

Intelligent detection of money laundering and other financial crimes

by Laurence Jacobs¹ and Ralph Wyss²

Introduction

One often hears that identifying a money-laundering event is like trying to find a needle in a haystack. This is not a good simile. For one thing, needles are all rather similar to each other; for another, needles don't look like hay at all. This is in stark contrast to what we know about money laundering and other financial crimes: the most serious instances of fraudulent financial activity are often characterized by the fact that they are like nothing seen before, and sophisticated criminals hide their activities by making them appear as normal as possible.

A somewhat better analogy, though still far from perfect, is that detecting a money-laundering event is like discovering a black hole in the far reaches of the universe. Black holes are very exotic objects predicted to exist by Einstein's cosmological equations. Among other bizarre properties, these objects are characterized by the fact that they are not directly visible. Indeed, any light that shines on a black hole is absorbed by it. How then are these objects *seen*? Black holes (and other exotic cosmic objects) are detected, indirectly, through the effects that they have on their surroundings. These effects, however, are very subtle, and are generally extremely difficult to measure. The conclusive discovery and detailed study of these strange cosmic objects has been possible only through the development of very clever and complex methods of data analysis.

In the past, it was possible for financial institutions to decide whether they would simply comply with existing regulations, often not very strict or carefully thought-out, or whether they should make the considerably more difficult and costly effort of seriously attempting to arrest financial crime. As regulators have become increasingly more sophisticated, and as the global risks and damage created through financial crime have become more widely appreciated, these two positions are gradually becoming indistinguishable from each other. Current regulations in Switzerland, the European Union, the United States, and elsewhere, are designed to impose the strongest possible obstacle to financial crime.

To a greater or lesser extent, all new regulations mandate that financial institutions investigate all *suspicious* activities of their customers and employees. This position is not completely new; what is new is exemplified by the ordinance published recently by the Swiss Federal Banking Commission. Unlike previous regulations, the new ordinance does not limit the meaning of suspiciousness to a fixed set of rules. This position follows from the understanding that a resourceful criminal can circumvent any finite set of published rules. The trick is, of course, to find practical ways in which suspiciousness can be measured in general.

What, exactly, is possible to combat financial crime? As it turns out, and as we discuss in this article, a great deal is possible. If the will is there, modern conceptual and technological means and, indeed, commercially available products, exist that can fundamentally change the world of financial crime.

¹ Dr. Laurence Jacobs is the Chief Technical Officer of kdlabs AG in Zurich, Switzerland. kdlabs provides products and services in the area of Knowledge Discovery and Application.

² Dr. Ralph Wyss, Attorney at Law in Zurich, Switzerland, advises financial institutions doing business in Switzerland in legal and regulatory matters.

Combating financial crime

Just as with biological disease, financial crime can be either prevented, by acting before it happens, or diagnosed once it has occurred. There are some similarities, as well as many differences, between these two actions.

Prevention is most often understood in terms of fixed compliance measures, either procedural, or on the basis of pre-established, quantifiable risks. An example of the former is the account-opening process in a bank, where certain steps, such as a formal verification of a new customer's identity, are mandated; an example of the latter is the assignment of risk measures to customer or transaction parameters, such as customer country of domicile, origin of a deposit, or other transaction-related attributes. What is common among the various prevention measures is that they generally apply to either static characteristics of an event, or to a single transaction.

Diagnosis generally addresses more complex aspects of financial activities, such as the relationships between a set of transactions associated with an account or customer, or activity-based relationships between different accounts and/or different customers. The more complex potential risk signals involved in diagnosis are generally referred to as *patterns*.

Some risks can be mitigated by a judicious choice of prevention rules, but the majority of cases of financial wrongdoing can only be detected after they have occurred.

A priori risks and risk groups

From the foregoing discussion, it is clear that the effectiveness of prevention and, to some extent, diagnosis, depends critically on a correct assignment of certain a priori risks. Regulatory and other agencies periodically review and update several risk measures associated with specific individuals or countries, and make these data publicly available.

Other mandated prevention measures are essentially the formalization of *best-practice* guidelines, such as the specification of the account-opening process at a bank.

These simpler, well-defined risks are amenable to real-time intervention. The implementation of measures to accomplish these real-time actions is generally of low technical complexity.

Other risks, however, are far too complex to formalize as fixed compliance rules. In addition, these risks are generally not static, but rather change over time. Most importantly, though, most of these complex risks are sensitive to the context in which they occur. Real-time actions based on these more complex risk measures are generally not possible or even sensible. In many cases, the technical complexity associated with these intricate risks can be very high.

Rule-based detection

There are strengths and weaknesses associated with fixed rules. Among the strengths are that fixed rules are technically simple to implement, can be very fast, and are generally easy to understand. The main weakness here is that only known risks can be implemented through fixed rules.

Apart from the basic compliance rules described above, fixed rules, in a sense summarizing past experience of regulators, are practically useless to detect criminal activity perpetrated by sophisticated criminals, but, then again, not all criminals are sophisticated.

In addition, many methods of money laundering are known, well understood, and some of these can be expressed as rules with a varying degree of complexity.

Erroneous detection

Any signal detection, however intelligently done, is prone to error. This basic, unavoidable fact must be understood, and safeguards against it must be implemented.

The effects of misdetection fall into two broad categories. The first, and most dangerous one, is associated with false-negative alerts, that is, with criminal activity that remains undetected. Fraudulent events that are subtle and difficult to detect can be the most damaging, since, when they surface, as many eventually do, the damage to the victim institution and principals is bound to be large. False-positives, on the other hand, affect mainly customer satisfaction, institution prestige, and, potentially, customer loyalty; they also generate great costs to the institution in that unnecessary human effort must be expended into the investigation of irrelevant events.

Understanding the fact that, by definition, no detection method can be perfect, makes real-time detection of activity-based signals impractical, at best, and extremely dangerous, at worst. The most powerful detector will always require human investigation and verification. This, of course, is not true of fixed compliance rules, such as customer identity certification, or transactions originating in blacklisted countries, or from blacklisted individuals, which can be detected in a fully automatic fashion.

Are rules enough?

As we have argued, straightforward ad-hoc rules have a central role to play in any detector of criminal financial activity. However, ad-hoc rules without a context are often foolish at best. Perhaps the most obvious example is a rule that triggers an alert whenever a deposit into an account exceeds a fixed amount. In general, the vast majority of instances of the application of such a rule will be either false-positives, or false-negatives. On the other hand, the very same rule, but with a dynamical threshold reflecting a *maximum expected deviation* measure associated with, for example, an individual customer, or a customer group, may make perfect sense.

Thus, when discussing whether a rule-based detector makes sense, one should carefully distinguish between fixed, context-independent rules, and the more general interpretation of a rule as discussed in the example of the preceding paragraph.

Any detector, however complex and powerful, can be implemented through a finite set of rules, but these rules may, in turn, act on predicates that are themselves complex, and include, in particular, dynamical contextual information associated with the event under investigation.

The predicates

Where do the complex predicates referred to above come from, and how can they be evaluated? This is an area of intense current research in which a great deal has been discovered in recent years. These predicates generally refer to sophisticated, and potentially very complex statistical models whose output, when acting on a given event, is the value of the predicate. In simple terms, given a model that measures some type of risk, and an event, such as a transaction made by a particular customer, the rule that generates alerts based on the application of the model to the event might look as follows,

```
IF model.risk[event] > dynamical.threshold[customer]  
THEN trigger.alert
```

Where `trigger.alert` is some process that initiates an investigation, either by another model, or by a human, into the suspect customer/event pair. The complexity of a rule such as the example above is thus not in the rule itself, but rather in the predicates (`model.risk` and `dynamical.threshold` in the example) and objects (`event` and `customer` in the example) that define the rule.

In the context of fraud detection, the most powerful signal detection and identification systems known today are based on one of two possible methods (or a combination of the two), known as *supervised learning* and *unsupervised learning*, which we will define in a moment. These two methods form the basis of what is popularly known as *data mining* (*inductive data mining*, to be a bit more precise).

These analytical methods rely on a wealth of algorithms and technologies developed in the last two or so decades which form the core of the field of *knowledge discovery*. The fundamental principle of inductive data mining is that knowledge, in the form of models, is derived directly from observation (data) through a process of learning, verification and generalization. The knowledge discovered directly from historical data can be used, for example, to classify new data.

Supervised learning leads to models that use one or more characteristics of cases in the historical data to build *classes*. A class is a subset of the data where one or more attributes are most likely to have a particular value. Classification models could, in principle, be used to score the risk that a transaction is fraudulent. In the case of money laundering this is, however, generally not possible. The main reason for this is that supervised learning requires that a relatively large proportion of the cases in the historical record be known to be fraudulent, which is not the case for any given financial institution. A related limitation is that the accuracy with which such a classifier would score new cases cannot be high enough to be practical in the context of money laundering or other equally relatively rare types of financial crime.

The process of unsupervised learning, on the other hand, looks for similarities between events (such as the activity profiles of customers in a bank), and groups these into distinct *segments*. These segments, generally known as *clusters*, are characterized by the fact that an event in one cluster is more similar to other events in the same cluster than it is to events in any other cluster. This is ideal for the problem at hand, where the goal is to identify *suspicious* events. Suspiciousness of an event is then the extent by which the event departs from expectation.

The core of an effective detector of suspicious behavior is a procedure that detects *anomalies* in the data. By definition, an anomalous event is one that does not fit its context. What remains, then, is to define the context in which an event occurs.

As the context changes, so should the measure of suspiciousness change. This is the meaning of *adaptive learning*; an effective method of detecting and measuring the degree of suspiciousness associated with an event must be able to adjust itself to changes in the context in which the event occurs, and it must do so in an unbiased, self-consistent manner.

Context

In the general area of financial crime detection, there are essentially two contexts in which to place an event, such as a transaction, in order to assess whether or not the event is anomalous.

The first context is the historical profile of activity of the individual customer (or account) associated with the event. In simple terms, normalcy in this context is measured by the degree to which the event under consideration is *expected* as compared with the rest of the events that define the activity profile of the customer in question. We refer to the analytical methods in this context as *self-history analysis*.

The second main context is the historical profile of activity of the *peer group* to which a customer or account belongs to. We shall discuss how to define and build peer groups in a moment. For now, it is sufficient to think of a peer group as sets of customers or accounts that somehow fit together. A simple example of a peer group might be the set of cinemas in the Zurich area. Normalcy in this context refers to the degree to which the event under consideration fits the expected patterns associated with the peer group. We refer to the analytical methods in this context as *peer-group analysis*.

In general, an effective anomaly detector must use both of the contexts we have described here. It is easy to imagine situations where truly suspicious activity may only appear to be so in one of these two complementary contexts.

It should be clear from the discussion above that contexts are dynamical; detectors that use contexts as we have defined them adapt to changes in customer or peer group behavior. This type of *adaptive learning* is crucial for effective detection of money laundering.

Peer Groups

How are peer groups defined and built? The example we mentioned above, that of a particular type of banking customer in a particular geography is an instance of one of two general kinds of peer group, namely a segmentation that is defined deductively using the values of one or more attributes (two in this case) associated with customers. The other general kind of peer group is directly induced from the data, and is thus called *inductive segmentation*. Inductive segmentation, which can, in principle, use all the attributes associated with customers, including their activity profiles is an example of unsupervised learning.

Most commonly, the highest quality segmentation used in peer-group analysis is built using a hybrid of inductive and deductive methods.

Profiles

Our use of the term *profile* in the previous paragraphs requires explanation.

In the case of a customer or account, a profile is a set of data that includes some of the static attributes associated with the customer, such as country of domicile, together with a set of metrics associated with the customer's activity over a given period of time. The process of determining such metrics from the raw transaction data involving the customer in question requires what is called *time-series abstraction* (or *summarization*). The series of transactions itself is seldom directly useful in this type of analysis. A very simple example of a summarization metric could be the mean value of incoming transactions associated with the customer. More complex metrics that are useful in this context may include things such as the spectral characteristics of the transaction time series.

Peer group profiles can be defined in a similar manner, using, for example, the most probable values of static attributes, or average values associated with the summarized time series of the dynamical attributes defining the peer group.

Profiles in the sense used here are dynamical structures that generally change over time. Thus, defining and using profiles in this sense are also instances of adaptive learning.

KYC

Most of the current views on compliance in the context of money laundering strongly emphasize, indeed, mandate, the principle of *know your customer*. How does one, in practical terms, do this? KYC is a short acronym that summarizes our discussion of the last several paragraphs. Our terms *profile*, *self-history*, and *peer groups* are non other than KYC in detail.

This brings us to an often-overlooked point. Financial institutions mostly view compliance as a necessary, but generally unattractive task, an often expensive, unprofitable, and time-consuming burden on their normal business operations. If KYC is understood correctly, however, its powers as a business *driver* become obvious. When looked at in detail, an unexpected change in a customer's profile, or a switch of peer groups, is not always a signal of malfeasance; it can also lead to a business opportunity for the financial institution. Examples abound. A simple one is the discovery of a customer's migration to a peer group that has a higher profitability measure, associated, for example, with more a profitable investment profile. Knowledge of this development allows the financial institution to react to the change in a way that profits both the customer and the institution.

In its pure form, KYC is the basis of customer relationship management (also in its pure form).

Link Analysis

Link analysis is the investigation of relationships between separate entities involved in financial transactions. Several known methods of money laundering are based on the use of networks of accounts that are not otherwise known to be legally related.

The analysis of networks of activity is often considered as a separate area of detection. This is not altogether reasonable.

From an analytical perspective, some of the mathematical methods traditionally used in link analysis are quite different from most other methods used in the detection of money laundering activities. This fact has somehow set this line of analysis apart from all others. However, from a conceptual perspective, separating link analysis from self-history and peer-group analyses leads to conceptual and practical difficulties and, eventually, to a weak solution of the problem. In fact, some of the most powerful aspects of link analysis are only possible when implemented as a downstream component of a multi-stage detector that includes self-history and peer-group components as well.

Understanding an anomaly

Detecting a suspicious event is but the first step in the process of identifying criminal activity. A fantastically powerful detector that operates like a black box makes succeeding steps, necessarily involving human investigators, difficult, expensive, and prone to error.

For a detector to be truly useful, it should also expose the origins of an alert that it triggers, identifying, for example, those aspects of the event that are most suspicious. While technically difficult, this is an essential characteristic of a useful detector.

Heuristics

Many aspects of the analytical methods discussed here are computationally very complex. In some cases, as in link analysis, for example, brute-force methods are available, but impossible to implement in any but the simplest of cases.

In some cases, however, it is possible to find good approximations to a mathematically impossible problem that have, if not the full power of the exact solution, sufficient value to make them useful. It is sometimes possible, for example, to find fairly simple rules that summarize the results of very complex analysis.

Implementation

The analytical methods we have discussed in this article are not simple. Given that the volume of daily transactions in a bank can reach into the millions, the demands on computational systems implied by sophisticated detection technologies can be enormous. A practical implementation of a detector that can deal with large volumes of data while meeting the most stringent constraints on accuracy must use state of the art technology in computing hardware and associated data organization and handling software. However, given the current state of development of these technologies, systems that scale gracefully to attack even the largest problems in this field are readily available today.

Limitations

The main limitation of the analytical picture sketched in this article is in the nature, quality, and amount of data that are available for analysis. The clearest example of this limitation is in the analysis of networks. Whereas a great deal of information is available to a financial institution in investigating characteristics of networks of its own accounts, it generally knows almost nothing of nodes in the network that correspond to accounts in different financial institutions.

Investigation support

An area not often mentioned in the context of money laundering, and not at all well represented by commercial solution providers, is the process of investigation performed by the various financial investigation units after a suspicious case is reported to them by a financial institution.

With modern methods of analysis, such as those discussed in this article, while the human effort for financial institutions is bound to decrease, the number of cases reported to financial investigation units is likely to increase as a result of an increasing number of reported cases of potentially higher complexity. Dealing with an increase in the number of potentially more complex cases will require new methods of analysis as support for these investigators.

Current technology in the areas of supervised and unsupervised learning in general, including text mining, could have a profoundly positive effect on how such investigations are conducted. In particular, issues of effectiveness, consistency, and speed could be substantially improved using modern analytical methods.

Of course, this area of application of modern analytical methods suffers from some of the same limitations we have just addressed, but in many ways, the problem is simpler than the initial detection of suspicious activity. Among other things, financial investigation units have access to many more suspicious cases than any single financial institution, making supervised learning methods potentially applicable. In addition, and most importantly, such agencies have access to data from all financial institutions in their jurisdiction, and are thus not affected by some of the limitations we have mentioned.

Conclusions

It might seem that the use of the most powerful analytical and computational methods necessarily carries with it a large price tag for financial institutions. In fact, the opposite is true. There are two reasons for this assertion. First, state of the art methods of analysis, and state of the art computing platforms and other support systems are not necessarily expensive. Second, the price paid by society, and financial institutions for the failure of implementing adequate safeguards to combat financial crime is extremely high; higher, without doubt, than the most expensive anti money-laundering systems available today.

The detection and identification of financial crime is not simpler than the detection of black holes. In some ways it is indeed more difficult. A successful attack on financial crime requires the most powerful modern methods of data analysis. As we have argued in this article, however, sophisticated data analysis is necessary, but not sufficient, to detect and identify instances of financial wrongdoing.

In addition to the powerful analytical techniques we have discussed in this article, a successful attack on financial crime requires a comprehensive approach, including deep domain knowledge, powerful support and delivery technologies, and collaborative, enlightened and sophisticated regulators.