 1$ if $B = 0$. p_0 is the audit probability that deters an individual with no moral cost ($\omega = 0$) from defrauding. We have $\phi(0)>0$, $\phi(p_0)=0$ and $\phi'(p)<0$ if $0<p<p_0$. $\phi\left(Q^f\right)$ is the critical value of the moral cost under which cheating overrides honesty as a rule of behavior. When $Q^f < p_0$ the higher the probability of being detected, the lower the threshold of the moral cost and thus the lower the frequency of fraud. When $Q^f \geq p_0$, there is no more fraud.[8]

When a policyholder files a claim — be it honest or fraudulent — the insurer privately perceives a *k*-dimensional signal $\sigma \in \{\sigma_1,\sigma_2,...,\sigma_\ell\} = \Sigma$. Hereafter, *k* will be interpreted as the number of fraud indicators (or red flags) privately observed by insurers. Fraud indicators are claim-related signals that cannot be controlled by the defrauder and that should make the insurer more suspicious[9]. When all indicators are binary, then $\ell = 2^k$ and σ is a vector of dimension k all of whose components are 0 or 1: component j is equal to 1 when indicator j is "on" and it is equal to 0 when it is "off".

Let p_i^f and p_i^n be respectively the probability of the signal vector σ taking on configuration σ_i when the claim is fraudulent and when it corresponds to a true accident (non-fraudulent claim), i.e.:

$$p_i^f = P\left(\sigma = \sigma_i \middle| F\right); \;\; p_i^n = P\left(\sigma = \sigma_i \middle| N\right)$$

with $i = 1,\ldots,\ell$, where F and N refer respectively to "fraudulent" and "non-fraudulent" and $P(\bullet)$ denotes probability[10]. Of course, we have:

$$\sum_{i=1}^{\ell} p_i^n = \sum_{i=1}^{\ell} p_i^f = 1 .$$

The probability distribution of signals is supposed to be common knowledge to the insurer and to the insureds. For simplicity of notations, we assume $p_i^n > 0$ for all $i = 1,\ldots,\ell$ and *w*.l.o.g. we rank the possible signals in such a way that[11]

$$\frac{p_1^f}{p_1^n} < \frac{p_2^f}{p_2^n} < \ldots < \frac{p_\ell^f}{p_\ell^n} .$$

This ranking allows us to interpret $i \in \{1,\ldots,\ell\}$ as an index of fraud suspicion. Indeed let $P\left(F \middle| \theta\right)$ be the proportion of fraudulent claims among the claims filed by type-θ individuals. Section 3 will show how $P\left(F \middle| \theta\right)$ can be deduced from the insurer's auditing strategy. Bayes law shows that the probability of fraud conditional on signal σ_i and type θ is:

$$P\left(F \middle| \sigma_i, \theta\right) = \frac{p_i^f \; P\left(F \middle| \theta\right)}{p_i^f \; P\left(F \middle| \theta\right) + p_i^n \left(1 - P\left(F \middle| \theta\right)\right)} \tag{1}$$

which is increasing with i. In other words, as index i increases so does the probability of fraud.

3. Auditing strategy

The insurer may channel dubious claims to a Special Investigative Unit (SIU) where they will be verified with scrupulous attention. Other claims are settled in a routine way. The SIU referral serves to detect fraudulent claims as well as to deter fraud. We assume for simplicity that an SIU referral always allows the insurer to determine beyond the shadow of a doubt whether a claim is fraudulent or not. In other words, the SIU performs perfect audits. An SIU claim investigation costs c to the insurer with $c < t$.

The insurer's investigation strategy is characterized by the probability of an SIU referral, this probability being defined as a function of individual-specific variables and claim-related signals. Hence, we define an investigation strategy as a function $q:\Theta\times\Sigma\to[0,1]$. A claim filed by a type-θ policyholder is transmitted to the SIU with probability $q(\theta,\sigma)$ when signal σ is perceived.

Let $Q^f(\theta)$ – respectively $Q^n(\theta)$ – be the probability of an SIU referral for a fraudulent –non-fraudulent – claim filed by a type-θ individual. $Q^f(\theta)$ and $Q^n(\theta)$ result from the insurer's investigation strategy through:

$$Q^f(\theta)=\sum_{i=1}^{\ell} p_i^f q(\theta,\sigma_i);\ \ Q^n(\theta)=\sum_{i=1}^{\ell} p_i^n q(\theta,\sigma_i) \tag{2}$$

In particular, a type-θ defrauder knows that his claim will be subjected to careful scrutiny by the SIU with probability $Q^f(\theta)$. The insurer knows that, given his investigation strategy, truthful claims filed by type-θ individuals are mistakenly channeled to the SIU with probability $Q^n(\theta)$.

The total expected cost of fraud includes the cost of investigation in the SIU and the cost of residual fraud (fraudulent claims not detected). A type-θ individual has an accident with probability π and in such a case his claim will be channeled to the SIU with probability $Q^n(\theta)$. If he has not had an accident, he may decide to file a fraudulent claim, and he will actually do so if his moral cost ω is lower than $\phi\left(Q^f(\theta)\right)$ which occurs with probability $H\left(\phi\left(Q^f(\theta)\right)\middle|\theta\right)$. Hence, the expected investigation cost is:

$$c\pi E_\theta \; Q^n(\theta) + c(1-\pi) \; E_\theta \; Q^f(\theta) \; H\Big(\phi\big(Q^f(\theta)\big)\big|\theta\Big) \tag{3}$$

where E_θ denotes the mathematical expectation operator with respect to the probability distribution of θ over the whole population of insureds. The cost of residual fraud is written as:

$$t\,(1-\pi) \; E_\theta\big(1-Q^f(\theta)\big) \, H\Big(\phi\big(Q^f(\theta)\big)\big|\theta\Big) \tag{4}$$

The total expected cost of fraud is obtained by adding up (3) and (4). An optimal investigation strategy minimizes this total expected cost with respect to $q(\cdot):\Theta\times\Sigma\to[0,1]$. Such a strategy is characterized in the following proposition.

Proposition 1: *An optimal investigation strategy is such that*

$$q(\theta,\sigma_i)=0 \text{ if } i<i^*(\theta); \;\; q(\theta,\sigma_i)\in(0,1] \text{ if } i=i^*(\theta); \;\; q(\theta,\sigma_i)=1 \text{ if } i>i^*(\theta)$$

where $i^*(\theta)\in\{1,...,\ell\}$ *is a critical suspicion index that depends on the vector of individual-specific variables.*

Proposition 1 says that an optimal investigation strategy consists in subjecting claims to an SIU referral when the suspicion index *i* exceeds the individual-specific threshold $i^*(\theta)$, or equivalently when the probability of fraud is larger than a type-dependent threshold. Hence the insurer plays a so-called red flags strategy: for some signals σ_i – those for which $i > i^*(\theta)$ – claims are systematically audited at SIU, while there is no audit when $i<i^*(\theta)$[12]. Furthermore the optimal red flags strategy is type-dependent: if σ_i is perceived and $i^*(\theta_0)<i<i^*(\theta_1)$, then the claim should be sent to SIU if it has been filed by a type-θ_0 individual but no special investigation should be performed for a type-θ_1 insured.

Let $\tau(Q,\theta)=(1-\pi)\;H(\phi(Q)|\,\theta)$ and $\eta(Q,\theta)=-Q\phi'(Q)h(\phi(Q)|\,\theta)/H(\phi(Q)|\,\theta)>0$. $\tau(Q,\theta)$ is the fraud rate, i.e. the average number of fraudulent claims for a type-θ insured, when the probability of being detected is equal to Q.[13] $\eta(Q,\theta_1)$ is the elasticity of the fraud rate (in absolute value), i.e. the percentage decrease in the fraud rate following a one percent increase in the probability of detection.

Proposition 2: *Assume that* $(1-Q)H(\phi(Q)|\theta)$ *is convex in Q. If* $\tau(Q,\theta_1)\geq\tau(Q,\theta_0)$ *and* $\eta(Q,\theta_1)\geq\eta(Q,\theta_0)$, *with at least one strong inequality, then* $Q^f(\theta_1)>Q^f(\theta_0), Q^n(\theta_1)>Q^n(\theta_0)$ *and* $i^*(\theta_1)\leq i^*(\theta_0)$.

$(1-Q)H(\phi(Q)|\theta)$ is the rate of undetected fraud among the type-θ individuals who have not suffered any accident. It is decreasing from $H(\phi(0)|\theta)$ > 0 to 0 when Q goes from 0 to p_0. The convexity assumption made in Proposition 2 conveys the decrease in the marginal deterrence effect when the fraud detection probability is increasing. Auditing will cut fraud costs all the more efficiently if the insured belongs to a group with a high fraud and/or elasticity rate, hence the statement in Proposition 2.

The conditional probability of fraud $P(F|\sigma_i,\theta)$ is given by (1) with

$$P(F|\theta)=\frac{(1-\pi)H\left(\phi\left(Q^f(\theta)\right)|\theta\right)}{\pi+(1-\pi)H\left(\phi\left(Q^f(\theta)\right)|\theta\right)}. \tag{5}$$

When signal σ_i is perceived, the expected benefit of an SIU investigation is $P(F|\sigma_i,\theta)t-c$.

Proposition 3 shows that the optimal investigation strategy involves transmitting suspicious claims to the SIU in cases where the expected benefit of such a special investigation may be negative and consequently it highlights the importance of the insurer's commitment.

Proposition 3: *If* $\left(p^f_{i^*(\theta)}/p^n_{i^*(\theta)}\right)-\left(p^f_{i^*(\theta)-1}/p^n_{i^*(\theta)-1}\right)<c^2\pi\eta(Q,\theta)/\tau(Q,\theta)(t-c)(t-c+c\eta(Q,\theta))$, *then the optimal investigation strategy is such that* $P\left(F\left|\sigma_{i^*(\theta)},\theta\right.\right)t<c.$

The condition introduced in Proposition 3 ensures that the increase in $P(F|\sigma_i,\theta)$ is small when we go from $i=i^*(\theta)-1$ to $i=i^*(\theta)$, which makes sense when ℓ is large. The Proposition means that there exists $i^{**}(\theta)$ larger than $i^*(\theta)$ such that:

$$P\left(F\left|\sigma_{i^{**}(\theta)},\theta\right.\right)t\leq c<P\left(F\left|\sigma_{i^{**}(\theta)+1},\theta\right.\right)t. \tag{6}$$

Forwarding the claim to the SIU is profitable only if the suspicion index i is larger than $i^{**}(\theta)$. Hence, it is optimal to channel the claim to the SIU when $i^*(\theta) \le i \le i^{**}(\theta)$, although in such a case the expected profit drawn from investigation is negative. Indeed the investigation strategy acts as a deterrent: It dissuades some insureds (those with the highest moral costs) from defrauding. Such a strategy involves a stronger investigation policy than the one that would consist in transferring a claim to the SIU when the direct monetary benefits expected from investigation are positive. A consequence of this result is that the SIU should not be organized as a profit center of the insurance company, for otherwise its objective would be in conflict with the implementation of the optimal auditing strategy.

In practice, when ℓ is large, p_i^n and p_i^f are very small and we can write $Q^f(\theta) = \lambda(i^*(\theta))$ and $Q^n(\theta) = \mu(i^*(\theta))$, where[14]

$$\lambda(i) = \sum_{j=i}^{\ell} p_j^f \text{ and } \mu(i) = \sum_{j=i}^{\ell} p_j^n . \tag{7}$$

$\lambda(i)$ and $\mu(i)$, which are decreasing functions, respectively denote the probability of channeling a fraudulent claim and a non-fraudulent claim to the SIU when the critical index of suspicion is i. Using Proposition 1 allows us to reduce the insurer's optimization problem to the choice of the type-dependent suspicion threshold: $i^*(\theta)$ minimizes the total expected cost of fraud

$$c\pi\mu(i) + (1-\pi)H\left(\phi(\lambda(i))|\theta\right)\left(c\lambda(i) + t(1-\lambda(i))\right) \tag{8}$$

with respect to $i \in \{1,...,\ell\}$. For a type-θ individual, the expected cost attributable to fraud is the sum of $C^n(i) \equiv c\pi\mu(i)$ which is the expected investigation cost of non-fraudulent claims that are incorrectly referred to the SIU (i.e. the cost of type-2 errors) and of

$$C^f(\theta,i) = (1-\pi)H\left(\phi(\lambda(i))|\theta\right)\left(c\lambda(i) + t(1-\lambda(i))\right)$$

which is the expected cost of fraudulent claims. This cost includes the investigation cost of the claim channeled to the SIU and the cost of paying out unwarranted insurance indemnities. $C^n(i)$ and $C^f(\theta,i)$

are respectively decreasing and increasing[15] with respect to i. The optimal investigation strategy trades off excessive auditing of non-fraudulent claims against inadequate deterrence and detection of fraudulent claims. The optimal critical suspicion index $i^*(\theta)$ minimizes $C^n(i) + C^f(\theta,i)$.

(Figure 1 about here)

The optimal auditing policy is illustrated by the full lines in Figure 1. When i^* decreases from ℓ to 1, $\mu(i^*)$ and $\lambda(i^*)$ go from 0 to 1. In the literature on classification techniques, the locus $\{(\mu(i^*),\lambda(i^*)),\ i^*=1,...\ell\}$ is known as the Receiver Operating Characteristic (ROC) curve (Viaene, et al., 2002). The monotonicity of p_i^f / p_i^n with respect to i implies that the ROC curve is concave. The optimal auditing procedure minimizes the expected cost of fraud with respect to (μ, λ), under the constraint that (μ, λ) is on the ROC curve. Figure 1 shows the dependence of the optimal solution on the agent's type. If the fraud rate and the fraud elasticity are larger for θ_1 than for θ_0, then isocost curves single cross as in Figure 1, with larger optimal audit probability for θ_1, which illustrates Proposition 2.

In practice (and particularly for the calibration on real data), we may assume that the activity of the SIU is budget-constrained: Antifraud expenditures should be less than some (exogenously given) upper limit K, which gives the following additional constraint:[16]

$$c\pi E_\theta Q^n(\theta) + c(1-\pi) E_\theta Q^f(\theta) H\left(\phi\left(Q^f(\theta)\right)\middle|\theta\right) \leq K . \tag{9}$$

An optimal investigation strategy then minimizes the total expected cost of fraud with respect to $q(\cdot):\Theta\times\Sigma\to[0,1]$ subject to (9). Corollary 1 shows that the qualitative characterization of the antifraud policy is not affected by the addition of this upper limit on possible investigation expenditures.

Corollary 1: *Propositions 1 and 2 are still valid when the investigation policy is budget constrained.*

Under the optimal decision rule, a claim with signal σ_i is audited when $i \geq i^*(\theta)$. For type-θ individuals, the hit rate (i.e. the proportion of successful audits for each investigator at SIU) is then:

$$\frac{1}{(1+X_\theta)} \equiv \frac{P(F|\theta)Q^f(\theta)}{P(F|\theta)Q^f(\theta)+\left[1-P(F|\theta)\right]Q^n(\theta)}. \tag{10}$$

We may check that this hit rate varies with θ. In other words, the optimal investigation policy does *not* equalize the probability of success (i.e. of catching defrauders) across individuals' types. See Appendix A2 for an illustration.

The fact that the optimal hit rate is type-dependent affects the incentive mechanism that should be used by the insurance company to stimulate the activity of its staffs at SIU. In particular, paying a constant bonus each time SIU staffs catch a defrauder whatever his type is *not* an optimal incentive mechanism since it would lead investigators to concentrate on the claims with the highest hit rates and to neglect the other claims. At equilibrium, the hit rate would be the same for all audited claims and it would be lower for non-audited claims, which would not be optimal. On the contrary, if SIU staffs receive a bonus $b_\theta = K(1+X_\theta)$ for any type-θ hit, with $K > 0$, and if they are risk neutral with respect to their global earnings, then their own financial interest will not be in conflict with the optimal audit policy of the insurance company. In other words, the lower the probability of success, the larger must be the bonus to the auditor when a defrauder is caught.[17]

4. Extensions

4.1 *Signal manipulations by defrauders*

So far we have assumed that fraud signals cannot be manipulated by policyholders. We now extend our model to a case where defrauders may affect signals perceived by the insurer. To keep matters as simple as possible, we assume that there are k binary fraud indicators, each of them taking value 0 or 1, with $\ell = 2^k$. In what follows, "switching off" indicator j means changing σ_{ij} from 1 to 0 in the case indicator j would have been equal to 1 if there were no signal manipulation activity.

Potential defrauders are heterogeneous with respect to their ability to manipulate signals and this ability cannot be observed by the insurer. The distribution of manipulation ability among policyholders is described by random vector $s = (s_1,...,s_k) \in [0,1]^k$, where the random variable s_j denotes the probability for the defrauder to be able to switch off indicator *j*. When a type-*s* individual knows in advance (i.e. before filing a fraudulent claim) whether he will be able or not to switch off indicator *j*, we have $s_j = 1$ or 0. However our formulation includes the case where defrauders are unsure as to whether they will be in a position to switch off some indicators. Hence an individual is now characterized by observable characteristics θ and two unobservable characteristics: the moral cost of fraud ω and the signal manipulation ability *s*. For notational simplicity, we assume that *s* is not correlated with θ and ω.[18] Let α_j^f be the probability that indicator *j* is on when the file is fraudulent, assuming that the defrauder does not engage in signal manipulation activity. Likewise indicator *j* is on with probability α_j^n when the claim is not fraudulent with $\alpha_j^f > \alpha_j^n$. Hence, in the absence of signal manipulation, indicator *j* is more frequently generated when the claim is fraudulent than when it is honest.

A high degree of manipulability may fully cut out some fraud indicators of their informativeness. In what follows we consider a simple case where this is not the case, which requires that Es_j is not too large for all indicators *j* = 1,…,*k*. We also assume that the probability distributions of the *k* indicators are independent conditional on the fact that the claim is fraudulent or non fraudulent. Under this conditional independence assumption, when defrauders do not manipulate signals, we have:

$$p_i^f = P\left(\sigma = \sigma_i \middle| F\right) = \prod_{j \mid \sigma_{ij}=1} \alpha_j^f \prod_{j \mid \sigma_{ij}=0} \left(1 - \alpha_j^f\right) \tag{11}$$

$$p_i^n = P\left(\sigma = \sigma_i \middle| N\right) = \prod_{j \mid \sigma_{ij}=1} \alpha_j^n \prod_{j \mid \sigma_{ij}=0} \left(1 - \alpha_j^n\right) \tag{12}$$

for all $i = 1,...,\ell$. Defrauders with manipulation ability *s* generate signal σ_i with probability

$$p_i^f\left(s\right) = \prod_{j \mid \sigma_{ij}=1} \alpha_j^f\left(1 - s_j\right) \prod_{j \mid \sigma_{ij}=0} \left(1 - \alpha_j^f + \alpha_j^f s_j\right). \tag{13}$$

The optimal audit strategy minimizes the total expected cost of fraud which may be written as

$$c\pi E_\theta Q^n\left(\theta\right) + \left(1-\pi\right) E_{\theta,s}\left(t - \left(t-c\right) Q^f\left(\theta,s\right)\right) H\left(\phi\left(Q^f\left(\theta,s\right)\right) \middle| \theta\right) \tag{14}$$

where $Q^n(\theta)$ is given by (2) and

$$Q^f(\theta,s)=\sum_{i=1}^{\ell} p_i^f(s)\,q(\theta,\sigma_i). \tag{15}$$

Proposition 4: *An optimal investigation strategy under manipulation of fraud signals is characterized by a threshold $L(\theta)>0$ and weights $\delta(\theta,s)>0$ for all θ, s such that*

$$q(\theta,\sigma_i)=0 \quad if \quad \frac{\overline{p}_i^f(\theta)}{p_i^n}<L(\theta)$$

$$q(\theta,\sigma_i)=1 \quad if \quad \frac{\overline{p}_i^f(\theta)}{p_i^n}>L(\theta)$$

where $\overline{p}_i^f(\theta)=E_s\left[\delta(\theta,s)p_i^f(s)\right]$ *for all i and* $E_s\delta(\theta,s)=1$.

Proposition 4 extends our results on the optimality of the red flags strategy to the case where fraud indicators can be manipulated by defrauders. Claims are not audited when the insurer perceives signals σ_i such that $\frac{\overline{p}_i^f(\theta)}{p_i^n}<L(\theta)$ and they are systematically audited when $\frac{\overline{p}_i^f(\theta)}{p_i^n}>L(\theta)$, where $\overline{p}_i^f(\theta)$ is a weighted average of the $p_i^f(s)$, with weight $\delta(\theta,s)$.[19] Equivalently a SIU referral should be triggered when the probability of fraud is large enough, under the assumption that the population of type-θ defrauders is similar to a representative defrauder who sends out signal σ_i with probability $\overline{p}_i^f(\theta)$.

For the sake of consistency, it remains to be checked that it is in the interest of defrauders to switch off fraud indicators when possible. Consider two vector signals σ_i and $\sigma_{i'}$, and indicator h such that $\sigma_{ih}=0$, $\sigma_{i'h}=1$ and $\sigma_{ij}=\sigma_{i'j}$ for all $j\neq h$. Let

$$A_{ii'}(\theta)\equiv\frac{\overline{p}_{i'}^f(\theta)}{p_{i'}^n}-\frac{\overline{p}_i^f(\theta)}{p_i^n}.$$

Proposition 4 shows that $q(\theta,\sigma_{i'}) \geq q(\theta,\sigma_i)$ if $A_{ii'}(\theta) \geq 0$. Hence a sufficient condition for switching off indicator h to be an optimal strategy of defrauders may be written as $A_{ii'}(\theta) \geq 0$ for all signals σ_i and $\sigma_{i'}$ with the above characterization.[20] A simple calculation gives

$$A_{ii'}(\theta) = E_s \left\{ \frac{\alpha_h^f (1-s_h) - \alpha_h^n}{\alpha_h^n \left(1-\alpha_h^n\right)} \; \delta(\theta,s) \prod_{\substack{j|\sigma_{ij}=1 \\ j\neq h}} \frac{\alpha_j^f \left(1-s_j\right)}{\alpha_j^n} \prod_{\substack{j|\sigma_{ij}=0 \\ j\neq h}} \frac{\left(1-\alpha_j^f + \alpha_j^f s_j\right)}{1-\alpha_j^n} \right\}.$$

When the s_j are independently distributed, we can check that

$$\alpha_h^f \left(1 - Es_h\right) \geq \alpha_h^n \frac{\underset{s}{\text{Max}}\, \delta(\theta,s)}{\underset{s}{\text{Min}}\, \delta(\theta,s)} \tag{16}$$

is a sufficient condition for $A_{ii'}(\theta) \geq 0$. The case where all defrauders have the same manipulation ability s is simpler. (16) is then replaced with $\alpha_h^f \left(1-s_h\right) > \alpha_h^n$, which simply means that indicator h is still informative about fraud despite the fact that a proportion s_h of defrauders are able to switch it off. The ROC curve shifts downward when the manipulation ability of the representative defrauder goes from 0 to s, hence an efficiency loss as shown by the dotted lines in Figure 1.

4.2 *Imperfect information on the auditing strategy*

So far we have assumed that potential defrauders can observe the audit probability $Q^f(\theta)$.[21] Although such an assumption is widely made in the literature on optimal auditing, one may argue that individuals have imperfect information about the insurer's auditing strategy: they just have beliefs about the probability of being spotted if they defraud. However beliefs may be affected by actual auditing decisions either when insurers spread information about their claims verification procedure or from other insureds by word of mouth. We can check that our conclusions still hold in such a setting.

Let us focus on the case where signals cannot be manipulated. Assume that policyholders do not know $Q^f(\theta)$. Prior to the current period, they just have subjective beliefs $\tilde{Q}(\theta)$ about the audit frequency. $\tilde{Q}(\theta)$ is a random variable which is distributed on [0,1]. Let T be the number of fraudulent claims an individual has heard of during the current period before making up his (her) mind to file a fraudulent-claim or not.[22] T is random with probability $P(T)$ for T = 0,1,2... and $\sum_{t=0}^{+\infty} P(T) = 1$. For each claim in this sample, the individual knows whether it has been audited or not. Let $X \in \{0,1,...,T\}$ denote the number of claims which have been audited among the T claims sample. X is distributed according to a binomial law $B\left(T, Q^f(\theta)\right)$. Type-$\theta$ individuals decide to defraud on the basis of the revised belief $E\left(\tilde{Q}(\theta)|X,T\right)$, which is the expected value of $\tilde{Q}(\theta)$ conditionally on X,T.[23] Since type-θ individuals file a fraudulent claim if $\omega \leq \phi\left(E\left(\tilde{Q}(\theta)|X,T\right)\right)$, their average fraud rate is written as

$$(1-\pi)\sum_{T=0}^{+\infty} P(T) \sum_{X=0}^{T} P\left(X \middle| T, Q^f(\theta)\right) H\left[\phi\left(E\left(\tilde{Q}(\theta)|X,T\right)\right)|\theta\right] \tag{17}$$

Substituting (17) to $(1-\pi)H\left[\phi\left(Q^f(\theta)\right)|\theta\right]$ in (3) and (4) lets our results qualitatively unchanged. Indeed the fraud rate remains a decreasing function of $Q^f(\theta)$ because, for any value of T, an increase in $Q^f(\theta)$ shifts the distribution of X in the sense of first order dominance and $E\left(\tilde{Q}(\theta)|X,T\right)$ is increasing in X. In particular, Proposition 1 on the optimality of red flags strategies and its consequences still hold. This approach can be easily extended to the case where individuals can manipulate signals by substituting $Q^f(\theta,s)$ to $Q^f(\theta)$ and possibly by conditioning prior beliefs $\tilde{Q}(\theta)$ on s.

5. Regression Analysis

Data are drawn from the automobile claims (theft and collision) of a large European insurer. Its audit strategy uses fraud indicators that are analyzed by an investigators team. Some of these claims had been spotted as fraudulent by the investigators, while other ones could be reasonably considered as truthful. The first group of files (A) comes from the company's SIU. This is the population of claims referred to this unit over a given period by claims handlers suspecting fraud. Of the 857 files referred to the SIU, 184 contained no fraud and 673 were classified as cases of either established or suspected fraud. As in Belhadji et al. (2000), we considered all these 673 files as fraudulent. The second group of files (B) was selected from the population of claims that the insurer did not think contained any type of fraud during the same period of time. We first chose to randomly select only about 1,000 files in the control group, because the cost of compiling information on fraud indicators is very high. Over the 1,000 selected claims in group B, 945 were classified as without fraud. The 55 remaining files contained some indication of fraud. The 184 files without any fraud in A were transferred to B, yielding two groups of files (A' with fraud and B' without fraud) and showing that 37% of the files contained established or suspected fraud. The 55 files with some suspicion of fraud in B were not included in A'.

The econometric analysis allowed us to identify relevant variables that are correlated with the frequency of fraudulent claiming, i.e. fraud indicators or individual characteristics. For that purpose, we used the standard Logit model for binary choice.

(Table 1 about here)

Table 1 reports the regression results. A detailed description of the variables is presented in Appendix A4. The first column (without θ variables) in Table 1 is limited to variables identifying fraud indicators (the so-called red flags σ^j)[24]. All these variables are significant in explaining (positively) the probability that a file may contain either suspected or established fraud at a level of at least 95% (with one exception at 90%). The second column yields similar results but takes into account two additional θ variables that represent characteristics of policyholders. In connection with the theoretical part of the paper, these variables are used to approximate the individual private cost of fraud which includes a pure moral cost component but also a monetary cost component. For the sake of brevity, we here restrict attention to two significant variables: θ_7 and θ_{16} indicate owners of vehicles whose value does not match the policyholder's income and which are not covered by damage insurance (i.e. have third party insurance only). Implicitly, it is suggested that such people have a lower moral cost, hence a higher probability for filing a fraudulent claim.[25]

6. Model Calibration

Our sample didn't include a realistic proportion of fraudulent claims. Having the true proportion of fraudulent claims is not important for regression analysis but it is essential for the calibration of the optimal auditing model. We used bootstrapping techniques in order to obtain a final sample representing the estimated proportion of fraudulent claims in the company's portfolio (8%). Details are in Appendix A5. We now tackle the innovative part of the empirical analysis related to the calibration of the theoretical model.

Let $\hat{\pi}(\theta)$ be the probability that a type-θ individual files a claim during a one-year time period when there is no auditing (which is supposed to correspond to the status quo situation in the insurance company) and let t be the average cost of a claim for the insurer (average amount paid above the deductible). For the time being, we do not distinguish the groups of insureds: a type-θ individual is thus a representative policyholder of the company. Data from the company give $\hat{\pi}(\theta)$ = 22% and t = €1,284. The audit cost c of a claim is equal to €280 (including investigation costs, lawyers fees, SIU overheads, …) and the current proportion of claims with fraud is $z(\theta) = P(F|\theta) = 8\%$. Since $\hat{\pi}(\theta)$ contains fraudulent claims, the true loss probability π is given by $\pi = \hat{\pi}(\theta)(1 - z(\theta)) = 0.2024$. From the above data $\tau(0,\theta) = \hat{\tau}(\theta)$ can be approximated by $\hat{\tau}(\theta) = \hat{\pi}(\theta) z(\theta) = 0.0176$, which amounts to assuming that the observed current anti-fraud policy of the company does not entail any deterrence effect.

Estimating the elasticity of fraud with respect to the audit probability can only be a matter of approximation: Indeed, the elasticity $\eta(Q,\theta)$ depends on the distribution of moral costs in the population of policyholders as well as on the relationship between the audit probability, the moral cost, and the decision to file a fraudulent claim. In short, the elasticity of fraud with respect to Q depends on $H(\phi(Q)|\theta)$ and θ. Such information is obviously unobservable. This is why we will content ourselves with an approximation of $\tau(Q,\theta)$. We will assume:

$$\tau(Q,\theta) = \hat{\tau}(\theta)(1-Q)^{\gamma(\theta)}$$

where $\gamma(\theta)$ is a parameter used to define $\eta(Q,\theta) = \gamma(\theta)Q/(1-Q)$, the elasticity of the fraud rate with respect to Q. Increasing $\hat{\tau}(\theta)$ and/or $\gamma(\theta)$ allows us to perform simple comparative static analysis in order to illustrate the effect of the fraud rate and/or the fraud elasticity on the optimal audit strategy. The expected cost of fraud can then be rewritten as:

$$c\pi\mu(i) + \hat{\tau}(\theta)(1-\lambda(i))^{\gamma(\theta)}(t-\lambda(i)(t-c)) \tag{18}$$

where $\lambda(i)$ and $\mu(i)$ are given by (7). We will assume that the fraud indicators are independent conditional on the fact that the file is *F* or *N*.[26] p_i^f and p_i^n are then given by (11) and (12) which allows us to compute $\lambda(i)$ and $\mu(i)$ for all *i*.[27] The optimal threshold i^* is obtained by minimizing (18) with respect to *i*.

7. Calibration results

The calibration results are summarized in Table 2. Column 1 presents the identification indexes *i* of the signals σ_i including the threshold *i**. Table 2 has 2^{13} = 8,192 lines because the regression analysis identified 13 significant binary indicators. Signals in Column 2 result from various combinations of 0 and 1 where 0 (resp. 1) means that an indicator is off (resp. on). For example, the first line in Column 2 indicates that no significant fraud indicator is on. Line 2 indicates that only the 9th fraud indicator is on. Signals are ranked in such a way that p_i^f / p_i^n is increasing, as shown in Column 3.

(Table 2 about here)

Columns 4 and 5 yield $\lambda(i)$ and $\mu(i)$. Column 6 provides $C^n(i) + C^f(\theta,i)$, the expected cost of fraud under the assumption $\gamma(\theta) = 0$, i.e. ignoring the deterrence effect. Column 7 gives the average audit rate. Finally, the hit rate is at column 8. The optimal solution is at line 194 = $i^*(\theta)$ where the expected cost of fraud reaches its minimal value at €13.21. We then have $\lambda(i^*)$ = 0.67 and $\mu(i^*)$ = 0.04, which means that 67% of the fraudulent claims are audited while only 4% of the non-fraudulent claims are audited. The optimal strategy entails auditing 9.23% of the files (column 7) with hit rate $P\left(F\left|i > i*\right.\right) = 57.8\%$.

Sensitivity analysis allows us to illustrate Proposition 2. Restricting our sample to the claims files by policyholders such that $\theta_7 = \theta_{16} = 0$ yields the fraud rate $z\left(\theta\right) = 5{,}83\%$, while $z\left(\theta\right) = 10.01\%$ when $\theta_7 = 0$ and $\theta_{16} = 1$. Still assuming $\gamma\left(\theta\right) = 0$ gives $i^* = 250$ with audit rate 6.44% and hit rate 62.10% when $\theta_7 = \theta_{16} = 0$ and $i^* = 157$, with audit rate 15.84% and hit rate 69.52%. when $\theta_7 = 0$ and $\theta_{16} = 1$. We can also check that optimal auditing intensifies when $\gamma\left(\theta\right)$ is positive and the fraud elasticity increases (see Table 3 in Appendix).

The monetary gains of auditing can be estimated as follows. The claims rate (over the whole portfolio) of the insurer is 22%, which represents about 500,000 claims for the corresponding time period. Without auditing, the total claim cost is 500,000 $\times$ €1,284 = €642 M, including 8% of fraudulent claims (€51 M). When the optimal auditing strategy of Table 2 is implemented, then 9.23% of the files are audited at average cost €280. Furthermore, 67% of the fraudulent claims are audited without any insurance coverage. However, 33% of the fraudulent claims would not be audited. The total cost (including paid claims and audit costs) will then be equal to 9.23% $\times$ 500,000 $\times$ €280 + €591 M + 33% $\times$ €51 M = €621 M. The optimal audit strategy thus entails a saving of €21 M, which represents 41% of the current cost of fraudulent claims. This reduction in cost would even be larger if the audit probability depends on the policyholders' type and if the deterrence effect were taken into account. Appendix A7 shows how our numerical estimates are affected by the possibility of fraud manipulation.

8. Conclusion

This article aimed at making a bridge between the theory of optimal auditing and the actual claims auditing procedures used by insurers. More generally, we have developed an integrated approach to auditing and scoring which is much more closely related to the actual auditing procedures used by insurers, bankers, tax inspectors or governmental regulatory agencies than the abstract costly state-verification modeling. A complete modeling has been developed for the detection of insurance fraud, but the same methodology could be adapted to other hidden information problems.

We have shown that the optimal auditing strategy takes the form of a type dependent red flags strategy which consists in referring claims to the SIU when some fraud indicators are observed. The classification of fraud indicators corresponds to an increasing order in the probability of fraud and such a strategy remains optimal if the investigation policy is budget constrained. It is robust to some degree of signal manipulation and to imperfect information of policyholders. Because of its deterrence function, the optimal investigation strategy may lead to an SIU referral in cases where the direct expected gain of such a decision is negative. A strong commitment of the auditor is thus necessary for such a policy to be fully implemented. Finally the optimal hit rate depends on the policyholder's type, which affects the optimal incentive mechanism of SIU staffs. The properties of the optimal audit strategy have been implemented by building a simple model of automobile insurance calibrated with data from a large European insurance company.

Among the various possible extensions of this model, it would be particularly interesting to consider a multiperiod setting in order to analyze the adaptative behavior of defrauders who may progressively learn how to better conceal evidence on fraudulent claiming, as well as the reaction of insurers who adapt the set of red flags to the change in defrauders' behavior. Another useful extension would consist in inserting the audit problem into a more global model of the insurance market. In particular some characteristics of the policyholders may then be affected by the audit strategy (e.g. an individual may choose to purchase an expensive car while thinking of cheating on his insurance contract). The interaction between the specification of the insurance contract and the audit strategy would then be of crucial importance.

References

Antle, R., Gordon, E., Narayanamoorthy, G, and Zhou, L., 2006. The Joint Determination of Audit Fees, Non-Audit Fees, and Abnormal Accruals. *Review of Quantitative Finance and Accounting* 27, 235-266.

Artis, M., Ayuso, M., and Guillén, M., 2002. Detection of Automobile Insurance Fraud with Discrete Choice Models and Misclassified Claims. *Journal of Risk and Insurance* 69, 325-340.

Baiman S., and Rajan, M.V., 1994. On the Design of Unconditional Monitoring Systems in Agencies. *The Accounting Review* 69, 217-229.

Beasley, M., 1996. An Empirical Analysis of the Relation Between the Board of Director Composition and Financial Statement Fraud. *The Accounting Review* 71, 443-465.

Belhadji, E.B., Dionne, G., and Tarkhani, F., 2000. A Model for the Detection of Insurance Fraud. *Geneva Papers on Risk and Insurance Issues and Practice* 25, 517-538.

Beneish, M., 1997. Detecting GAAP Violation: Implications for Assessing Earnings Management Among Firms with Extreme Financial Performance. *Journal of Accounting and Public Policy* 16, 271-309.

Caplan, D., 1999. Internal Controls and the Detection of Management Fraud. *Journal of Accounting Research* 37, 101-117.

Christensen, P.O., Demski, J.S., and Frimor H., 2002. Accounting Policies in Agencies with Moral Hazard and Renegotiation. *Journal of Accounting Research* 40, 1071-1090.

Crocker, K.J. and Morgan, R.J., 1997. Is Honesty the Best Policy? Curtailing Insurance Fraud through Optimal Incentive Contracts. *Journal of Political Economy* 106, 355-375.

Crocker, K.J. and Tennyson, S., 2002. Insurance Fraud and Optimal Claims Settlement Strategies. *Journal of Law and Economics* 45, 469-507.

Dechow, P., Ge, W., Larson, C.R., and Sloan, R.G., 2007. Predicting Material Accounting Misstatements. Mimeo, University of Michigan and University of Washington.

Dechow, P., Sloan, R., and Sweeney, A., 1996. Causes and Consequences of Earnings Manipulation: An Analysis of Firms Subject to Enforcement Actions by the SEC. *Contemporary Accounting Research* 13, 1-36.

Derrig, R.A. and Weisberg, H.I., 2003. Auto Bodily Injury Claim Settlement in Massachusetts. Document, Automobile Insurers Bureau of Massachusetts, 36 p.

Dionne, G. and Gagné, R., 2001. Deductible Contracts Against Fraudulent Claims: Evidence from Automobile Insurance. *Review of Economics and Statistics* 83, 290-301.

Dionne, G. and Gagné, R., 2002. Replacement Cost Endorsement and Opportunistic Fraud in Automobile Insurance. *Journal of Risk and Uncertainty* 24, 213-230.

Duda, R.O., Hart, P.E., and Stork, E.G., 2001. *Pattern Classification*.Wiley, New York.

Dunn, P., 2004. The Impact of Insider Power on Fraudulent Financial Reporting. *Journal of Management* 30, 397-412.

Dye, R.A., 1993. Auditing Standards, Legal Liability, and Auditor Wealth. *Journal of Political Economy* 101, 887-914.

Gale, D. and Hellwig M., 1985. Incentive-Compatible Debt Contracts: The One-Period Problem. *Review of Economic Studies* 52, 647-663.

Graetz, M.J., Reinganum, J.F., and Wilde, L.L., 1986. The Tax Compliance Game: Toward and Interactive Theory of Law Enforcement. *Journal of Law, Economics and Organization* 2, 1-32.

Knowles, J., Persico N., and Todd P., 2001. Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy*, 109, 203-229.

Kofman, F. and Lawarrée, J., 1993. Collusion in Hierarchical Agency. *Econometrica* 61, 629-656.

Melumad, M.D. and Mookherjee, D., 1989. Delegation as Commitment: The Case of Income Tax Audits. *The RAND Journal of Economics* 20, 139-163.

Mookerjee, D. and Png, I.P.L., 1989. Optimal Auditing Insurance and Distribution. *Quarterly Journal of Economics* 104, 399-415.

Morton, S., 1993. Strategic Auditing for Fraud. *The Accounting Review* 68, 825-839.

Narayanan, V.G., 1994. An Analysis of Auditor Liability Rules. *Journal of Accounting Research* 32 (Supplement), 39-59.

Newman, D.P., Patterson, E., and Smith, R., 2001. The Influence of Potentially Fraudulent Reports on Audit Risk Assessment and Planning. *The Accounting Review* 76, 59-80.

Pae, S., and Yoo S.-W., 2001. Strategic Interaction in Auditing: An Analysis of Auditors' Legal Liability, Internal Control System Quality, and Audit Effort. *The Accounting Review* 76, 333-356.

Persico, N., 2002. Racial Profiling, Fairness, and Effectiveness of Policing. *The American Economic Review* 92, 1472-1497.

Picard, P., 1999. On the Design of Optimal Insurance Contracts under Manipulation of Audit Costs. *International Economic Review* 41, 1049-1071.

Picard, P., 2000. Economic Analysis of Insurance Fraud. In *Handbook of Insurance*, G. Dionne (Ed.), Kluwer Academic Publishers, Boston, 315-362.

Shibano, T., 1990. Assessing Audit Risk from Errors and Irregularities. *Journal of Accounting Research* 28, 110-140.

Simunic, D.A., 1984. Auditing, Consulting, and Auditor Independence. *Journal of Accounting Research* 22, 679-702.

Skousen, C.J. and Wright, C.J., 2006. Contemporaneous Risk Factors and the Prediction of Financial Statement Fraud. Working paper, University of Texas at Arlington.

Summers, S. and Sweeney, J., 1998. Fraudulently Misstated Financial Statements and Insider Trading: An Empirical Analysis. *The Accounting Review* 73, 131-146.

Townsend, R., 1979. Optimal Contracts and Competitive Markets with Costly State Verification. *Journal of Economic Theory* XXI, 265-293.

Townsend, R., 1988. Information Constrained Insurance: The Revelation Principle Extended. *Journal of Monetary Economics*, 21, 411-450.

Viaene, S., Derrig, R.A., Baesens, B., and Dedene, G., 2002. A Comparison of State-of-the-Art Classification Techniques for Expert Automobile Insurance Claim Fraud Detection. *Journal of Risk and Insurance* 69, 373-422.

Notes

[1] The papers by Knowles et al. (2001) and Persico (2002) are noteworthy exceptions to this separation between the literature strands on audit and scoring. They are about auditing strategies (more specifically automobile control by police officers) in a setting where auditors can condition the audit probability on the agents' observable particulars and they focus on the detection of possible bias in the auditors' behavior (e.g. racial bias in vehicle searches).

[2] Crocker and Morgan (1997) have developed a *costly state falsification* approach to insurance fraud which has conceptual similarities with the models of costly state verification with audit cost manipulation. Other references on this issue are Crocker and Tennyson (2002), and Picard (1999). See Picard (2000) for an overview.

[3] See Dionne and Gagné (2001, 2002), Derrig and Weisberg (2003), Artis et al. (2002) and Crocker and Tennyson (2002) for different econometric applications, and volume 69, no 3, of the *Journal of Risk and Insurance*, September 2002, for a state-of-the-art presentation of claims fraud detection methods.

[4] A similar result is obtained by Persico (2002) in a different context. He analyses the behavior of police officers who choose whom to investigate when citizens of two groups may engage in crime. The equilibrium search is obtained when both groups have the same fractions of criminals, while maximal search effectiveness usually entails different hit rates between the two groups.

[5] The differences in moral cost may more generally reflect the heterogeneity in the private economic gains and losses derived from fraud. For instance, some defrauders may need to reach a more or less costly collusive agreement with service providers (e.g. car repairers) or, in a slightly different version of the model, some policyholders may face wealth constraints and those who have less to lose would choose to defraud. Our qualitative conclusions would be unchanged in such settings.

[6] θ gathers the whole information about the policyholders that may be available to the insurer, including civil status, occupation, urban or rural environment and possibly income and wealth. θ thus corresponds to a bundle of individual characteristics that may affect moral cost, hence the statistical link between ω and θ. There is no intrinsic ordering with respect to fraud propensity in the components of θ and moral cost may be affected by the interactions of these characteristics. For instance, in the case of car insurance, data may reveal that on average young urban drivers have less morality than older urban drivers (or conversely), while age does not affect the morality of rural drivers. A limitation of our study is that we consider the distribution of types among policyholders as well as their insurance contract as given. A most useful extension of the present paper would consist in inserting our analysis of

the optimal audit strategy in an insurance market model where insurers would interact and try to attract specific types of customers by making more or less advantageous offers to insurance seekers (and possibly by not offering any coverage to some types). Ideally, the (type dependent) insurance contracts, the distribution of policyholders' types among insurers and their auditing strategies could then be jointly determined at the equilibrium of such a more general model.

[7] The penalty B does not play any crucial role in the model and $B = 0$ is a possible case. In a dynamic setting the cost to individuals caught defrauding may include a non-monetary reputation cost when defrauders have difficulty to get insurance in the future.

[8] If W_0 and t were type dependent, then ϕ would be a function of p and θ without entailing any substantial consequence in the model. Note that laws prohibiting discrimination may prevent insurers from conditioning coverage or premium on type.

[9] See Section 4.1 on signal manipulation.

[10] We assume that p_i^f and p_i^n do not depend on θ. See Appendix A8 for tests of this assumption.

[11] Of course if $p_i^n = 0$ and $p_i^f > 0$ then the optimal investigation strategy involves channeling the claim to SIU (see the definition and the role of SIU hereafter) when $\sigma = \sigma_i$. Indeed the claim is definitely fraudulent in such a case.

[12] The only case where a random investigation may be optimal is for $i = i^*(\theta)$.

[13] If type θ_0 morality dominates type θ_1 in the sense of moral cost first-order stochastic dominance, then $\tau(Q,\theta_0) < \tau(Q,\theta_1)$ for all Q.

[14] Hence, we here disregard the fact that we may have $q(\theta, \sigma_i) < 1$ when $i = i^*(\theta)$.

[15] If $c > t$ then $C^f(\theta,i)$ may be (at least locally) decreasing in i and a corner solution at $i^*(\theta) = \ell$ is possible.

[16] Constraint (9) may also be the consequence of a fixed number of investigators at SIU, each of them being able to audit at most a certain number of claims. K/c then denotes the maximum number of audits per policyholder.

[17] See Appendix A3.1 for a simple model of auditors' incentives that leads to this result and Appendix A3.2 for a discussion on the revelation principle in relation with the model of this article.

[18] In a multiperiod setting, manipulation ability could be the outcome of a learning process. Since low moral cost individuals derive more advantage from learning how to cheat on their insurance policy, a negative correlation between manipulation ability and moral cost would then make sense.

[19] The proof of Proposition 4 shows that weights $\delta(\theta,s)$ are proportional to the decrease in the expected cost of fraud from type θ,s individuals following a unit increase in the audit probability for these individuals. Hence, when increasing the audit probability entails more beneficial effects for type θ,s than for type θ,s', the weight of group s should be larger than the weight of group s'.

[20] More precisely, under this condition it is optimal to manipulate indicators when all defrauders do so.

[21] When there is no signal manipulation, individuals do not need to know that the audit decision is made on the basis of fraud indicators: they just know the probability of being spotted if they file a fraudulent claim. Under signal manipulation, type θ,s individuals know that they will be spotted with probability $Q^f(\theta,s)$.

[22] In this modeling T cannot be chosen by defrauders, i.e. the defrauders' learning procedure is exogenous. The model could be extended to a more realistic setting in which the T claims sample includes unknown proportions of fraudulent and non-fraudulent claims.

[23] The probability distribution of the revised beliefs follows from Bayes' law by using the fact that X follows a binomial law $B\left(T,Q^f(\theta)\right)$. If $\tilde{Q}(\theta)$ has a beta distribution, then $E\left(\tilde{Q}(\theta)\middle|X,T\right)$ is a weighted average of the prior expected value and X/T.

[24] Some indicators were directly available in the data warehouse while others were in the paper files.

[25] See footnote 5 for alternative interpretations.

[26] The conditional correlations are very small in our data set. The correlation tables are in Appendix A5.

[27] This procedure allows us to estimate the ROC curve. It is known in the literature as the simple Bayes classifier method which is equivalent to the Bayes optimal classifier only when all predictors are independent in a given class. It has been shown that this simple Bayes classifier often outperforms more powerful classifiers (Duda et al., 2001).

Tables and Figures

Figure 1

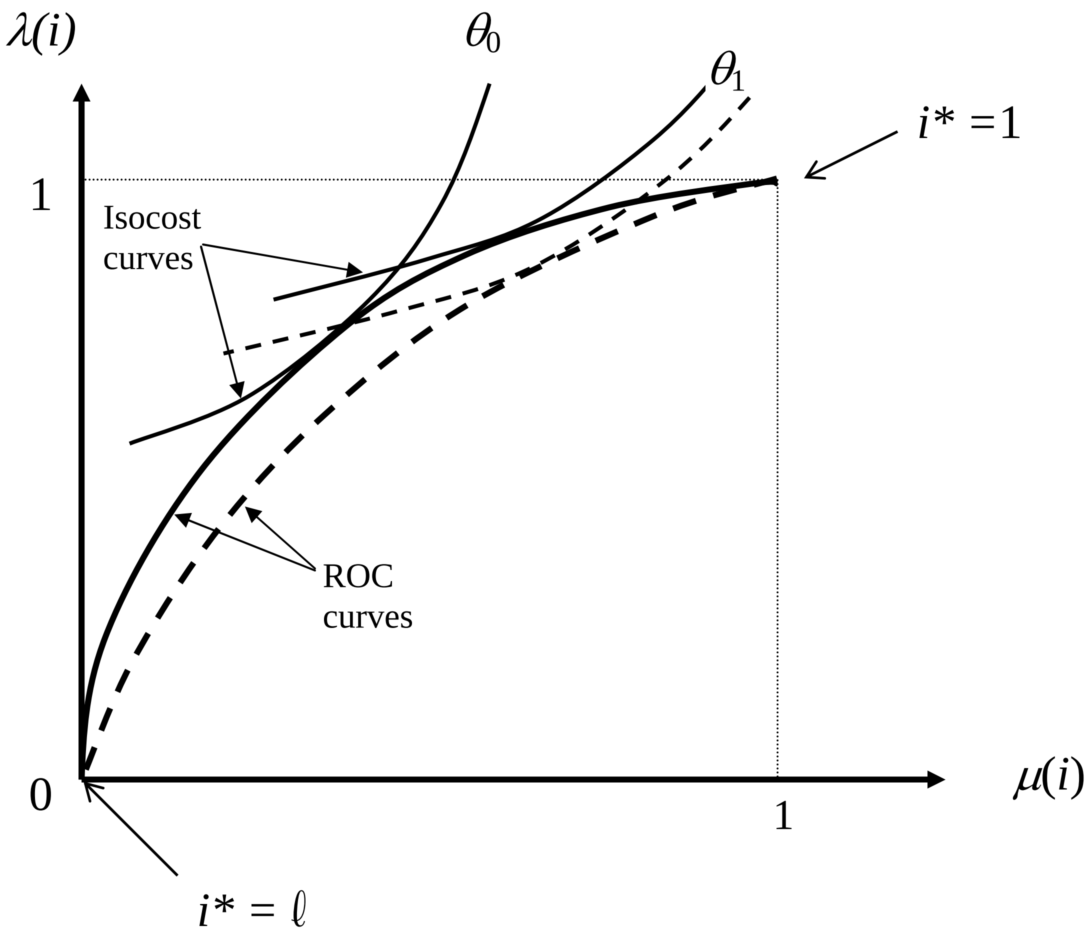

Table 1
Regression Results

		Without θ variables		**With θ variables**	
σ^j	Variable	Parameter	P	Parameter	P
	Intercept	5.8157	<0.0001	6.4780	<0.0001
σ^1	Fraud alert by expert	1.6365	<0.0001	1.6299	<0.0001
σ^2	Fraud alert by external organizations	1.4072	<0.0001	1.4387	<0.0001
σ^{12}	Retroactive effect of the contract	0.2103	0.0674	0.2329	0.0455
σ^{18}	Delay in filing accident claim	0.1808	0.0087	0.1567	0.0247
σ^{19}	Falsified documents or duplicated bills	1.2227	<0.0001	1.2287	<0.0001
σ^{20}	Reluctance to provide original documents	0.6099	<0.0001	0.5980	<0.0001
σ^{21}	Variation in the initial claim	0.5681	<0.0001	0.5218	<0.0001
σ^{22}	Harassment to obtain quick settlement	0.6197	0.0001	0.5933	0.0004
σ^{30}	Lack of clarity in the accident description	0.3590	<0.0001	0.3546	<0.0001
σ^{32}	Abnormal high frequency of accident	0.2352	0.0075	0.2404	0.0068
σ^{34}	Claim maker different to policyholder	0.2044	0.0246	0.1999	0.0297
σ^{35}	Too long lag between car purchase and guarantee	0.6714	<0.0001	0.6340	<0.0001
σ^{36}	Date of subscription to guarantee too close	0.4520	<0.0001	0.4264	<0.0001
θ_7	Too high value of vehicle			0.7742	0.0002
θ_{16}	No damage insurance			0.2302	0.0009
Log Likelihood		–694.72		–682.55	
Number of observations		1,802		1,802	

Table 2
Calibration Results

1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
i or $i*$	σ_i	p_i^f / p_i^n	$\lambda(i)$	$\mu(i)$	$C^n(i)+C^f(\theta,i)$	$z(\theta)\lambda(i)+(1-z(\theta))\mu(i)$	$P\ (F/i>i*)$
1	0000000000000	0.0779	1.0000	1.0000	61.600	1.0000	0.08000
2	0000000010000	0.1146	0.9816	0.7633	48.513	0.7808	0.10057
…							
10	1000010000000	0.2019	0.9382	0.4364	30.750	0.4765	0.15751
…							
…							
***i** = 194*	**0000110001000**	**3.1979**	**0.6672**	**0.0423**	**13.208**	**0.0923**	**0.57810**
…							
…							
8 192	1111111111111	63291108.94	0.0000	0.0000	22.598	0.0000	1.00000

Supplement

Table 3
Sensitivity of Optimal Solutions With Respect to $\gamma(\theta)$ and $z(\theta)$

Parameters					Optimal Threshold	Expected Cost of Fraud	Expected Audit Probability	Audit Probabiliy of Defrauders
$z(\theta)$	$\eta(\theta)$ [1]	$\gamma(\theta)$	π	$\tau(\theta)$	i^*	$C^n(i^*(\theta)) + C^f(\theta, i^*(\theta))$	$z(\theta)\lambda(i^*) + (1-z(\theta)) \times \mu(i^*)$	$\lambda(i^*)$
0.0800	0.00	0.00	0.2024	0.0176	194	13.21	0.0923	0.6672
0.0800	0.10	0.05	0.2024	0.0176	187	12.63	0.0937	0.6709
0.0800	0.21	0.10	0.2024	0.0176	183	12.08	0.0971	0.6798
0.0800	0.75	0.35	0.2024	0.0176	178	9.74	0.0979	0.6818
0.0583	0.00	0.00	0.2024	0.0128	250	10.15	0.0644	0.6210
0.0583	0.08	0.05	0.2024	0.0128	246	9.75	0.0646	0.6216
0.0583	0.17	0.10	0.2024	0.0128	240	9.37	0.0652	0.6238
0.0583	0.59	0.35	0.2024	0.0128	223	7.71	0.0669	0.6295
0.1001	0.00	0.00	0.2024	0.0220	157	15.84	0.1164	0.6952
0.1001	0.11	0.05	0.2024	0.0220	155	15.10	0.1170	0.6965
0.1001	0.23	0.10	0.2024	0.0220	153	14.39	0.1174	0.6974
0.1001	0.85	0.35	0.2024	0.0220	149	11.45	0.1232	0.7095

(1) $\eta(\theta) = -(\lambda(i^*))/1 - (\lambda(i^*))\gamma(\theta)$ in absolute value.

Appendix A1: Proofs

Proof of Proposition 1

If $Q^f(\theta) \geq p_0$, there is no residual fraud and reducing $q(\theta,\sigma_i)$ for any i decreases the investigation cost, hence a decrease in the total cost. Thus the optimal investigation strategy involves $Q^f(\theta) < p_0$, with $\phi'\left(Q^f(\theta)\right) < 0$. Using equations (2) to (4), pointwise minimization of the total expected cost of fraud with respect to $q(\theta,\sigma)$ gives:

$$c\pi p_i^n + (1-\pi)\ A_1\left(Q^f(\theta),\theta\right) p_i^f \begin{cases} \leq 0 & if \quad q(\theta,\sigma_i) = 1 \\ = 0 & if \quad 0 < q(\theta,\sigma_i) < 1 \\ \geq 0 & if \quad q(\theta,\sigma_i) = 0 \end{cases}$$

where $A(Q,\theta) = \left(cQ + t(1-Q)\right) H\left(\phi(Q) \middle| \theta\right)$ and A_1 denotes the partial derivative of A with respect to Q. Note that $\phi' < 0$ and $t > c$ give $A_1 < 0$, which proves the proposition, with $i^*(\theta)$ given by:

$$\frac{p^f_{i^*(\theta)-1}}{p^n_{i^*(\theta)-1}} < \frac{-c\pi}{(1-\pi)A_1\left(Q^f(\theta),\theta\right)} \leq \frac{p^f_{i^*(\theta)}}{p^n_{i^*(\theta)}} \tag{19}$$

Q.E.D.

Proof of Proposition 2

Let θ_0 and θ_1 in Θ such that $\tau\left(Q,\theta_1\right) \geq \tau\left(Q,\theta_0\right)$ and $\eta\left(Q,\theta_1\right) \geq \eta\left(Q,\theta_0\right)$ with at least one strong inequality. Assume moreover that $i^*(\theta_1) > i^*(\theta_0)$, which gives $Q^f(\theta_1) < Q^f(\theta_0)$. Let $i \in \{1,...\ell\}$ such that $i^*(\theta_0) \leq i < i^*(\theta_1)$. Writing optimality conditions as in the proof of Proposition 1 yields:

$$c\pi p_i^n + (1-\pi) A_1\left(Q^f(\theta_0),\theta_0\right) p_i^f \leq 0 \tag{20}$$

and

$$c\pi p_i^n + (1-\pi) A_1\left(Q^f(\theta_1),\theta_1\right) p_i^f \geq 0 \tag{21}$$

Using $Q^f(\theta_1) < Q^f(\theta_0)$ and the convexity of $Q \to A(Q,\theta)$ gives:

$$A_1\left(Q^f(\theta_0),\theta_1\right) > A_1\left(Q^f(\theta_1),\theta_1\right). \tag{22}$$

(21) and (22) give:

$$c\pi p_i^n + (1-\pi) A_1\left(Q^f(\theta_0), \theta_1\right) p_i^f > 0\,. \tag{23}$$

(20) and (23) then imply:

$$A_1\left(Q^f(\theta_0), \theta_1\right) > A_1\left(Q^f(\theta_0), \theta_0\right). \tag{24}$$

We have

$$A_1(Q,\theta) = (c-t) H\left(\phi(Q) \mid \theta\right) + \phi'(Q) h\left(\phi(Q) \mid \theta\right)\left(cQ + t(1-Q)\right)$$

which may be rewritten as:

$$A_1(Q,\theta) = -\frac{\tau(Q,\theta)}{1-\pi}\left(t - c + \frac{cQ + t(1-Q)}{Q}\eta(Q,\theta)\right).$$

The assumptions on $\tau(Q,\theta)$ and $\eta(Q,\theta)$ give

$$A_1\left(Q^f(\theta_0), \theta_1\right) < A_1\left(Q^f(\theta_0), \theta_0\right)$$

which contradicts (24). Hence, we may conclude that $i^*(\theta_1) \le i^*(\theta_0)$, which completes the proof. Q.E.D.

Proof of Proposition 3

We have

$$A_1\left(Q^f(\theta), \theta\right) < (c-t)\; H\left(\phi\left(Q^f(\theta)\right) \mid \theta\right).$$

Hence

$$\frac{-c\pi}{(1-\pi)\; A_1\left(Q^f(\theta), \theta\right)} < \frac{-c\pi}{(1-\pi)(c-t) H\left(\phi\left(Q^f(\theta)\right) \mid \theta\right)} \tag{25}$$

which gives:

$$\frac{p_{i^*(\theta)}^f}{p_{i^*(\theta)}^n} < \frac{-c\pi}{(1-\pi)\; (c-t)\; H\left(\phi\left(Q^f(\theta)\right) \mid \theta\right)} \tag{26}$$

when $\left(p_{i^*(\theta)}^f / p_{i^*(\theta)}^n\right) - \left(p_{i^*(\theta)-1}^f / p_{i^*(\theta)-1}^n\right)$ is small enough. We can check that this is the case under the condition stated in the Proposition.

Using (6), (26) and

$$\frac{p_{i^*(\theta)}^f}{p_{i^*(\theta)}^n} = \frac{\left(1-P\left(F \mid \theta\right)\right) P\left(F \;\middle|\; \sigma_{i^*(\theta)},\theta\right)}{P\left(F \mid \theta\right) \left(1-P\left(F \;\middle|\; \sigma_{i^*(\theta)},\theta\right)\right)}$$

gives $$P\left(F \;\middle|\; \sigma_{i^*(\theta)},\theta\right) t < c.$$ Q.E.D.

Proof of Corollary 1

Let χ be a (non-negative) Kuhn-Tucker multiplier associated with (9) when the total expected cost of fraud is minimized with respect to $q(\theta,\sigma)$ subject to (6) and (9). Pointwise minimization gives:

$$c\pi\left(1+\chi\right)p_i^n + \left(1-\pi\right)\tilde{A}_1\left(Q^f\left(\theta\right),\theta\right)p_i^f \begin{cases} \leq 0 & \text{if } \; q\left(\theta,\sigma_i\right)=1 \\ =0 & \text{if } \; 0<q\left(\theta,\sigma_i\right)<1 \\ \geq 0 & \text{if } \; q\left(\theta,\sigma_i\right)=0 \end{cases}$$

where $\tilde{A}\left(Q,\theta\right)=\left(c\left(1+\chi\right)Q+t\left(1-Q\right)\right)H\left(\phi\left(Q\right)\middle|\,\theta\right)$ and $\tilde{A}_1$ is the partial derivative of $\tilde{A}$ with respect to Q. Corollary 1 can then be proved in the same way as Propositions 1 and 2. Q.E.D.

Proof of Proposition 4

Using equations (2) and (15), pointwise minimization of total expected cost of fraud with respect to $q\left(\theta,\sigma\right)$ gives

$$c\pi p_i^n + \left(1-\pi\right)E_s A_1\left(Q^f\left(\theta,s\right),\theta\right)p_{is}^f \begin{cases} \leq 0 & \text{if } q\left(\theta,\sigma_i\right)=1 \\ =0 & \text{if } 0<q\left(\theta,\sigma_i\right)<1. \\ \geq 0 & \text{if } q\left(\theta,\sigma_i\right)=0 \end{cases}$$

Let

$$\delta\left(\theta,s\right)=\frac{A_1\left(Q^f\left(\theta,s\right),\theta\right)}{E_{s'}A_1\left(Q^f\left(\theta,s'\right),\theta\right)}.$$

We then get

$$c\pi p_i^n + \left(1-\pi\right)\overline{p}_i^f\left(\theta\right)E_s A_1\left(Q^f\left(\theta,s\right),\theta\right) \begin{cases} \leq 0 & \text{if } q\left(\theta,\sigma_i\right)=1 \\ =0 & \text{if } 0<q\left(\theta,\sigma_i\right)<1. \\ \geq 0 & \text{if } q\left(\theta,\sigma_i\right)=0 \end{cases}$$

Proposition 4 then follows from

$$L(\theta) = -\frac{c\pi}{(1-\pi) E_s A_1\left(Q^f(\theta, s), \theta\right)} > 0. \qquad \text{Q.E.D.}$$

Appendix A2: Illustration of a Type-dependent Investigation Policy

Let $\lambda_\theta \equiv \lambda\left(i^*(\theta)\right)$ and $\mu_\theta \equiv \mu\left(i^*(\theta)\right)$. When (9) is binding, the optimal audit strategy minimizes the expected cost of residual fraud

$$E_\theta\left\{tH\left(\phi(\lambda_\theta)|\theta\right)(1-\lambda_\theta)\right\}$$

with respect to λ_θ and μ_θ, $\theta \in \Theta$, subject to the upper limit on the number of audits:

$$E_\theta\left\{\pi\mu_\theta + (1-\pi)H\left(\phi(\lambda_\theta)|\theta\right)\lambda_\theta\right\} \leq \frac{K}{c} \tag{27}$$

with $(\lambda_\theta, \mu_\theta)$ on the ROC curve. For type-θ policyholders, the hit rate is $1/1+X_\theta$ with:

$$X_\theta = \frac{\left[1-P\left(F|\theta\right)\right]Q^n(\theta)}{P\left(F|\theta\right)Q^f(\theta)} = \frac{\pi\mu_\theta}{(1-\pi)H\left(\phi(\lambda_\theta)|\theta\right)\lambda_\theta}.$$

Writing the first-order optimality condition yields

$$X_\theta = \frac{1}{\xi(1-\pi)\nu_\theta}\left[\left(t-\xi(1-\pi)\right)(1-\eta_\theta) + \frac{t\eta_\theta}{\lambda_\theta}\right]$$

where ξ is a Lagrange multiplier associated to (27), $\eta_\theta \equiv \eta(\lambda_\theta, \theta)$, and ν_θ is the elasticity of the ROC curve at $(\lambda_\theta, \mu_\theta)$. For illustrative purpose, assume that the ROC curve equation is $\mu = \lambda^\nu$ with $\nu > 1$, which gives $\nu_\theta = \nu$ for all θ. Assume also $\eta_\theta = \eta$ for all θ. We then have:

$$X_\theta = \frac{1}{\xi(1-\pi)\nu}\left[\left(t-\xi(1-\pi)\right)(1-\eta) + \frac{t\eta}{\lambda_\theta}\right] \tag{28}$$

for all θ. If $\lambda_{\theta_1} > \lambda_{\theta_0}$, which holds under the conditions stated in Proposition 2, then $X_{\theta_1} < X_{\theta_0}$ from (28). The hit rate is then larger for type θ_1 than for type θ_0.

Appendix A3: Incentives of SIU's Staffs and Revelation Principle

Appendix A3.1: Incentives of SIU's Staffs

Assume that scrutinizing claims allows the SIU auditors to discover with certainty whether these claims are fraudulent or truthful. The non contractible effort of auditors defines the number of claims that they actually scrutinize. In fact auditors may pretend to have carefully examined some claims without having found anything fraudulent, while they just have validated those claims without making any careful inquiry. The only verifiable outcome of their activity is the number of fraudulent claims they discover.

Assume that auditors' earnings include a fixed wage w_0 and a bonus $b(\theta)$ for each type θ catch. Let $n(\theta)$ be the number of type θ scrutinized claims: $n(\theta)$ cannot be verified.

Let $R-C(N)$ be the auditor's net earnings, where *R* denotes the monetary earnings and *N* is the number of scrutinized files with effort cost $C(N)$, $C'>0$, $C''>0$.

Let us assume that the auditors are risk neutral. We have

$$ER = w_0 + \sum_{\theta\in\Theta} \frac{n(\theta)b(\theta)}{1+X_\theta}$$

$$N = \sum_{\theta\in\Theta} n(\theta)$$

Each auditor chooses $n(\cdot):\Theta \to R^+$ in order to maximize $ER-C(N)$. His behavioral function $n(\cdot)$ will be such that $n(\theta)>0$ for all θ only if $b(\theta)=K(1+X_\theta)$, with $K=C'(N)$.

Appendix A3.2: Revelation Principle

Using the ROC curve allows us to highlight the specificity of scoring-based audit procedures by referring to the Revelation Principle. Let $x\in\{A,N\}$ by the insured's type for given observable data. Here type is defined at the interim stage, i.e. after an accident has occurred $(x=A)$ or not $(x=N)$. A type-dependent allocation may be written as $\{y_A, y_N, Q^f, Q^n\}$ where $y_x\in[0,1]$ is the probability that a type-*x* individual files a claim and Q^f, Q^n are the audit probabilities for types N and A respectively, which are linked by the ROC curve $Q^n = f\left(Q^f\right)$.

For illustrative purpose let us focus on the simplest case where $\omega = 0$ with probability 1 and where the insurer targets the allocation where fraud is deterred at the lowest cost, i.e. $y_A = 1, y_N = 0, \quad Q^f = \phi^{-1}(0), \quad Q^n = f\left(Q^f\right)$. Consider a direct revelation game where (at the interim stage) (1) the insured announces a type $\hat{x} \in \{A, N\}$, (2) the allocation is a function of $\hat{x}$, i.e. $y = y^*\left(\hat{x}\right)$ and $Q = Q^*\left(\hat{x}\right)$, where y is the probability of filing a claim and Q is the audit probability.

Implementing the desired allocation at a truthful equilibrium of this revelation game would require $y^*(A) = 1, \quad y^*(N) = 0, \quad Q^*(N) = \phi^{-1}(0)$ and $Q^*(A) = f\left(\phi^{-1}(0)\right)$. In this case, type-$N$ individuals are deterred from announcing $\hat{x} = A$ if $Q^*(A) \geq \phi^{-1}(0)$, which is not the case since $f(Q) < Q$ for all Q.

Hence the Revelation Principle does not apply in this setting. Intuitively this comes from the fact that scoring allows the insurer to use auditing as a type-dependent threat. Indeed if a type-N individual announces that he is a type-A, he would not be subject to the same audit probability than if he actually were a type-A. This takes us out of the costly state verification environment contemplated by Townsend (1988) where the outcome of strategic interactions can be replicated through a truthful revelation mechanism.

Appendix A4: Detailed Description of Variables in Regression Analysis

σ^{1}	Fraud alert by expert (expert warns that damages to vehicle do not correspond to those claimed by policyholder).
σ^{2}	Fraud alert by ARGOS or ALFA professionals (ARGOS alerts by mail, telephone, or via ARVA, a telematic management link with experts).
σ^{12}	Retroactive effect of the contract or guarantee (when the date of accident report is close to the guarantee or contract date).
σ^{18}	Delay in filing accident claim (refer to periods stated in contract: more than 5 business days for all types of accidents and more than 2 days after the accident occurs and/ or the theft is noticed.)
σ^{19}	Production of questionable or falsified documents means that the policy holder is submitting fraudulent invoices. The anti-fraud agent is alerted when the policyholder submits photocopies or duplicates of bills.
σ^{20}	Refusal or reluctance to provide original documents.
σ^{21}	Variations in or additions to the policyholder's initial claims (additional reports…).
σ^{22}	Harassment from policyholder to obtain quick settlement of a claim (abnormally frequent letters or calls from the policyholder).
σ^{30}	Description of circumstances surrounding the accident either lack clarity or seem contrived (accident report is too perfect; policyholder's descriptions of accident are too detailed or too vague).
σ^{32}	Abnormally high frequency of accidents (more than 3 accidents a year).
σ^{34}	Claimant not the same as policyholder (this can be detected in the claim or police report).
σ^{35}	Too long a lag between date of purchase and date of guarantee: policyholder takes out insurance a week after purchasing the vehicle whereas, logically, this should be done the same day.
σ^{36}	Date of guarantee subscription and/or date of its modification too close to date of accident (< 1 month).
θ_{7}	Vehicle whose value does not match income of policyholder (ex: an unemployed person or a welfare recipient with a Porsche).
θ_{16}	Policyholder with no damage insurance and/or one who would suffer harm if found at fault (for example, a policyholder who is insured only for third party damages may collude with a friend to stage an at-fault accident in order to split the total payment from his insurer).

Appendix A5: Calibration

To obtain the estimated fraud rate of the insurer (8% according to the views of the company managers), we first replicated the non fraud sub-sample six times, yielding 6,774 observations ($6 \times 1{,}129$). Then we took a random sample (with replacement) from these 6,774 observations in order to obtain the additional 953 observations needed to produce a fraud rate of 8%. The final sample contains 8,400 files, 673 files with fraud and 7,727 files without fraud. For a matter of robustness, we did also select randomly (with replacement) 6,598 non-fraud files from the original population of 1,129 to obtain the 7,727 non-fraud files. The results are very similar and are available from the authors.

Appendix A6: Correlation tables

Table A – Correlations conditional on *F*

	σ^{12}	σ^{18}	σ^{19}	σ^{21}	σ^{30}	σ^{32}	σ^{36}	σ^{35}	σ^{34}	σ^{22}	σ^{20}	σ^{1}	σ^{2}
σ^{12}	1	-0.02	-0.05	0.02	-0.04	0.02	0.05	-0.02	-0.01	-0.03	-0.00	0.01	-0.00
σ^{18}		1	0.04	0.02	-0.01	0.01	-0.02	0.06	0.09	0.01	-0.04	-0.11	0.04
σ^{19}			1	0.16	0.12	-0.02	0.02	-0.01	-0.02	0.06	0.01	0.07	0.01
σ^{21}				1	0.20	0.05	-0.05	-0.01	0.09	0.03	0.12	0.11	0.07
σ^{30}					1	0.03	0.04	-0.00	-0.01	0.03	0.12	0.08	0.02
σ^{32}						1	-0.06	-0.05	0.12	0.04	0.04	0.00	0.02
σ^{36}							1	0.19	-0.06	0.08	0.10	-0.11	0.02
σ^{35}								1	-0.07	0.10	0.10	-0.04	-0.00
σ^{34}									1	0.05	-0.03	-0.06	-0.05
σ^{22}										1	0.14	-0.02	-0.02
σ^{20}											1	0.05	0.05
σ^{1}												1	-0.15
σ^{2}													1

Table B – Correlations conditional on *N*

	σ^{12}	σ^{18}	σ^{19}	σ^{21}	σ^{30}	σ^{32}	σ^{36}	σ^{35}	σ^{34}	σ^{22}	σ^{20}	σ^{1}	σ^{2}
σ^{12}	1	-0.04	0.06	-0.00	-0.02	-0.02	0.06	0.00	-0.03	0.01	0.00	0.07	-0.02
σ^{18}		1	0.04	-0.02	0.06	-0.06	0.08	0.04	0.08	-0.03	0.03	0.02	0.03
σ^{19}			1	0.13	0.07	-0.01	0.04	0.02	0.01	0.10	0.05	0.00	-0.01
σ^{21}				1	0.03	0.00	-0.03	0.03	0.03	0.08	0.03	0.05	-0.01
σ^{30}					1	-0.02	0.02	-0.04	0.02	0.01	0.05	0.01	0.01
σ^{32}						1	-0.10	-0.07	0.02	-0.02	0.01	0.04	0.01
σ^{36}							1	0.13	0.02	0.12	0.09	0.00	-0.03
σ^{35}								1	0.04	-0.02	0.01	-0.02	-0.01
σ^{34}									1	-0.03	0.04	0.03	0.01
σ^{22}										1	0.02	0.14	-0.01
σ^{20}											1	0.07	-0.01
σ^{1}												1	-0.01
σ^{2}													1

Appendix A7: Fraud indicators manipulation

Section 4.1 of the paper has extended our theoretical model to the case where insureds can manipulate fraud signals. Here we show how our numerical estimates are affected by the possibility of fraud manipulation. The results are summarized in Table A7.1.

We first consider five indicators that are likely to be manipulated (18, 21, 22, 30, 34). In the first row, the manipulation rate s —the same for all individuals— ranges from 1/20 to 1/5 and is assumed, for the moment, to be the same for these five indicators. Technically, we cannot go much further under this uniformity assumption because s_{18} must be lower than 0.22 in order to have $\alpha_{18}^{f}(1-s_{18}) \geq \alpha_{18}^{n}$. Indicator 18 would not be informative any more for higher manipulation rates. As the results clearly show, these levels of fraud manipulation do not significantly affect the benefits of the audit policy. They drop from \$21.1 millions to \$19.9 millions when the manipulation rate increases from 0 to 1/5.

If now we use different manipulation rates and we move the five indicators near their informative limit, the benefit of the audit policy drops to \$18.2 millions, which is not yet a large decrease. This result suggests that the five indicators subject to manipulation are not the most important in terms of informativeness since their manipulation does not strongly affect the outcome of the auditing process.

The last row of the table shows a much more radical manipulation case that moves all indicators near to the limit of their informative value. In such a case the benefit from the audit policy dramatically declines and only 24.1% of audited claims are fraudulent. This is of course much lower than 57,8%, the hit rate without manipulation.

Table A7.1 – Fraud indicators manipulations

Manipulation rate	1	3	4	5	10	12	Benefit of audit policy (thousand \$)
	i^*	p_i^f / p_i^n	$\lambda(i)$	$\mu(i)$	$z(\theta)\lambda(i) + (1-z(\theta))\mu(i)$	$P(F / i > i^*)$	
0	194	3.2138	0.6672	0.0423	0.0923	0.5781	21,119.2
$s = 1/20$	195	3.2120	0.6547	0.0417	0.0910	0.5781	20,649.7
$s = 1/15$	198	3.2165	0.6541	0.0415	0.0905	0.5779	20,689.1
$s = 1/10$	199	3.2215	0.6477	0.0412	0.0897	0.5775	20,474.7
$s = 1/8$	199	3.2214	0.6450	0.0416	0.0898	0.5743	20,323.0
$s = 1/5$	209	3.2078	0.6493	0.0459	0.0942	0.5514	19,926.3
$s_{18} = 0.20; s_{21} = 0.80;$ $s_{22} = 0.80; s_{30} = 0.60;$ $s_{34} = 0.25$	483	3.2167	0.6200	0.0497	0.0954	0.5200	18,264.0
$s_1 = 0.90; s_2 = 0.90$ $s_{12} = 0.35; s_{18} = 0.20;$ $s_{19} = 0.80; s_{20} = 0.80;$ $s_{21} = 0.80; s_{22} = 0.80;$ $s_{30} = 0.60; s_{32} = 0.25;$ $s_{34} = 0.25; s_{35} = 0.75;$ $s_{36} = 0.50$	2,666	3.2074	0.0077	0.0021	0.0025	0.2414	42.7

$\alpha_j^f (1 - s_j) \geq \alpha_j^n \quad \Leftrightarrow \quad s_1 \leq 0.93; \quad s_2 \leq 0.95; \quad s_{12} \leq 0.40; \quad s_{18} \leq 0.22; \quad s_{19} \leq 0.90; \quad s_{20} \leq 0.83;$

$s_{21} \leq 0.83; \quad s_{22} \leq 0.86; \quad s_{30} \leq 0.63; \quad s_{32} \leq 0.27; \quad s_{34} \leq 0.28; \quad s_{35} \leq 0.76; \quad s_{36} \leq 0.52.$

Appendix A8: Distribution of p_i^f / p_i^n

In the paper it has been assumed that the probability distributions of signals (i.e. the p_i^f and p_i^n) are the same for all individuals, i.e. they do not depend on θ.

We did take this assumption very seriously into account because implementing an optimal audit strategy with type-dependent signal distributions would require a much larger database. The practical implementability of the optimal audit strategy could then be jeopardized. We proceeded in two steps in order to test the relevance of our assumption. First we tested whether α_j^f and α_j^n statistically differ for all indicators *j* between the three different samples we have contemplated in our numerical analysis. Secondly, we analyzed the distributions of the ratio p_i^f / p_i^n obtained from these alpha values over the same three samples and we tested whether they statistically differ.

The second regression in Table 1 has allowed us to identify two variables (θ_7 and θ_{16}) that are correlated with the fraud rate. So our tests compare α_j^f and α_j^n for all *j* and the distributions of the ratio p_i^f / p_i^n for three different samples.

In Table A8.1 below we test whether the α_j^f and α_j^n are type dependent. For each indicator *j* we first present the α_j^f and α_j^n values obtained from the original sample with all files. We then present their corresponding values for each new sample. The test indicates that when the *p*-value is lower than 5%, we reject the hypothesis $H: p_1 = p_2$ where p_1 is the proportion of claims with indicator *j* in the original sample and p_2 is the proportion in the additional samples. It turns out that all the computed *p*-values are higher than 5% which means that the alpha values, in the two additional samples, do not significantly differ from those in the first sample.

In Table A8.2 we present the three distributions of p_i^f / p_i^n obtained from the three samples. The last column (sample A) is the distribution from the original sample and the two other columns come the two additional samples (sample B: $\theta_7 = 0$ and $\theta_{16} = 0$; sample C: $\theta_7 = 0$ and $\theta_{16} = 1$. We applied the

Kolmogorov-Smirnov test in order to determine if the distributions differ significantly. The two p-values are 0.9334 and 0.3050. They permit to conclude to a non rejection of the null hypothesis that the distributions are identical for the two populations when each additional sample is compared with the original sample.

For illustration we present the three histograms for values ≤ 500 which are those where the optimal values for i* are obtained (Figure A8.1). Since the potential values for the ratio range from 0 to more than 6 millions (see Table 2) it is difficult to draw histograms over all values. The histograms show clearly that the three distributions do not differ significantly. We obtained results from Kernel density estimation of the three distributions and we plot them in Figure A8.2. This figure confirms the Kolmogrov-Smirnov test that the three distributions are not different. Finally, Figure A8.3 represents the Kernel densities along with the corresponding histograms.

Table A8.1 – Dependence of α_j^f and α_j^n

θ_7	θ_{16}	α_j^f (%)	INF	SUP	p-value	α_j^n (%)	INF	SUP	p-value
		σ^{12}							
		12.1842	9.6386	14.7298		7.35	5.7835	8.9165	
0	1	13.068	9.4048	16.7312	0.6844	6.93	4.5059	9.3541	0.7690
0	0	10.830	6.9898	14.6702	0.5564	7.79	5.6611	9.9189	0.7340
		σ^{18}							
		50.8172	46.9658	54.6686		39.77	36.8708	42.6692	
0	1	48.295	42.9326	53.6574	0.4432	44.59	39.9492	49.2308	0.0762
0	0	52.708	46.6479	58.7681	0.5962	35.73	31.9837	39.4763	0.0906
		σ^{19}							
		46.3596	42.5177	50.2015		4.78	3.4912	6.0688	
0	1	43.750	38.4255	49.0745	0.4256	5.41	3.2390	7.5810	0.5994
0	0	49.458	43.3896	55.5264	0.3846	4.27	2.6453	5.8947	0.6199
		σ^{21}							
		21.8425	18.6465	25.0385		3.63	2.4947	4.7653	
0	1	21.591	17.1506	26.0314	0.9262	4.33	2.3658	6.2942	0.5092
0	0	20.578	15.6366	25.5194	0.6662	3.05	1.6567	4.4433	0.5157
		σ^{30}							

θ_7	θ_{16}	α_j^f (%)	**INF**	**SUP**	***p*-value**	α_j^n (%)	**INF**	**SUP**	***p*-value**
		33.7296	30.0833	37.3759		12.58	10.6013	14.5587	
0	1	37.784	32.5768	42.9912	0.1966	12.34	9.2327	15.4473	0.8955
0	0	28.159	22.6817	33.6363	0.0946	12.82	10.1834	15.4566	0.8832
		σ^{32}							
		22.5854	19.3519	25.8189		16.39	14.1863	18.5937	
0	1	21.591	17.1506	26.0314	0.7162	15.15	11.7724	18.5276	0.5406
0	0	23.466	18.2948	28.6372	0.7689	17.40	14.4203	20.3797	0.5820
		σ^{36}							
		41.7533	37.9531	45.5535		19.84	17.4695	22.2105	
0	1	44.3182	38.9866	49.6498	0.4304	21.43	17.5800	25.2800	0.4742
0	0	37.9061	32.0122	43.8000	0.2725	17.71	14.7101	20.7099	0.2695
		σ^{35}							
		9.8068	7.4855	12.1281		2.39	1.4548	3.3252	
0	1	9.9432	6.6750	13.2114	0.9445	2.60	1.0407	4.1593	0.8058
0	0	8.6643	5.1709	12.1577	0.5847	1.83	0.7272	2.9328	0.4354
		σ^{34}							
		19.3165	16.2595	22.3735		13.99	11.9223	16.0577	
0	1	19.0341	14.7909	23.2773	0.9132	14.29	10.9905	17.5895	0.8759
0	0	20.5776	15.6362	25.5190	0.6568	13.44	10.7515	16.1285	0.7453
		σ^{22}							
		12.4814	9.9100	15.0528		1.77	0.9566	2.5834	
0	1	10.227	6.9195	13.5345	0.2864	2.16	0.7262	3.5938	0.6034
0	0	12.996	8.8555	17.1365	0.8282	1.37	0.4034	2.3366	0.5193
		σ^{20}							
		13.5215	10.8637	16.1793		2.30	1.3813	3.2187	
0	1	16.193	12.2025	20.1835	0.2479	3.46	1.6852	5.2348	0.1899
0	0	10.108	6.3777	13.8383	0.1486	1.53	0.5137	2.5463	0.2648
		σ^{2}							
		34.0267	30.3727	37.6807		1.68	0.8860	2.4740	
0	1	33.807	28.7231	38.8909	0.9438	1.52	0.2961	2.7439	0.8192
0	0	34.657	28.8723	40.4417	0.8524	1.83	0.7272	2.9328	0.7911
		σ^{1}							
		6.2407	4.3388	8.1426		0.44	0.0096	0.8704	
0	1	7.670	4.7479	10.5921	0.3858	0.00	-0.1082	0.1082	0.1533

θ_7	θ_{16}	α_j^f (%)	**INF**	**SUP**	***p*-value**	α_j^n (%)	**INF**	**SUP**	***p*-value**
0	0	5.054	2.2938	7.8142	0.4803	0.76	0.0186	1.5014	0.2998

Table A8.2 – Distribution of p_i^f / p_i^n

	Sample B		**Sample C**		**Sample A**	
p_i^f / p_i^n	$\theta_7 = 0$ N	$\theta_{16} = 0$ %	$\theta_7 = 0$ N	$\theta_{16} = 1$ %	N	%
(0 – 15]	574	7.01	601	7.34	567	6.92
(15 – 50]	544	6.64	528	6.45	528	6.45
(50 – 120]	534	6.52	520	6.35	526	6.42
(120 – 250]	546	6.67	521	6.36	542	6.62
(250 – 500]	571	6.97	555	6.77	573	6.99
(500 – 1 000]	628	7.67	598	7.30	608	7.42
(1 000 – 1 500]	374	4.57	358	4.37	372	4.54
(1 500 – 3 000]	657	8.02	631	7.70	657	8.02
(3 000 – 5 000]	476	5.81	448	5.47	477	5.82
(5 000 – 10 000]	613	7.48	602	7.35	608	7.42
(10 000 – 20 000]	569	6.95	561	6.85	574	7.01
(20 000 – 40 000]	504	6.15	505	6.16	504	6.15
(40 000 – 100 000]	547	6.68	572	6.98	561	6.85
(100 000 – 350 000]	527	6.43	569	6.95	543	6.63
> 350 000	528	6.45	623	7.60	552	6.74
Total	8,192	100.00	8,192	100.00	8,192	100.00
KS-test	*p*-value = 0.9334		*p*-value = 0.3050			

Figure A8.1: Histogram of p_i^f / p_i^n

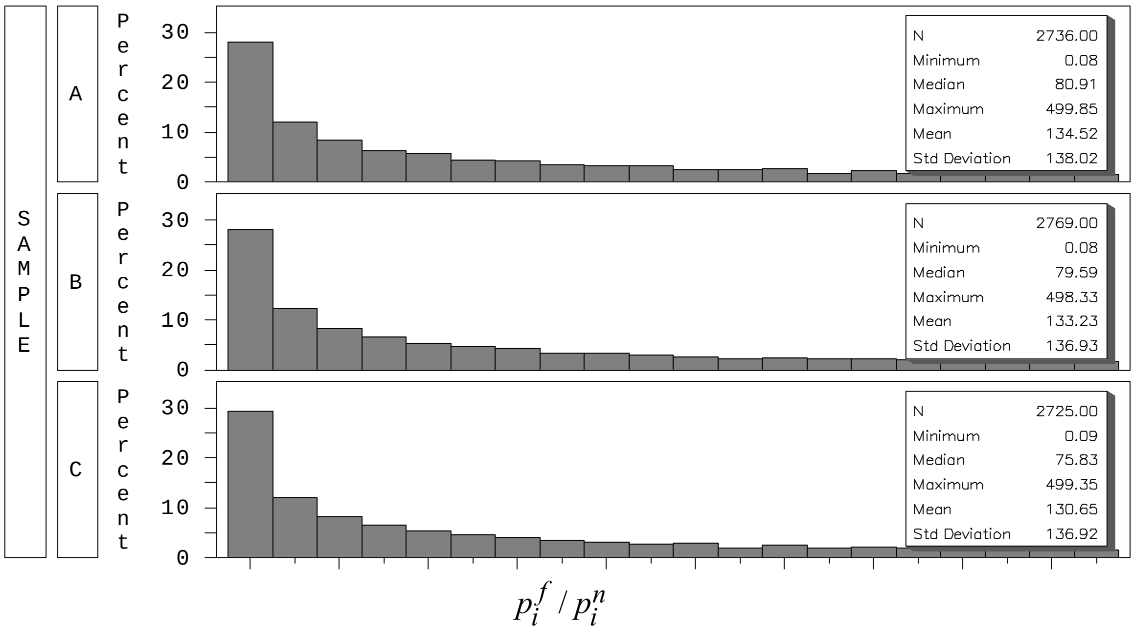

Figure A8.2: Kernel Density Estimate of p_i^f / p_i^n

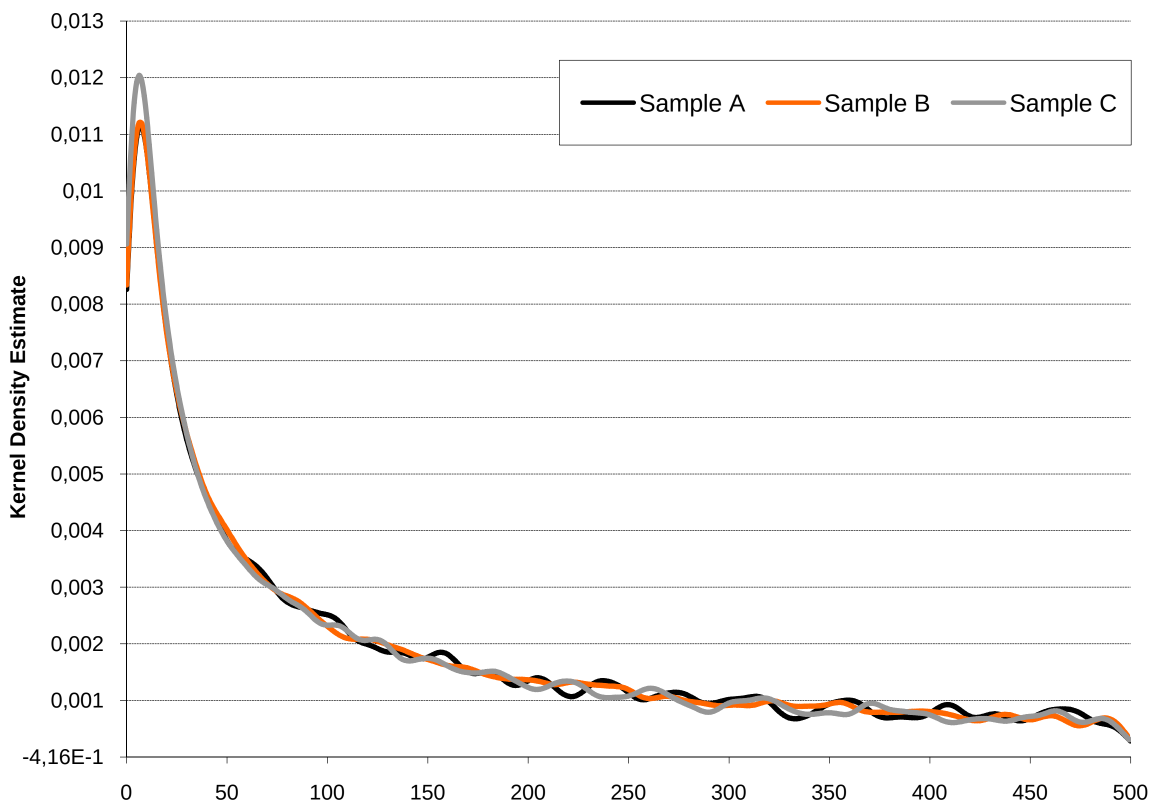

Figure A8.3: Histograms

Sample A

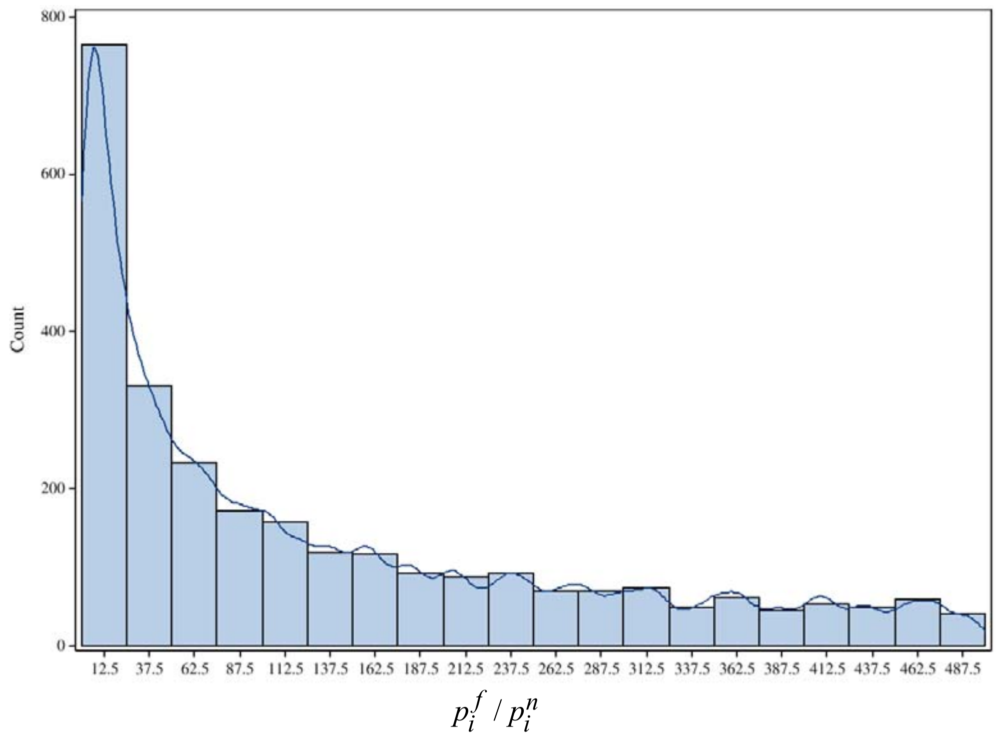

Sample B

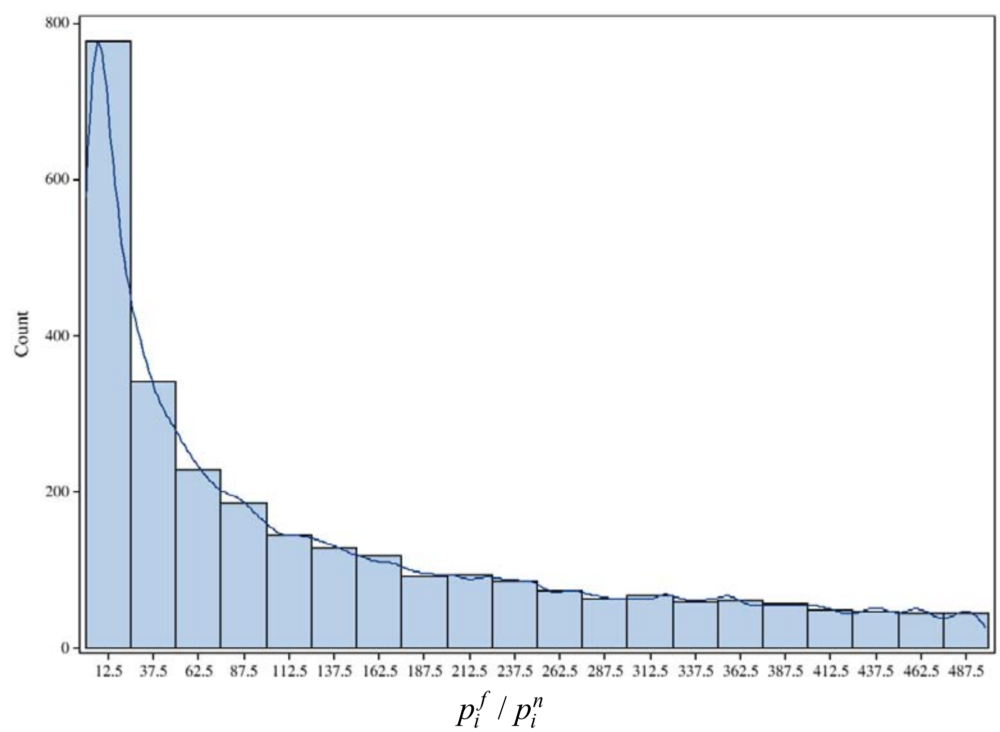

Sample C

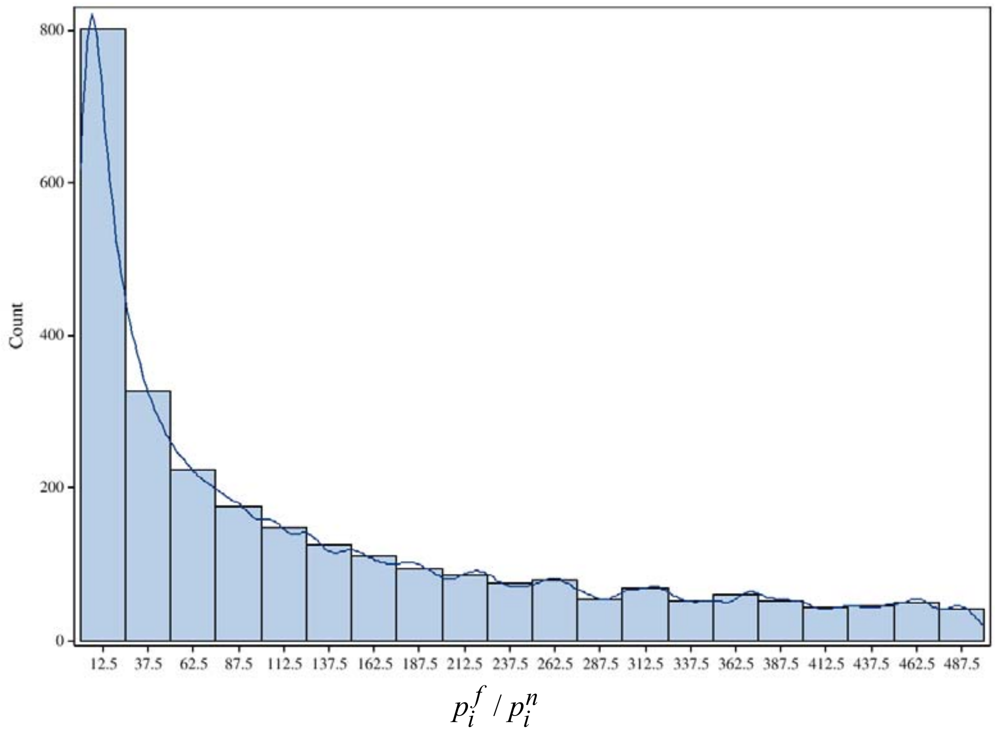

800
600
400
200
0
Count
12.5
37.5
62.5
87.5
112.5
137.5
162.5
187.5
212.5
237.5
262.5
287.5
312.5
337.5
362.5
387.5
412.5
437.5
462.5
487.5
p_i^f / p_i^n