<http://www.diva-portal.org>

This is the published version of a paper presented at *20th Nordic Conference, NordSec 2015 Stockholm, Sweden*.

Citation for the original published paper:

Lopez-Rojas, E A., Axelsson, S. (2015)

Using the RetSim Fraud Simulation Tool to set Thresholds for Triage of Retail Fraud.

In: Sonja Buchegger • Mads Dam (ed.), *Lecture Notes in Computer Science - Secure IT Systems*

Stockholm, Sweden: Springer

<http://dx.doi.org/10.1007/978-3-319-26502-5>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:bth-10868>

Using the RetSim Fraud Simulation Tool to set Thresholds for Triage of Retail Fraud

Edgar Alonso Lopez-Rojas¹ and Stefan Axelsson²

¹ Blekinge Institute of Technology, `edgar.lopez@bth.se`

² Gjøvik University College, `stefan.axelsson@hig.no`

Abstract. The investigation of fraud in business has been a staple for the digital forensics practitioner since the introduction of computers in business. Much of this fraud takes place in the retail industry. When trying to stop losses from insider retail fraud, triage, i.e. the quick identification of sufficiently suspicious behaviour to warrant further investigation, is crucial, given the amount of normal, or insignificant behaviour.

It has previously been demonstrated that simple statistical threshold classification is a very successful way to detect fraud [15]. However, in order to do triage successfully the thresholds have to be set correctly. Therefore, we present a method based on simulation to aid the user in accomplishing this, by simulating relevant fraud scenarios that are foreseeing as possible and expected, to calculate optimal threshold limits. Our proposed method gives the advantage over arbitrary thresholds that it reduces the amount of labour needed on false positives and gives additional information, such as the total cost of a specific modelled fraud behaviour, to set up a proper triage process. With our method we argue that we contribute to the allocation of resources for further investigations by optimizing the thresholds for triage and estimating the possible total cost of fraud. Using this method we manage to keep the losses below a desired percentage of sales, which the manager considers acceptable for keeping the business properly running.

1 Introduction

The economic impact of fraud by staff can be substantial in several types of business. Thus the detection and management of fraud is an important topic. In the retail store the cost of fraud is of course ultimately transferred to the consumer, and finally impacts the overall economy. For example; in one recent case the major US home improvement chain *Home Depot* was the target of a fraudulent return scam where two inside members of staff netted several thousand dollars before being caught. They perpetrated the fraud by abusing their knowledge of the processes of the shipment and return of products [8]. Retail fraud was estimated to cost US retailers about 42 billion dollars in 2013 from which 43% was committed by dishonest staff [2]. Due to the seriousness of this type of fraud, both EU and US recently started to mandate the use of fraud detection as one part of the minimum security requirements for financial services [7, 6].

The constant change of criminal behaviour and patterns, and the introduction of new fraud schemes makes this an important topic for researchers and practitioners alike. For a multitude of reasons (e.g., privacy related, legal, financial, or contractual) the state of practice in fraud research is to work with sensitive and hence secret data [12]. This difficulty hinders researchers to develop methods to detect, prioritize investigations (triage process), and finally share the results with other researchers without problem. We name this problem the *data secrecy problem*.

The *data secrecy problem* has previously been addressed by the use of synthetic data, generated by the RetSim simulator [14]. However, a simulator has other benefits aside from being able to share relevant data for research. One major advantage is that different scenarios can be tested. In these scenarios parameters such as the number of fraudulent staff, their propensity to perpetrate fraud, cost of merchandise etc. This enables the testing of e.g. new fraud detection schemes. These detection schemes can later be applied to real data, so we can prioritize and allocate resources for performing further investigation of fraud.

In this paper we address the problem of how to apply statistical threshold detection to perform triage by proposing a technique which uses synthetic data, generated by the RetSim simulator [14], to test different fraud scenarios so that thresholds can be set and the resulting performance studied. These methods can later be applied on real data.

In our study we target fraud caused by fraudulent refunds performed by sales staff. This is a common type of fraud and accounts for around 28% globally of the total fraud in a retail store, with this situation being more critical in North America [2]. This fraud scheme takes advantage of the lack of security controls inside the store and the difficulty to perform inventory control often in most of the retail business settings. Once the missing inventory is noticed by the inventory control officers, a digital forensic investigation can be performed over the available evidence to identify the people responsible, which in this particular case happens to be associated in most of the cases with the salesperson staff.

However, it is difficult to prioritise which of the staff members should be investigated, especially when we are dealing with a chain of stores with multiple branches and the corresponding number of staff. When investigating losses from theft and fraud in the retail setting, we are most often not interested in finding every last instance of fraud but rather of limiting our overall loss to an acceptable level, often put as a set percentage of overall turnover. Spending time and resources on the investigation of the pettiest of thefts of office supplies is counter productive, as the investigative resources are both scarce and expensive. It can also negatively affect the workplace atmosphere. Thus, being able to focus the investigative effort on the cases that can affect the bottom line is vital.

Thus there is an evident need to prioritize and allocate personal for performing such investigations of staff fraud in the retail sector. We do this by using a triage process model and categorising the fraud threat into critical, important and low impact. However, the financial and transaction data available is large

and hence we need an effective way of performing triage, to focus efforts where they may make the most impact, and hence keeping the total loss to fraud at a set, acceptable, level.

2 Related Work

Simulations in the domain of retail stores have traditionally been focused on finding answers to logistics problems such as inventory management, supply management, staff scheduling and customer queue reductions [5, 19, 4]. We find no research focusing on simulations generating fraud data to be used for fraud detection in retail stores besides the RetSim simulator [14].

One of the reasons data is simulated instead of taken directly from the original source is the *data secrecy problem*. In our experience, the privacy of the customers has always been the main concern when disclosing any transactional data. This can be seen by the lack of any kind of public transactional data set that reflects financial statement of individual persons. Many anonymization techniques have been used to preserve the privacy of sensitive information present in data sets. But de-anonymizing data sets is not an insurmountable task, far from it [17]. This is one of the reasons why we have decided to use simulation techniques to keep specific properties of the original data set, such as statistical and social network properties, and at the same time providing an extra layer of insulation that pure anonymization does not provide.

However, using a simulator also has many other benefits, the main one being that the experimenter is in total control of the environment and can vary parameters to try different scenarios; increasing and decreasing the intensity and severity of fraud, for example.

There are tools such as IDSG (IDAS Data and Scenario Generator [11]) that were developed for the purpose of generating synthetic data based on the relationship between attributes and their statistical distributions. IDSG was created to support data mining systems during the testing phase, and it has been used to test fraud detection systems. The RetSim approach differs in that it is implementing an agent-based model which is based on agent micro behaviour rather, than a fixed statistical distribution of macro parameters.

Other methods to generate the necessary fraud data have been previously proposed [20, 16, 10]. The work by Yannikos et al. [20] lets the user specify the assumptions about the environment at hand; i.e., there is no need for access to real data. However, this will certainly affect the quality of the synthetic data. The work by Lundin et al. [16] makes use of a small sample of real data to generate synthetic data. This approach is similar to the one in RetSim. However, the direct use of real data to prime the generation of synthetic data is limited in that it makes it harder to generate realistic data with other characteristics than those of the original real data [20]. The work by Kargupta et al. [10] focused on privacy-preserving methods for data mining. However, that method also does not have the possibility of generating realistic data with other characteristics than those of the original data. RetSim, uses social simulation, which makes it

Fig. 1. Overlap of Two Runs of RetSim vs Real Data

possible to change the parameters of the agents in the model to create realistic synthetic data, potentially producing emergent behaviour in the logs which is hard to produce in other ways.

Previous research on fraud detection algorithms has showed that data mining and machine learning algorithms can identify novel methods of fraud by detecting those records that are different (anomalous) in comparison with benign records, e.g., the work by Phua et al. [18]. This problem in machine learning is known as *novelty detection*. Furthermore, supervised learning algorithms have been used on synthetic data sets to prove the performance of outlier detection [1, 16]. More particularly in retail stores the use of pattern discovery to address retail fraud has been used by Gabbur et al. [9] with many limitations to train a classifier due to the lack of reliable fraud data. However none of these studies made use of synthetic data from retail stores. To our knowledge, there has been no investigation of what the limits of effectiveness of e.g. simple threshold based monitoring are over other complex techniques such as machine learning and pattern recognition.

3 RetSim: the Simulator for Retail Store Data and Fraud

Since we have access to several years worth of transaction data from one of the largest Scandinavian retail shoe store chains, we made use of *RetSim*[14], a *Retail shoe store Simulation*, built on the concept of Multi-Agents Based Simulation (MABS). *RetSim* is intended to be used in developing and testing fraud scenarios at a retail shoe store, while keeping business sensitive and private personal information about customers consumption secret from competitors and others.

Table 1. Triage Threshold Limits

Data Set-Triage	Red 98%	Blue 95%	Green 90%
rs5125-No Fraud	0.136	0.133	0.128
rs3712-Moderate Fraud	0.191	0.182	0.147
rs3302-Aggressive Fraud	0.209	0.185	0.170

RetSim uses the relevant parameters that govern the behaviour in and of a retail store to simulate *normal* behaviour. The output of this process is a synthetic data set that contains similar properties as the original data set. We also model the malicious behaviour of staff and simulate this behaviour together with our normal behaviour to produce a rich data set useful for fraud research.

One of the main advantages of simulating data for fraud over real data is that it can quantify the loss due to the identification of malicious agents since the activities of these are known [12, 13]. Due to this capability, one of the main results of previous research with RetSim is that in many cases a proper setting of a threshold detection control can be enough to keep the loss of a business to fraud, at a desired level, and at the same time avoiding the cost and complexity of more advance methods, such as data mining and machine learning.

Fraud in the retail setting is in many cases perpetrated by the staff so we have decided to focus on that. A common example of such fraud is *Refunds due to Fraudulent Returns*. In this paper we make use of the *RetSim* tool to study this specific fraud scenario that includes agents defrauding the store and performing known fraud behaviour patterns.

With the help of *RetSim*, we produced a simulation that results in data comparable to our real data set. The generated synthetic data set contains 36 salesmen and around 45,000 receipts and 81,500 articles sold. The simulation was seeded with a subset of about 11,000 articles from the real store (that we named *Store 1*). One of the challenges when simulating data is to evaluate how realistic the data is in comparison with the source. In Figure 1 there is evidence of the similarities from an overlapping plot of the generated distribution of sales by price of both: original data set (*Store 1*) and simulated data sets (*rs3658* and *rs5125*). In this paper we make use of *rs5125* as a reference data set for a *No Fraud* scenario. This evidence is part of previous work using this tool to generate a realistic data set for research [14].

Now that we have a proper data set simulated that resemble the original data, the next step is to inject malicious behaviour that can be used for research. In the chosen scenario of fraudulent returns we include cases where a salesman creates fraudulent refund slips, keeping the cash refund for him- or herself. In terms of the object model used in RetSim the refund scenario can be implemented by: Estimating the average number of refunds per sale and the corresponding standard deviation. Use these statistics for simulating refunds in the RetSim model. Fraudulent salesmen will perform normal refunds, as well as fraudulent one. The volume of fraudulent refunds can be modelled using specific parameters

Fig. 2. Triage cut off using as reference no fraud behaviour

that determine the *aggressiveness* of the fraudulent behaviour. The “red flag” for detection will in this case be a high number and value of refunds in average divided by the total sales for a salesman.

4 Triage Process in a Retail Store Scenario

When investigating the instance of fraud in a retail store, it is important to quickly eliminate all the normal background behaviour that is not indicative of fraudulent behaviour. This is of course (hopefully) the overwhelming majority of the transactions. So we need to perform some form of triage, where we quickly identify the abnormal behaviour and single that out for further investigation.

So inspired by the original triage,³ we have chosen to divide the studied behaviour into three categories, bins, to classify suspicious activity of staff members: the first category requires *critical* or urgent investigation, Category one - red line in figures), the second category is *important* to detect significant loss to the business (Category 2 - blue line in figures), the last category is the category where investigation would probably not be fruitful and have *low impact* on business (Category 3 - green line in figures). The idea being an investigator ought to focus on the red category (Critical), maybe keeping an open mind regarding people in the blue category (Important), and disregard the green (Low Impact) as a cost of doing business if indeed any problematic behaviour should lurk in that category.

³ From the French *trier* (v): to separate, sift, or select.

Fig. 3. Triage cut off using moderate fraud behaviour as reference

Fig. 4. Triage cut off using aggressive fraud behaviour as the reference

To illustrate this triage process in a fictitious retail store we use the return fraud scheme as an example, and we wish to set statistical thresholds to identify the limits of the categories. Staff that process many refunds in comparison to their sales are subject to investigation, therefore our definition of suspicious behaviour is: *An unusually high fraction of the total value of refunds to the total value of sales for the individual salesman in question, in comparison with the average value for the sales staff as a whole.* I.e. the fraud score for the individual salesman is:

$$FraudScore = \frac{ValueOfRefunds}{ValueOfSales}$$

However the task of finding the limits for each of the categories is not trivial. If we set the limits to high, much fraud will be unidentified. On the other hand if we set a lower limit we will experience many false positives. False positives are always a problem when doing any form of detection, and must be avoided [3].

To start formulating the problem we assume that all staff members have a similar probability of performing returns, and that returns of valuable articles are more interesting than lower priced articles from the criminal perspective. We chose, somewhat arbitrary, threshold cut-off limits of (90%, 95% and 98%) to cover at least 10% of the staff that are of particular interest due to suspicious fraud behaviour (higher amount of refunds). The calculated values for the thresholds limits are shown in table 1. We do this in order to show how a simulation tool can be used for setting these limits.

Figure 2 shows the values for both the normal behaviour, and two simulations with injected *return fraud*. This figure shows the total value of refunds divided by the total sales for each salesman as percentages, for the three simulations *rs5125*, *rs3712* and *rs3302* explained below.

To model a normal sales scenario we used information about relevant parameters describing normal behaviour without any fraud and selected *rs5125* as the

No Fraud reference data set. As we said before in section 3, this data set was one of the two evaluated against the original data set to verify that we are using a realistic synthetic data set that resembles the original, and maintains interesting properties of sales without revealing specific details of particular customers.

In a normal situation the only information available is the normal behaviour of the refund process (no fraud). If we set up our threshold limits for triage using this data as a reference (as shown in figure 2), we notice that the limits are perhaps too close to each other due to the assumption, as perceived by studying the real data from the store, that each salesman has a similar probability of processing refunds.

The resulting triage processes using an arbitrary threshold for each of the simulations is not optimal for detecting fraud in most cases due to the possible high number of false positives needed to have sufficient effectiveness. For instance, if we set up the thresholds of triage using the data set that contains no fraud (First case Triage 1 in tables 2 and 3), we have many salesmen to investigate that are flagged as the top priority (red). This could lead to an overwhelming effort to investigate all the fraud without the necessary resources.

5 Tuning the Parameters of the Triage set up

In this section we propose a method to set up a triage process that fits the business expectations and limitations. To start, in section 4 we analysed the consequences of setting arbitrary thresholds due to lack of information of possible fraud. Our proposed method works for either detecting new fraud as it takes place or for the processing of historical data of refunds in order to perform a forensic digital investigation.

To begin, we make use of the RetSim simulator to generate a synthetic data set that contains information about an expected fraud behaviour scenario. In this study we started modelling two possible scenarios, one with moderate fraud behaviour and another one with aggressive fraud behaviour.

Table 2. Triage of moderate fraud data set with rs3712-5% 800u (Top Fraud Score)

ID Salesman	Fraction of Sales	Refunded	Total Sales	Total Stolen	Triage1	Triage2	Triage3
S836051140	0.191	-20234	106026	3874	Red	Blue	Blue
S1068592722	0.191	-12774	67000	0	Red	Blue	Blue
S1408212765	0.160	-10168	63409	0	Red	Green	
S1948780723	0.149	-42865.86	287783.1	7702.857	Red	Green	
S1568033761	0.129	-406122.8	3155567	119662	Green		
S1434682851	0.128	-53809.7	418757.1	0	Green		
S193026137	0.124	-14805	118980.4	0			
S24105143	0.123	-35499.56	288014	0			
S705613182	0.120	-99021	822423.2	16449			

The first fraud simulation (*rs3712*) shows a conservative fraud behaviour agent where each of the fraudsters will attempt to commit fraud only if the sales value is worth more than 800 units in the fictitious currency, and the frequency with which he/she commits this fraud is 5% of all sales. The total amount pilfered by all fraudulent agents in a year is 161,630 units in this scenario, which is around 0.43% of total revenue (39,085,000 units).

The second fraud simulation (*rs3302*) represents an aggressive fraud agent behaviour where the threshold to commit fraud is 600 units and the frequency is 10% of sales. The total amount defrauded by all agents is 400,451 units per year, which is around 1.09% of total revenue (36,584,000 units).

The percentage of returns per salesman for the three generated synthetic data sets are plotted together in figures 2, 3 and 4 for comparison purposes. One of the many benefits of using a simulator is that we can flag all sales refunds that are fraudulent. In a real data set this is unknown unless someone has already vetted the entire data set, which is difficult both from a practical and theoretical standpoint. In table 2 and 3 we can partially see the information concerning the plots for those salesmen with higher percentage of refunds per the total sales, more specifically information about total value of sales, value of refunds and the total value stolen.

Now that we have all information required of the expected fraud scenarios, we can make use of these two simulations to set up the same arbitrary thresholds limits of (90%, 95% and 98%). The thresholds limits are shown in figures 3 and 4 and we will name them Triage 2 and 3.

After looking more in detail into the second case (Triage 2 in table 3), we see that inside the categories that are detected as red and blue, there are only 2 salesmen detected. Finally the last case (Triage 3 in table 3), does not flag any salesman in the blue category. But we notice that the amount of red flags for investigation is considerably lowered in comparison to Triage 1.

The evaluation of the fraud detection methods using different triage processes is presented in table 4 and 5. From these tables we see that when the fraud is

Table 3. Triage of aggressive fraud data set with rs3302-10% 600u (Top Fraud Score)

ID Salesman	Fraction of Sales	Refunded	Total Sales	Total Stolen	Triage1	Triage2	Triage3
S836051140	0.214	-3597	16783	1199	Red	Red	Red
S1068592722	0.189	-11087	58619.38	0	Red	Blue	Red
S1568033761	0.175	-553671.7	3160122	267760.9	Red	Green	Green
S1948780723	0.170	-10965	64498.8	5343	Red	Green	Green
S1063661000	0.166	-41420.67	249302.5	0	Red	Green	
S705613182	0.151	-185363	1225409	68065	Red	Green	
S1884511064	0.148	-55632	374763	29095	Red	Green	
S944780329	0.132	-51983.12	392897.4	28989	Blue		
S1888626692	0.123	-66295.58	537627.2	0			

Table 4. Fraud Detection Results for Triage of moderate fraud using rs3712

Statistic	Triage 1	Triage 2	Triage 3
True Positives	3	2	1
False Positives	3	2	1
False Negatives	3	4	5
Detected	131239	11577	3874
Not Detected	30392	150054	157757
Precision	50%	50%	50%
Recall	50%	33%	17%

Table 5. Fraud Detection Results for Triage of aggressive fraud using rs3302

Statistic	Triage 1	Triage 2	Triage 3
True Positives	6	5	3
False Positives	2	2	1
False Negatives	0	1	3
Detected	399253	371463	274302
Not Detected	0	28989	126149
Precision	75%	71%	75%
Recall	100%	83%	50%

more aggressive, the triage process has a higher precision and recall than in any of the other triage set ups. However if the goal is to minimize the false positives we should carefully chose a threshold limit that minimize this value.

If we aim to investigate moderate fraud we should carefully set up triage thresholds somewhere in between the settings for Triage 1 and 2. Triage 3 is very inefficient in this scenario as seen in table 4. From our results in table 5 for the aggressive fraud, we can see that the higher recall is of course with a very low threshold as in Triage 1, but the effort to investigate 8 members of the staff is higher than using Triage 3, where we are still able to detect most of the fraud but lower the number of staff to investigate by half and still detect about 68% of the fraud committed with just one false positive.

Finally, after having all the information available, a fraud investigator can decide on the goal of the investigation and calculate the effort needed to investigate each of the categories and establish new thresholds according to business strategy to detect and prevent fraud, with the certainty that the chosen triage fraud detection strategy will cover a specific fraud behaviour, and therefore minimise the risk of a big loss.

For example if the goal is to detect as much as possible without considering cost, we can set up very flexible thresholds to achieve this goal. If on the other hand the goal is to catch only substantial amounts of losses we can filter out the staff that has not had enough total refund value from the flagged staff to not spend any resources investigating them. However, if the goal is to deter thieves

Fig. 5. Percentage of fraud divided by total sales grouped by number of fraudsters

from stealing, then we can think about prosecuting even minor fraud that become evident and hence easy to flag due to a higher proportion of refunds versus sales.

6 Using the Triage Setup

Now we are ready to start using the threshold for triage. In the previous section we set the triage limits using 3 different ways (Table 1): using no fraud data (rs5125), moderate fraud behaviour (rs3712) and aggressive fraud behaviour (rs3302). In this section we make use of the RetSim simulator to simulate enough probable scenarios where the fraud can vary from none to aggressive fraud that rises up to nearly 9% of total sales as shown in fig. 5.

Using similar settings as the previous simulation we simulated scenarios for 48 different fraud behaviours each one 10 times randomly, which in total were 480 data sets. We did this by changing 3 different parameters of the fraud behaviour: the number of fraudsters, the minimum article value they are willing to defraud, and finally the frequency with which they steal and keeping the same parameters of sales for each store.

Each simulation is based on a store with 36 salesmen that according to different situations could work full time, part time or seasonally. The different circumstances vary the total amount of refunds considerably and the quantities sold by each salesman. This is the main reason for using *the fraction of total value of refunds divided by total value of sales per each staff* as a Fraud Score

Fig. 6. Percentage of loss divided by Total Sales vs Detected

Fig. 7. Percentage of loss divided by Total Sales vs Undetected

in our fraud detection model. We assumed that all the salesmen should perform at or below approximately the same value with this indicator. Any deviation is due to possible abnormal or fraudulent behaviour.

With these scenarios the goal is to study how much fraud we can catch on average, and if we can keep the loss of the business below a certain level.

Fig. 8. False Positives Frequency on Different Triage Models

We will primarily focus on the red flags identified by the triage process since in most of the cases this would be the main focus of investigations, and also where most of the fraud investigation resources are placed.

Figure 6 compares the effectiveness of the different methods. A perfect detection method should detect all possible cases of fraud, and as expected a low threshold such as the one for no fraud can detect almost all frauds. However there is a cost of detecting all fraud, which is the presence of false positives (see figure. 8).

A false positive in the context of a retail store might incur cost and difficulties to investigate all possible salesmen that matches the low suspicious criteria and perhaps if they notice it, an internal conflict due to the submission of staff to unwarranted suspicion.

One important fact about fig 6 is that the moderate and aggressive models used to calculate the triage limits are sufficient in most of the cases, to keep the loss to a maximum of 3% of total sales in the worst case, and most often below 2% of total sales, which are commonly accepted figures for fraud risk (see figure 7).

7 Discussion

The biggest threat to the validity of our research is concerning the synthetic data. Many would think that using a synthetic data set is not the same as

using the original one. That is true to some extent. A synthetic data set is an abstraction of the original data set. The main goal in our case is to preserve the privacy of customers and business by keeping the original source secret but allowing third parties to interact with the essence of the data in order to provide access to the business through a simulation tool that aids to develop a layer of fraud protection.

Now moving to the retail store business, one of the biggest questions when inventory is missing is: where did it go? Missing inventory directly affects the revenue of many retail operations, especially in markets with lower margins. The causes can vary from customer theft, to staff fraud. In this paper we focus on the study of losses due to fraudulent refunds by the staff.

Gathering evidence of this fraud is a difficult task, specially when there is not a clear starting point. The triage process came about in wartime where medical first responders needed a way to prioritise how to allocate resources to wounded soldiers on the battle field. By analogy, using a similar triage process in the retail setting, when presented with evidence of loss due to fraud in the form of missing inventory, can be a useful way to allocate scarce investigative resources.

The goal of the investigation is an important variable when setting up the thresholds for triage. Often the investigation of small losses is costly for the business and we can reduce the number of investigations to those where the amount is substantial enough to affect the bottom line.

After the triage process is properly calibrated to fit the goals of the retailer, then a proper investigation can be carried out. There are many tools available for performing such an investigation, and discussing these go beyond the scope of this paper, but one could for example match the time stamps of refund receipts to video surveillance records of the cash register, to see if the receipts match the expected transaction.

The main concern of the retail fraud executives is to keep the business loss as low as possible. Reducing the cost of fraud is a constant process that requires business resources. Since the final cost of the fraud and the investment in fraud detection ends with the customer, managers have an important role to play in this process. They can either go for the detection of the minimum case of fraud with big investments, or accept that part of the business is to keep running at a certain low level of fraud which does not deeply affect the business end customer. The total percentage defrauded by all agents is in our aggressive simulation around 1.09% of total revenue. Most businesses would consider this an acceptable loss rate because it will not severely affect the final customer.

8 Conclusions

The RetSim simulator is a useful tool to implement a triage process based on a suspected fraud detection scenario. Without knowing the loss for a specific fraud scenario, retail loss prevention executives are basically left to set up arbitrary thresholds for the triage process. When using the RetSim simulator, they can

model the expected cost of fraud by simulating and analysing a synthetic data set with already identified instances of fraud.

Developing a proper and effective triage process without knowing all the underlying details of the expected fraud is difficult. By using the RetSim simulator we can gather enough information for starting a digital forensic investigation since we can model and investigate how different parameters affect the situation.

For threshold detection to be effective, the reference data set that is used to calculate the cut off should contain enough fraud data. Otherwise the range of the categories for a triage process are in risk of being so slim that they can not detect any fraud or perhaps the fraud is outside of the region delineated by the thresholds.

The generation of synthetic fraud is necessary in many cases, to properly set up the triage categories. Our triage set up method substantially reduce the scope of investigation by using triage optimization based on simulated data generated from expected fraud scenarios. It is a management decision to properly set up the threshold according to the resources available for investigating more or fewer cases of suspicious staff returns, given the accepted overall loss to the business. It is off course a task for the managers to decide on the different thresholds to correspond with their business goals, which could be i.e. to maximise the fraud detection while avoiding the cost of invested resources for investigating many false positives, or keeping the total losses below 2% of turnover annually etc. In our examples above, we meet that level.

Further research on this topic will be focused on simulating new types of fraud and at the same time apply these results on the real business to measure the effectiveness of detecting simulated fraud versus real fraud *in vivo*. Another direction of research is to identify other domains where similar methods as those presented in this paper can be applied to circumvent the *data secrecy problem* such as when studying bank transactions, mobile payments and similar financial services.

Acknowledgements

This work is part of the research project "Scalable resource-efficient systems for big data analytics" funded by the Knowledge Foundation (grant: 20140032) in Sweden.

References

1. Abe, N., Zadrozny, B., Langford, J.: Outlier detection by active learning. Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD 06 p. 504 (2006)
2. Arora, R., Khan, A., Deyle, E.: The global retail theft barometer 2014. Thorofare, NJ USA: Checkpoint Systems, Inc.(1-81). Acedido em p. 60 (2014)

3. Axelsson, S.: The base-rate fallacy and the difficulty of intrusion detection. *ACM Transactions on Information and System Security (TISSEC)* 3(3), 186–205 (2000)
4. Bovinet, J.: *RETSIM: A Retail Simulation with a Small Business Perspective*. West Pub. Co., Minneapolis/St. Paul (1993)
5. Chaczko, Z., Chiu, C.: A smart-shop system - Multi-agent simulation system for monitoring retail activities. *20th European Modelling and Simulation Symposium* pp. 20–26 (2008)
6. Council, F.: Supplement to Authentication in an Internet Banking Environment. URL: <http://www.ffeec.gov/pdf/Auth-ITS-Final> pp. 206–222 (2011)
7. ECB, E.C.B.: Recommendations for the Security of Internet Payments. Tech. Rep. January, European Central Bank (2013)
8. FBI: Ticket Switch Fraud Scheme at Home Deopt (2013)
9. Gabbur, P., Pankanti, S., Fan, Q., Trinh, H.: A pattern discovery approach to retail fraud detection. *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD 11* p. 307 (2011)
10. Kargupta, H., Datta, S., Wang, Q.: On the privacy preserving properties of random data perturbation techniques. *Third IEEE International Conference on Data Mining* pp. 99–106 (2003)
11. Lin, P., Samadi, B., Cipolone, A.: Development of a synthetic data set generator for building and testing information discovery systems. In: *ITNG 2006*. pp. 707–712. IEEE (2006)
12. Lopez-Rojas, E.A., Axelsson, S.: Money Laundering Detection using Synthetic Data. *The 27th workshop of Swedish Artificial Intelligence Society (SAIS)* pp. 33–40 (2012)
13. Lopez-Rojas, E.A., Axelsson, S.: Multi Agent Based Simulation (MABS) of Financial Transactions for Anti Money Laundering (AML). *The 17th Nordic Conference on Secure IT Systems* pp. 25–32 (2012)
14. Lopez-Rojas, E.A., Axelsson, S., Gorton, D.: RetSim: A Shoe Store Agent-Based Simulation for Fraud Detection. *The 25th European Modeling and Simulation Symposium (2013)*, (Best Paper Award)
15. Lopez-Rojas, E.A., Gorton, D., Axelsson, S.: Using the RetSim Simulator for Fraud Detection Research. *International Journal of Simulation and Process Modelling* 10(2) (2015)
16. Lundin, E.: A synthetic fraud data generation methodology. In: *Information and Communications Security*, pp. 265–277. Springer, Berlin Heidelberg (2002)
17. Narayanan, A., Shmatikov, V.: De-anonymizing Social Networks. *2009 30th IEEE Symposium on Security and Privacy* pp. 173–187 (May 2009)
18. Phua, C., Lee, V., Smith, K., Gayler, R.: A comprehensive survey of data mining-based fraud detection research. *Arxiv preprint arXiv:1009.6119* (2010)
19. Schwaiger, A., Stahmer, B.: SimMarket: Multiagent-based customer simulation and decision support for category management. *Multiagent System Technologies* pp. 74–84 (2003)
20. Yannikos, Y., Franke, F., Winter, C., Schneider, M.: 3LSPG : Forensic Tool Evaluation by Three Layer Stochastic Process-Based Generation of Data. In: *4th International Workshop in Computational Forensics*. pp. 200–211. Tokyo, Japan (2010)