December 1997

Compliance Monitoring in a Complex Environment: An Overview

Peter Goldschmidt
The University of Western Australia

Compliance Monitoring in a Complex Environment: An Overview

Peter Goldschmidt
The Department of Information Management and Marketing
The University of Western Australia.
Email pgold@ecel.uwa.edu.au

Executive Summary
This paper discusses concepts of compliance monitoring for anomaly detection [CMAD] in the business environment in terms of the functionality of these systems, the computational approaches, the classes of CMAD environments, the decision making requirements and the agents involved in the CMAD decision making. It includes a review of relevant literature on CMAD, reported problems and proposed solutions.

1.0 Introduction

With the increase in electronic commerce using the internet, and the proliferation of national and organisational intranet structures, issues of security and compliance are essential to ensure and maintain the integrity of transactions conducted via this new technology. For the purposes of risk management, governments and commercial organisations typically monitor transactions that may impact on their operations.

The electronic monitoring of data related to individuals and groups of individuals is described as "dataveillance" by Clarke (1988). He highlights the inherent dangers of drawing conclusions resulting from this data, and points out that a major problem in "dataveillance" is the high noise to signal ratio which may result in misleading conclusions.

Subsequent advances made in improving the quality of this data have, in general, reduced the problem of misleading results produced due to this "noisy data". These advances include improvements in data processing and the increased use of sophisticated computational techniques such as statistical, knowledge-based and artificial neural computational methods.

These systems are typically centred on the events being monitored and the events' source agents. The results of these systems however may still require human judgment to determine their validity. A more detailed discussion of the statistical and knowledge-based approaches is presented in section 2.[1]

1.1 The Function of CMAD Systems

The process of categorising an event by its deviation from some predetermined pattern or theory is termed anomaly detection. An anomaly as defined by Garner and Chen (1994) is a subjective, post-data manifestation. This could be a data point which is a member of a set of data points and is determined to be an outlier, or the behaviour of an event's source agent that is manifest by non-compliant, such as fraudulent, behaviour.

Automated data analysis techniques identifying these variances typically depend on data that identifies source agents and their relationships, and is used to draw a compliance agent's attention to a particular event or group of events that indicate possible anomalies. Clarke (1988).

Examples of CMAD range from standard data processing routines that ensure internal control, such as data input, processing and output compliance,[2] to the monitoring of events transacted in more complex environments via sophisticated statistical, artificial intelligent and neural computing techniques, or hybrid combinations. We describe these devices as primary monitoring systems.

---

1 Artificial neural nets are not dealt with in this paper.

2 Weber (1988) provides a comprehensive discussion on this type of CMAD.

In the business environment, primary CMAD systems traditionally function as a process by which the integrity of transaction data as well as the entire transaction, or event, are examined to ensure that both comply with predetermined conditions. If the event complies, it is accepted for future use; if it does not, it is rejected from further processing, pending some remedial action.

The compliance process compares some predetermined conditions of acceptance with the actual data or event, which is a matching process. If any variance is detected between the conditions and the actuals, an exception report is produced, identifying the variance. This identification of the variance either fulfills the conditions of necessary and sufficient evidence and thus determines an instance of non-compliance, or if not, it may be only an indicator of possible non-compliance. In the latter case further evidence may then be sought to fully substantiate the hypothesis of non-compliance.

The CMAD decision process helps to determine if there has been an instance of non-compliance, based on the evidence of an occurrence of a variance between the preset conditions and the actual data or event. The function of a CMAD system is therefore twofold, namely identifying a variance, and producing and accumulating (if required) supporting evidence. When both conditions are met, the evidence points to the detective, corrective or preventative actions as required.

The observed variance takes the form of an exception report, produced by the primary monitoring system, indicating why the exception was triggered. The detective function is fulfilled by the recognition of the variance; the correction function identifies the changes to be made to the data or the event, which can then be re-processed; and the preventative function is fulfilled by recognising and reporting a variance that will result in the suspension or rejection of similar, future events.

Where the evidence or the accumulation of evidence does not directly indicate what action is required, or indicates only the possibility of non-compliance, it is then incumbent on human agents to interpret this evidence to determine what action is required, or to determine if the non-compliant indicator is a true or a false positive directive.

1.1.1 Computational Approaches to CMAD

O'Leary (1991) discusses two approaches to the procedural and declarative CMAD techniques that can provide support for the identification of events in an automated environment. They are demons, and objects.

Demons are defined as computerised routines that are instantiated by data or events received, as opposed to being requested by some program. One of the reasons for their use is that "demons add knowledge to a system without specification of where they will be used ... like competent assistants they do not need to be told when to act" Winston (1977, p. 380). This allows them to be data or event dependent, rather than program dependent, and to provide intelligent self activation for monitoring data when appropriate. The compliance threshold levels would be the appropriate triggers for this activation. O'Leary points out that demons have been developed to monitor patterns for the purpose of auditing activities conducted on computer-based systems. Examples include O'Leary's (1992a) intrusion detection systems which he defines as "those systems which are designed to monitor an agent's activities to determine if the agent is acting as expected or if the agent is exhibiting unexpected behaviour." These systems protect against unauthorised use of computer systems and the protection of the system when transactions are entered. Correspondingly, the demons are activated by an unauthorised user attempting a transaction, or when an authorised user attempts an unauthorised transaction. Vasarhelyi and Halper (1991) describe an alternate audit approach called CPAS, a Continuous Process Audit System. CPAS allows for the continuous audit of on line systems by monitoring transactions to determine the presence of a variance between the monitored information and expected information. Threshold levels are used for the compliance metrics which trigger the appropriate demons. "CPAS monitors key operational analytics, compares these with standards, and calls the auditor's attention to any problems that may exist. Ultimately, this technology will utilise system probes that will monitor the auditee system and intervene when needed" Vasarhelyi and Halper (1991, p. 1).

Objects are the basis of the object oriented paradigm, whereby data is combined with knowledge. An object is defined as an entity with its attributes attached by means of the object's properties. Actions or procedures (such as demons) can also be associated with the object. Attributes and actions form one integrated object. A unique characteristic of this technique is that objects can be a member of a class of

Applying the concept of objects to the CMAD construct, we identify an event as an object whose attributes are defined by the characteristics of the event, and whose actions are demons that are instantiated when the event or class of events occur. We use the notion of a class template, which specifies the common features of a collection of objects. A class is a set of objects which possess the common features specified by the class template. Objects which are members of a class are created by instantiating a class template.

Each attribute can have an assigned value. We call this an object_attribute_value [O_A_V] triplet.

2.0 The CMAD Environment

CMAD can be classified by the level of complexity characterised by the environment in which it operates, and by the decision required to determine instances of non-compliance. The environments are either simple or complex. In practical terms, CMAD systems could fall anywhere on the simple-complex continuum.

Within these environments, decision makers are confronted with problems of different levels of complexity. Figure 2.1 illustrates the breadth and depth of the problem domains within the simple and complex environment.

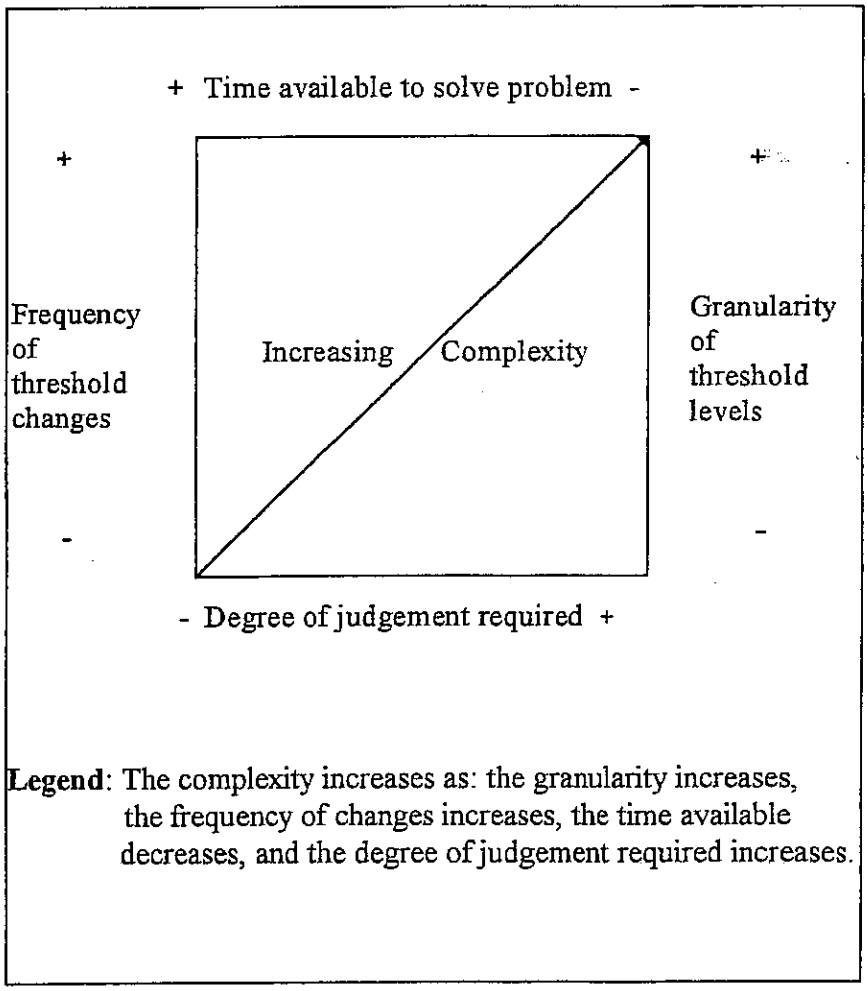

Figure 2.1 Simple / Complex CMAD Problem Domain

2.1 CMAD in a Simple Environment [$CMAD_S$]

The constraints used may take the form of an organisation's predetermined policies and procedures, predetermined constraints needed to ensure data and event integrity, contractual agreements, and statutory requirements. These constraints are not necessarily mutually exclusive and can be seen as bounds or threshold levels.

The various parameters used to construct these levels may change over the longer term. Changes occur because of changes in threshold requirements, such as evolutionary changes in policies and procedures, statutory regulations and changes in data and event requirements.

This environment is called simple for three reasons: 1) the threshold levels either seldom change or change only over the longer term, 2) the identification of the variance fulfills the conditions of necessary and sufficient evidence to determine an instance of non-compliance, and 3) the decisions needed to determine if an event complies lies on the structured to highly structured portion of the decision making continuum.

The degree to which the bounds of the threshold levels are set, very narrow to very broad, determines the type of decision required. Under a simple environment the bounds are narrow, characteristic of structured decisions, such as data input integrity and customer credit checks. Decision making in this environment is ex ante, is made in a single step, and the constraints are all predetermined. Typical examples of $CMAD_S$ would be record, field and batch checks which validate input data to a system. Figure 2.2 shows a macro model of $CMAD_S$ and in section 3 we discuss the concepts of problem and decision structure.

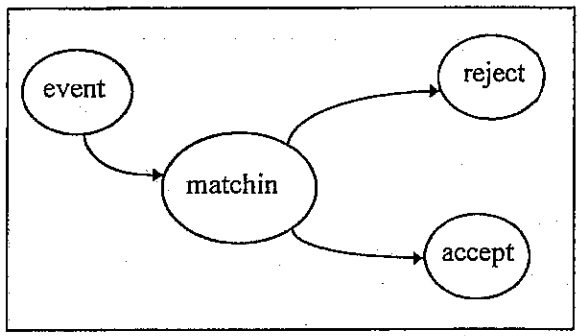

Figure 2.2 A Macro Model of $CMAD_S$

2.2 CMAD in a Complex Environment [$CMAD_C$]

In a complex environment, the decision making is ex post, more complex and may require multiple steps. The event monitoring and decision making is in a domain where the initial monitoring uses a priori thresholds broader than in a simple environment, i.e. it is more granular. This initial monitoring produces exceptions that identify suspected non-compliant events [SNCEs]. Once these exceptions have been produced, it is then the task of the decision maker to substantiate true positive exceptions. True positives are those exceptions that the decision maker has determined are indeed anomalous and have the evidence to support this assertion. To obtain this supporting evidence the decision making uses the results of the initial monitoring as well as important information, related to the event, and characterised by its interpretive nature, requiring judgmental expertise.

This task must be broken down into smaller components and sub-goals must be developed (Simon 1973), namely to identify, categorise and discard any false positive exceptions. These are exceptions that have signaled suspect events that require further scrutiny, and are subsequently rejected by the decision maker, for various reasons. On the other hand, false negatives are events for which the current monitoring facilities do not generate an exception, and allow possible suspect events to slip through the CMAD sieve. If the threshold limits are stringent enough, it can be argued that the marginal false negatives could be subsumed and later considered. Nevertheless, this would not necessarily reduce the occurrences of *true* false negatives as their characteristics may not be known. Figure 2.3 shows a macro model of $CMAD_C$.

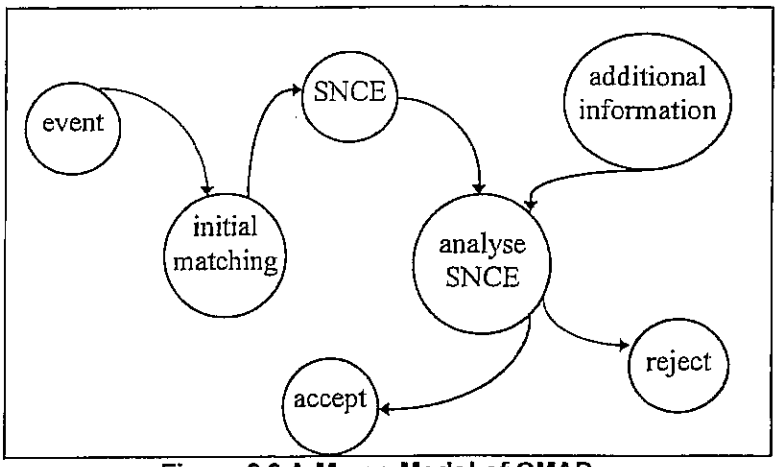

Figure 2.3 A Macro Model of $CMAD_C$

A typical example of $CMAD_C$ would be the monitoring of insurance claims to determine instances of insurance fraud. Operationally, the matching process uses predetermined tolerance metrics to identify possible instances of fraud, such as multiple claims over a period of time. However, in order to substantiate the non-compliant hypothesis, additional information may be required to determine if the identified claim is legitimate or not.

A sample of reported $CMAD_C$ systems includes Lecot (1988), who describes the use of procedural and declarative techniques to assist in the detection of debit card fraud at the Security Pacific National Bank; Byrnes, et al. (1990), who use statistical and declarative techniques to monitor worldwide foreign exchange events conducted by Manufacturers Hanover Bank, to ensure transaction compliance; Major and Riedinger (1992), who describe the use of combining statistical methods with declarative (knowledge based) systems to detect medical insurance fraud; and Senator et al. (1995), who describe the U.S. Treasury Department's Financial Crimes Enforcement Network AI System (FAIS), which identifies potential money laundering activities. FAIS uses statistical methods, database search techniques, declarative systems and graphical users interfaces [GUIs].

The above examples of $CMAD_C$ decision making environments fall on the lower to middle range of the $CMAD_C$ complexity continuum. CMAD in the data intensive capital market [$CMAD_{CM}$], due to the temporal and context sensitive nature of the information relating to the events, tends to the extreme complex end. CMAD in this environment is discussed below.

2.3 CMAD in a Data Intensive Capital Market [$CMAD_{CM}$]

A data intensive capital market [cm] is characterised by its complex and dynamic nature, multiple participants, reliance on timely information flows, and its impact on national and international economies. In this environment, compliance monitoring for anomaly detection [CMAD] differs from other business environments.

$CMAD_{CM}$ is conducted by the regulatory authorities to detect unusual trading behaviour.[3] It is an important tool for building market confidence, thereby increasing market liquidity, and ultimately decreasing the cost of capital to business, and has potentially far-reaching implications for the economy as a whole. Research in $CMAD_{CM}$ has focused on improving the initial identification of anomalous events. This has lead to the introduction of various techniques that have, in general, been an improvement on the preceding ones. The improvements have been partly motivated by a desire to improve the accuracy of the results generated by these compliance systems; to reduce the costs associated with $CMAD_{CM}$ and to increase its reliability, consistency, productivity and effectiveness; to free up resources so that they can be redirected to activities with greater payoffs; and to improve overall risk management, in general, in an increasingly complex and competitive global market.

[3] These may be market specific, regional, national and international authorities.

$CMAD_{cm}$ systems are used by institutions, such as financial intermediaries and capital market providers, to identify anomalous events that have occurred and have influenced, or may potentially influence, subsequent agent behaviour.[4]

The capital market can be conceptualised as objects representing subjects and the relationship between them. Figure 2.4 illustrates this conceptualisation.

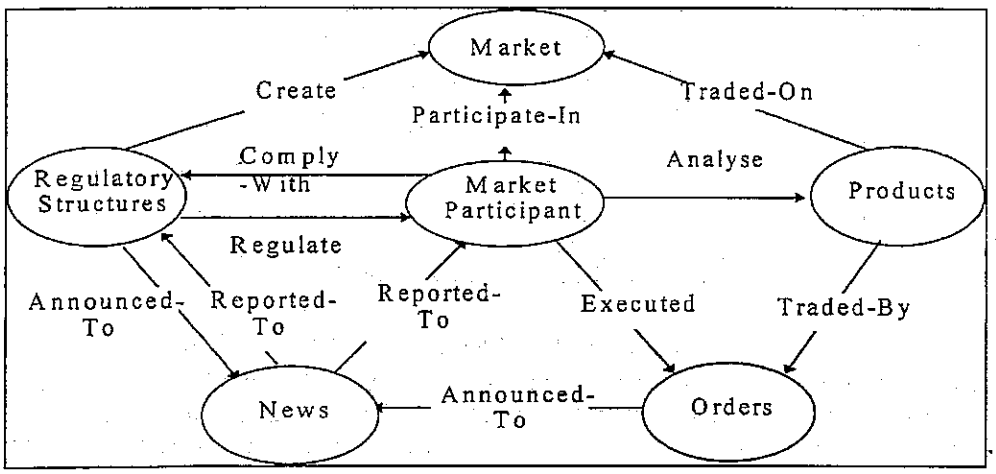

Figure 2.4 Object Oriented Representation of the Data Intensive Capital Market. Source: Freedman and Mathia (1995)

Figure 2.5 shows an expansion of the regulatory component of this capital market object approach.

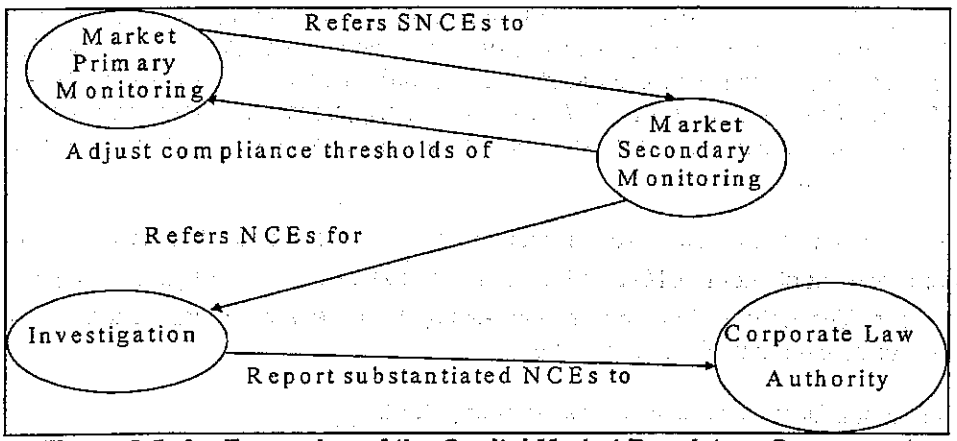

Figure 2.5 An Expansion of the Capital Market Regulatory Component

The $CMAD_{cm}$ systems comprise agents, both human and machine based, with the function of identifying the possible breaches of various rules and regulations pertaining to events conducted on the exchanges. The breaches may be instances of insider trading or market manipulation.

The techniques used for $CMAD_{cm}$ are procedural, declarative or a combination of both. A sample of reported $CMAD_{cm}$s is as follows. Berry and Yanko (1990) and Aitken (1991) describe the current surveillance operation at the Australian Stock Exchange (ASX), which uses an analytical model based on the statistical matching approach to $CMAD_{cm}$. It combines computer-based decision support systems to analyse market events, with communications software, text retrieval and graphics. SOMA primarily uses statistical methods such means, variances, moving averages, days since last traded, etc., to identify SNCEs.

4 "At least on an 'exception' basis, every firm should have a daily capacity for surveilling principal and customer securities activities. If, for no other reason, ... to satisfy regulatory and self-regulating organisations requests for trade information". (Pessin, 1990, p. 415).

Davis and Ord (1990) discuss a sophisticated statistical approach using a form of the Capital Asset Pricing Model to produce a set of compliance indicators or threshold levels. This system is proposed for use by The National Securities Dealers Association's [NASD] surveillance operation. Davis and Ord acknowledge the general problem of setting these threshold levels in an ever changing environment.

> With any set of tolerance levels, deviant (even fraudulently motivated) behaviour may escape detection. Tightening tolerance levels limits increases the likelihood that these exception conditions will trigger an alert. However this increased detection capability does not come without cost. Tightening these limits also increases false positive alerts since the number of instances that fall outside the tolerance must necessarily increase as the limits become more restricted. The cost for the analyst (the decision maker) to review the additional non-exception condition alerts must be assessed in relation to the imputed value of identifying the additional true exceptions detected by more stringent limits. Davis and Ord (1990, pp. 39 - 40).

This cost is due to the temporal and context sensitive nature of the information required to evaluate each exception and to confirm or revoke the evidence supporting the assertion of non-compliance in this complex environment.

Buta and Barletta (1991) describe the Intelligent Market Monitor [IMM], a procedural system supplemented by a case-based reasoning [CBR] approach to $CMAD_{cm}$ built for The Toronto Stock Exchange's [TSE] surveillance operation.[5]

They contend that CBR has the advantage over both the statistical and the knowledge based approaches as CBR "learns" from past cases, whereas the other two approaches do not.[6] Additionally they note that statistical pattern recognition approaches "attempt to observe patterns of stock price and volume measurements, and lack the fundamental view of the companies" and that "even if a reasonable representation of the data can be developed, ... it can only provide a black and white decision. With respect to a declarative approach, they point out the well known "bottle-neck" problem of knowledge acquisition, and the cost of updating existing rules or adding new ones. At this point, it should be pointed out that this CBR is built as an adjunct to a SOMA type primary system. The CBR's function is to supplement the alerted transaction with case details relating to the traded stock, such as recent trading details, macro economic data, ex-dividend dates and the like. The resulting information is then sent to the surveillance analysts for study.

Freedman (1991) and Freedman and Mathai (1995) discuss generic $CMAD_{cm}$ in terms of risk management and regulations. Freedman led the development of the New York Stock Exchange's Integrated Computer Assisted Surveillance System [ICASS]. ICASS incorporates a procedural primary system supplemented by an off line, CBR approach, to assist in identifying instances of insider trading after the initial identification of SNCEs. This CBR system ranks all source agents and uses a suspicion level based on the likelihood of activities being associated with insider trading.

In discussing problems of detecting insider trading, they conceptualise market participating agents, and events, in terms of subjects and relationships between subjects. The subjects are the traders and market providers and the relationships are between these traders and the possessors of proprietary knowledge. They also point out that "Difficulties arise because in general, (1) the unusual subjects are not known - they must be discovered or inferred from the data; (2) the definition of unusual pattern of behaviour is subjective and possibly changes with every analysis and over time; and (3) the quantity of the data in an analysis is overwhelming." (p. 321).

Insider trading is one of two basic types of NCEs. The other is market manipulation. The problems encountered with analytical models used or proposed in the highly complex capital market, including

---

5 The approach is similar to Slade's (1991) CBR approach to support financial decision making.

6 Buta and Barlletta do not indicate that the "learned" cases are automatically available for use by the IMM system, but need a knowledge engineer to manually update the case-base. It should be noted that the author has been advised by TSE that, in practice, the "learning" aspect of this CBR is not as yet automated.

models for predicting market behaviour as well as models used for $CMAD_{cm}$ includes: Incomplete model theories, models often contain incomplete theories as well as incomplete data; Incomplete model inputs, even the best models occasionally produce decisions much worse than a human analyst would, because they do not include some important factors; Incomplete model outputs, the analyst's risk preference in dealing with uncertain outcomes might differ from that of the model; Conversely, the analyst's role is trivialised if the model makes all the decisions; Incomplete explanations, models provide precision at the expense of intuition and common sense.

These analytical, predictive and compliance models are often rejected by the decision makers. Consequently, to compensate for these limitations, some analysts "tune" these results by making heuristic adjustments to the analytical model. This tuning produces a model forecast that is more consistent with intuitive expectations, and maintains the detail and structure of the analytical model. However, as Pindyck et al. (1976) and Freedman et al. (1991) show, tuned forecasts can easily be misused. Alternatively, a cognitive model of an analyst, implemented as an expert system, might perform better at predictive tasks than an analytical model. However, cognitive models fail in domains where there is too much reliance on judgment. In these domains, judgments are dynamic and their representation is difficult to quantify and verify.

Goldschmidt (1995, 1996) and Brown and Goldschmidt (1996) present an extension to the $CMAD_{cm}$ construct. This extension adds an extra dimension to the $CMAD_{cm}$ construct, namely the agents responsible for the interpretation, analysis and classification of SNCEs. This includes the modeling of an aspect of $CMAD_{cm}$ agent's [$CMAD_{cm}$A's or A's] cognitive processes in the form of a cognitive computational model. This model facilitates the accumulation of the (post initial identification) evidence supporting the assertion of non-compliance, and includes the introduction of a multi-agent infrastructure architecture supporting the $CMAD_{cm}$A's review process.

The construct takes the form of an intelligent decision support system using multi-agent technology [IDSS-MAT], supporting a team of analysts whose task is to evaluate the exceptions produced by the $CMAD_{cm}$ primary monitoring systems. The multi-agent components include a relational database which contains, (1) the output from the primary monitoring system, the current SNCE details under scrutiny, (2) links to reference databases, and the SNCE's details and subsequent classification supported by evidence, and (3) control rules, including the coordination knowledge; expert systems [ESs], based on fuzzy set theory and appropriate for each level of expertise in the team hierarchy, and containing the knowledge of the lower level ESs plus (if required) the knowledge specific to that level. A blackboard approach, Hayes-Roth (1983), is used to record control rules, and meta rules controlling part of the heuristic level knowledge: for example, the rules governing which hypotheses to consider given the SNCE alert type presented. The linguistic variables [LVs], each agent's results and the accumulated evidence may also be on the blackboard, depending on the status of the diagnostic process.

This design allows for the complete review of the agents' assumptions, in the form of the LVs, and of their decisions based on the accumulated evidence. It also provides for a more complete CMAD audit trail.

The review process is supported by this construct to assist in the $CMAD_{cm}$ problem solving and the decision making process.

3.0 $CMAD_{cm}$ Problem Solving and the Decision Making Process

Secondary monitoring (SM) problem solving is the human evaluation of the exceptions produced by the primary monitoring system, a process of determining if a generated exception is feasible. This is similar to the analytical review (AR) conducted by auditors and characterised by Libby (1985) as a diagnostic-inference process. Koonce (1993), reviewing past research of cognitive studies of AR, defines AR as the diagnostic process of identifying and determining the cause of unexpected fluctuations in account balances and other financial relationships. Similarly, the SM problem solving is the $CMAD_{cm}$ diagnostic process of identifying and determining the cause of the unexpected variances determined by the primary monitoring facility.

Blocher and Cooper (1988) found that auditors performing AR typically follow four distinct diagnostic inference components: mental representation - the accumulation and evaluation of the relevant problem information; initial recognition of unusual fluctuations in a company's financial statements; subsequent hypothesis generation - the generation of potential causes of the observed fluctuations; and finally, information search and hypothesis evaluation - the search for and evaluation of the information relevant to the causes.

With $CMAD_{cm,}$ the mental representation component is guided by the results of the primary monitoring facility which accumulates and evaluates the relevant compliance problem information leading to the initial recognition of a variance. This is followed by the subsequent hypothesis generation of the potential causes of the observed variance based on the search for and evaluation of the information relevant to its causes.

The diagnostic approach to $CMAD_{cm}$ takes the form of defeasible logic, which means that any inference made may be only tentative, as the inference may require revision if new information is presented. This is due to the default assumption that there is a legitimate cause of the observed variance. It is the task of the decision maker to evaluate all possible legitimate reasons for its occurrence. If none are found, the hypothesis of non-compliance is strengthened.

3.1 $CMAD_{cm}$ Problem Structure

In section 2 we saw that the function of $CMAD_c$ is twofold, the identification of a variance and the accumulation of supportive evidence. Correspondingly, following Sol's (1982), p. 5) definition of problem structure, the structuredness of the $CMAD_c$ problem is also twofold; the identification of a variance is the structured component, and the accumulation of evidence supporting or refuting the NCE hypothesis is the ill-structured component. This is because the variance is typically the product of some algorithm indicating a possible occurrence of NCEs, but in order to substantiate a true NCE, the required accumulation of evidence involves using judgment of agent behaviour. The agents include the source of the event, the identification of the source agents' possible motivations, the environment in which the source agent is operating and the impact this event may have on the environment. These agents are termed source agents [A_s], who may be traders, brokers, and the like.

Additionally, the behaviour of the agent whose judgment is required to evaluate the SNCE is another component that impacts on the accumulation of evidence. These agents are termed evaluating agents [$CMAD_{cm}A_{ei}$], or [A_{ei}], where the i denotes the level of the ith agent's expertise.

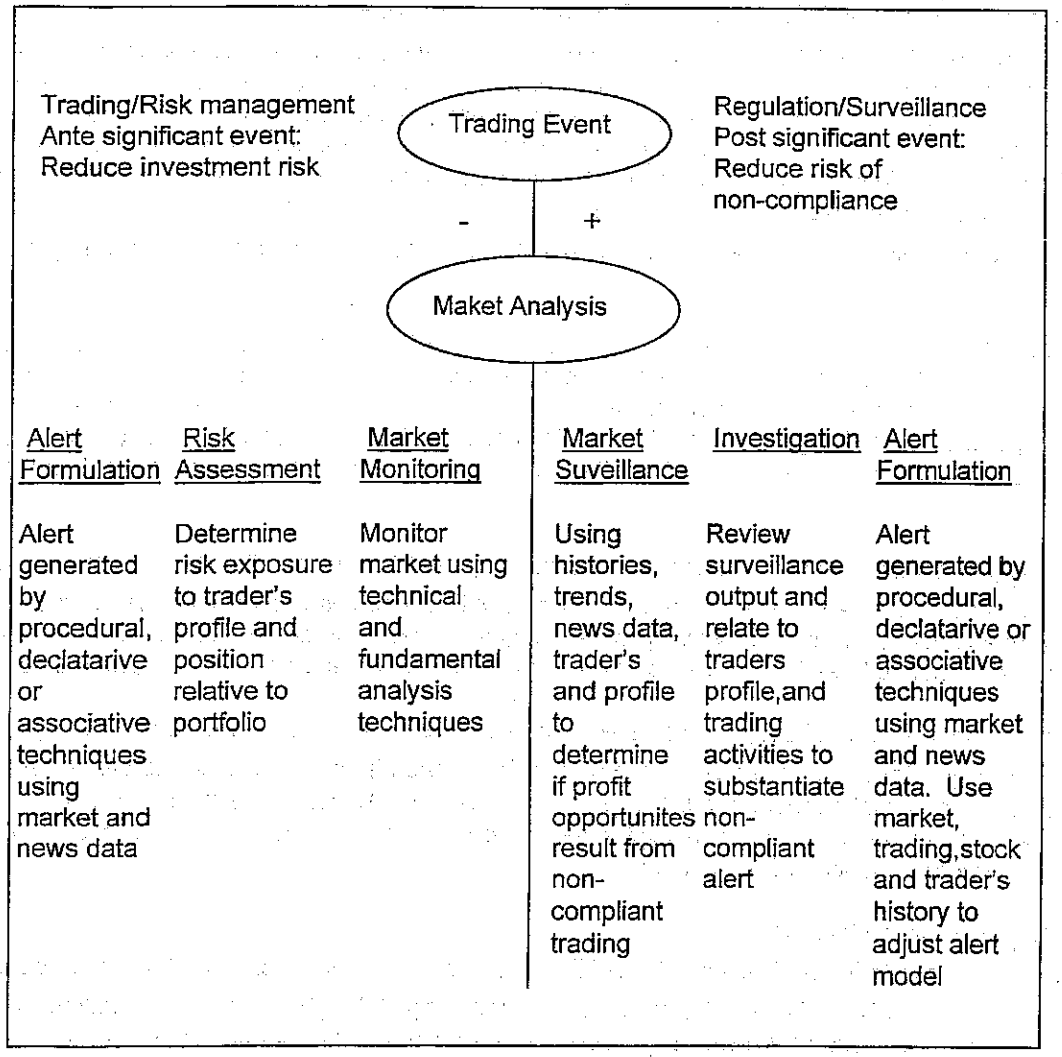

Figure 3.1 Trading versus Regulation: an Inverse Relationship [Adapted from Freedman and Mathai (1995)]

Figure 3.1 shows the inverse relationship between trading and regulation.
Dataveillance techniques for $CMAD_{CM}$ and the subsequent tasks of the A_{ei} depend on data that identifies A_Ss and their relationships if any, Freedman and Mathai (1995).

The contextual component may be determined by the timing of the event, the A_S, the possible motivation for the event, and the market environment in which the event was transacted. Intuitively an A_S's motivation would be to reduce risk. However this may not be immediately transparent if the event is one in a series of events conducted for risk reduction. Therefore the scope of the A_{ei}'s knowledge is an important factor in the coordination of the $CMAD_{CM}$ problem solving task. Coordination refers to the managing of interactions between multiple agents cooperating in some collective task. Pete et al. (1993) show that the optimal organisational design depends on the task environment and, as with an audit team or group, a typical A_{ei} organisational structure is hierarchical, with the A_{ei+1}'s evaluation of the variance depending on the event in question and A_{ei}'s evaluation.

The objective of this cooperation is to reduce the problems discussed by Freedman (1991), to reduce any potentially redundant activities conducted by the A_{ei}s, and to render the $CMAD_{cm}$ process more efficient and effective.

With $CMAD_{cm}$, in order to ensure assumption truth maintenance [ATM], de Kleer (1986), it is expedient for the judgments of less experienced A_{ei}s to be reviewed by more senior agents.

The process of review when evaluating judgments made on accounting data and information is well established in the auditing literature. "The hierarchical structure of the audit team and the sequential and iterative review processes which dominate interactions among its members are well-recognised characteristics of the audit decision-making environment (see e.g. Ashton et al., 1988; Mautz and Sharaf, 1961). Their role in quality assurance is deeply ingrained in the audit standards and firm policies" Libby and Trotman (1993, p. 559).

To facilitate an efficient process, the audit task is typically reduced, by subdivision into subtasks, each of which is conducted by individual auditors who coordinate their findings. The auditors, organised in teams involve a process in which the auditors are hierarchically organised and make interdependent judgments, usually in a sequential and iterative manner, resulting in a decision that has been reviewed by more senior auditors. An alternate audit organisation is the audit group. The group is differentiated from the team in that a group is non-hierarchical and decision making is done collectively in a simultaneous rather than strictly sequential mode, Chang et al. (1993). To facilitate this coordination, it is necessary for the A_{ei}s to communicate their findings via a communication protocol.

The communication protocol establishes the means and modes of the information communication between agents. This information exchange can be either via an implicit communication mechanism such as a common memory or blackboard, or via explicit communication mechanisms, such as message sending. The former is simpler in that the agents do not communicate directly with one another, but their communications are posted on the blackboard for use by others.

The blackboard approach allows for the posting of the SNCE's details plus the $CMAD_{cm}$A's assumptions and results. This is analogous to an electronic document which is communicated to the different agent levels and facilitates the more senior agents imposing their criteria on lesser agents' results, as well as using their task specific criteria to further refine the classifications. More specific details of this communication protocol follow.

The review process for ATM essentially involves the communication of the $CMAD_{cm}$A's belief structure, supporting both the intermediate and final beliefs of a SNCE, from the $CMAD_{cm}$A at node n_i to each n_{i+k} in the hierarchy. The communication is facilitated using a communication protocol and the blackboard. The protocol is governed by a set of rules contained in the blackboard which also records the SNCE generated by the primary monitoring system. These rules post the belief structures from the n_is to the blackboard or from the blackboard to the n_{i+k}s.

4.0 Conclusion

$CMAD_{cm}$ is presented as an example of CMAD operating in a highly complex environment. However other domains, where the threshold granularity is high and the decision making time factor is short, may benefit from the decision support discussed. It is essential for accountability that these organisations ensure transactions identified as SNCE are scrutinized and substantiated. This will assist in minimising false positive conclusions that may result because of the speed, volume and increased complexity of the transactions, and the information used to analyse them.

References

Aitken, M. "An Evaluation of ASX's Surveillance Of Market Activity Program". Unpublished manuscript, Australian Stock Exchange Surveillance Division, 1991.

Ashton, R. A., Kleinmuntz, D.N., Sullivan, J.B., and Tomassini, L.A. "Audit Decision Making". In (Abel-khalik, A. and I. Sullivan eds.): Research Opportunities in Auditing: The Second Decade, American Accounting Association, Saratosa FL, 1988, pp. 95-132.

Berry, J. and Yanco, G. "Enter the ASX Computer Police". Journal of the Securities Institute of Australia, Volume 1, March 1990, pp. 2-5.

Blocher, E. and Cooper, J.. "A Study of Auditors' Analytical Review Performance". Auditing: A Journal of Practice and Theory, Volume 7 (2), 1988, pp. 1-28.
Brown, P. and Goldschmidt, P. "ALCOD IDSS: Assisting The Australian Stock Market Surveillance Team's Review Process". Applied Artificail Intelligence: An International Journal, Special Issue, Volume 10(6), 1996, pp. 625-641
Buta, P. and Barletta, R. "Case-Based Reasoning for Market Surveillance". Unpublished manuscript, Cognitive Systems, Inc., Boston, MA, 1991.
Byrnes, E.; Thomas, C.; Henry, N; and Waldman, S. (1990). "INSPECTOR An Expert System for Monitoring World-wide Trading Activities in Foreign Exchange". AI Review, Volume 3, 1090, pp. 9-13.
Chang, A.; Bailey Jr., A.; and Whinston, A. "Multi-Auditor Cooperation: A Model of Distributed Reasoning". IEEE Transactions on Engineering Management, Volume 20(4), 1993, pp. 346-59.
Clarke, R. "Information Technology and Dataveillance". Communications of the ACM, Volume 31(5), 1988, pp. 498-512.
de Kleer, J. "An Assumption Based Truth Maintenance System". Artificial Intelligence, Volume 28(2), pp. 127-62.
Davis, S. and Ord, K. "Improving and Measuring the Performance of a Security Industry Surveillance System". INTERFACES, Volume 20(5), pp. 31-42.
Freedman, R. S. "AI on Wall Street". IEEE Expert, Volume 6(2), 1991, pp. 2-7.
Freedman, R. S. and Mathai, J. "Market Analysis for Risk Management and Regulation: An Artificial Intelligence Approach". In R. Freedman, R. Klein and J. Lederman, (eds.) Artificial Intelligence in the Capital Market. Probus Publishing, Chicago, Illinois, 1995, pp. 315-26.
Freedman, R. S. and Stuzin, G. J. "Knowledge-Based Methodology for Tuning Analytical Models". IEEE Transactions on Systems, Man, and Cybernetics, Volume 21(2), 1991, pp. 347-58.
Garner, B. and Chen, F. "Case-Based Interaction for Fraud Detection in EDP". Proceedings of International Conference on Information Processing & Systems, October: Beijing, China, 1992.
Goldschmidt, P. "ALCOD: An IDSS for Stock Market Surveillance". Proceedings of The Third International Conference on Artificial Intelligence Applications on Wall Street, New York, June: 1995, pp. 24-35.
Goldschmidt, P. "Compliianced Monitoring for Anomaly Detection in a Complex Environment using Multiple Agents: Supporting Capital Market Surveillance". Ph.D. dissertation, The University of Western Australia, Australia, 1996.
Hayes-Roth, "The Blackboard Architecture: A General Framework for Problem Solving". Technical Report HPP-83-30, Stanford University, Stanford, CA, 1983.
Koonce, L. "A Cognitive Characterisation of Audit Analytical Review". Auditing: A Journal of Practice and Theory, Volume 12(Supplement), 1993, pp. 57-76.
Lecot, K. "Using Expert Systems in Banking: the Case of Fraud Detection and Prevention". Expert Systems Review for Business and Accounting, Volume 1(3), 1988, pp. 17-20.
Libby, R. "Availability and the Generation of Hypotheses in Analytical Review". Journal of Accounting Research, Volume 23(2), 1985, pp. 648-67.
Libby, R. and Trotman, K. "The Review Process as a Control for Differential Recall of Evidence in Auditor Judgments". Accounting, Organisations and Society, Volume 18(6), 1993, pp.559-74.
Major, J. and Riedinger, D. "EFD: A Hybrid Knowledge / Statistical-Based System for the Detection of Fraud". International Journal of Intelligent Systems, Volume 7(7), 1993, pp. 687-703.
Mautz, R. K. and Sharaf, H. A. The Philosophy of Auditing. American Accounting Association, Sarasota Fl., 1961.
O'Leary, D. "Artificial Intelligence and Expert Systems in Accounting Databases: Survey and Extensions". Expert Systems with Applications, Volume3(1) 1991, pp. 143-52.
O'Leary, D. "Case-Based Reasoning and Multi-agent Systems for Accounting Regulation Systems with Extensions". Journal of Intelligent Systems in Accounting, Finance and Management, Volume 1(1), 1992a, pp. 41-52.
O'Leary, D. "Intrusion Detection Systems". Journal of Information Systems, Volume 6(1) 1992b. pp. 63-73.
Pessin, A. H. Securities Law Compliance: A Guide for Brokers, Dealers, and Investors. Dow Jones-Irwin, NY, New York, 1990.
Pete, A.; Kleinman, D.; and Pattipati, K. "Tasks and Organisational: Signal Detection model of Organisational Decision Making". Intelligent Systems in Accounting, Finance and Management, Volume 2(4), 1993. pp 289-303.
Pindyck, R. S. and Rubinfeld, D. L. Econometric Models and Economic Forecasts. McGraw-Hill, New York, 1976.
Senator, T.; Goldberg, H.; Wooten, J.; Cottini, M.; Khan, A.; Klinger, C.;W. Llamas, C.; M. Marrone, M.; and Wongs, R. "Financial Crimes Enforcement Network AI System (FAIS)". AI Magazine, Volume16(4), 1995, pp. 21- 39.
Simon, H. "The Structure of Ill Structured Problems". Artificial Intelligence, Volume 4(3-4), 1973, pp.181-201.
Slade, S. "Case-Based Reasoning for Financial Decision Making". Proceedings of the First International Conference on Artificial Intelligence Applications on Wall Street. IEEE Computer Society Press, 1991, 232-7.
Sol, H. G. "Simulation in Information Systems Development", Ph.D. dissertation, University of Groningen, The Netherlands, 1982.
Trotman, K. T. and Yetton, P. W. (1985). "The Effect of the Review Process on Audit Judgement". Journal of Accounting Research, Volume 23(1) 1985, pp 256-67.
Trotman, K. T. "The Review Process and the Accuracy of Auditor Judgements". Journal of Accounting Research, Volume 23(2), 1985, pp. 740-52.
Vasarhelyi, M.A. and Halper, F .B. (1991). "The Continuous Audit of Online Systems". Auditing: A Journal of Theory and Practice. Volume 10(1) 1991, pp. 110 -25.
Weber, R. EDP Auditing: Conceptual Foundations and Practice. McGraw-Hill, Tokyo, 1988.
Winston, P. Artificial Intelligence. Addison-Wesley, Reading, MA, 1977.