



Fighting money laundering with technology: A case study of Bank X in the UK

Dionysios S. Demetis

Centre for Systems Studies, Hull University Business School, Cottingham Rd, Yorkshire HU6 7RX, UK



ARTICLE INFO

Article history:

Received 4 December 2016

Received in revised form 29 October 2017

Accepted 21 November 2017

Available online 2 December 2017

Keywords:

Money laundering

Systems theory

Transaction monitoring systems

Profiling

ABSTRACT

This paper presents a longitudinal interpretive case study of a UK bank's efforts to combat Money Laundering (ML) by expanding the scope of its profiling of ML behaviour. The concept of structural coupling, taken from systems theory, is used to reflect on the bank's approach to theorize about the nature of ML-profiling. The paper offers a practical contribution by laying a path towards the improvement of money laundering detection in an organizational context while a set of evaluation measures is extracted from the case study. Generalizing from the case of the bank, the paper presents a systems-oriented conceptual framework for ML monitoring.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Despite the critical role of technology in tackling money laundering (ML), this socially important phenomenon has been examined sparingly in Information Systems (IS) research [1, p.115], typically from technical perspectives that focus on algorithmic optimizations alone. By and large, these approaches have failed. While algorithmic optimizations remain important and technology is central to Anti-Money Laundering (AML) efforts [2], a richer nexus of intertwined human- and technology-generated decisions supports the identification of ML-behaviour. This paper is motivated by the desire to find a way forward in improving money laundering detection (true positive rates), exploring ML-detection challenges within a real organizational context and extracting evaluation metrics for AML that can assist the key decision makers. In an attempt to uncover what elements affect ML-profiling beyond algorithmic techniques, this paper explores these challenges in the richer context of a case study. This paper focuses on banks, which remain the main avenue for ML activity [3]. For example, in the UK, banks file more than 85% of potential ML-activity to the authorities by submitting *Suspicious Activity Reports (SARs)* [4].

By presenting an in-depth longitudinal case study of a UK bank in the interpretivist tradition [5,6], this paper follows the bank's efforts to improve ML-detection. This research is based on two different but connected EU-funded projects spanning approximately 2.5 years in total (JAI/2004/AGIS/182 and SEC6-PR-205800). The case is enriched by data from other key project stakeholders (including asset recovery agencies, financial intelligence agencies, central banks, police, and other banks from the UK and abroad that were project partners).

The paper uses the concept of *structural coupling* from systems theory to depict the dynamic relationship between *computer profiling* and *human profiling* in the targeting of ML; structural coupling also helps us organize the interactions and reflect on the way the balance between information/redundancy is expressed. The paper provides a practical contribution to the improvement of the True Positive Rate (TPR) of Transaction Monitoring Systems (TMS) and extracts AML evaluation measures that can be considered by key decision makers within AML systems. The development of a systems-theoretical conceptual framework for ML detection synthesizes the theoretical contribution (Fig. 1). The structure of the paper is as follows: a) in the literature review, the substantive problem (ML) is reviewed briefly in tandem with IS-related influences; second, the core theoretical concept of structural coupling drawn from systems theory; b) the methodology; c) the description of the case study and its analysis; d) the discussion, where the findings and the concepts are used in light of the theory to outline the contributions; e) research limitations, suggestions for future research; f) conclusions.

2. Literature review

2.1. (Anti-)money laundering and information systems research

ML constitutes the process of *masking* monetary gain resulting from any type of criminal activity (e.g. drug & human trafficking, fraud, extortion, kidnap-for-ransom); the totality of such criminal activities associated with ML are labeled as *predicate offences*, but ML is also considered a crime on its own [7,8]; thus, someone can be charged with drug trafficking & money laundering in a combined prosecution (or indeed, separately if the laundering activity is facilitated by an associate). The purpose of ML is to create the appearance that money generated by

E-mail address: d.demetis@hull.ac.uk.

predicate criminal offences has a legitimate source [9]. Ultimately, the conversion into legal money assists criminals to benefit from the formal financial system, thereby gaining them more flexibility in the movement of funds that may subsequently be used to fund further crime. Despite difficulties in estimating the size of underground markets [1,p.13, 10], ML is estimated to account for between \$1.5 trillion to \$2.85 trillion annually [11,12], making it the world's third largest market [13,14].

In the IS literature, ML is significantly under-researched despite early calls by Mumford [15,p. 258] who considered it an important research problem and Demetis [1] who studied the phenomenon by looking at the AML risk-based approach and technology. Furthermore, the semantic ambiguity in how ML has been framed within IS has created some confusion in how it is approached. For example, Chung, Chenb, Changc and Chouc [16,p.671] view it in the context of cybercrime, even though money laundering rarely involves a security breach or a hardware/software compromise of financial systems. Ngai, Hu, Wong, Chen and Sun [17,p. 561] see it as part of financial *fraud*; as we shall see in the case study, that remains equally problematic from both organizational and detection perspectives. Of course, ML itself rarely involves deception that results in personal or institutional loss of a financial nature; quite the contrary: most launderers are happy to pay tax if that means legitimizing their profits and banks usually profit from ML activities. Even in the highest-recorded financial fine levied (against HSBC), the bank remained relatively unscathed as it settled for \$1.9bn (when accused of laundering money for Mexican drug cartels) and turned a net profit in the same year. While ML is related to cybercrime, fraud, and risk [18,p. 609], this paper calls for it to be treated as a stand-alone phenomenon, characterized by distinct stages; it aligns with Pramod and Gao who emphasize the importance of working on conceptual frameworks related to AML itself [19].

The limited number of IS articles dealing with ML has focused mostly on ML-detection through clustering, classification [20], sequence analysis [16,21], outlier detection [17,22], and *real-time detection* for ML-monitoring [23]. Attributing *ML-risk scores* to individual customers [24] has also been explored in the context of a risk-based approach to AML [25]. Meanwhile *network recognition* [26] is increasingly emphasized, as terrorist financing has become an integral part of the AML-regime [27]. But while there seems to be agreement that ML-detection is very challenging, scholars in this domain describe possible techniques *decontextualized* from an organizational setting. Also, they do not discuss the True Positive Rates (TPR) of ML-monitoring systems and the challenges that specific financial institutions face in improving those. For example, Drezewski, Sepielak and Filipkowski [22] describe the broader computational architecture that would be required to tackle ML while Khan [24] describes how Bayesian networks can assist detection. However, no TPRs are recorded or described in an applied banking environment. Ngai [15] does mention the importance of looking at false positives and false negatives in a business context but this is not investigated at an empirical level. Canhoto and Backhouse [31] mention how most financial institutions experience very low True Positive Rates in a general profiling context of customers (often <1%) due to: the underground nature of the problem, the complex behaviour of launderers in their efforts to obfuscate the money trail, and the sheer volume of banking transactions [28], but they do not relate or explore AML profiling strategies and options for TPR development. Just how persistent this problem of low TPRs in ML-detection has been is evident in the results of a wide survey conducted by PwC to its AML-customers where “PricewaterhouseCoopers analysis indicates that 90 percent to 95 percent of all alerts generated by AML alert engines are false positives. These high false-positive rates lead to significant monitoring costs” [29,p. 1] while further liabilities also occur from non-compliance and the actual costs can be significantly higher. In summary, previous work is deficient as TPR improvements are not reviewed against a real organizational setting so the challenges that decision makers face (and the impact) is decontextualized from technology applications in AML. More recent research reinforces the concern that serious problems lurk in ML detection, and calls for far more emphasis to be placed on personnel-training [2]. Indeed, while technology remains essential for

filtering financial transaction data for ML [30], banks have largely failed to implement AML-specific IS [18]; these have been unsuccessful, overloaded users with white noise [31] and largely failed to support the decisions of ML analysts and the key decision makers who corroborate/dismiss ML-suspicions. This has had a negative impact in both banks and regulators. Yet, despite technology's central role in ML detection, technology monitoring in AML has not been advanced theoretically or practically. This paper is motivated by the aforementioned gaps. It deconstructs how TPR improvements can be achieved, introduces AML evaluation measures and develops a systems theoretical-oriented conceptual framework for ML detection. In order to put the ensuing discussion into context, we turn to a few core AML concepts. First, ML can be considered to have three distinct stages: the *placement* stage where illegally earned money enters the financial system; the *layering* stage, where the money trail is deliberately complicated in order to avoid detection; and the *integration* stage, where money becomes part of the legitimate economy and used to fund other illegal and even legal activities. While technology supports AML-related work across all three stages (and the ensuing discussion applies to all three), the most critical stage for detecting ML is considered to be placement as it poses the highest risk for the launderers before breaking up the transactions and moving funds [32]. In order to put the role of technology into perspective, it is important first to review the lifecycle of reporting of suspicious transactions. Customers transact with financial institutions in a number of ways (ATMs, branches, e-banking, etc.). Checking for ML activity typically takes place through Transaction Monitoring Systems (TMS) that profile ML behaviour through the batch processing of 3-months of data at regular intervals (e.g. daily, weekly, monthly). The output is *alerts* (suspects on ML) and is subsequently evaluated manually by a dedicated team of ML analysts who decide whether the transaction(s) are indeed suspicious enough to be escalated internally or put on file. If suspicion is corroborated, the analysts will escalate the case to the Money Laundering Reporting Officer (MLRO). The MLRO's role is central; she is the only designated individual within the bank that is authorized to file a formal Suspicious Activity Report (SAR) to the Financial Intelligence Unit (FIU), the national aggregator of the country where the institution is based (the FIU also forwards SARs to prosecution authorities). Thus, the MLRO is a key decision maker. Admittedly, different theoretical approaches can be used to study AML in relation to IS, but systems theory is selected here for four main reasons. First, ML and AML co-evolve; systems theory (and structural coupling in particular) can enable us to conceptualize this co-dependency. Second, we follow Mumford's [15] call for using systems theory to deconstruct ML. Third, though the IS community has focused around social theorists like Giddens [33] and Latour [34], and more recently socio-materiality [35,36], systems theory is built into the very roots of information systems [37]; with it, IS research can expand its ways of theorizing. Finally, society's excessive reliance on technology [38,39] places significant demands on IS research for integrating theoretical perspectives that help us understand increasingly complex phenomena [40]. This paper aligns with the aspiration that IS research should take *systems theory* seriously [37] and the study of AML offers a fruitful ground of exploration within IS.

2.2. The concept of structural coupling in systems theory

While it is beyond the scope of this paper to delve into the history of systems concepts [41] or deconstruct the exceptions of its use in IS research (e.g. [42]), it is recognized that systems concepts have been used in work [44], user-centered design [45,46], design theory [47], information growth [48] and others [43,49]. Here we focus on the core concept of *structural coupling*. This helps us reflect on the relationship between computer and human profiling and organize the different AML structural couplings. Other systems concepts presented in this section, help us illustrate – and are in support of – the concept of structural coupling.

All branches of systems theory agree that we can distinguish between a *system* and its *environment* [50]. A *boundary* lies between system/environment, and is part of both. The connection and the co-evolution

between a system and its environment, is captured in the concept of *structural coupling*, introduced by Maturana and Varela [51]. Any structural coupling is a highly selective set-up; not everything that exists in the environment will be structurally coupled with the system [52]. What is selected by a system, specifies *what* can be channeled into it. This is labeled a “*canalisation of causalities*” [50,p.85]: a necessary selection of limited pathways that structure how system/environment interact. Through such pathways, the system becomes sensitive to specific irritations of the environment while it ignores others [51]. For example, “the brain is structurally coupled with the external environment via eye and ear. It thus possesses a very narrow bandwidth of sensibilities that produces what can be seen, limits the spectrum of colours, and equally reduces what can be heard. It is only because things are this way that the system is not overburdened with external influences” [50,p. 86]. The irritations received by the system are interpreted as information and prompt it to adjust [50], leading it to either: a) structural *development*, where irritations from the environment will be exploited to the system’s benefit, b) structural *drift*, where such irritations undermine the system’s ability to control itself.

An information *system* tackles environmental complexity by transmuting data to information. Data are channeled (from environment to system), internalized and transformed into “*information (that) is a product of the system itself and not something which exists out there*” [53,p.258]. An illustrative example is the use of search engines (e.g. Google). After the user’s search query, the – typically – tens of millions of results become irrelevant. In order for the user to transform the data into information, he/she is forced to engage in *human profiling*: a selection of will be internalized (i.e. read by the user) and what will be ignored. Thus, the top 30 results from a search get over 90% of traffic [54]; everything else is dismissed. Eventually, in information systems, structural coupling can be perceived to expand and reframe the notion of computer/human interaction into another expression: the distinction (and structural coupling) between *information* and *redundancy* [55]. Redundancy is equally (if not more) important in eliminating the demand for (and interest in) further information; it protects the system from being overburdened with external influences (similar to our sensory example). In this regard, information systems are also redundancy systems. Structural coupling then plays a pivotal role in conceptualizing how an (information) system interacts with its environment, and sets up its own structures in order to exploit the (data) resources of the environment. In summary: structural coupling between a system (e.g. AML department) and its environment implies a ‘canalisation of causalities’ (from environment to system). Environmental stimuli are internalized by the system as data and the system’s internal structures convert them to (an output of) information. The concept of *structural coupling* illustrates the sensitivity (of the system) to the environment in the process of absorbing environmental data through specified canals, as well as optimizing these selections based on whether they result in information or redundancy. This also applies to subsystems of the system; for example, if we consider the AML department of a bank as a system, then its environment will include other departments within the bank, as well as the external – to the bank – environment.

3. Research methodology

A combination of data collection methods was used during the 2.5 years of the EU-funded projects on which this research is based (JAI/2004/AGIS/182 and GATE SEC6-PR-205800). This includes individual interviews, project meetings (typically a full-day event every two months), interviews during on-site visits to Bank X (as well as other institutions), an experiment with anonymized raw transaction data, as well as feedback from dissemination seminars from high-profile events in London, Rome, Athens, Dublin, and Barcelona. While there are research design challenges in using a mixed-method approach, in the context of a domain like AML where critical information is held by different stakeholders, a mixed-method approach provides a richer picture [56]. Due to both the criticality of the domain and confidentiality agreements, at no point was the

researcher given access to any personally identifiable information from Bank X or from other sources. Also, as the project involved other researchers, the author did not include on-site visits to all project partners, but focused mostly in Bank X (which is the main case study presented in this paper). The concept of structural coupling guided theoretical development and data collection by exploring different types of system (i.e. Bank X)/environment interactions (different forms of *structural coupling* are expressed in Table 2). Structural coupling also helped reframe interactions (e.g. between computer/human profiling) as a structural coupling that is expressed by the distinction between *information/redundancy*. This has implications for information systems as information is seen as structurally coupled with, and depended on, redundancy; this reframing is of particular interest in domains like AML where such structural couplings are shaping the key constructs of the domains in which they are applied (e.g. shaping who will be considered as a suspect and who will be forwarded for prosecution). Details on the data used for the analysis and the profiles of interviewees of Bank X are presented in the table below (Table 1). These are accompanied by important ethical considerations and confidentiality issues.

All interviews were conducted in a semi-structured format, allowing the flexibility to explore any interesting emergent lines of enquiry. The objective of the interviews was to explore the approach of Bank X to its difficulties with ML-monitoring. In analyzing mostly qualitative data, an unavoidable process of categorization takes place. In the context of this research, and as scholars have analyzed [57–61], this resembles the following 3-stage process: A) open coding (of the broader research domain): AML and the role of technology in identifying suspicious behaviour; B) axial coding (themes selected): problems of Bank X in the use of transaction monitoring systems, identification of different data sources for integration into ML-profiling (presented in Table 2); and C) selective coding (last pass of incorporating further data elements): developing profiling strategies for the improved detection of money laundering.

4. Case study of Bank X in the UK

4.1. Background to Bank X and its use of transaction monitoring (relating to phases 1 & 2)

In order to avoid confusion in terminology, we will treat the AML department of Bank X as our *system* (in the systems theoretical sense); any software installations (e.g. *Transaction Monitoring System*) will be labeled by their abbreviations (e.g. TMS) or referred to as software.

Bank X is an important financial institution in the UK, focusing mostly on retail banking with branches throughout the country. The bank handles almost half a billion transactions in a quarter; due to the volume, the ML-monitoring of any form of transacting (e.g. e-banking, use of credit/debit/pre-paid cards, ATM, face-to-face) is down to filtering/monitoring data by automated means. Transactions are profiled for spotting potential ML behaviour but then, due to regulatory requirements and the conduct of due diligence, ML-analysts have to scrutinize the technology-generated reports and decide whether suspicion can be corroborated or dismissed (a number of IS support ML-analysts in this process but the TMS is the most important). For conducting such manual checks, the AML department employs several analysts (~10) that review suspicious ML-cases on a daily basis. The ML-analysts evaluate the overall financial position and transacting behaviour of technology-flagged suspected customers.

Bank X had implemented an off-the-shelf TMS that would: a) deal with the volume of raw transaction data (past 3 months of transactions), b) model ML-behaviour in a set of SQL rules, c) apply the profiles onto transactions in a batch-processing mode every night so that, d) flagged ML-suspects would be investigated further by ML-analysts the following day. Unfortunately, according to a ML analyst, when the bank implemented the TMS, it entered a “*prolonged period of havoc*” with a severe negative impact to the MLRO as the key decision maker (the MLRO was experiencing an extreme imbalance between information/redundancy due to the

Table 1
Types of data analyzed, ethical and confidential considerations.

Partner Institution	Data types used/collected by the researcher	Ethical considerations and Confidentiality issues (these prohibit disclosing actual names of individuals and masking institutions' names where possible)	
Financial Regulator	Personal notes during project meetings	Not applicable as any regulatory representation was fundamentally disclosing information that was already publicly available.	
Law Enforcement	Personal notes during project meetings	Access to sensitive information, which at the time could not have been disclosed, involved: techniques for terrorist financing and the cost of attacks in the London Underground on 7/7.	
Bank X (focus of research in how it becomes structurally coupled with other entities in the environment for AML improvement)	Personal notes, e-mails, profiling schemes, raw transaction data for experiment/visualization, models for monitoring ML, semi-structured interviews.	1) Names of customers could not have been disclosed and full anonymization was necessary in the context of the bank releasing 250 million financial transactions (corresponding to three months of data). 2) While the general model and principles of modeling ML-behaviour were discussed and created the basis for a number of developments/considerations, the actual thresholds being applied could not have been disclosed for all parameters.	
	Profiles of interviewees		
	Title		Job Description
	Money Laundering Reporting Officer (MLRO)		Manages the AML department, Files SARs to the FIU and reports to the board of directors. The MLRO is the key decision maker in the bank in the domain of AML.
	AML Manager		Responsible for SQL queries in the transaction monitoring system; manages a small team of ML-analysts who go through the submitted reports and check them for due diligence.
	Deputy AML Manager	Assists the AML Manager in his duties and works mostly on the scorecard system of risk	
	ML-analyst	Conducts due diligence and checks transaction history/behavior of suspected customers	
Banks Y,Z	Personal notes during project meetings, profiling models, TPR measurements in AML	These Bank cases are not used in the context of this paper.	
Financial Intelligence Unit (FIU)	Personal notes during project meetings	High-level issues with Financial Intelligence Units (FIUs), guidance on potential monitoring techniques were discussed frequently. ML-monitoring application consequences at FIU-level were the main focus.	
Central Bank	Personal notes during project meetings, e-mail communications, profiling structures at national level, ML-modeling frameworks	In one case, the central bank acted as the FIU as well with a unique scope of pro-active analysis and modeling of ML in the national context. Sensitive techniques that would combine different types of behaviours for ML and that have led to the identification of suspects have been redacted.	

skyrocketing number of internal ML-related reports from the technology). Having experienced a reasonably straightforward implementation of credit-card anti-fraud software, the bank was surprised by the extremely low True Positive Rate (TPR). The TPR was calculated by the bank as the following percentage: $TPR_{TMS} = [(CR_{ML-Analysts}/A_{TMS}) * 100]\%$ where A_{TMS} was the total number of Alerts generated by the TMS and scrutinized by staff whereas $CR_{ML-Analysts}$ was the number of Corroborated Reports after ML-analysts conducted a thorough analysis those alerts. So, if for a given month, the ML-analysts had work-capacity to scrutinize 100 alerts from the TMS and only 5 would be labeled as 'truly suspect' then the TPR_{TMS} would be 5%. Both lack (at first) and then scarcity of feedback from the FIU, meant that the bank relied on this internal metric for gauging its efforts in ML detection. A different metric TPR_{staff} was taken for ML-suspicions forwarded to the AML team by staff members from the branch network, substituting A_{TMS} for A_{staff} . The latter fluctuated from about 40% to 65% for well trained staff – performing a lot better than technology as subtle face-to-face interactions often gave rise to suspicious behaviour (a window of detection that is *en route* to becoming marginalized as e-banking becomes dominant). Also, the final decision on whether a 'truly suspect' SAR from the TMS would be sent to the FIU is taken by the MLRO, the key decision maker in the bank.

Evidently, in the context of the TPR_{TMS} , the percentage was influenced by what queries were used by the bank. The AML-software vendor had preconfigured around 100 SQL queries (system rules), from which the bank initially activated 7–8 (those they deemed to be capturing important ML techniques). Automated alerts from the TMS would be listed out in a queue each morning and indicate the particular system rule which has been hit, together with a finite amount of static data from the main legacy system that held the details of the product and customer. These alerts would then be allocated to an analyst and reviewed/worked in the same way as a SAR from a member of staff. The choice of 7–8 rules was the first important reduction in complexity: a highly selective conditioning through which specified ML behavioural indicators aligned the system with part of its environment (money launderers). More specifically, a very limited subset of ML-behaviour was framed within the 7–8 queries

alone; this selection channeled potential ML-behaviour and enabled the system to react (support the determination of suspicion and forwarding to the FIU for further analysis/prosecution). In effect, the bank, or the AML department to be more specific, reconfigured its structural coupling with its environment (the money launderers) by re-arranging the mechanism through which the distinction between information/redundancy would be instantiated.

Although the 7–8 queries were selected with caution by the bank in the hope of uncovering truly suspicious cases, the TPR_{TMS} was $< 0.1\%$. ML-analysts would check 1000 suspicious alerts manually before finding one worthy of escalation. After two years of trial-and-error in query-adjustment, the bank raised the TPR_{TMS} to 1.42%, while the cost to the bank for checking false positives alone was around £450,000/year (with other costs for non-compliance/fines creating additional risk for the bank). The very low TMS success rate had a serious impact in key stakeholders. ML-analysts were demotivated and demoralized from having to check very large numbers of false positives and their decisions were influenced by their TMS-use (they were more inclined to dismiss a technology-based case). Also, MLROs were adversely affected by the volume of suspicions but had to defend the bank's use of AML technology in regulatory audits. Based on several interviews, the difficulties were attributed to several factors: a) with ML connected to serious criminal activities (e.g. drug trafficking), launderers exercise caution when placing money into the financial system; b) banks operate as 'silos', and software typically monitors ML-behaviour based on *raw transaction data*; while convenient for operational reasons, ML behaviour is far more complex; c) databases degenerate and profiling is based on data that is incomplete, unsound, polluted by mistaken manual entries, etc.; d) there is minimal feedback from the FIU to the bank.

To put it in the words of another MLRO, "...we pass on almost every case that is generated by the software - let the authorities worry about it; besides, how are we supposed to know whether in a national context the case does not make sense? We don't have access to the transacting data of other banks, or tax information from HMRC (Her Majesty's Revenue & Customs)" That approach that led to a staggering increase of SARs in the UK (up by

almost 150,000 SARs from 20,000 SARs within a decade). Technology-supported decisions by MLROs created massive ambiguity in AML with information becoming intertwined with *redundancy (and noise)*. Bank X sought to refine its approach as the very low TPR_{TMS} was perceived as structural drift; it also raised management concerns of how much money can actually be spent on AML compliance. The bank tried to find ways through which the *computer profiling* of the large volume of transacting data would assist the subsequent *human profiling* (by ML-analysts) achieve better TPR_{TMS} results. The sections that follow describe these efforts towards structural development in brief.

4.2. Visualizing transactions and 'ML transactions' (relating to phase 3 – Step 1)

One technique that was explored was the dynamic visual exploration of 250 million *anonymised* financial transactions (with Datadesk software – Appendix 1). While visual exploration became difficult as the number of categories increased (not computationally impossible, just user unfriendly), some promising reflections were gained through it. Instead of visualizing all transactions, Bank X isolated from its AML Case Management System, only those accounts that corresponded to *already-flagged and suspicious customers* (but only where the suspicion originated by staff members that had a high success rate in detection; this set out to exploit a structural coupling between A_{TMS} and A_{STAFF}). By applying that *recursively* to transaction data, it extracted and then visualized the corresponding transaction data of (potential) money launderers; this led to some surprising results: while the AML department was expecting that laundering would be spread out across a spectrum of categories, the visualization indicated otherwise and led to the decision to restructure the information/redundancy dynamics of the AML department.

For example, out of a total of 99 transaction categories recorded in the transacting databases (e.g. deposit of cash would be one category, ATM withdrawal another, ATM deposit another, and so on across all products), it was only 14 transaction categories that were eventually identified as relevant to ML for the bank (this was based on *ten years of recorded suspicious cases*). From the 14 categories, it surfaced that the biggest transacting category for money launderers was cash-based transactions. In fact, the top two categories would be 60% for cash-based transactions, and 12% for wire transfers. This led the bank to apply a risk scoring system on top of the transactions for identifying the probability of ML-transacting (more on that in Section 4.4.). This implied a stricter selection of pathways and entailed the *risk* of ignoring suspects that fell out of the scope of what those choices reflected. The narrower selection recognized the need to find a better structural coupling between human and computer profiling.

As Bank X concentrated in these 14 transacting categories for profiling ML, the TPR of the TMS increased to about 5–7%, while a model was created that sought to enhance profiling by integrating other sources of data. Based on project workshops and consolidating ML indicators, three other types of behaviours were identified as important: lifestyle, business, and criminal. It is beyond the scope of the present paper to deconstruct the details of each one, however, lifestyle and business behaviour are most closely related to the following section.

4.3. The breakthrough in ML profiling of Bank X (relating to phase 3 – Step 2)

With transacting behaviour being viewed as just one relevant type to ML activities, the question that emerged was how could the profiling of money laundering be enhanced by using additional data, from either within the system or the environment? In one of the project meetings that led to interesting insights, the MLRO of Bank X said that: "...we realized that the marketing department of the bank had access to a (*demographics*) database that they use all the time... It took us quite a bit of time because at first they could not understand what we were saying (implying here about ML modeling). Until we put it in their language and said: How would you market a product to a money launderer? What sort of characteristics would you be looking for?"

Even though the exact profile was not disclosed for confidentiality concerns, the data categories that were used are listed in Table 3 (phase 3, Step 2) and described here in brief. Bank X integrated these categories of data from the database of an off-the-shelf product that the marketing team had bought. The marketing team was using that database in order to target customer groups for particular products, for example, loans, credit cards, etc. While that database contained in excess of 60 micro-categories through which they were classifying the totality of the UK population, after several workshops between AML/Marketing teams, Bank X decided to focus on: i) key *demographics* (household, population, background) as these would yield additional information – at postcode level – for several different segments of the population, ii) *socio-economics and consumption* data (occupation structure, employment status, socioeconomic status, spending, etc...) that would yield lifestyle and other characteristics that would be compared against transacting behaviour; for instance, if a customer was matched against a category of a region where the socioeconomic indicators pointed to a deprived region (whereas the transacting behaviour of the customer gave a very different footprint) then along with the combination of his/her personal details, occupation, spending patterns, he/she could be flagged as a suspect; iii) *property value and location* (value, tax band, accessibility, urbanization) and *financial measures* (directorships, bad debt, credit behaviour) were also used routinely as an information pathway towards investments and business behaviour. Infiltrating businesses, investing in real estate (particularly in the London property market – a favorite of Mafiosi around the world), setting up cash-intensive businesses with high turnover within short periods of time from their establishment (e.g. restaurants), were just a few of the ML typologies that the bank was attempting to target from these last two categories.

A combination of the above elements was used to inform the profiling queries further and as disclosed by the MLRO, increased the TPR_{TMS} to 17.3%. Here, it is worth mentioning that in AML technology use, TPR_{TMS} is the key metric used widely by banks and the key decision makers in AML; through its monitoring, banks try to improve and justify their AML compliance efforts against central bank & regulatory audits on AML technology. But it is not just a performance measure. Banks change their algorithmic detection practices routinely in order to attempt to improve their TPR and by doing so, they change their operational focus on determining suspicious behaviour. In a sense, most AML-operations are guided by the need to improve the TPR but any such changes imply a shift in who is actually targeted for ML suspicious behaviour. The core findings/reflections in the limitations of the traditional transaction monitoring approach and the addition of marketing-oriented data are summarized below:

- a) Peaks and troughs in transactional activity do not give high enough returns of true positive identification of ML. There is an attempt to use any data that differ greatly from the norm. However, as the head of ML analysts always maintained: "good money launderers try to look like the norm". This complicates matters considerably as 'outlier detection' techniques are not as useful as they are in fraud; profiling is conducted under conditions of considerable ambiguity.
- b) Bank X found – accidentally – that high-risk customers for ML also matched a marketing demographics group that had greater propensity to apply for a loan; through that comparison they identified key postcode areas as high-risk for ML; these later corresponded with some of the postcodes of the suspects that the Police arrested for terrorism after the 7/7 bombings in London.
- c) A simple profile that was originally developed and targeted ML-suspects that were: *male, 38–45 years of age, with an association to the bank between 3 and 5 years, and living in specific postcodes around the UK*. This has changed considerably in profiling complexity.
- d) The initial model of Bank X looked at specifics such as: product holdings, income, balance of account, debit values compared to credit values, transaction timestamps and amounts overdrawn, number of transactions, types of transactions, number of debit card transactions,

merchant categories where transactions take place, international location of transactions, time of association between customer and bank, etc. By approaching this as a problem of the system, its subsystems, and its environment, the bank included data related to the behaviour of money launderers, but turned to the structural development of the system itself (instead of merely modeling the behaviour of money launderers). Following the relative success of the approach, Bank X sought to extend its profiling activities by *expanding the variety of data that would support the modeling of ML-behaviour*. This included additional marketing data, referrals from staff members where the narrative part of the SARs would be analyzed for additional profiling indicators, Law Enforcement Agencies enquiries (e.g. when the police would request transacting data for an individual the bank would look at the indicators for this request, potential associates of the individual that could be involved in ML, etc.), fraud data (including internal fraud), data from media sources (TV, Newspapers, etc), industry forums, etc. The rationale of Bank X behind integrating these sources is described in the table below (Table 2). A particular disconnect came with fraud data, described separately in Section 4.5. However, what did make a difference and led to another step-change in the improvement of TPR_{TMS} was the introduction of a scorecard system.

4.4. Risk-based customization and scorecard development (relating to phase 4)

Due to the introduction of the risk-based approach in AML monitoring, banks in the UK (and globally) were encouraged to create risk-profiles for ML and tailor their own AML handling as they saw fit. Bank X chose to do this by attributing a risk score for ML to every single customer in the bank. First, the bank reverse-engineered the general categories from already-

submitted SARs (corroborated by analysts) so that a general risk-score-card-weighting system could be developed. The scale used was 0 for the lowest risk weight and 5 for the highest. For instance, W1 was the assigned risk-code given to gender. If it was found that based on all former submissions (based on 10 years of SARs-data), men had a higher propensity to be reported further in a case due to their suspicious behaviour and engagement in ML, then a higher risk weight would be established based on gender (in this case $W1_{Male} = 1.52$ and $W1_{Female} = 0.62$). A number of different weights were issued and implemented (disclosed in Table 3 along with their numerical values). These risk weights would be assigned automatically to each customer and this would provide the analysts with an automated risk-score for any given customer that was forwarded to them for further investigation (either by the TMS or by staff). That risk-score was also using (and building upon) the marketing database; customer demographics, socio-economics, etc., were bundled into classifications (i.e. “buckets of risk” as the head of AML analysts called them) and these became variations of another risk weight (W7). For example, customers that – based on the marketing data – resided in certain high-risk postcodes (for ML purposes) were grouped into ‘classification1’ and were given a score of 3.48 (out of 5) for that weight.

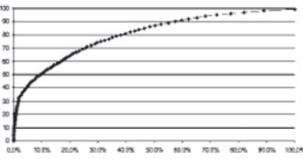
By applying these weights in an automated manner, the ML-analysts were given an additional indication whether the technology-generated alert was truly suspicious. A web based interface allowed analysts to get an instant customer ML-risk score at any time. However, on top of the quantitative metric that this risk provided, ML-analysts were asked to add an additional risk-metric following their manual assessment (low-medium-high). The application of both created a composite risk and a further information/redundancy consideration; the case management system of the AML-team was scripted to combine the two ratings and produce the final risk factor. By using the risk weights as listed in Table 3, the TPR of the TMS was increased further to 20–22%. An important

Table 2
Internal, external and self-referential structural couplings from the AML system.

ML-profiling sources (S/E)	Rationale for ML-profiling integration and restrictions	Identified structural coupling
Marketing data (System)	The use of data from the marketing department has been one of the fundamental successes of Bank X in improving its TPR for money laundering profiling, as the marketing department held important demographic data that could be used/juxtaposed against transaction data.	Internal between AML/Marketing departments.
Suspicious Activity Reports (SARs) submitted from members of staff (System)	The use of SARs intelligence internally has been critical in improving automated profiling as a feedback loop was essentially created between the two. As staff-initiated SARs for money laundering had a much higher TPR (following manual analysis from the AML team), isolation of raw transaction data that corresponds to customers who have been reported through this route, yielded insights about the suspicious clientele of the bank. This process is described in the context of the case study in four steps.	Self-referential, reflexive/recursive where the AML dept., seeks to improve based on the data it owns
Fraud data (System)	As fraud is a very different phenomenon to ML, Bank X did not have much success in using fraud data for money laundering purposes with the MLRO arguing that these departments (i.e. fraud and AML) should remain separate and exchange financial intelligence when it makes sense. However, the conditions for how such an exchange could shape up ML-profiling efforts would demand further research. At a minimum, the stance of Bank X in this matter illustrates that the transferability of techniques from fraud to AML is not straightforward. Outlier detection used in fraud techniques is not easily transferable in AML (as in ML the suspects do their very best to look like the norm), fraud in banking often involves taking control of individual accounts (the owners of which have critical interests in reporting them).	Internal between the AML department and the Fraud department – on exception and not a regular selection of conditions for interaction
Enquiries from Law Enforcement Agencies (LEAs) (Environment)	When an LEA makes a request to the bank for information (this is typically a request to hand-over transaction data for a specific period of time), then the bank may use the request as an information resource itself. In other words, if the request is related to the provision of information for a suspect that has an account with the bank, then the bank can explore further suspects by its own initiative (e.g. potential associates, family members).	A classic system/environment structural coupling between Bank X and data triggered by LEAs
Data from media and social Networking Services (Environment)	Using information from public media is typically highly-unstructured and not systematic in any way. Banks can use XML feeds from news agencies to filter negative news on potential suspects and trigger alerts in the TMS through those. Also, a top international bank is using public data from Facebook for ML profiling purposes – clearly there are serious implications for privacy here that need to be taken into account.	A classic system/environment structural coupling between Bank X and data by media sources

Table 3

Types of data used along the various profiling phases towards Bank X's structural development towards AML; approach and True Positive Rate (TPR).

AML phases	(Types of) Data used by the bank	Approach taken in Bank X	TPR _{TMS}
Phase 1	Raw transaction data that were filtered through the TMS based on applied queries	The vendor had provided 100 predefined SQL queries that were 'profiling' ML behaviour. Upon implementation the bank activated only 7–8 queries.	0.1%
Phase 2	Raw transaction data that were filtered through the TMS based on applied queries	Trial-and-error optimization of the 7–8 queries. Basic model contained at least 4 main elements that characterized a money-launderer: a) male, b) 25–40 years old, c) certain postcodes, d) account holder 2–5 years with Bank X	1.42%
Phase 3	Step 1: Extraction of account numbers of truly suspect launderers from Case Management System data (10 year period) after ML-analyst due diligence; application of account numbers to raw transaction data and subsequent visualized data of the launderers' transactions (Datadesk – Appendix 1) Step 2: Introduced marketing data : <i>Demographics</i> (household, population, background), <i>socio-economics and consumption</i> (occupation, employment status, socio-economic status...), <i>property value and location</i> (value, tax band, accessibility, urbanization) and <i>financial measures</i> (directorships, bad debt, credit behaviour)	see Appendix 1 for screenshots of Step 1	Step1: 5–7% Step 2 17.3%
Phase 4	Scorecard weighting Data* (0 low – 5 highest): W1 Gender: M (1.52), F (0.62), W2 Age Group risks: 25–34 (1.71), 18–24(1.66), 35–44 (1.45), Under 18 (1.04), 45–54 (1.04), 55–64(0.71), 65+ (0.33), W3 Marital status: Single (1.90), Other (0.92), W4 Identification of high risk postcode groups: G1 (2.82), G2(1.3), G3 (0.83) ..., W5 UK Residency: Non-resident (2.7), Resident: (1.0), W6 Directorship: Director (1.92), Non-Director (0.95), W7 AMLMARK (Marketing Database) Code: Classification1 (3.48), Class2 (1.32), ... * Existing SARs cases (10 years of data) were used to calculate the weights and W7 was based on the marketing database classifications.	Introduced customized scorecard for all customers  10% of customers generated half of all SARs	22–23%

finding that the bank stressed was the retrospective fit of SARs cases as they were scored out by the AML risk-model. This gave the bank an additional indication of their 'ML-suspicious customers'. By applying the risk model to SARs cases, the bank found that about 10% of its customers (based on the characteristics assigned) had generated 50% of all SARs ([Table 3](#)).

4.5. The AML/fraud disconnect

The expansion of the profiling elements to a wider variety of data sources created a number of difficulties for the bank. From the case, it became clear that simply including more data to target ML behaviour would not work. An important finding was the comingling of fraud/ML in the context of behavioural modeling and the difficulties that it created. In most cases, integrating insights from fraud data to ML profiling gave rise to suspects that were meaningless from an investigative perspective. According to the MLRO: "As phenomena, these are very different so it is difficult to make the claim that we can take data from one and use it to model the other. We are still looking into this area, but so far, we have failed". While fraud rings were on the rise and the supply of fake ID documents (e.g. through www.confidentialaccess.com) resulted in large scale mortgage fraud, the bank found the convergence between fraud/ML to be particularly problematic. The reasons are summarized here: a) *fraud precedes ML*: fraud generates the illegally earned money that subsequently needs to be laundered; in this regard, fraud-detection techniques do not apply to the ML phenomenon per se, but only to stopping one of the many avenues for generating illegally-earned money, b) while VAT fraud (e.g. carousel fraud or Missing trader Intra-Community Fraud), or other types of fraud constitute *predicate offences* and may be flagged as such in the SARs submissions to assist further investigations (or pass onto relevant authorities like the HMRC if it involves tax evasion), from a behavioural profiling perspective ML is the intersection where all criminal activities meet; thus, the focus on detection should be on the process itself and finding mechanisms through which computer profiling and human profiling can be structurally coupled for the improvement of TPR; c) while the bank had implemented a TMS that was developed from a company with very strong anti-fraud solutions, the transferability of fraud-oriented transaction queries did not work well for Bank X. The exact zone of congruence in terms of detection could not be determined but this would be a subject for future exploration – if not a research agenda on its own.

Similar difficulties were raised with the integration of other data sources, even though it became clear that the agenda and strategy for such integration was much broader. For example, the head of the ML analysts mentioned that they assigned one analyst to check news reports, newspaper articles, police announcements, etc., for information that they could check/integrate for targeting ML. This was expanded later to include news feeds in XML format while, over time, profiling became more aggressive in how personal data were being used for monitoring purposes. Another bank in our project considered using Facebook data to profile customers' lifestyle characteristics and juxtapose with transacting patterns. While out of scope for the present paper, this raises substantial privacy concerns; indeed, the European Commission has recently moved to suggest that Data Protection Authorities should be granted the right to veto AML legislation.

4.6. Evaluating AML: lessons for AML management in banks, MLROs and regulators

The study indicates that Money Laundering Reporting Officers (MLROs) that remain the key decision makers within a financial institution must make a clear separation between fraud and anti-money laundering if they seek to improve AML detection. As MLROs and Heads of AML departments use AML-technology systems to support their decisions in determining truly suspicious cases before these can be forwarded to FIUs, they need to approach money laundering detection systemically and look into a much more structured approach between technology and human co-improvement; for instance, training of staff and the improvement of the TPR_{staff} index can be fed back for improvements in detection. MLROs need to steer their departments from simply monitoring ML to *deconstructing the transacting footprint of their own suspicious customers*. The one-size-fits-all integration of SQL queries that capture ML typologies algorithmically has not worked. Furthermore, thus far, MLROs' decisions to forward technology-based suspicions for ML to the FIU have had an overall negative impact on national AML systems with skyrocketing SARs of poor quality (the vast majority of which are technology-generated). If this situation is to be (partly) reversed, MLROs need to steer their AML departments and integrate technology monitoring for ML differently. More specifically, the specific sequence of phases 1–4 as described above (based on raw transaction data filtering, reverse engineering of transaction from staff-identified suspicious cases – and their

Table 4
Evaluation metric considerations.

AML evaluation metrics	Implications for key decision makers (MLROs & Head of AML)
$TPR_{TMS} = [(CR_{ML-Analysts}/A_{TMS}) * 100]\%$ and $TPR_{Staff} = [(CR_{ML-Analysts}/A_{Staff}) * 100]\%$ where: CR: Corroborated reports by analysts A_{TMS} : Alerts from the TMS A_{Staff} : Alerts from branch staff	Both of these measures are internal to the bank and MLROs should continuously support the improvement of both the TPR_{Staff} (e.g. by enhancing AML training activities) and managing TPR_{TMS} improvements. The two are structurally coupled. As key decision makers that rely on both staff and technology in <i>establishing</i> and <i>forwarding</i> suspicion for ML to the FIU, MLROs should seek to balance these measures with forwarding rates.
$R_w(Analyst) = \frac{W_{staff-alerts}}{W_{technology-alerts}}$ where $W_{staff-alerts} = (Allocated(A_{Staff})/A_{Staff})$ $W_{tech-alerts} = (Allocated(A_{TMS})/A_{TMS})$ Where W is the workload for the ratio of ML-Alerts created by staff reports from branches and the TMS.	Staff demotivation from false positives in managing A_{TMS} needs to be managed by MLROs carefully so that their decisions for forwarding SARs can be based on high quality reports. A balance of the workload of ML-analysts should be maintained in how allocations are being conducted between the higher-quality alerts from staff (A_{Staff}) and the much lower-quality alerts from technology (A_{TMS}). Rotations of staff between the two would assist in managing this process. Technology should be optimized so that the workload ratio (R_w) for each analyst takes into account additional weights (probability of suspicion for A_{TMS} and A_{Staff})
$(F_{Rate})_{TOTAL} = \frac{(MLRO_{decisions-to-forward})}{(A_{TMS} + A_{Staff})}$	The decision for an internally corroborated report to be forwarded to the FIU rests with the MLRO. The <i>forwarding</i> rates (representing this decision) should be monitored closely in relation to both A_{TMS} and A_{Staff} . Very high forwarding rates where low TPRs exist have a negative impact to FIUs that are struggling to cope with the volume of reports. MLROs should safeguard the process of corroborated reports by analysts instead of amplifying any ambiguities and forwarding everything for consideration. In addition, FIUs can provide feedback for the adjustments of this, by informing MLROs of criteria/cases that are helpful in prosecutions. Some FIUs are doing this in a handful of countries but this is neither structured nor part of an AML evaluating strategy that involves both staff and technology.
$(F_{Rate})_{STAFF} = \frac{(MLRO_{decisions-to-forward}/A_{Staff})}{(MLRO_{decisions-to-forward}/A_{Staff})} \rightarrow$ indicates how many alerts from branch staff are eventually forwarded to the FIU following an MLRO decision	Delineating the suspect transacting categories for a bank focuses the scope of profiling and can also be used to monitor the variety of transacting categories exploited over time by money launderers. This gives an indication of the variety of the suspects' transacting footprint. The closer V is to zero, the more focused the selected categories become based on A_{Staff} . A result closer to the number one would imply a very wide variety and no real reduction in transacting complexity.
$(F_{Rate})_{TECH} = \frac{(MLRO_{decisions-to-forward}/A_{TMS})}{(MLRO_{decisions-to-forward}/A_{TMS})} \rightarrow$ similarly for technology	With the daily capacity of members of staff to corroborate reports (CR) being limited, the alerts of the TMS (A_{TMS}) needs to be optimized so that $A_{TMS} + A_{Staff}$ does not exceed the capacity of staff to analyze reports.
$V_{Suspect\ categories} = \frac{Y_{selected}(A_{Staff})}{X_{categories}}$ where X the total number of transacting categories (99 for Bank X), Y the selected subset that originates from staff alerts and V the variety of transacting categories.	
$C_{Staff} = (A_{TMS} + A_{Staff})/CR$ C_{Staff} : staff capacity to analyze alerts	

AML management of the bank as it can demonstrate compliance, reduce costs of checking false positives, avoid financial fines and contribute to the identification of illegal behaviour. Furthermore, regulators ought to encourage this approach and work with other key decision makers both in banks and other reporting entities to reach a consensus of how AML technology can be evaluated and audited in a more comprehensive way. In addition to the TPR that remains central and indicates how banks progressively customize their AML technology solutions to achieve better detection, the case shows that regulators must enhance their feedback mechanisms so that MLROs can develop their decisions and deconstruct what suspicious cases were fruitful for prosecution and asset recovery and which ones could be improved by changing internal processes. In this context, based on the insights from the case of Bank X, the following measures can be used in tandem with the TPR. These can be useful in managing AML systems, evaluating them, auditing them and considering their impact on MLROs and other key decision makers (both within the system and its environment). They can also be considered in light of future AML-IS technological developments (e.g. when designing analytics dashboards for MLROs so that there's comprehensive AML monitoring). These are discussed briefly in the table below (Table 4).

5. Discussion

5.1. Implications for practice & contributions

Abstracting from the case, we can look to implications for institutions that try to improve in their handling of AML. One important aspect is that the (AML) system, in its efforts to reduce the complexity of the environment, is forced to succumb to a default *two-step reduction of environmental complexity*. The first step involves complexity reduction via technology: data from the environment are internalized by the system, which trigger algorithms based on what phenomenon is being modeled. The second step is a follow-up complexity-reduction by human activity systems. As shown in Table 1, the (AML) system develops three types of structural couplings for TPR improvements: i) *internal* (between itself as a subsystem of the bank and other departments like marketing), ii) *self-referential* (with AML recursive explorations like that in phase 3 – step 1), iii) *classic/external* types of structural couplings with the environment of the bank. These should not be thought of as distinct but as intertwined, affecting ML modeling efforts in complex ways. They are depicted in the conceptual model as [A], [B], and [C].

The concept of structural coupling and the findings from the case study alert us to the critical dynamics between computer profiling (CP) and human profiling (HP). These impose restrictions on each other: CP is essential to filter the massive volume of transactions and flag suspected cases by *reducing* behaviour (lifestyle, business) into models that can be monitored. While human profiling is typically perceived as an after-thought to CP, the case of AML illustrates that HP places severe restrictions that need to be built into the profiling efforts. This re-emphasizes the relationship between CP and HP while structural coupling helps us reconsider HP, not as a residual category to CP, but as constitutive to it. This change in emphasis implies that profiling a phenomenon must be seen as an exercise far beyond the computational modeling of a domain (like ML) as HP reconstitutes the way CP is considered. In this context, structural coupling also helps us reframe the traditional distinction between computer/human profiling into a highly asymmetric relationship between *information/redundancy*. In domains like AML where we have high-ambiguity phenomena (e.g. where we try to model, detect and evaluate suspicious behaviour), information systems must be considered as redundancy systems first: systems where the reduction of environmental complexity takes precedence; systems where data (e.g. about suspicious behaviour in ML) is at the environment of the system and it is only through highly selective conditionings that reduce its potentiality that it can become information. The following conceptual framework synthesizes the aspects brought forward (and abstracted) from the case study and offers a

subsequent visualization, as well as the inclusion of specific marketing data in tandem with risk-based scorecards for which indicative weights have been provided), could yield considerable improvements in the TPR_{TMS} . In turn, this can have a considerable positive impact for the

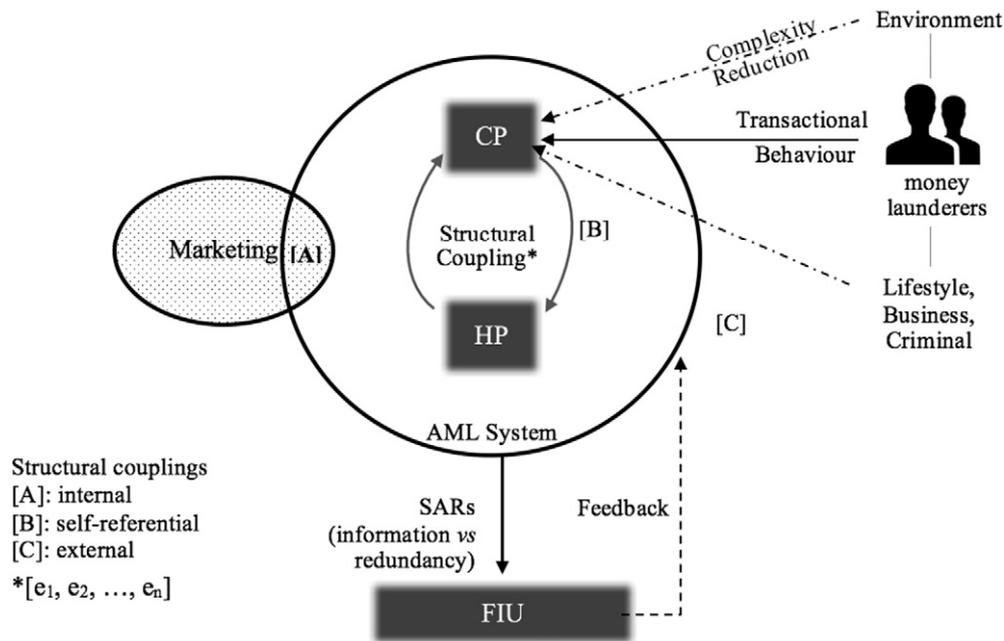


Fig. 1. Systems theoretical framework for ML monitoring.

systems theoretical model of ML monitoring, grounded on the concept of structural coupling.

In ML, the identification of suspicious behaviour and its consequences are more challenging and technology constructs suspicious behaviour (and its identification). The norm against which outliers are constructed is very different and much more challenging to fraud (where the victims are the customers). This has implications for profiling, for how information systems are designed, developed, implemented, and also audited. An important implication for IS is that the 'normal' interaction between 'computer' and 'human profiling' mutates into a massive information-redundancy structural coupling in high-ambiguity domains. IS scholars exploring such high-ambiguity domains will benefit from the framework by rendering any computer/human profiling interactions onto an information/redundancy structural coupling and are prompted to: a) explore and integrate lifestyle, business and criminal behaviour as a different set of behaviours in tandem with transacting behaviour, b) separate AML from fraud as discussed in Section 4.6., and explore what internal structural couplings would be most beneficial for the AML system in the context of profiling (e.g. marketing), c) take into account the regulating elements [illustrated as e_1, e_2, \dots, e_n] that influence the process of structural coupling between CP/HP and build them into any future profiling; these can also be taken into account in IS development, implementation, audit, and can also double as evaluation metrics presented to a dashboard to the two key decision makers in a bank (i.e. Head of AML analysts, Money Laundering Reporting Officer); acknowledging the structural coupling between the two 'turns the tables' on AML profiling and prompts for the integration of HP-oriented indicators in the refinement of CP (e.g. adjusting algorithms so that human demotivation, workload, forwarding rates are taken into account alongside the modeling of the ML-behaviour), d) explore how feedback from the FIU should be enhanced so that decision makers can adjust the structural coupling between CP/HP within financial institutions, e) reflect and explore further the conditions under which the asymmetry between information and redundancy shapes the construct of suspicion in AML; f) consider how profiling decisions create redundancy (e.g. in eliminating some suspects over others) and, g) what additional evaluation metrics can be developed in the context of AML.

In this context, the structural coupling between CP and HP in the theoretical model is considered to be the core of the conceptual model and alerts us to a number of systems theoretical conditions that are important for ML monitoring. First, the way CP considerations inform any AML

behavioural modeling effort can now be reversed in light of the structural coupling. For example, instead of handling CP on its own and ending up with a cost of £450,000 for checking false positives, the *financial cost of checking for false positives* can become part of the modeling effort (though as we've said, liabilities can be far greater in cases of financial fines for non-compliance). Thus, the TMS that is supposed to 'detect ML behaviour and generate alerts must be optimized to produce a number of alerts that would be financially viable for the AML system. Second, any restrictions posed by HP that are usually ignored in the context of behavioural modeling, or treated as an after-thought, must be factored into the profiling dynamics. For instance, ML-analysts typically become demotivated by having to check such a large volume of false positives and this affects their own performance in conducting checks for real suspicion. Thinking about this problem systemically – and through the structural coupling between CP/HP, turns the tables (so to speak) on this problem. Thus, the identification of staff demotivation (as an AML problem) when struggling to evaluate false positives might be considered as a prerequisite to adjust computer profiling by *rotating the allocation of technology-generated alerts* to different members of staff based on the probability for suspicion. While human profiling is typically treated as an after-effect in information systems and a residual category, structural coupling tells us that it must be built into profiling. In this manner, the scorecard development at phase 4 could include an additional weight, W_8 , which would assign a probability of demotivation to a ML-analyst if a large number of low-risk alerts had been assigned to him/her that – following analysis – turned out as false positives. In this way, the distribution of high-risk TMS alerts would be conducted in a way that would be conducive to increased vigilance. Structural coupling alerts us to the need to consider not only the behavioural modeling of money launderers but also of ML-analysts.

The systemic rationale for this applies to many different elements (indicated as general conditions $[e_1, \dots, e_n]$ in the model). Individual considerations for either side of the CP/HP structural coupling (say the cost of checking false positives from a HP perspective), transcend one side of the CP/HP distinction and affect the other side as well. For example, abstracting from the bank, these include: e_1) cost of checking false positives, e_2) capacity for human profiling (i.e. how many alerts can be manually investigated per day), e_3) balancing the percentage of analyzing technology-generated alerts that have a low TPR vs staff-generated alerts from branches that have a much higher (up to 60%) TPR, e_4) staff demotivation, etc. It can also include environmental restrictions; for

example, the FIU capacity to investigate SARs could be fed back to individual AML systems based on what FIUs perceive as *information vs redundancy*; in turn, the AML system (Bank X) would optimize its SARs output. Unfortunately, FIUs generally give *minimal* or often no feedback and without this type of (environment-to-system) coupling, initiated by the FIU, banks operate as black-boxes that output noise (i.e. redundancy) instead of information. This affects FIUs adversely. This appears to be the case for most AML-national systems around the world and constitutes one of the most critical obstacles to developing AML practice further.

At the same time, any potential TPR improvements should consider all types of structural couplings: internal with other subsystems, external and self-referential/recursive improvements. Phases 1–4 as discussed in the case study would relate to internal and self-referential improvements in the TPR and outline the practical contribution of the study. Of particular interest could be social media data and how they might impact CP/HP. In the context of the framework, the argument can be made that within a bank, despite the fact that ML-detection is a legal requirement, the details of who exactly is being ‘detected’, monitored, and produced in a list of suspects remains an emergent phenomenon, based on different forms of structural coupling as identified above. Here, a significant concern must be raised over aspirations to engage in real-time detection [23]. In that case, algorithms alone would largely define criminality.

While the concept of the structural coupling denotes the form of the interaction between the system and its environment [50,51,62], it is not very clear in what way data is transformed into information other than the claim that a system is capable of effecting this transformation [48].

Viewed from this outlook, structural coupling does capture the dependency between system/environment [51] but we find that it is structured via the system’s predisposition to engage with selected ‘data sources’ and externalize them as information, as well as the evolving dynamics between information/redundancy. The transition then from data to information is facilitated through the dynamic demands that computer profiling and human profiling place on each other: this interaction between CP and HP guides the distinction between information/redundancy so that meaningful decision-making can take place within an organization. While we are in agreement with Vanderstraeten [53, p.258], that information emerges as a product of the system, we find that this is instantiated in information systems through a negotiation between CP/HP. Applied to AML and based on the developed conceptual framework, suspicious behaviour then is the product of the AML system, based on the multiplicity of structural couplings (internal, external, self-referential). In fact, we can observe that all of the profiling efforts of Bank X were geared to reduce the complexity of the ML-characteristics and reach a more manageable subset of suspicious cases for ML-analysts. In this way, the information-value of data, and what can be recognized as useful information (instead of redundancy), is only realized after it is internalized within the system. For example, the exact same transaction *data* for customer A can be *information* for the fraud department but *redundant* for the AML department. Subsystems will exercise their own internal structural couplings between CP and HP.

While the key theoretical contribution of this paper is the systems-theoretically-oriented conceptual model, the identification of different types of structural couplings and recognition of their interactions on an information/redundancy level also contributes to systems theory (the concept of structural coupling in particular). In this context, structural coupling can now be perceived as the totality of inter-related system/environment couplings, even when the environment is the system itself (in case of a self-referential, recursive application of the concept). On the basis of how the system evaluates its own structural development or drift, the system can adjust the sensitivity of its boundary and modify its structural couplings with the environment (e.g. adopt data from other sources, ask additional data from customers, etc). As we have seen the system can even mitigate its structural drift by accepting its own redundancy as a legitimate form of information output (e.g. not scrutinizing all SARs carefully in order to avoid the cost of analysis and passing a volume of useless SARs to the FIU). In this regard, systems do not always

optimize their structural coupling based on information and structural development, but also based on accepting redundancy and structural drift. The unavoidable consequence is that carving such pathways through which the (AML) system internalizes data runs an important risk in AML: criminal suspects may be informed of how their behaviour is being modeled algorithmically and so avoid detection altogether.

5.2. Study limitations

A number of limitations are acknowledged. First, the UK banking environment and the national regulatory practices create a unique set of contingencies. While a degree of generalizability is claimed in the development of the framework in Fig. 1, it is recognized that ML techniques often change from one country to another. Also, the transacting behaviour of ‘ML suspects’ is dependent on the economic landscape of different countries and the institutionalization of specific modes of transacting (e.g. other economies are much more cash-based in transacting within the EU, the use of cheques in the US is much more widespread, and so on). A second limitation within the context of bank X was posed by the difficulty in getting access to additional departments other than the AML group (part of the compliance department).

5.3. Suggestions for future research

AML is under-researched in IS so there are multiple opportunities for further research. First, a wide range of different technologies influence AML and these require further examination [63]. These include case management systems, sanction control systems, intranets, etc. Also, industry matters within IS research considerably [64] and taken that the securities industry, exchange bureaux, insurance companies, casinos, credit card companies, the energy sector, retail, football, real estate, non-profit organizations, and so on, are amidst a large number of industries affected/exploited by money launderers (and where technology plays a critical role in detection/monitoring/prevention), the scope is truly vast. Finally, following up on the findings of our project, we tried to expand our research agenda and conduct benchmarking of TPRs across multiple financial institutions within the City of London. All of the banks approached, declined to participate and cited confidentiality concerns in case of possible leaks (the implication is that if a very low – relative – score became known, it might signal a bank’s AML weaknesses and attract active money laundering or cause reputational harm). This is an area that we prompt others to explore; developments in distributed forms of anonymized tools could facilitate this process.

6. Conclusion

The present study highlights the critical dynamics between computer profiling and human profiling in an AML organizational context. The rich context of the case study and the longitudinal nature of this work, advances our understanding on AML detection, provides an approach to key decision makers (e.g. MLROs) that can be implemented towards AML improvements, delineates a series of steps that a bank can take and provides reflections on key AML evaluation metrics (TPR and others in Section 4.6) that are useful to both banks and regulators. This new study contextualises AML detection within the structural coupling of computer/human profiling, re-orientes the rationale of AML profiling by emphasizing contingent human profiling elements that are often ignored in profiling and disentangles AML from fraud, thereby prompting IS scholars to study AML separately. The ensemble of these contributions offers practical guidance for banks that seek to develop their AML monitoring strategies and prompts IS researchers for additional work in this space: not only from a purely technical-orientation but also from a deeper exploration of the organizational dynamics that affect ML-monitoring and the key decision makers in an AML system and its environment.

Appendix 1. (Relating to Phase 3 – step 1)

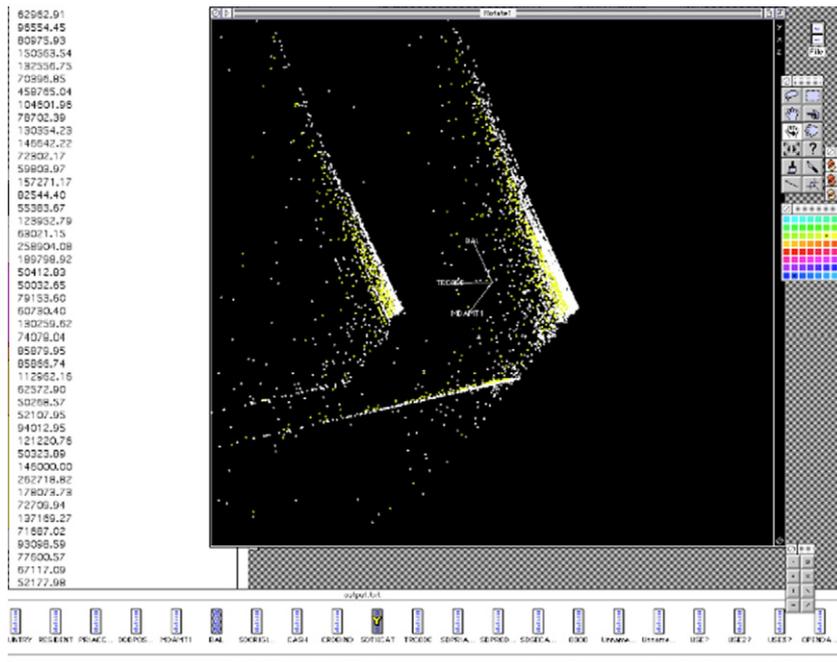


Fig. A1. Visualization of a dataset that money laundering analysts would explore/interact within the context of the project in order to reflect about the qualitative characteristics of suspicious ML-behaviour.

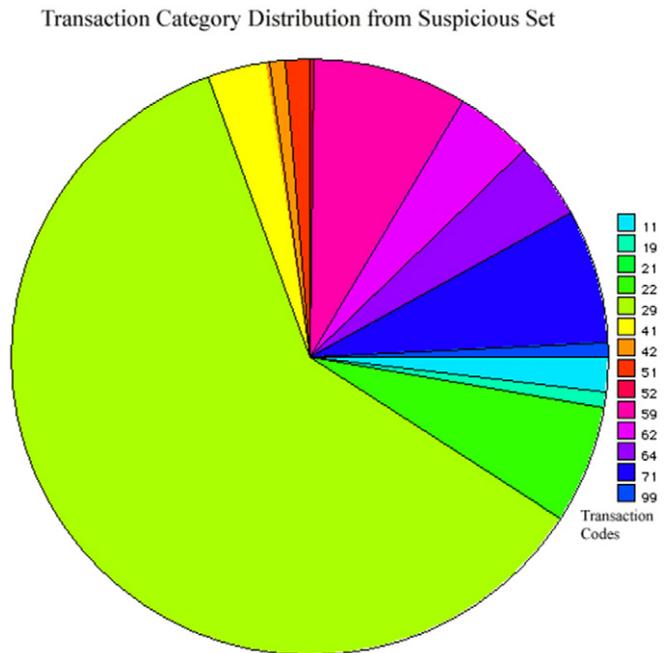


Fig. A2. Reduction of the suspicious complexity. From the 99 transaction codes recorded in the legacy system, the isolated suspicious transactions from the money launderers would indicate that only 14 categories were dominant. The largest area (cash-based transactions) would be present in >60% of suspicions.

References

[1] D. Demetis, Technology and Anti-Money Laundering: A Systems Theory and Risk-Based Approach, Edward Elgar, 2010.

[2] S. Vaithilingam, M. Nair, T. Thiyagarajan, Managing money laundering in a digital economy, *J. Asia Pac. Bus.* 16 (1) (2015) 44–65.

[3] A. Chong, F. Lopez-De-Silanes, Money laundering and its regulation, *Econ. Polit.* 27 (1) (2015) 78–123.

[4] NCA, Suspicious Activity Reports (SARs) Annual Report 2013, 2013.

[5] M. Myers, M. Newman, The qualitative interview in IS research: examining the craft, *Information & Organization* 17 (1) (2007) 2–26.

[6] G. Walsham, The Emergence of Interpretivism in IS Research, *Information Systems Research*, 6: (4), 1995.

[7] UN, United Nations Convention Against Illicit Traffic in Narcotic Drugs and Psychotropic Substances, 1988.

[8] UN, The United Nations Convention against Transnational Organized Crime, 2000.

- [9] F. Calderoni, The analysis and containment of organized crime and transnational organized crime: an interview with Ernesto U. Savona, *Trends in Organized Crime* 18 (1) (2015) 128–142.
- [10] V. Tanzi, Uses and abuses of estimates of the underground economy, *Econ. J.* 109 (1999) 338–347.
- [11] D. Smith, *Black Money: The Business of Money Laundering*, 2011.
- [12] Walker, How big is global money laundering? *Journal of Money Laundering Control* 3 (1) (1999).
- [13] C. Chambers, Can you ever regulate the virtual world against economic crime? *Journal of International Commercial Law and Technology* 7 (4) (2012) 339.
- [14] R. Stokes, Virtual money laundering: the case of Bitcoin and the Linden dollar, *Information and Communications Technology Law* 21 (3) (2012) 221–236.
- [15] E. Mumford, Problems, knowledge, solutions: solving complex problems, *J. Strateg. Inf. Syst.* 7 (1998) 255–269.
- [16] W. Chung, H. Chenb, W. Changc, S. Chouc, Fighting cybercrime: a review and the Taiwan experience, *Decis. Support. Syst.* 41 (2006) 669–682.
- [17] E.W.T. Ngai, Y. Hu, Y.H. Wong, Y. Chen, X. Sun, The application of data mining techniques in financial fraud detection: a classification framework and an academic review of literature, *Decis. Support. Syst.* 50 (2011) 559–569.
- [18] J. Goldstein, A. Chernobai, M. Benaroch, An event study analysis of the economic impact of IT operational risk and its subcategories, *J. Assoc. Inf. Syst.* 12 (9) (2011) 606–631.
- [19] V. Pramod, J. Li, P. Gao, A framework for preventing money laundering in banks, *Information Management and Computer Security* 20 (3) (2012) 170–183.
- [20] X. Luo, Suspicious transaction detection for anti-money laundering, *International Journal of Security and its Applications* 8 (2) (2014) 157–166.
- [21] T. Senator, H. Goldberg, J. Wooton, A. Cottini, A. Umar, C. Klinger, W. Llamas, M. Marrone, R. Wong, *The FinCEN Artificial Intelligence System: Identifying Potential Money Laundering from Reports of Large Cash Transactions*, 7th Conference on Innovative Applications of AI, Menlo Park, CA, 1995.
- [22] R. Drezewski, J. Sepielak, W. Filipkowski, System supporting money laundering detection, *Digit. Investig.* 9 (1) (2012) 8–21.
- [23] S. Fong, Z. Luo, B.W. Yap, S. Deb, *Proceedings of the IASTED International Conference on Artificial Intelligence and Applications*, 2014.
- [24] N.S. Khan, A.S. Larik, Q. Rajput, S. Haider, A Bayesian approach for suspicious financial activity reporting, *Int. J. Comput. Appl.* 35 (4) (2013) 181–187.
- [25] D.S. Demetis, I.O. Angell, The risk-based approach to AML: representation, paradox, and the 3rd directive, *Journal of Money Laundering Control* 10 (4) (2007) 412–428.
- [26] D. Chakrabarti, C. Faloutsos, Y. Zhan, Visualization of large networks with min-cut plots, A-plots and R-MAT, *Int. J. Hum. Comput. Stud.* 65 (5) (2007) 434–445.
- [27] W.A. Tupman, Ten myths about terrorist financing, *Journal of Money Laundering Control* 12 (2) (2009) 189–205.
- [28] A. Canhoto, J. Backhouse, Profiling under conditions of ambiguity - an application in the financial services industry, *J. Retail. Consum. Serv.* 14 (2007) 408–419.
- [29] PWC, *From Source to Surveillance: The Hidden Risk in AML Monitoring System Optimization*, 2010.
- [30] S. Ramage, Information technology facilitating money laundering, *Information and Communications Technology Law* 21 (3) (2012) 269–282.
- [31] M. Gill, G. Taylor, Can information technology help in the search for money laundering? The views of financial companies, *Crime Prevention and Community Safety* 5 (2) (2003) 39–47.
- [32] CS, *A Model of Best Practice for Combating Money Laundering in the Financial Sector*; Economic Paper 43, Commonwealth Secretariat, 2001.
- [33] A. Giddens, *The Constitution of Society: Outline of the Theory of Structuration*, Polity press, Cambridge, 1984 174–175.
- [34] B. Latour, *Science in Action: How to Follow Scientists and Engineers Through Society*, Harvard University Press, 1988.
- [35] W. Orlikowski, Sociomaterial practices: exploring technology at work, *Organ. Stud.* 28 (9) (2007) 1435–1448.
- [36] W. Orlikowski, S. Scott, Sociomateriality: challenging the separation of technology, work and organization, *Acad. Manag. Ann.* 2 (1) (2008) 433–474.
- [37] D. Demetis, A. Lee, Crafting theory to satisfy the requirements of systems science, *Inf. Organ.* 26 (4) (2016) 116–126.
- [38] J.K.B.O. Friis, S.A. Pedersen, V.F. Hendricks, *A Companion to the Philosophy of Technology*, Wiley-Blackwell, Chichester, UK, Malden, MA, 2013.
- [39] N. Postman, *Technopoly: The Surrender of Culture to Technology*, 1st Vintage Books, Vintage Books, New York, 1993.
- [40] D. Demetis, A. Lee, When Humans Using the IT Artifact Becomes IT Using the Human Artifact, 50th Hawaii International Conference on System Sciences HICSS, Hawaii, United States, 2017.
- [41] K. Bausch, Roots and branches: a brief, picaresque, personal history of systems theory, *Syst. Res. Behav. Sci.* 19 (5) (2002) 417–428.
- [42] J. Mingers, Can social systems be autopoietic? Assessing Luhmann's social theory, *Sociol. Rev.* 50 (2) (2002) 278–299.
- [43] R. Ackoff, Towards a system of systems concepts, *Manag. Sci.* 17 (11) (1971) 661–671.
- [44] S. Alter, Are the fundamental concepts of information systems mostly about work systems? *Communication of AIS* 5 (11) (2001) 1–67.
- [45] P.J. Lewis, *Information Systems Development: Systems Thinking in the Field of Information Systems*, Pitman, London, 1994.
- [46] E.J. Garrity, Synthesizing user centered and designer centered IS development approaches using general systems theory, *IS Frontiers* 3 (1) (2002).
- [47] M.L. Markus, A. Majchrzak, L. Gasser, A design theory for systems that support emergent knowledge processes, *MIS Q.* 26 (3) (2002).
- [48] J. Kallinikos, *The Consequences of Information: Institutional Implications of Technological Change*, Edward Elgar, Cheltenham, 2006.
- [49] M. Mora, O. Gelman, F. Cervantes, M. Mejía, A. Weitzenfeld, A systemic approach for the formalization of the information systems concept: why information systems are systems? in: J.J. Cano (Ed.), *Critical Reflections on Information Systems: A Systemic Approach*, Idea Group Publishing, Hershey, PA 2003, pp. 1–29.
- [50] N. Luhmann, *Introduction to Systems Theory*, Polity Press, 2012.
- [51] H. Maturana, F. Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*, Shambhala, Boston & London, 1998.
- [52] N. Luhmann, System as difference, *Organization* 13 (1) (2006) 37–57.
- [53] R. Vanderstraeten, The social differentiation of the educational system, *Sociology* 38 (2) (2004) 255–272.
- [54] B. Jansen, A. Spink, *An Analysis of Web Documents Retrieved and Viewed*, the 4th International Conference on Internet Computing, Las Vegas, Nevada, 2003 65–69.
- [55] N. Luhmann, *Law as a Social System*, Oxford University Press, Oxford; New York, 2004.
- [56] S. Sawyer, Analysis by long walk: some approaches to the synthesis of multiple sources of evidence, in: E.M. Trauth (Ed.), *Qualitative Research in IS: Issues and Trends*, IDEA Group Publishing 2001, pp. 163–189.
- [57] H.M. Cooper, L.V. Hedges, *The Handbook of Research Synthesis*, Russell Sage Foundation, New York, 1994.
- [58] A. Coffey, P. Atkinson, *Making Sense of Qualitative Data: Complementary Research Strategies*, Sage Publications, Thousand Oaks, 1996.
- [59] B.L. Berg, *Qualitative Research Methods for the Social Sciences*, 4th ed. Allyn and Bacon, Boston, 2001.
- [60] C.F. Auerbach, L.B. Silverstein, *Qualitative Data: An Introduction to Coding and Analysis*, New York University Press, New York, 2003.
- [61] L. Richards, *Handling Qualitative Data: A Practical Guide*, SAGE Publications, London, Thousand Oaks, CA, 2005.
- [62] N. Luhmann, *Risk: A Sociological Theory*, Transaction Publishers, New Brunswick, 1993.
- [63] D.S. Demetis, *Technology and Anti-Money Laundering: A Systems Theory and Risk-Based Approach*, Edward Elgar, Cheltenham, 2010.
- [64] M.W. Chiasson, E. Davidson, Taking industry seriously in information systems research, *MIS Q.* 29 (4) (2005) 591–605.



Dr. Dionysios Demetis holds a PhD in Anti-Money Laundering (AML) and Information Systems from the London School of Economics. He has taught on Information Systems Security and Information Systems (at the LSE) while his research on AML and Risk has been featured in the select bibliography of the United Nations. He has taught extensively on International Compliance, Anti-Money Laundering and Crime/Bribery for the PhD programme of the Thomas Jefferson School of Law, in San Diego, California. For his work on big data and its consequences for financial crime, he has won the Emerald Highly Commended Award from the *Journal of Money Laundering Control*. He has worked on large-scale research projects for the European Commission and acted as a convener of the steering group for the Future of Identity for the Information Society (FIDIS); this was a large scale consortium of 25 academic/industry participants across the EU (including Microsoft and IBM) that explored several social, legal, technical, economic repercussions of technology and its consequences for the concept of identity in the information age. He was an active member of the Computer Security Research Centre at the LSE and is currently Co-Chair of the International Security Conference (in its 16th year now). Additionally, he is the author of two books: *Science's First Mistake*, published worldwide by Bloomsbury Academic, and *Technology and Anti-Money Laundering: A Systems Theory and Risk-Based Approach*, published worldwide by Edward Elgar.