The Geneva Papers on Risk and Insurance Vol. 25 No. 4 (October 2000) 517–538

A Model for the Detection of Insurance Fraud*

by El Bachir Belhadji, Georges Dionne** and Faouzi Tarkhani

The aim of this article is to develop a model to aid insurance companies in their decision-making and to ensure that they are better equipped to fight fraud. This tool is based on the systematic use of fraud indicators. We first propose a procedure to isolate the indicators which are most significant in predicting the probability that a claim may be fraudulent. We applied the procedure to data collected in the Dionne–Belhadji study (1996). The model allowed us to observe that 23 of the 54 indicators used were significant in predicting the probability of fraud. Our study also discusses the model's accuracy and detection capability. The detection rates obtained by the adjusters who participated in the study constitute the reference point of this discussion. As shown in the Caron–Dionne (1998), there is the possibility that these rates underestimate the true level of fraud.
JEL numbers: D81, G14, G22.

1. Introduction

The purpose of this research is to develop a tool to help insurance companies to detect fraud. The application is limited to property damages in the automobile sector. The work consists in setting up a simple model allowing the adjuster to calculate the probability of fraud in each file studied. After having determined the most significant fraud indicators and their respective weights, we formulate a function linking these indicators to a probability of fraud.

The article is divided into six sections. In the second section, we describe the procedure for sampling the files. In section 3, we review the literature on fraud indicators and we justify our choice of the selected indicators. We then calculate the parameters allowing us to pick out which indicators are significant in predicting the probability that a file contains fraud. Section 3 uses the statistically significant parameters to calculate the probabilities of fraud in different files. This section also explains the procedures to follow in constructing a sample of files to be reviewed, while keeping within the budget allotted to fraud by the company. Finally, this section looks at the prevalence of fraud, relying exclusively on the evidence provided by our model. Section 5 briefly discusses the costs of investigations, and section 6 presents avenues of research.

2. Method and sampling procedure

In this section, we will examine in turn the representativeness of our sample, the different stages of the study, the notions of suspected and established fraud, as well as the methods used to select the actual sample.

* This research was made possible thanks to funding from the Insurance Bureau of Canada (BAC), FCAR Québec, CRSH Canada and the FFSA in France. We would like to thank Raymond Medza, General Manager, BAC, for his collaboration at different stages of the study, as well as Gilles Chevrier and Gilles Beaulieu of Axa (Canada) for their help in searching out certain useful information.

** Risk Management Chair, HEC-Montreal (dionne@crt.umontreal.ca).

Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK.

2.1 Representativeness of the sample

After having studied the market shares of insurers in collaboration with the Insurance Bureau of Canada, we selected 20 insurance companies among the biggest in Quebec's automobile insurance sector. According to the data in the *Rapport annuel sur les assurances* (1994) from the Inspector General of Financial Institutions, the twenty companies selected held 78.5 per cent of the automobile insurance market in Quebec. Eighteen of the 20 firms selected agreed to participate in the study: a 90 per cent participation rate. Taking into account the relative market shares of these 18 firms, we thus obtained the participation of firms representing 70 per cent of the market, which is very satisfactory. So, we can conclude that our study is representative of the market studied.

The participating firms agreed to make a detailed analysis of their closed claims files, in order to provide us with the different kinds of information needed to understand the fraud phenomenon. Their collaboration also allowed us to evaluate the level of fraud and to study the main indicators of automobile insurance fraud.

2.2 Stages of the study

The files were selected at random from among all the files closed during the period from 1 April 1994 to 31 March 1995. The number of questionnaires assigned to each company was

Table 1:
Number of questionnaires received from participating companies

Company	Number of questionnaires completed	% of sample
1	100	4
2	89	4
3	166	7
4	108	4
5	274	11
6	48	2
7	370	15
9	45	2
10	39	2
11	60	2
12	275	11
13	346	14
15	69	3
16	39	2
17	100	4
18	40	2
19	161	6
20	180	7
Total	2,509	100

proportional to its share of the market. We suggested a method of random sampling (described in Dionne and Belhadji, 1996) to the firms. The study took place in the spring of 1995. The insurers chose which adjusters would answer the questionnaire. We had asked them to designate the adjusters who had already handled the files in the sample. Strict procedures were implemented to protect the identity of the adjusters and the companies. The 18 companies returned 2,509 completed questionnaires to us, corresponding to a response rate of 98 per cent for the 18 participating companies. Table 1 gives the breakdown of the files received from the participating companies.

2.3 Established and suspected fraud in the data of the sample

Table 2 gives the number of established and suspected cases of fraud over the 2,509 closed files. These numbers reflect the opinion of the adjusters who answered the questionnaires. By established fraud, we mean that there was effectively fraud in a particular file, whether or not the case was brought to court. Fraud is suspected if the adjuster had suspicions of fraud while handling a file which was never investigated further.

Table 2:
Classification of files

Classification	**N No fraud**	**S Suspected fraud**	**E Established fraud**	**(S + E)%**	**E%**	**N + S + E**
Total	2,375	116	18	5.34%	0.72%	2,509

3. Fraud indicators

The survey carried out in spring 1995 had a dual purpose: it was aimed, on the one hand, at evaluating the extent of automobile insurance fraud[1] and, on the other hand, at developing an expert system for the automatic detection of such fraud. This second part of the study consists in determining a set of indicators sensitive to the detection of fraud and the suspicion of fraud. In this section we shall justify the choice of the indicators selected as well as the criteria forcing us to eliminate certain others.

3.1 Choice of indicators

We presented a list of 50 fraud indicators to participating adjusters. We asked them to include one or more indicators in any file containing one or more correspondent character-

[1] This part of the study has been completed and a copy of the report can be obtained from the Insurance Bureau of Canada (Dionne and Belhadji, 1996). It has also been published in broad outline in the October 1996 issue of *Assurances*. An extension of this fraud evaluation was made by Caron and Dionne (1998).

istics. The 50 indicators are enumerated and categorized in the Appendix.[2] The indicators recorded in bold are the ones which are significant in predicting the probability that the file is fraudulent.

In the literature on indicators, there is very little empirical evidence which could scientifically demonstrate why some fraud indicators should be more relevant than others. It is for this very reason that we were attracted to the subject. The list we drew up (see Appendix) and which was included in the questionnaire sent to the claims adjusters was based on the literature in print at the time when we compiled our questionnaire.

Our list was based on two types of literature: applied literature concerned with the attempt to pick out relevant indicators and writings based on the experience of insurance professionals. In the category of applied literature, we note the work done by the University of Florida's Institute Research Center. Their study dealt with automobile insurance fraud in Florida. Its authors calculated conditional probabilities of fraud for each of the 90 indicators proposed in their list. Some of them applied to very few files (less than ten) and we had to eliminate them. We selected all the indicators relevant to the coverage (chapters) offered in Quebec and whose conditional probability of fraud was higher than 10 per cent. One should, however, note that this American study did not study all the indicators simultaneously: it only calculated probabilities of fraud based on the presence of one particular indicator in the file (one indicator at a time).

Another study (Weisberg and Derrig, 1993) was used to complete our list of significant indicators. To ensure that our list would contain a professional component, we drew on descriptive studies conducted for the Insurance Bureau of Canada and SACA and on internal information on indicators published in the March 1993 issue of *Property and Casualty Claims Services.* The idea behind the selection process was to retrieve indicators common to several different studies.

More recently, Derrig and Weisberg (1998) have developed a methodology that can identify claim suspicion. This methodology is based on a linear regression model where the dependent variable is the level of suspicion on a ten-point scale from 0 to 10. Here, since we are more interested in estimating the probability that a file contains some fraud, we will use a Probit model, where the dependent variable is equal to one when the file has been judged fraudulent (suspected or established fraud) and equal to zero otherwise. Our results will be compared to those of a linear regression model with a ten-point scale dependent variable (since we also have this information). Finally, in order to add a component closely connected with the Quebec automobile insurance market, we consulted company executives to find out which indicators were most used. The list of indicators in the Appendix was thus drawn up based on these studies and these consultations.

3.2 Criteria for limiting the number of indicators to be included in the regression

Since our list contained a very large number of indicators, we wanted to find a way of reducing them to a reasonable number for inclusion in our regressions. It seemed that the most effective means of doing this (while also eliminating non-relevant indicators) would be to calculate the conditional probabilities of fraud for each of the indicators. Table 3 indicates these conditional probabilities. Column 1 reports the indicator's number as it appears in the

[2] In fact, the Appendix contains a list of 54 indicators since four of the 50 primary indicators were used in a second category.

Table 3:
Conditional probability of fraud (E or S) based on the presence of indicator "x"

Number	Indicator	Total number	Number (E + S)	Rate (%)
1	Police	363	34	9.37
2	Minor	74	6	8.11
3, 4	Incoh (1, 2)	57	18	31.58
5	Theft	32	12	37.50
6	Recent	80	7	8.75
7	Bill	30	3	10.00
8	Interest	71	13	18.31
9	Commer	24	3	12.50
10	Scope	41	5	12.20
11	Agent	148	9	6.08
12, 13	Diffic (1, 2)	43	16	37.21
14	Occup	7	3	42.86*
15	Receives	71	12	16.90
16	Small	15	8	53.33*
17	Rapid	56	15	26.79
18, 19	Jargon (1, 2)	48	13	27.08
20	Travel	19	7	36.84
21	Eager	17	8	47.06
22	Proxim	87	13	14.94
23	Act68	3	0	0.00*
24	Taxi	17	7	41.18
25	Proof	19	12	63.16
26	Guilty	15	5	33.33*
27	History	14	7	50.00*
28	Third	22	2	9.09
29	Documen	21	5	23.81
30	Samegar	21	2	9.52
31	Represen	4	1	25.00*
32	Repair	20	3	15.00
33, 34	Witness (1, 2)	97	19	19.59
35	Denial	13	4	30.77*
36	Sole	142	19	13.38
37	Unident	69	7	10.14
38	Guarantee	9	3	33.33*
39	Rental	3	2	66.67*
40	Contact	8	3	37.50*
41	Sign	34	4	11.76
42	Settlem	6	3	50.00*
43	Cash	33	8	24.24
44	Unemploy	18	7	38.89
45	Found	6	3	50.00*
46	Overassur	8	0	0.00*

continued overleaf

Table 3:
(continued)

Number	Indicator	Total number	Number (E + S)	Rate (%)
47	Prem	1	1	100.00*
48	Motel	1	0	0.00*
49	Agress	26	9	34.62
50	Refuse	15	1	6.67*
51	Nervous	42	17	40.48
52	Numerous	129	12	9.30
53	Title	5	2	40.00*
54	Preced	7	2	28.57*

* Since they contain 15 observations or less, these results must be interpreted with caution.

Appendix. Each indicator has been given a name, which is shown in the second column. The third column displays the total number of files where the indicator is recorded. The fourth column is similar to the third except that the number recorded is that of files marked "E" (established fraud) or "S" (suspected fraud). Finally, the last column reports on the rate of fraud in a file given the presence of a specific indicator (example: police).

Note that the confidence intervals can be calculated based on the data in Table 3. For all the estimates of the table, the standard deviation can be calculated as follows:

$$\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

where $\hat{p}$ is the value of the estimate and "n" is the frequency of that particular indicator (third column). The confidence interval will thus be:

$$\left[\hat{p} - Z_{a/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}};\ \hat{p} + Z_{a/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}\right]$$

where $Z_{a/2}$ is the $1 - a/2$ percentile of a normal distribution.

The standard deviation of the estimate for the number 3 indicator, for example, is calculated as follows:

$$\sqrt{\frac{(0.316)(0.684)}{57}} = 0.062$$

The confidence interval at 95 per cent of the estimate for this indicator would thus be:

$$[0.316 - 1.96(0.062);\ 0.316 + 1.96(0.062)]$$

or equivalently:

$$[19.5\%;\ 43.7\%]$$

This table, which served only as a guideline in the choice of indicators for the econometric analysis, allowed us to eliminate all the indicators present in only 15 files or less.[3] So, the indicators marked with an asterisk (fifth column) should not be interpreted as being non-significant, since the reason for their exclusion is merely insufficient data. Any future study dealing with a large number of files (many more than 2,500 files) should include and test these indicators.

Note that, because the different indicators are viewed in isolation, Table 3 is not sufficient evidence of the possibility of fraud. A regression taking into account all the indicators present in a file is preferable. This is what we will do in the Probit regressions below. These regressions will allow us to determine which are the most significant indicators.

3.3 Regression model and results

The Probit model that we use supposes a y_i^* response variable as defined by the following relation:

$$y_i^* = b'x_i + u_i$$

where y_i^* is not observable. The x_i vector represents the indicators present in the file, whereas the b' vector includes the values of their respective parameters.

We do, however, observe the binary variable y_i which is defined by:

$$y = 1 \quad \text{if } y_i^* > 0$$

$$y = 0 \quad \text{otherwise}$$

In our specific case:

$$y = 1 \text{ if the file has been judged fraudulent (suspected or established fraud)}$$

$$y = 0 \text{ otherwise}$$

It follows that:

$$\begin{aligned} \text{prob}(y_i^* = 1) &= \text{prob}(u_i > -b'x_i) \\ &= 1 - F(-b'x_i) \end{aligned}$$

where F is the cumulative distribution function of u. The likelihood function will thus be:

$$L = \prod_{y_i=0} F(-b'x_i) \prod_{y_i=1} [1 - F(-b'x_i)]$$

In the Probit model, u_i follows a normal distribution $N(0, \sigma^2)$. In this case:

$$F(-b'x_i) = \int_{-\infty}^{-b'\frac{x_i}{\sigma}} \frac{1}{(2\pi)^{1/2}} \exp\left(-\frac{t^2}{2}\right) dt$$

The results of the Probit regression are recorded in Table 4. In this table only the significant indicators are presented. These are the indicators which will be used to calculate

[3] Technically, these indicators can be used in the analysis. However, for a matter of stability of the coefficients, it is better to not use them.

Table 4:
Results of Probit model estimate

No.	Indicators	Coefficients	t
	Constant	−2.41746	26.32507
	W/Ocol	0.30366	2.35227
1	Police	0.41268	3.15259
2	Minor	0.52156	2.21195
5	Theft	1.26335	4.63594
8	Interest	0.47338	1.95909
10	Scope	0.78794	2.19894
13	Diffic2	0.79851	2.51028
19	Jargon2	1.10237	3.88710
21	Eager	1.49044	4.57367
22	Proxim	0.60740	3.12840
24	Taxi	0.80453	2.28934
25	Proof	1.68070	5.11908
29	Documen	0.96332	2.69495
34	Witness2	1.18144	4.26807
36	Sole	0.64804	3.92314
37	Unident	0.53696	2.22246
43	Cash	0.82919	2.92983
49	Agress	1.02128	3.58915
51	Nervous	1.02193	4.58424
52	Numerous	0.32338	1.69918
3	Incoh1	1.59655	6.72115
12	Diffic1	1.20468	3.58305
18	Jargon1	1.21734	3.22767
33	Witness1	0.91227	4.19262

Log likelihood: −350.70.
Number of observations: 2,509.

the probabilities of fraud in files. The first and second columns give, respectively, the numbers and names of the indicators selected. The third column presents the estimated coefficients, whereas the last column indicates their t-statistics. It is to be noted that the indicators followed by the number 1 are those considered the most important in the file for the investigator. In contrast, the indicators followed by a 2 represent those whose order of importance range between 2 and 12 in their classification. In Table 4′, we present results of Ordinary Least Squares estimates when the dependent variable is a ten-point scale.

Results in Table 4′ are closed to the results obtained from the Ordinary Least Squares model in Derrig and Weisberg (1998). Few indicators are significant. The main difference between the two models is that the Probit model is a qualitative model that estimates the probability that a file contains some fraud while the Ordinary Least Squares model establishes a quantitative linear relationship between different indicators and the different values of the ten-point scale. The results from the two tables confirm that a qualitative model dominates a

Table 4′:
Results of Ordinary Least-Squares Estimate

No.	Indicators	Coefficients	t
3	Incoh1	0.174129	2.817
5	Theft	0.413884	6.046
17	Rapid	0.201829	3.784
19	Jargon2	0.156729	2.210
21	Eager	0.266572	2.981
25	Proof	0.255759	2.948
49	Agress	0.452627	6.072
52	Numerous	0.077020	2.371

Adjusted R^2: 0.06.
Number of observations: 2,509.

quantitative one for the purpose of getting subjective information from the indicators listed by the adjusters and to give a probability that a file contains fraud.

4. Model where only adjusters' predictions are relevant

We will first compare the predictions of our model with the results of the investigation; we will then propose examples of decisions as to whether or not further investigation should be conducted.[4]

4.1 Comparison of regression results with those of the investigation

The predictions of our model are compared with the decisions of adjusters regarding their suspicions – these suspicions being considered as complete and that there is no other fraud (or suspicions) in the sample examined.

The model that we have used has generated probabilities of fraud for each file. These vary between 0.067 per cent (no indicator) and 99.817 per cent. In setting the probability threshold for a closer investigation of files, we implicitly determine the detection rate for fraud as well as the model's level of accuracy in spotting suspected cases of fraud. We know that out of the 2,509 claims investigated, adjusters found 134 suspected or established fraudulent cases. This gives us a fraud rate of 5.34 per cent.[5]

Take for example the 10 per cent threshold of probability. At the 10 per cent level, the model will generate 296 cases of fraud. Among these 296 cases, the adjusters suspected (and/or established) 93 cases of fraud. The rate of accurate fraud detection is thus equal to 31.42 per cent (or 93/296). The cases classified by the model as "non-fraud" number 2,213 (or 2,509 − 296). Among these, 2,172 cases are classified as "non-fraudulent" by investigators: an accuracy rate of 98.15 per cent (or 2,172/2,213). The fraud rate in the industry is the basis

[4] For a different methodology, see Brockett, Xia and Derrig (1998).

[5] Here it is supposed that the adjusters will detect all the cases of fraud contained in the files. This hypothesis has been questioned by Caron and Dionne (1998) who show that the adjusters observe only one-third of the fraud.

for the proportion of the sample to be selected.[6] The company will select "X" percentage of the files to be reviewed, as calculated by solving the following equation:

$$\alpha\ .\ X\% + (1 - \beta)(1 - X\%) = 5.34\%$$

where:

α: is the rate of accuracy (in %) for fraud cases captured by the model. If the threshold is 10%, this rate is equal to 31.42%.

β: is the rate of accuracy (in %) for non-fraudulent cases generated by the model. If the probability threshold is 10%, this rate is equal to 98.15%.

Note that the equation above contains two parts: the first part, $\alpha\ .\ X\%$, expresses the percentage of fraud (as predicted by the model) to be found among the cases selected (cases where the probability of fraud exceeds the threshold). The second part, $(1 - \beta)(1 - X\%)$, expresses the proportion of fraud detected by the investigators, but a proportion lower than the threshold chosen (cases not selected for review). For example, solving the equation above for a 10 per cent threshold, we will obtain a proportion of the sample to be re-examined, "X", equal to 11.80 per cent.

To sum up: for a 10 per cent threshold of fraud probability, we should select 296 cases (according to the regression model); among these cases, adjusters found that 93 cases were fraudulent and 203 non-fraudulent, which means:

$$296 \rightarrow \begin{cases} NF = 203 \\ F = 93 \end{cases}$$

The proportion of the sample to be re-examined = X = 11.80%.

The accuracy rate for cases of fraud (F): 93/296 = 31.42%.

The number of files below the 10 per cent threshold is equal (according to the model) to 2,213. Among these, 2,172 cases are classified as non-fraudulent by adjusters:

$$2{,}213 \rightarrow \begin{cases} NF = 2{,}172 \\ F = 41 \end{cases}$$

The rate of accuracy for "non-fraud" (NF) = 2,172/2,213 = 98.15%.

4.2 Decision to investigate

We now propose a model for the decision to investigate. *At a probability threshold of* $P = 10$ *per cent*, the company should select 11.80 per cent of its claims for re-examination, if it feels its fraud rate is equivalent to the industry's 5.34 per cent level.

---

[6] This percentage can be replaced by another if the company feels its fraud rate differs from the average for the industry.

Thus, for 100 claims, we will get the following results:

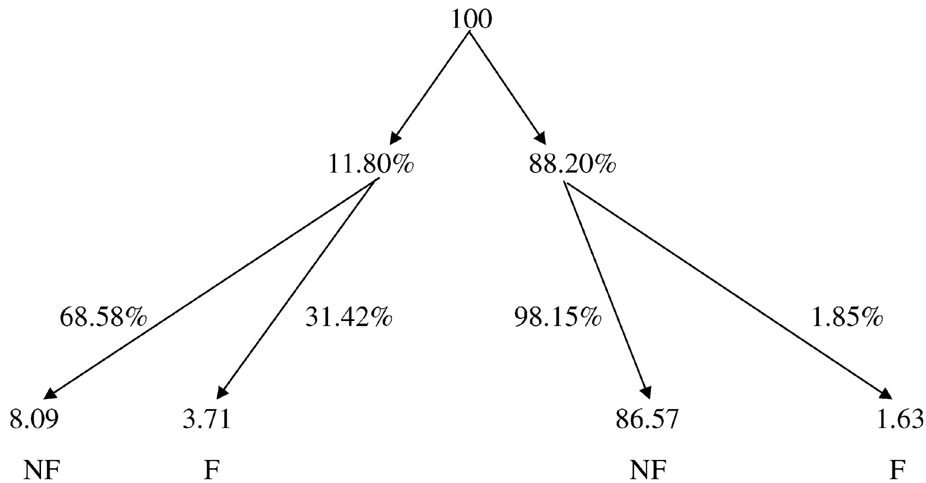

Total fraud rate $= 3.71\% + 1.63\% = 5.34\%$[7]
Rate of accuracy for fraud cases (F) $= 31.42\%$ (see the calculation above)
Rate of detection $= 3.71\%/5.34\% = 69.40\%$.

Similarly, for a *$P = 20$ per cent probability threshold*, the model detects 168 cases of fraud. Of these, 72 were classified "fraudulent" by the investigators. The model's accuracy rate is thus 42.86% (or 72/168). When the threshold is dropped below 20 per cent, the model generates 2,341 "non-fraud" cases. Adjusters determined that 2,279 of these 2,341 cases, were not fraudulent. The "non-fraud" accuracy rate is thus 97.35 per cent (or 2,279/2,341). To sum up, a probability threshold of 20 per cent will give us:

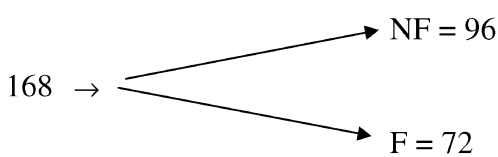

Proportion of sample to be re-examined $= X = 6.70\%$.
Rate of accuracy for fraud cases (F): $72/168 = 42.86\%$.

The number of files below the 20 per cent threshold is (according to the model) 2,341. Among these, 2,279 cases are classified as non fraudulent by investigators:

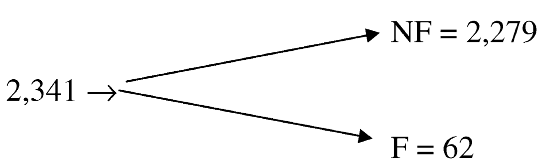

[7] Note that this figure may not be exactly equal to 5.34% due to the rounding off of percentages.

The rate of accuracy for "non-fraud" (NF) = 2,279/2,341 = 97.35 per cent.

At a P = 20 per cent probability threshold, the company should select 6.70 per cent of its claims for re-examination, if it feels its fraud rate is equivalent to the industry's 5.34 per cent level.

So, out of 100 claims, we will obtain the following results:

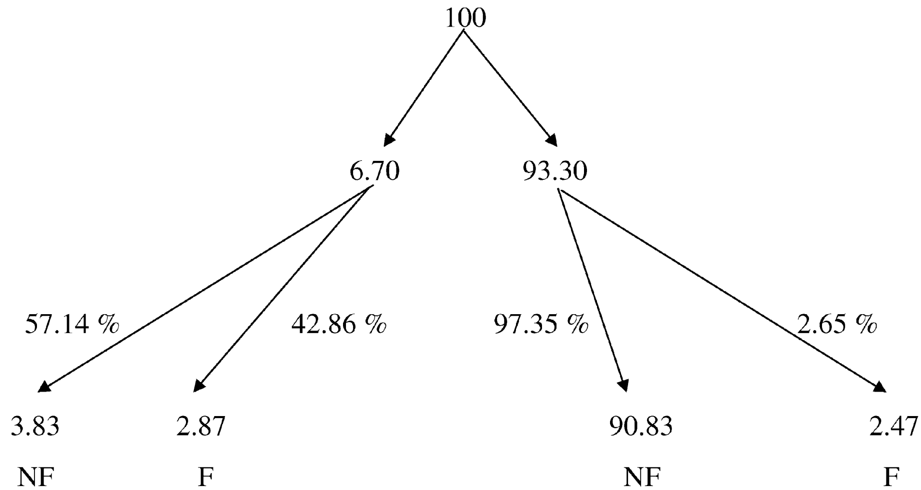

Total fraud rate = 2.87% + 2.47% = 5.34%
Rate of accuracy for fraud cases (F) = 42.86%
Rate of detection = 2.87%/5.34% = 53.73%.

Similarly, calculating for a P = 90 per cent probability threshold, we obtain the following results:
Proportion of sample = X = 0.20%
Rate of accuracy = 80.00%
Rate of detection of fraud = 2.99%

If these same calculations are repeated for different thresholds (P), we obtain the next summary table (Table 5).

Table 5 dictates the following conclusions: If the goal is detection, a large number of files must be sampled. Note that a sample where P > 10% contains more files than a sample where P > 20%. This methodology has the advantage of ferreting out a maximum number of suspect cases. It does, however, suffer from two shortcomings. The first is that reviewing a large sample of files is costly for the company. An average company which decides to re-examine every file with a fraud probability of over 10 per cent will have to pay for the review of 1,797 files[8] (automobile accidents excluding broken windows). The second shortcoming is that re-examination of a broad range of files will entail some "unfairness" towards clients who are not cheaters but who will be closely investigated by their insurance company. For example, if the 10 per cent threshold is chosen, we know that there will only be a 31.42 per cent level of

[8] This figure is equivalent to 11.80 per cent of the 15,224 claims files an average automobile insurance company handles annually. These figures are drawn from the Dionne and Belhadji study (1996).

Table 5:
Rates of accuracy and detection for different thresholds of probability of fraud

Probability	**Percentage of sample (%)**	**Rate of accuracy in %**	**Rate of detection in %**
10%	11.80	31.42	69.40
15%	8.61	38.43	61.94
20%	6.70	42.86	53.73
25%	4.90	50.41	46.27
30%	14.11	53.40	41.04
35%	13.07	61.04	35.07
40%	2.75	62.32	32.09
45%	1.99	62.00	23.13
50%	1.79	60.00	20.15
55%	1.63	58.54	17.91
60%	1.16	58.62	12.69
65%	0.80	55.00	8.21
70%	0.60	53.33	5.97
75%	0.48	58.33	5.22
80%	0.44	54.55	4.48
85%	0.32	62.50	3.73
90%	0.20	80.00	2.99
Average 20.75%	5.50	47.10	48.51

accuracy (see Table 5) and, thus, over two-thirds of the sample selected will not be cheaters (according to the adjusters in our study).

If, by contrast, the only concern is rate of accuracy, a very small number of files will be selected for re-examination. The advantage of applying such a method is the relatively low total costs of reviewing such a small sample. Indeed, an average company having decided to re-examine all files with a fraud probability of over 90 per cent will have to pay for the review of only 31 files[9] (automobile accidents excluding broken windows). The other advantage is the level of accuracy: whereas the accuracy of the 10 per cent threshold was 31.42 per cent, the 90 per cent threshold is 80.00 per cent accurate. In this case, more than three out of four cheaters are detected. The shortcoming of this decision threshold (90%) is that only 2.99 per cent of fraud is actually detected; most cheaters (97.01 per cent of cases) will slip pass the company.

Having examined these extreme cases ($P = 10\%$ and $P = 90\%$), we can see that there is a trade-off between detection and accuracy (see Table 5): the higher the fraud probability

[9] This figure is equivalent to 0.20 per cent of the 15,224 claims files an average automobile insurance company handles annually.

threshold the greater the accuracy and the weaker the detection. A very conservative company with no wish to get too involved in fighting fraud will probably opt for very high thresholds. An aggressive company with a strong wish to fight fraud and reduce its costs will opt for rather low thresholds. Finally, note that a threshold of about 25% (average threshold for the industry) will allow us to detect roughly half the fraud while achieving about 50 per cent in accuracy.

This table can also be used by companies wishing to know their rates for accuracy and for detection of fraud. If a company knows its budget constraints and decides to allow a certain sum (Z) to in-depth investigations, it will thus know the number of files to select in conducting its investigations. To each sample size will correspond a rate of detection and accuracy (if the fraud rate is the same as that in the industry).

4.3 Cases where our model predicts correctly

According to this approach, our model for predicting fraud is used to estimate the overall incidence of fraud. We consider all the cases with a probability above a given threshold as fraudulent. However, for the proportion of cases below the chosen probability threshold, the opinions of adjusters do come into play. This thinking is in some sort motivated by the fact that our model (with its indicators) fails to explain all aspects of fraud. We are here thinking in particular of the essentially subjective aspects of a routine investigation.

Investigators may also fail to observe all the elements that point to fraud. The total number of estimated fraudulent cases may consequently be greater than the number evaluated as fraudulent by adjusters.[10] All the fraud detected by the model but not detected by investigators is added in. We should remember that in the preceding case, fraud amounted to 5.34 per cent (equivalent to a threshold of $P = 20.75\%$). In this case, where all the files whose probability of fraud exceeds this threshold are considered fraudulent, fraud would climb to 8.25 per cent. The rate differential (i.e. 2.91%) could be due to hidden fraud which went undetected by adjusters. We are in fact faced with three scenarios:[11]

(1) A portion of fraud is detected simultaneously by our model and the investigators. These are the files with a probability of fraud higher than the threshold and which were spotted by the adjusters (65 cases).
(2) Another portion is detected by the model but not by investigators. These are the files whose probability of fraud is higher than the threshold but which were not detected by adjusters. This is what we have called hidden fraud ($138 - 65 = 73$ cases).
(3) Finally, a third portion is made up of files detected by adjusters but not detected by the model ($2,371 - 2,302 = 69$ cases). We mentioned above the subjective aspects at play in spotting these fraudulent files. Their probability of fraud fell below the threshold. These three cases (1, 2, and 3) are reproduced in the following figure:

[10] See Caron and Dionne (1998) for further details. Their model's best estimator generates a 10 per cent fraud rate or $113.5 million in the industry.

[11] The model would select 138 cases from which 65 would also be selected by the adjusters. Moreover, the number of non-fraudulent cases from the model would be 2,371 from which 2,302 would be by the adjusters. Consequently, the total number of fraud cases detected is given by: $138 + (2{,}371 - 2{,}302) = 207$ and $207/2{,}509 = 8.25\%$.

Threshold

Cases of non-fraud (model) Certain cases are considered fraudulent by investigators (3)	Cases of fraud (model): (1 + 2) Certain cases are considered non-fraudulent by investigators

5. Some notions about the costs of in-depth investigations

In deciding to pursue an in-depth investigation, the insurance company is faced with the following choice: if the cost of settlement without in-depth investigation is lower than the expected cost with investigation, there will be no investigation; otherwise, the investigation will be pursued.[12]

Given the following costs:

C_1: Cost of settlement without in-depth investigation
C_2: Cost of settlement after in-depth investigation
C_c: Cost of conducting an in-depth investigation[13]
$C_1 = R_1$ (value of the claim net of deductible)
$C_2 = C_c + p \,.\, R_2 + (1 - p) \,.\, R_1$

where:

R_2: amount of the claim paid if the investigation succeeds
p: probability investigation will succeed

Note that the probability of the investigation's success will generally depend on several factors such as the probability of fraud (calculated and noted as P), experience and training of investigators (noted respectively as EX and T), and the organization of the firm (noted as SIU). For example, the firm's organization may or may not include Special Investigation Units. This probability of success may thus be represented by the following function:

$$p = f(P, EX, T, SIU).$$

It should, however, be mentioned that the probability of fraud is a determining factor in the above function: a file with a very weak probability of fraud (P) should not generate a high probability of success (p) no matter what the values of the other factors (EX, T, SIU).

Finally, the firm's decision will take the following form:

If $C_1 \geqslant C_2 \Rightarrow$ The investigation is conducted.
If $C_1 < C_2 \Rightarrow$ The investigation is not conducted.

By replacing the values C_1 and C_2 in the first inequality above, we find that the company will conduct the investigation in the case where:

---

12 This deduction is valid only in short-term static thinking. A company which bases its decisions on more long-term strategies may wish to pursue the investigation even if, on average, its costs with investigation exceed its costs without investigation. In so doing, this company will earn a reputation for being active in the fight against fraud and cheaters will switch to other companies more tolerant of fraud. In the end, this will mean lower total costs (discounted) for claims and investigations.

13 This amount may be estimated by the average cost of conducting an in-depth investigation. It is however more accurate to consider each cost as being specific to a particular file. Some files are more difficult to handle than others.

$$p(R_1 - R_2) - C_c \geqslant 0.$$

In other terms, the investigation will be conducted if the expected claim–cost differential will at least cover the cost of the investigation.

6. Avenues of research

The model developed is based on claims files randomly selected from the population of files belonging to participating insurers. It thus treats all possible cases of fraud for any conceivable claim. The model generates a probability of fraud for an ordinary file (see Figure 1).

Insurance company executives may not only be interested in the probability of fraud in an ordinary file, but also in the probability an in-depth investigation will be successful (see Figure 2). Indeed, these two complementary steps should be executed simultaneously: Figure 3 sums up the procedure for this.

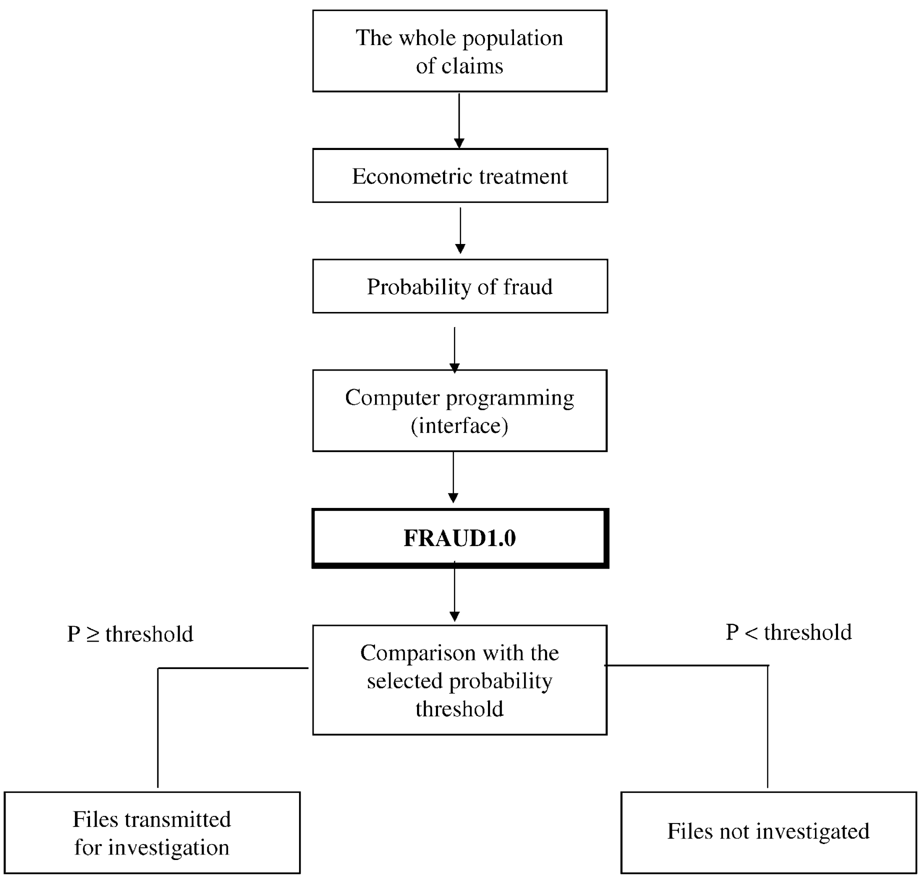

Figure 1: Determination of the probability of fraud (decision of adjusters)

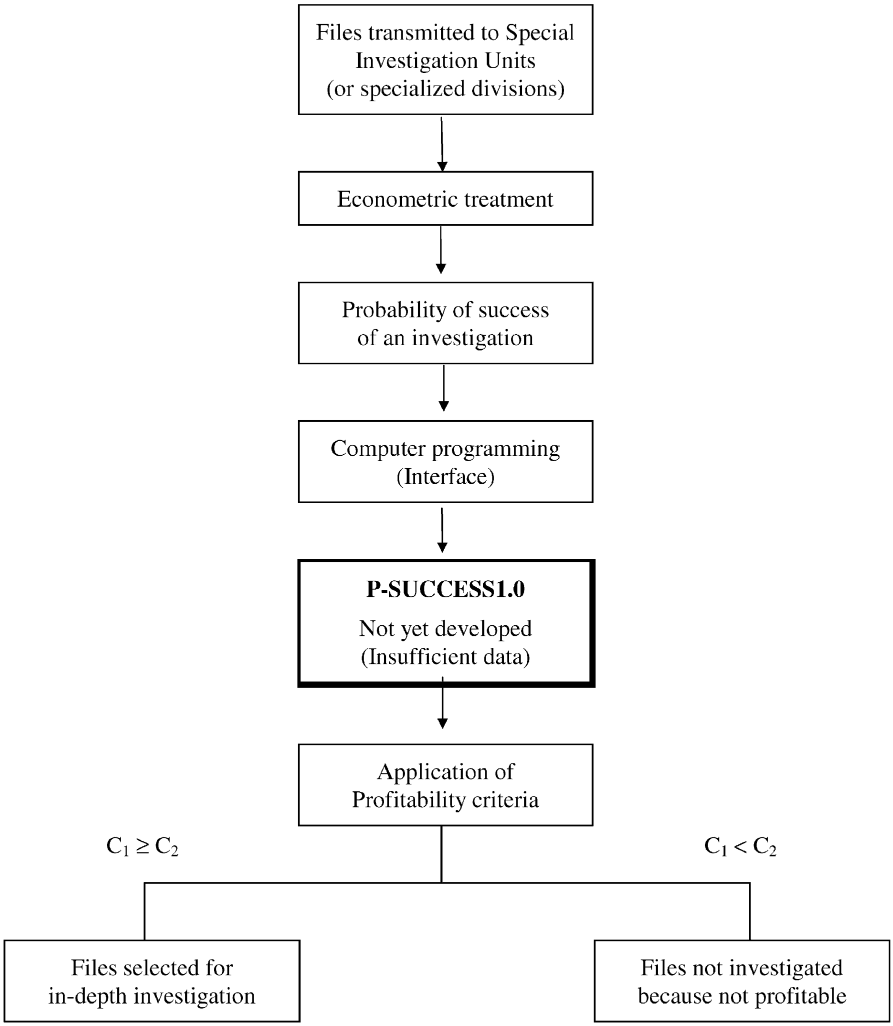

Figure 2: Determination of the probability of success of an investigation (decision by Special Investigation Units)

It is first a matter of deciding whether or not an ordinary file should be investigated: this is the goal of the present work where we have used a model to generate a probability of fraud. Our sample in this work was made up of 0.64 per cent of all the industry's automobile claims files (excluding broken windows).

Next, the so-called serious cases which have received special investigation must be studied. In these cases, the files having undergone in-depth investigation would need to be sampled. The ideal would be to sample a portion of the files handled by the Special Investigation Units (or any other division charged with conducting special investigations). To

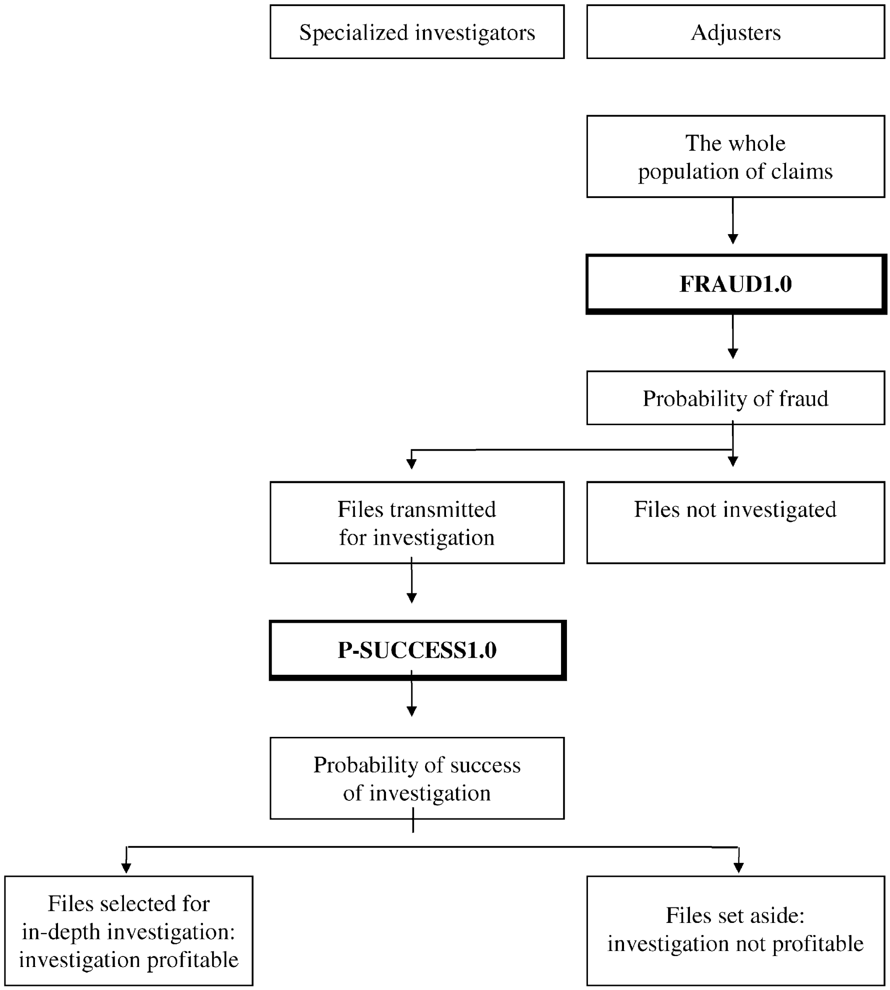

Figure 3: Probability of fraud and probability of success of an investigation (decision at two different levels)

the indicators in the first stage of the investigation could be added indicators specific to the second stage (in-depth investigation). This stage of investigation would take into consideration the indicators themselves as well as the training and experience of specialized adjusters. The goal of this stage would be to generate a probability of success for the in-depth investigation. After this, the investigation's profitability could be calculated in monetary terms. In our view, it is this second stage which must be achieved in order to complete the process of automatic investigation.

Finally, we have developed a user-friendly software with a graphic interface linking the probability of fraud function developed in this article. The procedure is explained in Belhadji and Dionne (1997). The user will enter the information required for the different indicators. The fraud (probability) score will appear at the bottom of the "questionnaire sheet". Based on this probability score, supervisors can decide whether or not to investigate certain files in depth. A complementary feature was conceived to measure the profitability of pursuing an investigation in each particular case. However, for lack of certain data, this second part is still in the exploratory stage.

It is important to stress that the model cannot be directly applied to a specific firm, since the parameters used are derived from calculations made with data from the industry as a whole (see Dionne and Belhadji, 1996). The potential user will have to conduct his own study – based either on all his files for a specific time period or on a random sample – and then compute the corresponding parameters.

REFERENCES

AUTOMOBILE INSURANCE FRAUD STUDY, 1991, University of Florida: Florida Insurance Research Center.

BELHADJI, E.B. and DIONNE, G., 1997, "Development of an Expert System for the Automatic Detection of Automobile Insurance Fraud", Working Paper 97-06, Risk Management Chair, HEC-Montreal.

BROCKETT, P.L., XIA, X. and DERRIG, R.A., 1998, "Using Kohonen's Self-Organization Feature Map to Uncover Automobile Injury Claims Fraud", *Journal of Risk and Insurance* no. 65, 2, pp. 245–274.

CARON, L. and DIONNE, G., 1998, "Insurance Fraud Estimation: More Evidence from the Quebec Automobile Insurance Industry", in Dionne G. and Laberge-Nadeau, C. (eds), *Automobile Insurance: Road Safety, New Drivers, Risks, Insurance Fraud and Regulation*, Kluwer Academic Publishers, pp. 175–182.

DERRIG, R.A. and WEISBERG, H.I., 1998, "AIB Claim Screening Experiment Final Report", Insurance Bureau of Massachusetts.

DIONNE, G. and BELHADJI, E.B., 1996, "Évaluation de la fraude à l'assurance automobile au Québec", *Assurances* no. 64, October, pp. 365–395.

INSPECTEUR GÉNÉRAL DES INSTITUTIONS FINANCIÈRES, 1994, "Rapport annuel sur les assurances".

JOHNSTON, D.J., 1997, "Combating Fraud: Handcuffing Fraud Impact Benefits", *Assurances*, no. 65, July, pp. 175–184.

PROMOTERS OF CREATIVE CLAIMS SOLUTIONS, 1993, "Fraud Indicators", Nationwide Insurance. Property and Casualty Claims Services, March.

WEISBERG, H.I. and DERRIG, R.A., 1993, "Fraud and Automobile Insurance", Insurance Bureau of Massachusetts.

Appendix

Fraud indicators

Indicators

The indicators below are reproduced exactly as they were presented to adjusters. Note, however, that adjusters were given detailed explanations of these indicators along with the list. The indicators in bold have been flagged as significant in explaining the probability of fraud in a file.

Accident/damage

1) **No police report when there should have been one (or report requested at the police station whereas the accident took place in sectors where the police usually respond quickly).**
2) **A minor collision has led to excessive costs.**

3) Existence of damage not related to the loss or inconsistent with the facts reported about the accident (1).[1]
4) Existence of damage not related to the loss or inconsistent with the facts reported about the accident (2).[2]
5) The vehicle is reported stolen and found shortly after with heavy damage.

Vehicle

6) Recent and expensive model.
7) Insured cannot produce bills for maintenance of the vehicle.
8) The vehicle is not attractive to thieves (i.e. ordinary old car)
9) Vehicle stolen in a shopping centre.

Agent/insurer

10) Shortly before the loss, the insured checked the extent of coverage with his or her agent.
11) The insurance agent has never seen the insured vehicle.

Financial

12) The insured is having personal and business-related financial difficulties (1).[3]
13) The insured is having personal and business-related financial difficulties (2).[3]
14) The insured's occupation does not justify the high value of his or her vehicle (and its accessories).

Settlement

15) The insured (claimant) is **too** eager to receive monetary compensation in lieu of repairs on his or her vehicle.
16) The insured is ready to accept a relatively small settlement rather than produce all the documents linked to the loss.

Claimant/insured

17) The insured is very insistent on a quick settlement.
18) The insured is extraordinarily familiar with the insurance and vehicle repair jargon (1).
19) The insured is extraordinarily familiar with the insurance and vehicle repair jargon (2).
20) The insured offers to come to the claims office for payment.
21) The insured (or claimant) is too eager or too frank to accept blame for the accident.
22) The accident (or loss) took place shortly after the vehicle was registred and insured or in the months preceding the end of the policy (or of coverage).

---

[1] This indicator is significant when it is ranged as most important indicator by the adjusters.
[2] In contrast, the same indicator is not significant when it is ranged between 2 and 12 in the adjusters order of importance.
[3] This indicator is significant not only when it is ranged as most important indicator, but also when it is ranged between 2 and 12 in the adjusters order of importance. See also nos 18, 19, 33 and 34.

23) Appealing to Act 68, the insured refuses to give his or her consent (for an in-depth investigation).
24) Numerous taxi receipts or bills for rental of vehicle from a body shop.
25) Bills or proofs of payment which seen phony or forged.
26) Insured's record: he or she has already been convicted of fraud or has already committed misdemeanours suggesting a potential for fraud.

Title and ownership

27) The history of the ownership of the vehicle cannot be established.

Other indicators
Accident/damage

28) The claim (and/or loss) is reported by a third party.
29) Documentation of the estimate and repairs is not available.
30) All the damaged vehicles are sent to the same garage owner.
31) The claims representative (adjuster) is not allowed to examine the vehicle.
32) The vehicle was repaired before being checked by the claims representative (adjuster).
33) Contradictory witness reports concerning the circumstances of the loss (1).
34) Contradictory witness reports concerning the circumstances of the loss (2).
35) The insured denies the versions of witnesses concerning the accident.
36) Accident involving a single vehicle.
37) Accident involving an unidentified third party.

Vehicle

38) Vehicle with a history of mechanical problems; the manufacturer's guarantee has expired.
39) Rented vehicle with high mileage.
40) Ignition block intact after vehicle has been recovered.
41) No signs of forced entry (door lock intact . . .).

Financial

42) Loan payment on vehicle late.
43) Vehicle purchased with cash.
44) Insured is unemployed; works in a depressed industry; lives in a poor region.

Claimant/insured

45) Vehicle found by insured.
46) Several types of coverage; several policies; over-insured loss.
47) Premium paid in person and in cash.
48) Problems with address: post office box; motel; false address; insured absent; lives with friends, etc.
49) Claimant is very aggressive (threatens to call a lawyer, contact the government, etc).
50) Claimant refuses to answer questions about the accident.

51) During the investigation, insured is nervous and seems confused.
52) Numerous claims filed in the past.

Title and ownership

53) Title recently transferred from another province (or another state).
54) Title of ownership is still in name of the previous owner.