



DEREK CHAU  
MAARTEN VAN DIJCK NEMCSIK

# ANTI-MONEY LAUNDERING

TRANSACTION MONITORING  
SYSTEMS IMPLEMENTATION

Finding Anomalies

WILEY

# **Anti-Money Laundering Transaction Monitoring Systems Implementation**

---

---

# Wiley and SAS Business Series

The Wiley and SAS Business Series presents books that help senior-level managers with their critical management decisions.

Titles in the Wiley and SAS Business Series include:

*Analytics: The Agile Way* by Phil Simon

*The Analytic Hospitality Executive* by Kelly A. McGuire

*The Analytics Lifecycle Toolkit: A Practical Guide for an Effective Analytics Capability* by Gregory S. Nelson

*Analytics in a Big Data World: The Essential Guide to Data Science and Its Applications* by Bart Baesens

*A Practical Guide to Analytics for Governments: Using Big Data for Good* by Marie Lowman

*Bank Fraud: Using Technology to Combat Losses* by Revathi Subramanian

*Big Data Analytics: Turning Big Data into Big Money* by Frank Ohlhorst

*Big Data, Big Innovation: Enabling Competitive Differentiation through Business Analytics* by Evan Stubbs

*Business Analytics for Customer Intelligence* by Gert Laursen

*Business Intelligence Applied: Implementing an Effective Information and Communications Technology Infrastructure* by Michael Gendron

*Business Intelligence and the Cloud: Strategic Implementation Guide* by Michael S. Gendron

*Business Transformation: A Roadmap for Maximizing Organizational Insights* by Aiman Zeid

*The Cloud-Based Demand-Driven Supply Chain* by Vinit Sharma

*Connecting Organizational Silos: Taking Knowledge Flow Management to the Next Level with Social Media* by Frank Leistner

*Data-Driven Healthcare: How Analytics and BI Are Transforming the Industry* by Laura Madsen

*Delivering Business Analytics: Practical Guidelines for Best Practice* by Evan Stubbs

*Demand-Driven Forecasting: A Structured Approach to Forecasting (Second Edition)* by Charles Chase

*Demand-Driven Inventory Optimization and Replenishment: Creating a More Efficient Supply Chain* by Robert A. Davis

*Developing Human Capital: Using Analytics to Plan and Optimize Your Learning and Development Investments* by Gene Pease, Barbara Beresford, and Lew Walker

*The Executive's Guide to Enterprise Social Media Strategy: How Social Networks Are Radically Transforming Your Business* by David Thomas and Mike Barlow

*Economic and Business Forecasting: Analyzing and Interpreting Econometric Results* by John Silvia, Azhar Iqbal, Kaylyn Swankoski, Sarah Watt, and Sam Bullard

*Economic Modeling in the Post Great Recession Era: Incomplete Data, Imperfect Markets* by John Silvia, Azhar Iqbal, and Sarah Watt House

*Enhance Oil & Gas Exploration with Data-Driven Geophysical and Petrophysical Models* by Keith Holdaway and Duncan Irving

*Foreign Currency Financial Reporting from Euros to Yen to Yuan: A Guide to Fundamental Concepts and Practical Applications* by Robert Rowan

*Harness Oil and Gas Big Data with Analytics: Optimize Exploration and Production with Data Driven Models* by Keith Holdaway

*Health Analytics: Gaining the Insights to Transform Health Care* by Jason Burke

*Heuristics in Analytics: A Practical Perspective of What Influences Our Analytical World* by Carlos Andre Reis Pinheiro and Fiona McNeill

*Human Capital Analytics: How to Harness the Potential of Your Organization's Greatest Asset* by Gene Pease, Boyce Byerly, and Jac Fitz-enz  
*Implement, Improve and Expand Your Statewide Longitudinal Data System: Creating a Culture of Data in Education* by Jamie McQuiggan and Armistead Sapp

*Intelligent Credit Scoring: Building and Implementing Better Credit Risk Scorecards (Second Edition)* by Naeem Siddiqi

*JMP Connections: The Art of Utilizing Connections in Your Data* by John Wubbel

*Killer Analytics: Top 20 Metrics Missing from Your Balance Sheet* by Mark Brown

*Machine Learning for Marketers: Hold the Math* by Jim Sterne

*On-Camera Coach: Tools and Techniques for Business Professionals in a Video-Driven World* by Karin Reed

*Predictive Analytics for Human Resources* by Jac Fitz-enz and John Mattox II

*Predictive Business Analytics: Forward-Looking Capabilities to Improve Business Performance* by Lawrence Maisel and Gary Cokins

*Profit-Driven Business Analytics: A Practitioner's Guide to Transforming Big Data into Added Value* by Wouter Verbeke, Cristian Bravo, and Bart Baesens

*Profit from Your Forecasting Software: A Best Practice Guide for Sales Forecasters* by Paul Goodwin

*Project Finance for Business Development* by John E. Triantis

*Retail Analytics: The Secret Weapon* by Emmett Cox

*Social Network Analysis in Telecommunications* by Carlos Andre Reis Pinheiro

*Statistical Thinking: Improving Business Performance (Second Edition)* by Roger W. Hoerl and Ronald D. Snee

*Strategies in Biomedical Data Science: Driving Force for Innovation* by Jay Etchings

*Style & Statistic: The Art of Retail Analytics* by Brittany Bullard

*Taming the Big Data Tidal Wave: Finding Opportunities in Huge Data Streams with Advanced Analytics* by Bill Franks

*Too Big to Ignore: The Business Case for Big Data* by Phil Simon

*The Value of Business Analytics: Identifying the Path to Profitability* by Evan Stubbs

*The Visual Organization: Data Visualization, Big Data, and the Quest for Better Decisions* by Phil Simon

*Transforming Healthcare Analytics: The Quest for Healthy Intelligence* by Michael N. Lewis and Tho H. Nguyen

*Using Big Data Analytics: Turning Big Data into Big Money* by Jared Dean

*Win with Advanced Business Analytics: Creating Business Value from Your Data* by Jean Paul Isson and Jesse Harriott

For more information on any of the above titles, please visit [www.wiley.com](http://www.wiley.com).

# Anti-Money Laundering Transaction Monitoring Systems Implementation

---

*Finding Anomalies*

**Derek Chau**

**Maarten van Dijck Nemcsik**

**WILEY**

Copyright © 2021 by SAS Institute Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at [www.copyright.com](http://www.copyright.com). Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at [www.wiley.com/go/permissions](http://www.wiley.com/go/permissions).

**Limit of Liability/Disclaimer of Warranty:** While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993, or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit [www.wiley.com](http://www.wiley.com).

### ***Library of Congress Cataloging-in-Publication Data***

Names: Yip, Derek Chau Chan, 1978- author. | Nemcsik, Maarten van Dijk, 1972- author.

Title: Anti-money laundering transaction monitoring systems implementation : finding anomalies / Derek Chau Chan Yip, Maarten van Dijk Nemcsik.

Description: First Edition. | Hoboken : Wiley, 2020. | Series: Wiley and SAS business series | Includes index.

Identifiers: LCCN 2020032916 (print) | LCCN 2020032917 (ebook) | ISBN 9781119381808 (cloth) | ISBN 9781119381969 (adobe pdf) | ISBN 9781119381938 (epub)

Subjects: LCSH: Money laundering—Prevention.

Classification: LCC HV8079.M64 Y57 2020 (print) | LCC HV8079.M64 (ebook) | DDC 363.25/968—dc23

LC record available at <https://lccn.loc.gov/2020032916>

LC ebook record available at <https://lccn.loc.gov/2020032917>

Cover image: Derek Chau

Cover design: Wiley

10 9 8 7 6 5 4 3 2 1

---

# Contents

**About the Authors** xiii

**Acknowledgments** xv

**Preface** xvii

Chapter 1	An Introduction to Anti-Money Laundering	1
	The Emergence of AML	2
	AML as a Compliance Domain	5
	The Objectives of AML	9
	Regulatory Reporting	9
	Corporate Citizenship versus Profitability	10
	About True and False Positives and Negatives	11
	The Evolution of Automated Transaction Monitoring	15
	From Rule-Based to Risk-Based	17
	From Static to More Dynamic Transaction Monitoring	22
	Latest Trends: Machine Learning and Artificial Intelligence	26
	Latest Trends: Blockchain	29
	Risk-Based Granularity and Statistical Relevance	34
	Summary	36
Chapter 2	Transaction Monitoring in Different Businesses	39
	Banking	43
	Correspondent Banking	46
	Banking – Trade Finance	49
	Banking – Credit Card	60
	Insurance	60
	Securities	63
	Stored Value Facilities (SVFs)	66
	Casinos and Online Gambling	68
	Lottery and Jockey Club	70

	Other Businesses	72
	Summary	72
Chapter 3	The Importance of Data	75
	ETL: Extract, Transform, and Load	76
	Extract: Data Availability and Sourcing	77
	Transform: Data Quality, Conversion, and Repair	80
	Data Load and Further Processing	89
	Loading of the Data	89
	Data Lineage	92
	Multiple ETLs	92
	Summary	93
Chapter 4	Typical Scenario Elements	95
	Transaction Types	96
	Actionable Entity	100
	Scenario Parameters	106
	Use of Maximum Instead of Minimum Value	
	Threshold	108
	Threshold per Customer	109
	Pre-Computing Data	110
	Timeliness of Alerts	112
	Use of Ratios	114
	Ratio as Degree of Change/Similarity	117
	Ratio as Proportion	119
	Other Common Issues	120
Chapter 5	Scenarios in Detail	121
	Large Aggregate Value	122
	Unexpected Transaction	123
	High Velocity/Turnover	129
	Turnaround/Round-Tripping	132
	Structuring	136
	Early Termination/Quick Reciprocal Action	141
	Watchlist	141
	Common Specifications across Unrelated Entities	142

	Involving Unrelated Third Party	144
	One-to-Many	144
	Transacting Just below Reporting Threshold	145
Chapter 6	The Selection of Scenarios	147
	Selecting Scenarios	148
	Regulatory Requirements	148
	Business Drivers	150
	Data Quality and Availability of Reference Data	152
	Maintenance of the Scenario Repository	152
	How Specific should a Scenario Rule Be?	153
	Overlapping Scenario Rules	155
	Summary	156
Chapter 7	Entity Resolution and Watchlist Matching	157
	Entity Resolution	158
	Watchlists	161
	Summary	184
Chapter 8	Customer Segmentation	185
	The Need for Segmenting Customers	186
	Approaches to Segmentation	188
	Overview of Segmentation Steps	191
	Organizational Profiling	193
	Common Segmentation Dimensions	195
	Considerations in Defining Segments	197
	Check Source Data for Segmentation	199
	Verify with Statistical Analysis	200
	Ongoing Monitoring	205
	Change of Segmentation	205
	Summary	207
Chapter 9	Scenario Threshold Tuning	209
	The Need for Tuning	210
	Parameters and Thresholds	210
	True versus False, Positive versus Negative	212
	Cost	213

Adapting to the Environment	214
Relatively Simple Ways to Tune Thresholds	215
Objective of Scenario Threshold Tuning	216
Increasing Alert Productivity	216
Definition of a Productive Alert	219
Use of Thresholds in Different Kinds of Scenario	
Rules	220
Regulation-Driven Rules	220
Statistical Outlier	221
Insignificance Threshold	225
Safety-Blanket Rules	225
Combining Parameters	226
Steps for Threshold Tuning	228
Preparation of Analysis Data	234
Scope of Data	234
Data Columns	234
Quick and Easy Approach	237
Analysis of Dates	238
Stratified Sampling	239
Statistical Analysis of Each Tunable Scenario Threshold	
Variable	239
Population Distribution Table by Percentile (Ranking Analysis)	244
Distribution Diagram Compressed as a Single Line	245
Multiple Peaks	246
Zeros	246
Above-the-Line Analysis and Below-the-Line Analysis	247
Above-the-Line Analysis	247
Below-the-Line Analysis	249
Use of Scatter plots and Interactions between	
Parameter Variables	251
Binary Search	258
What-If Tests and Mock Investigation	260
What-If Tests	260

Sample Comparisons of What-If Tests	261
Qualifying Results of What-If Tests	262
Scenario Review Report	263
Scenario Review Approach	268
Scenario Review Results	268
Summary	274
<b>Index</b>	<b>277</b>



---

# About the Authors

Chau Chan Yip (Derek), is a principal consultant in SAS Hong Kong. He has twenty years of IT system integration and implementation experience. He has participated in several complex and large-scale projects. He leads the SAS AML consulting team in Hong Kong overseeing the delivery and support of over 30 SAS AML sites. Derek designed a tool for analyzing and tuning transaction monitoring scenario thresholds. This tool was deployed in six financial institutes in Hong Kong, and he led the team in providing scenario review service recommending the thresholds, customer segmentation, and other adjustments to the AML transaction monitoring systems. He has also held advisory roles for some foreign delivery of this AML solution. Derek has an honors bachelor degree in computer science and computer engineering from the Hong Kong University of Science and Technology, and a master of science degree in computer science from the Chinese University of Hong Kong.

Maarten van Dijk Nemcsik, LL.M, PhD, has worked for SAS since 2012 as a financial crimes and (tax) compliance domain expert and solution lead. Over the years he has implemented SAS AML, Customer Due Diligence, Enterprise Case Management, and SAS Visual Investigator. Maarten is currently part of the SAS Global Fraud & Security Business Intelligence Unit, being responsible for internal and external training courses for the implementation of the SAS financial crimes and compliance products and project implementation advisory and support. After obtaining a master's and PhD degree in the field of criminal law and procedure and working as a post doctorate in the field of Organized Crime research at the University of Tilburg, the Netherlands, Maarten worked as a security fraud intelligence officer for ABN AMRO Bank NV in Amsterdam and The Royal Bank of Scotland Group in London. Maarten has published in the field of organized and economic crime, anti-money laundering, and compliance in the diamond sector.



---

# Preface

Over the past years, we have been supporting quite a significant number of AML and other financial crime-related implementation projects. We have been meeting with many customer project teams and end users, mainly across Asia Pacific, Europe, and the Middle East. From working with small financial institutions that operate strictly at a domestic level, to implementing our software in some of the largest financial groups worldwide . . . and almost everything in between. As each customer is different, so is each project and we have had to find solutions for a wide range of, sometimes peculiar, customer requirements. We are thankful to our customers for coming up with such a wide range of wishes and demands, as it challenges us not only as consultants but also as software developers, to continuously improve our software, but also our personal skills and our way of delivering tailor-made solutions to satisfaction.

SAS has a relatively low attrition rate, which is testament to SAS being one of the best employers worldwide (don't take our word for it, Google it!). But even then, as time passes, we see team members come and go and may ourselves move on to explore different pastures. With each skillful consultant that moves on in their career path, insights, expertise, and experience is bound to get lost. Back in 2016, Derek began thinking about preserving these insights, knowledge, and experience via writing and publishing. This initiative was well received within SAS, and a foundation was laid for this book . . . Derek's first one.

This was initially an individual effort, and not an easy task for a "technical" guy who, until now was happy conversing within his own comfort zone . . . the language of data and computers! After a while Maarten joined Derek and we have been working on the book ever since, in whatever time we could spare from both our busy jobs and, more importantly, family life.

Even at the time when social distancing wasn't even "a thing," we kept well apart from each other. Maarten working from Spain

and Derek from Hong Kong. Never on a project together, or even a training. As technically savvy as we believe ourselves to be, (why else would we work in this field and for a company like SAS!) we were fine communicating strictly through electronic means. The time zone difference only became a bit of a hindrance when our SAS editorial support, namely Lauree Shepard, was based at the SAS mothership in Cary, North Carolina.

So, this book is about transaction monitoring implementation, which is, at the end of the day, inescapably a technical subject. In the world of today, technology is still very much in the fast lane. Four years have elapsed since 2016 and during the process we were worried about some of the content becoming outdated. That worry has still not dissipated entirely, but we are modestly confident that most of the information and guidance that we set out to share will still be relevant for those entering the field of AML software implementation. We hope that most of our insights as captured in the following chapters will hold some truth for now and in the foreseeable future. And we would like to thank you, as our (aspirant) reader . . . for placing your trust in us and this book.

CHAPTER **1**

# **An Introduction to Anti-Money Laundering**

## THE EMERGENCE OF AML

Money laundering is generally understood as the concealment of an illegitimate source of assets, providing an apparent legal origin. People have various reasons to whitewash assets: they might want to conceal the original crime and not let their wealth become a whistleblower, or they may simply want to build up a reputation of being a successful and respectable member of the community. Since the late 1980s, there has been another reason to launder money. Law enforcement, initially as part of the US-driven war on drugs, started to clamp down on the financial aspect of crime and new laws were enacted globally to criminalize the laundering of assets itself, whilst at the same time making it much easier for law enforcement to seize assets and for the courts to include confiscation of assets, both as a penalty and a measure of redistributive justice. Against the backdrop of a publicly perceived rise in profit-driven crime, it was generally felt that criminals should be hit where it hurt the most: in their pockets. Criminalization of the act of money laundering and the emphasis on the law enforcement effort to go after the money are a natural extension of the age-old adage that no man should profit from his own crime. Consequently, money laundering touched upon the core beliefs about a just society, where advancing oneself by evading the rules is felt as unfair towards those citizens who abide by the law. As such, money laundering (as any criminal offence) is a crime against society, against the public . . . and there is a public duty to fight and prevent.

By the late 1990s, another dimension was added: the counteracting of the financing of terrorism, and this was further fueled by the US terrorist attacks on September 11, 2001. This dimension worked as a catalyst for the development of ever more stringent anti-money laundering regulations, adopted across the globe, and the emergence of what arguably can be called an entire new industry: Compliance. To a large degree, Compliance became a synonym for *AML* Compliance, *AML* standing for Anti-Money Laundering (and we will use this well-established acronym throughout the rest of this book), but even that is a *pars pro toto*, as it commonly also encompasses counter-terrorist financing (CTF).

There were two main factors that contributed to the emergence of *AML* Compliance. First, there is the down-to-earth fact that law

enforcement simply did not and does not have the capability or the capacity to do what is needed to detect money laundering. This is why financial institutions were recruited to partake in law enforcement as *gatekeepers*. Second, there was the shift in public perception about the role of private companies as *Corporate Citizens* and the intrinsic notion of *good citizenship*, linked to widespread notions on corporate moral responsibility, sustainability, and good standing and reputation. This is why the financial institutions, to date, accept the operational and cost burden of AML. Obviously, there is a clear financial incentive for financial institutions to be compliant: the fines imposed by the regulatory watchdogs for non-compliance are enormous. But beneath this mundane motivator there seems to be a genuine acceptance by the financial industry of the role they have to play, as members of society at large, to disallow and prevent the abuse of their systems.

Accepting this role is one thing, it is quite another to live up to it. Being a money laundering prevention gatekeeper imposes all kind of practices that need to be established in order to keep compliant with all regulatory requirements. There is the practice of “know your customer” (KYC), which in a nutshell means establishing that a customer is who he claims to be. Then there is the practice of enhanced due diligence: risk assessing a financial institution’s entire customer base on a continuous basis, specifically for the purpose of AML and CTF, and stepping up the monitoring effort or even reconsidering the relationship with the client in the case of high risk. Lastly, there is the practice of transaction monitoring: looking at behavior on the account to identify any suspicious<sup>1</sup> or unusual activity. When such activity is deemed to be found and cannot be refuted by further analysis or investigation, then the financial institution has the obligation to report this to the local Financial Intelligence Unit, which in most countries acts as the

<sup>1</sup>Whilst some jurisdictions allow for financial institutions to look for suspicious activities, other countries define this as an activity specifically for law enforcement and against the rule of law to task the banking sector with law enforcement tasks typically associated with the public domain. Whereas the semantic distinction between suspicious and unusual seem to point in the direction of the former compared to the latter, requiring a greater investigation effort by the financial institution, in reality, there does not seem to be much difference other than in the wording.

conduit between the financial institutions and law enforcement, and quite often acts in an investigative capacity itself.

It goes without saying that complying with AML along the lines of these three practices impose a huge administrative burden on the financial institutions, requiring significant investment in front, middle, and back office operations. This applies in particular to transaction monitoring where volumes of customers, accounts, and transactions are significant, and meaningful analysis cannot be done by human labor alone.

And this is where AML software enters the scene. Financial institutions not only deploy electronic means to detect suspicious or unusual activity because of the sheer scale of the data, but also because regulatory watchdogs require them to apply computerized forms of analysis to avoid inconsistency and too much reliance on the (frailty of) human capacity to do so. Whilst computer systems carry out the tasks they can do more efficiently than humans, there is still a role for human analysts and investigators to further verify the validity of the initial electronic analysis. Thus AML transaction monitoring is typically divided into three practices: electronic analysis of transactions and the subsequent generation of *alerts*, the assessment by a human analyst with regard to the validity of the alert(s), and the subsequent filing of a regulatory report if one or more related alerts cannot be refuted as false positives.

The end-to-end process of data collection involves electronic and then human analysis; further investigation of more complex cases; and reporting to the regulators and being able to explain to the regulator how risks have been assessed and mitigated appropriately. All of this constitutes a complex operation, driving up the (manufacturing) cost of financial services delivery and potentially upsetting customers, especially when these customers find themselves unjustifiably subject to a financial investigation, often with the added penalty of not being able to execute transactions and do business. Striking the balance between satisfying both the regulator (compliance), banking customers (services delivery), and shareholders (keeping operational costs down whilst maintaining a good reputation) is of utmost importance for financial institutions today. It is the AML transaction monitoring software that enables financial institutions to do this.

The aim of this book is to explain various aspects of AML transaction monitoring software and will mainly focus on the electronic analysis, both from a perspective of the logic applied in this analysis and the challenges faced during implementation of that software in terms of data integration and other technical aspects.

We, as authors of this book, share a history with SAS, representing a combined 17 years of experience in the implementation, configuration, and redesign of the SAS Compliance Software. SAS has been considered market leader in a number of significant areas relevant for the practice of AML: data integration, analytics, transaction monitoring, case analysis, and reporting.

For any AML software to be taken seriously, is by definition complex. This is because AML transaction monitoring is a multifaceted process. This book aims to provide further clarity mainly on the logic applied to the rules. We hope to provide a good starting point for the beginning AML scenario analyst and administrator. We also hope this will benefit AML alert analysts and investigators, so that they may understand a bit better the output of what we call the alerting engine in terms of the AML and CTF alerts. To that aim this book will put the technical and analytical detail in its business context.

One could legitimately ask if this would not play into the hands of those whose actions we try to detect and allow them to improve their ability to evade detection. We think this not to be the case. Whilst much of the rule-based approach is already in the public domain, the keys are in the actual thresholds financial institutions set as part of these rules for their specific customer segments.

## AML AS A COMPLIANCE DOMAIN

Money laundering is commonly defined as the act of providing a seemingly legitimate source for illegitimate funds. In order to conceal the illegitimate origin, the proceeds of economic crimes need to be white-washed. This will do two things for the criminal beneficiary. Firstly, it will cut the follow-the-money trail leading to the criminal acts, thus avoiding financial assets giveaways of the underlying crimes. But, secondly, even when it does and the criminal is successfully prosecuted and convicted for the criminal proceeds, when successfully laundered,

the proceeds will not be subject to confiscation, since no link between these assets and the convicted can be proven. Even if the criminal is imprisoned, the criminal or their family will still be able to dispose of the assets and wealth build from a criminal life.

Anti-Money Laundering, therefore, is, in its widest sense, a worldwide framework of legislation and regulations to prevent abuse of the financial system for the purpose of money laundering, and the practices and processes to comply with this framework. AML attempts to put in barriers to prevent abuse of the financial system by those seeking to conceal the criminal origins of assets and/or the link to their (criminal) ultimate beneficial owners.

What the legislation and regulations have in common is that they require private sector organizations, defined as Financial Service Providers, to put in place a mechanism to prevent the abuse of the financial system by perpetrators of financial crimes. Underpinning this effort is the adage that criminals should not be allowed to profit from their (economic) crimes and respectable financial service providers should not facilitate criminals in doing so. Some hold the view that AML regulations seek to recruit the financial service providers in the active detection of economic crime through the detection of money laundering that is often associated with it.

Compliance and AML cover a wide domain that includes many subdomains, such as customer acceptance (KYC, name screening), customer due diligence, handling of politically exposed persons, ongoing risk assessment, real-time screening on remittance, transaction monitoring, sanctions screening, regulatory reporting, and more. Although there are tools and technologies available for each and all of those aspects, of equal importance is the awareness and understanding of the institutes' personnel about AML principles and regulations and how to apply these in everyday practice. The software is there to assist responsible staff to efficiently and effectively work through the vast amount of customer transaction data, sanctions, and watch lists, and to detect links and patterns that are impossible to see with a human-driven and/or naked eye approach.

Within the various subdomains, transaction monitoring often incurs the highest spending by financial institutes. A transaction monitoring system basically looks at the financial transactions of customers

over time and identifies suspicious patterns. Virtually any monitoring rule can be implemented, given the relevant data available. It is an after-the-fact control, instead of preventive. At the time when unusual, suspicious, or otherwise alertable behavior is detected, the money has most likely already been transferred and beyond reach of the financial institution. Reporting to the authorities may, however, still provide useful indications of abnormal behavior for investigation and, thus, be a starting point for an investigation, especially when the alert from one bank can be triangulated with alerts by other banks. In most countries the alerts will end up with the Financial Intelligence Unit (FIU). This may lead to the unravelling of money laundering schemes and rings and possibly lead to prosecution and conviction of those responsible.

Central to AML is the concept of alerting. This has its origin in the gatekeeper function imposed on financial institutions by AML regulations. Counteracting money laundering and its predicate crimes is a task for law enforcement, governed by the rule of law. But in most jurisdictions, they lack the capability, capacity, and, more importantly, the legal backing to directly monitor the transactions of citizens. Financial institutions do, and one important driver of the AML regulations is to mobilize the financial institutions to contribute to the fight against crime by disallowing money launderers and criminals to reap the benefits of their criminal proceeds through the abuse of the financial institutions' infrastructure, products, or services. The main mechanism through which financial institutions contribute is by reporting suspicious or unusual behavior.<sup>2</sup>

Before they can submit a report, financial institutions must have identified a transaction or some behavior as reportable. This internal identification is commonly referred to as an alert. An alert, in this context, can be defined as something to draw attention to indications of possible money laundering or terrorist financing. There is no need

<sup>2</sup>Whilst some jurisdictions task the financial institutions under local AML legislation to detect *suspicious* behavior, in other jurisdictions the viewpoint is held that this belongs to the realm of law enforcement and therefore the financial institutions' involvement goes as far as identifying unusual behavior. This conceptual difference is often reflected in the naming of the reports: e.g. in the US, SAR stands for *Suspicious Activity Report*, and the name of the report in the Netherlands, literally translated, is *Unusual Transaction Report*.

to prove this (as that would bring the financial institution into the realm of law enforcement); it suffices to have assessed the likelihood of the alerted behavior pertaining to money laundering (or terrorist financing) as higher than average and worthy of further scrutiny.

Alerts are therefore indications that something may require further attention and those indications can both be human generated and machine generated. A typical scenario of human-generated alerts is when a front office, customer-facing bank employee observes something suspicious and decides to send an e-mail, pick up the phone, or submit an online form to make the Compliance department aware of his or her suspicions. Such notifications may come from employees, but also from customers, other financial institutions, law enforcement, or any other party. To further serve this kind of human alerting, many AML case management systems provide the option for end users to create alerts manually.

However, regulators these days require the majority of transactions to be analyzed individually and in the context of other transactions on the same account and by the same customer or customer group, and the sheer volume of data to be processed means that we can no longer rely on human vigilance alone. This is where software starts to come into the picture. For the past two decades, dedicated software has become increasingly sophisticated in how it analyses the bulk of transactions processed by the financial institutions on a daily basis.

Most commonly, this scrutiny takes place in two phases. The first is in the form of a high-level analysis. If reasons or conditions can be found that refute the hypothesis of money laundering (ML) or terror financing (TF), then the alert will be discarded as a false positive. If, however this cannot be done, the financial institution is obliged to further investigate the alert until a corroborated decision can be made to either report the activity to the regulator or to close the investigation without reporting.

Striking the balance between hit rate and alert productivity is vital to any AML operation driven by a commercially operating financial institution. It is here where a transaction monitoring system supporting a risk-based approach with a more advanced analytical approach, demonstrates its value. A robust system with a proper model in place,

and backed by a capability to change the model and scenario thresholds based on analysis, keeps both the regulator satisfied (hit rate) and minimizes the operational cost of alert handling (alert productivity).

## THE OBJECTIVES OF AML

Let us give some thought to the objectives from a financial institution's perspective in AML transaction monitoring. We take as an example a private sector, profit-driven financial institution, to whom compliance with legislation is merely a cost of doing business and has to be balanced with customer experience and satisfaction as one of the main drivers for being competitive and generating revenue by financial service delivery.

### Regulatory Reporting

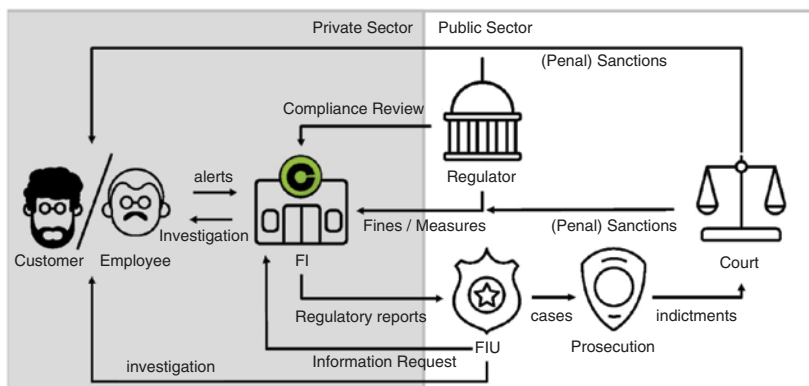
If you ask people working in the domain of AML transaction monitoring software about the purposes of such a system, many will respond along the lines of generating AML alerts. Others, who keep in mind the overall business context will say to report to the regulator. Indeed, the entire AML chain culminates in that one final objective: submitting regulatory reports. The financial institution's effort starts with the sourcing of the data and continues with the analysis of the data and the creation of alerts. During alert triage and the case investigation stage the financial institution will decide whether the findings are reportable in terms of the regulatory requirements imposed by the AML regulatory framework. If so, the relevant data must be submitted to the local regulatory authority, often in a format prescribed by the authority. Once submitted, feedback may be received on whether the report has been received correctly (this is technical and can be automated feedback from the recipient system talking to the sending system); other feedback will be around the content (or payload) of the report and may comprise a response from the FIU in terms of requests for further information. Evidently, this is an ongoing, never-ending process, but conceptually speaking many financial institutions consider their job done once the report has been submitted. The submission of the report represents the end of a cycle. At the surface level,

when someone says the ultimate objective of the AML transaction monitoring system is to submit regulatory reports required by law to the relevant regulatory recipient, then technically that person would not be wrong. Of course, there is more to it.

## Corporate Citizenship versus Profitability

AML regulations put the financial institution between a rock and hard place. Being compliant with legislation and regulations may be intrinsically important for a self-respecting financial institution that is serious about its role as a responsible corporate citizen, but abiding by the law itself does not generate revenue. Those with a more Machiavellian view on things may argue that financial institutions only comply with regulations from a perspective of sustainability and that any fight with the regulatory bodies is doomed to fail, and fines are so high that it endangers profitability, let alone continuity of the financial institution. In that case, satisfying the regulators, whilst avoiding costly fines, does not generate revenue. Revenue is generated by good customer service that customers can appreciate. Being reported to the regulator usually is not. Any financial institution that will see itself as a good corporate citizen is more than willing to comply with the law and even to accept some spillover of customers that are mistakenly or unnecessarily reported to the FIU. This becomes more of a problem if too many legitimate customers end up being reported.

Then there is the operational side of things. As most financial institutions are essentially profit-driven entities in a sometimes highly competitive market it is important to keep operational costs down. Most financial institutions are also publicly registered companies and shareholder value sometimes even trumps profitability in terms of key performance indicators. Shareholder value, by itself, has two sides: in most situations there is a strong correlation with profitability and economic performance, providing one more reason for financial institutions to keep operational costs, including those associated with AML operations, down. Regulatory fines by the regulatory watchdogs for being found insufficiently compliant can be huge and turn a profitable quarter or even year into a non-profitable one. On top of that, shareholder value may shrink upon the news of a major fine, let alone a multitude of fines. See Figure 1.1.



**Figure 1.1** Interactions around a financial institution on AML.

To avoid fines, financial institutions, these days, are putting a high emphasis on getting their houses in order and having a robust Compliance program in place to satisfy the regulator. However, and somewhat ironically, the more stringent the scrutiny of transactions, the higher the operating costs, especially when proper scrutiny and avoiding the unnecessary reporting of legitimate customers, at some point in the process, entails human involvement. AML operations, both from a human resource and an IT perspective, come with high costs. All of which are eating into the bottom line and may affect both profitability and shareholder value. So, when it comes to the true objectives behind AML operations, from a publicly registered, profit-driven financial institution's perspective, it all comes down to getting the balance right between overreporting and underreporting and to do so in the most economical way. And *this is the* key objective of any AML automation.

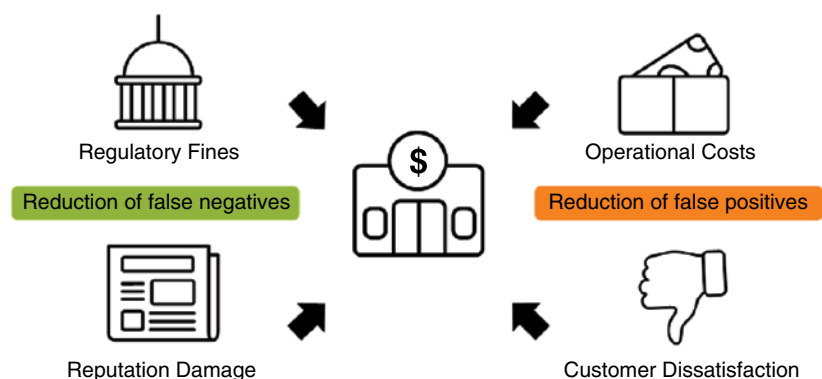
## About True and False Positives and Negatives

Automated AML transaction monitoring is a form of predictive analysis. Based on a limited set of data that, in itself, will never be conclusive on the subject of whether a transaction constitutes money laundering or is part of it, the system tries to filter those situations to point the analyst in the right direction. Ultimately, a human analyst

will look at the generated alerts and determine whether the system correctly singled out that situation. This could be the financial institution's compliance analyst, his/her peer or manager, an internal audit team, a regulatory watchdog audit team, or the analysts of the FIU to whom the truly suspicious events are reported. If an alert is being discarded as not being relevant or not (sufficiently) indicative of money laundering, then we call this a *false positive*. A false positive can mean different things at various stages of the process. Let us follow through the most common process and see what constitutes a false positive at these various stages. For most financial institutions, after alert generation, the process starts with an AML analyst who performs initial *triaging*. This term is borrowed from military practices: after a mass casualty event the sparse group of medical aids will have to prioritize their efforts to maximize the number of lives they will save. They must quickly assess on-site who are beyond help (either dead or so heavily wounded that with the time and effort to stabilize and help them, a lot of others could have been helped but will die because efforts are concentrated on this more badly damaged person). Triage in AML has nothing to do with saving lives, but the term is used to express that an analyst must first discard the obvious false positives and separate out those that require further investigation. Many Compliance departments (and the same goes for fraud investigation) distinguish between an alert and a case: the latter meaning that one or more alerts are singled out for further scrutiny and more time and effort is put into them to determine whether these alerts should or should not be reported to the FIU. When an analyst discards the alert, then from his/her perspective this will be a false positive, a false alarm so to speak. Both in AML and fraud it is widely accepted that 9 out of 10 alerts are false positives. At this stage, a true positive is when the alert might be potentially reportable and therefore either reporting will be decided or, as in most situations, a case will be launched to further investigate. To simplify: true positive = case. The case investigation can be done by the same analyst who now has decided to devote more time to it, but many financial institutions have created a separate role for that: the investigator or senior analyst. At this stage, the same question is again asked: are the alerts under investigation sufficiently out of line with what one would expect under the same circumstances to raise

enough suspicion so that the authorities should be informed? Here a true positive means that the details of the transaction(s) of the account and account holders will have to be shared with the FIU. Here: positive alert = regulatory report. From one perspective, the better the analyst does his job, the higher the ratio between cases created and reports submitted. This process will repeat itself at the level of the FIU, which will research the incoming reports, tie them together where relevant and hand over to law enforcement, which in many countries are specialist financial investigations units. Here, true positive means a report worth further investigation by the FIU, and ultimately it will mean producing a case that prosecution will take to court. Ultimately a true positive alert in its truest sense means that the ultimate money launderers and other who benefit from the money laundering scheme are convicted in court and that the ruling is upheld in appeal. See Figure 1.2.

Obviously financial institutions are usually not concerned with this legal follow-up. At most, records are formally requested by the prosecution to serve as evidence in their case, the financial crimes unit works together with the bank's fraud prevention department in investigation, and/or bank employees are summoned to court to give testimony. For financial institutions, whether an alert is a true or false positive is determined at the early stages of triage, case investigation or submission of the reports. However, it is important to understand that the notion of a false positive and the false positive rate, as key



**Figure 1.2** Financial Institutions have incentives to reduce both false positive and false negative alerts.

performance indicators for automated transaction monitoring, are multi-interpretable concepts and the team concerned with optimization of the system analytical performance will have to decide which measurement(s) they adopt as their key metric.

A high false positive rate indicates overreporting. While this may not necessarily be an issue for the regulator, as they would want financial institutions to stay on the safe side of things, it is a concern for the AML reporting chain, from financial institutions to FIUs, as too many false alerts may clog the system and actually draw away precious resources and capacity from more relevant investigations. From an operational and cost-efficiency perspective there is a big incentive for financial institutions to keep the false positives down and have a high true positive ratio. A key metric to express the quality of the system is detection rate: how many alerts are produced? But this detection rate must always be offset with the true and false positive rates. A high detection rate means nothing if the true positive alerts are still the proverbial needles in a haystack of false positives. At the moment, a 90% false positive rate, meaning that only 1 in 10 alerts is considered for further investigation, seems to be acceptable and is considered within the risk tolerance of many financial institutions and regulators alike.

Optimizing the analytical systems performance by looking at true or false positive is referred to as above-the-line analysis. This seems to imply there is also a below-the-line equivalent. Below-the-line analysis is concerned with the false *negatives*. A false negative is when the system failed to generate an alert in a situation where (often in hindsight) the Compliance analyst or investigator would have deemed the alert a true positive in the hypothetical case that was generated by the system.

Whereas true and false positive ratio are produced by the normal process of reviewing alerts by Compliance analysts and investigators, the analysis of the negatives requires a special effort and is usually only done periodically (or as a one-off) during a system optimization exercise. Below-the-line testing requires a lowering of the thresholds and then a review of the resulting additional alerts, to assess how many of these alerts *would* have been true positives, which makes them, in fact, false *negatives*. This exercise can be repeated with different threshold values. Depending on the number of false negatives, compared to

the number of overall additional alerts, the financial institution may decide to lower the thresholds to have a better hit rate.

From a regulatory viewpoint, false negatives are an issue. The financial institution has to defend why it has set its scenario parameters so high as to knowingly miss out on a number of valid alerts. One reason is that catching those few extra true positives may come with significant overhead for the handling of the additional false positives. This may not be of concern to the regulator, but it certainly is for the financial institution, who seeks to minimize its operational costs.

Analytical software supports the activity of above- and below-the-line testing and helps to run what-if scenarios to optimize the threshold parameters and strike the right balance between effectiveness and efficiency, meaning the best ratio between true and false positives, and true and false negatives. This is the primary, if not sole purpose of any AML transaction monitoring system: getting the balance right in having the most cost-efficient tool that satisfies the regulator in terms of risk tolerance.

## THE EVOLUTION OF AUTOMATED TRANSACTION MONITORING

In the field of AML transaction monitoring, technology is applied to support front and back office business processes with a specific business objective: being compliant with relevant legislation and regulations. Inherently these business objectives are predominantly driven by regulatory developments. These regulatory developments are, in turn, partially driven by a learning process on the part of the regulatory supervisors, which is enabled by the technological advancements and the increase in analytical capability. Regulatory developments are also driven by what regulators and law enforcement observe with who they target: the money launderers. Technology and the application of (advanced) analytics has the effect of tightening the net; however, a consequential effect is that those who seek to launder money through the system will adapt and survive. One can truly speak of an evolution whereby the various actors in this monetary and legislative eco-sphere react to each other's progress.

This evolution is increasingly determined and driven by advancements in technology: the introduction of electronic payments, online, and mobile banking on the one hand, and the increasing importance of data and data processing on the other hand. As can be seen in virtually all areas of commerce and corporate activity, everything is electronic-data driven. Transactions in financial institutions, even the ones that seemingly involve only cash, will be registered in a ledger, which these days is an electronic database.

This brings opportunities, but it also brings challenges. Big challenges. The first and foremost challenge concerns the volume of data, which is so big, that human, naked-eye scrutiny of this data is impossible and highly ineffective. Whilst this may seem an open door, it should be realized that even today some financial institutions rely, for parts of their transaction monitoring, on analysts who have as a day job to cast their eye on hardcopy printed lists of transactions, merely to see if they can catch any irregularities. These lists are at worst random samples of transactions and at best a tiny subset of the overall volume of transactions filtered out by very simplistic and crude rules, sometimes built in as a formula into the spreadsheet used for this purpose. Needless to say, this is a highly insufficient approach to transaction monitoring and most probably will not bear a regulatory approval, not today at least. But, apparently, it had done so in the past.

From a technology perspective we can distinguish between three areas of expertise that have enabled the evolution of AML transaction monitoring. These are data integration, analytics, and data visualization. The evolution of AML transaction monitoring followed a path along the lines of these areas, in this particular order. Whilst initially the concern was with how to link into the electronically available data and get all the data into one place for the use and analysis of the one specific goal of AML, the focus has now shifted to the analytics applied to that data to refine models and increase alert productivity. Whereas automated alerting has taken up an increasingly important role, most regulatory frameworks will not allow for a fully automated process, and the human element is still considered an essential component of any AML operation. To make sense out of the massive volumes of data and the information that can be drawn from it, data visualization is claiming its place.

Here, we also find an educational cycle between human and computer. While it is human expertise and knowledge that created the rules that instructed the computers, it is then with the aid of computers, databases, and electronic data analysis tools that humans learned much more about the data. The new insights are then reapplied in ever more elaborate and sophisticated models. Models are then trained by or with the help of electronic data, often enriched by human investigation, and gradually we move into the direction of self-learning machines that only need minimum input from humans to further train their own models and increase their own capability of separating the wheat from chaff.

## From Rule-Based to Risk-Based

AML and, in particular, transaction monitoring are, on the one hand, driven by regulatory developments which, in a strongly simplified way of speaking, represent how the world should be, the realm of desirable. On the other hand, the transaction monitoring is also driven by the progress the world has made in computing. Computing power and the ability to analyze big data in increasingly shorter amounts of time represents the realm of the possible. The increase in sheer computing power has had and will have big implications for AML transaction monitoring, especially the sophistication in the detection of unusual or suspicious behavior.

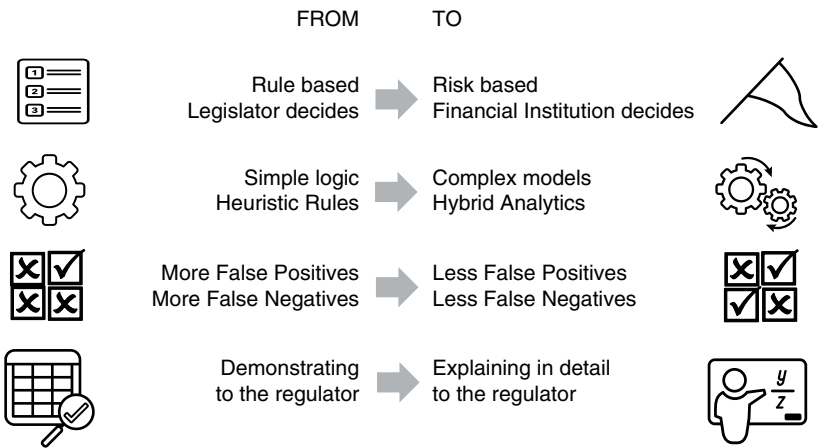
We have discussed the move from a rule-based to a risk-based system. In a nutshell, this move was the result of the system being clogged by a vast volume of false positives because the initial rule, prescribed by the regulators, were far too crude. This to the dismay of the FIUs, which lacked the capacity to review hundreds of thousands of incoming reports. Financial institutions saw operational cost towering without any evident benefit, and customers found themselves the subject of alerting for no good reason at all. The financial institutions jointly indicated that the regulators should trust them and allow them to replace or at least augment the initial hardcoded rules with more sophisticated scenarios.

The name of the game became find the odd one out. If a bank's customer base predominantly consisted of oranges, and not so many

apples, then under a rule-based regime, the apples were the ones reported on. Continuing this metaphor, under the risk-based regime, this shifted from finding the rotten orange amongst the healthy oranges and the rotten apple amongst the healthy apples. It further seeks to separate less rotten from severely rotten fruit. Of course, this is a simplification of the difference, but it serves to make clear that the ongoing evolution of transaction monitoring is one of continuing refinement of the way we can analyze data.

The risk-based approach incorporates a number of risk concepts that are utilized in different ways in transaction monitoring. Ultimately, transaction monitoring is looking at transaction risk: what is the likelihood that *this* transaction is concealed money laundering and therefore concealing criminal funds? The risk of the transaction is calculated on the basis of algorithms that incorporate other notions of risk. See Figure 1.3.

One of these notions is the risk that the parties to the transactions represent. If these are internal parties, i.e. *customers*, then a customer risk level, most commonly in terms of high, medium, or low, is established for the party. This customer risk may itself be based on a calculative score card or be the result of an investigation into the customer. Obviously, the latter is not done for all customers – that would be impractical – but for the high-value private and corporate customers,



**Figure 1.3** Evolution of transaction monitoring.

it is not uncommon to pinpoint a precise risk level if they are teetering on the edge of high risk. This is where the notions of Enhanced Due Diligence (EDD) and Special Due Diligence (SDD) come into play. The majority of customer risk is, however, calculated based on information gathered during account opening and throughout the lifespan of the customer relationship.

Most financial institutions are also risk assessing their own products. Product risk is often a separate category of risk that feeds into customer risk or directly finds its way to the transaction monitoring risk assessment. For example, purely domestic and common products, like a normal current or tax exempted savings account are assessed to represent low risk (for money laundering) and more exotic products with a foreign currency component or products that have been known to be favored by money launderers in the past may be allocated a higher risk level.

This brings us to a third risk dimension, that of country risk. Most regulatory frameworks require financial institutions to look at country risk separately. Most, if not all, financial institutions with cross-border activity maintain risk scores or levels for each country or territory. Sometimes these are simply copied from a subscription, but the larger financial institutions commonly have processes in place to assess or even calculate the risk level for all countries and territories worldwide. Under the rule-based system the non-governmental Financial Action Task Force (FATF) maintained a list of non-cooperative countries and territories and local regulations incorporated this as a blacklist. The list still exists, but over the years, countries featuring on this FATF list made sure they improved their AML regime and were subsequently taken off.

Countries may still appear on the so-called watchlists, alongside individuals, entities, and registered vessels. Watchlists arguably can be categorized as one of the risk factors. They feed into the customer risk, although most commonly they are used during the onboarding process of a customer and appearance on the list will prevent accounts being opened for the individual or entity. But also, transaction monitoring scenarios commonly include a number of rules that specifically look at parties and jurisdictions involved in a payment that are on watchlists.

Apart from these potential high-risk components of a transaction, there is a different notion of transaction risk that ties in with the assessment of the transaction as being unusual. The underlying premise is that the more unusual a transaction is for this customer or this account, the higher the ML risk. Determining whether a transaction is unusual has become an increasingly sophisticated process, in which ever more data is drawn into the equation. A transaction can be unusual from different angles and most transaction monitoring systems look at it from multiple perspectives.

First there is the perspective of being unusual compared to the usage of the account. Many algorithms look at how the account has been used and whether one or more transactions are (suddenly) deviating strongly from that pattern. One can look at the size of the transaction and the number of transactions in a given interval (velocity), but one has to account for more gradual increases or decreases of income and spending patterns, as these are in fact quite ordinary. One must account for seasonal variations: many individual customers and most retailers use their accounts more actively during the festive season, and sudden slumps or bursts are associated with vacation periods.

Further, the behavior on the account can, and often is, assessed in comparison with similar accounts to avoid comparing apples and oranges. High-net-worth customers will have a significantly different pattern from middle- and low-income customers. Retail business accounts may show very different patterns than those of manufacturing businesses. And even within retail, a low-end retailer in a busy shopping area will show a completely different income and outgoings stream than a more exclusive high-end retailer. Most financial institutions therefore divide their customer base into segments and either have separate rules in place specific to each segment or have set threshold parameters different across the thresholds. Those financial institutions that go a step further introduce the concept of peer groups: each customer or account is allocated to a group that is expected to behave similarly. A group profile is built, dynamically or statically, and transactions and summary profiles for the account are compared to the group profile. If there is a significant deviation, then an alert will be produced.

As we have mentioned above, the move from rule based to risk based brought about a shift in responsibility and accountability. In a

rule-based system, there is hardly any responsibility for the financial institutions beyond simply ensuring the rules are applied. The price to pay was a tsunami of alerts and reports, most of them being false positives. A risk-based approach may stem the flood of false positives and reports to be filed, but at the cost of increased liability and the investment into a framework that justifies the (risk) approach taken. As we have seen above, the risk-based approach opens the field wide, as there are endless variations of the approach and countless decisions to be made. At any given time, the regulatory watchdog may review that system and assess its robustness. Financial institutions will have to be able to explain why they set up transaction monitoring the way they did. If country risk is part of the process, then it may have to explain separately how country risk levels are arrived at and what is done with that knowledge further downstream. And so this goes for all the other concepts of risk that together make up the risk-based approach. These approaches and the decisions need to be documented in a way that they explain to a relative outsider (regulatory auditors) how it all works and that it is sufficiently robust.

A lot of the mathematical logic is built into the transaction monitoring software. It is, however, the financial institutions who will need to understand the software as AML accountability cannot be deferred to the software supplier. Financial institutions are free to hire consultants from the software vendor to assist with regulatory review. However, from the perspective of regulatory scrutiny, it is always a good thing if the solution is white box as opposed to black box. Black box software conceals most or parts of its built-in logic, its inner mechanisms, most often to protect its intellectual property. The software vendor might fear that making their solution fully transparent exposes them to IP theft. For example, the customer itself may decide to end the contract and build a solution themselves based on what they have learned from the vendor's software. Not all software vendors have the same fear and, like SAS, take a white box approach and are willing to give the customer access to (virtually) any component of the system. This means that analysts, consultants, and coders can dive deep into the system to understand how and explain why the transaction monitoring system yields the result it does. Our common experience is that most financial institutions feel more at ease

with a system that hides no secrets in terms of how risk is being calculated and used and how alerts are generated.

## **From Static to More Dynamic Transaction Monitoring**

In the quest for reducing false positives and the detection of relevant, i.e. significant deviations from behavioral patterns on one or multiple accounts, the aim is to refine the mechanism by way of recognizing (more complex) patterns.

To put it simply, if every transaction is measured against the same fix threshold, such as any cash deposit of (the equivalent of) \$10,000, this would be a very crude and most likely ineffective way to identify suspicious or unusual transactions. First, such a straightforward rule seems to assume that cash deposits are more likely to be done by money launderers than by law abiding citizens. It is true that, to a certain extent, crime economies are largely cash based. This is, for example, typically the case in drug trafficking, where across the entire chain from manufacturing through wholesale to retail distribution, cash seems to be the preferred method of payment. Much of that cash is cycled in the underworld drug economy, upstream in the supply chain, but many operational costs, such as pay-out of staff, purchasing of vehicles for transport, and hiring of warehouses, are also paid for in cash. From the net profit for the criminal entrepreneur some will be stashed as cash; some of it will be lent to others in cash, and paid back in cash; and for only part of it there will be a need to put it in a bank account. But many forms of crime are not predominantly cash based, such as many forms of fraud, and there is no need to siphon off profits held in cash and make large cash deposits. Also, cash deposits may be common in some perfectly legal trades, such as for new and second-hand cars. Putting up a fixed threshold will erect a type of barrier for cash-rich money launderers, but that will also affect a bank's legitimate customers in a cash-based industry and so will come at an operational and/or commercial cost. And to what effect? This barrier will be (and has been) quickly recognized and circumvented with structuring techniques, such as smurfing. Smurfing is the breakdown of large sums of money into (well) below threshold chunks and distributing it across multiple depositors who use multiple accounts to avoid

unwanted attention. These days transaction monitoring systems will not only look for above-the-threshold deposits, but also for amounts that are just below it. This does not make the rule more sophisticated or more dynamic; it is still a fixed threshold, except that the one for internal alerting is a different one than the one that is communicated publicly. But such sub-rules are often combined with an additional threshold on a minimum number of such below-the-threshold deposits. By doing so the financial institution has taken its first step towards a more dynamic rule.

The simplest form of a dynamic rule is where the threshold is manually set and reviewed either periodically or even on a driven basis. The form brings relatively high operational costs and seems inefficient. The setting of the threshold needs to be described in some type of process to avoid arbitrariness and regulatory exposure. The second form of threshold setting is relatively rare for AML transaction monitoring and more often seen in the context of fraud and abuse incident response, and heavily driven by alert analysis and case investigation. A good example is the activation or refocusing of a geographic rule if it is suspected that a certain postcode area has become the habitat of a fraud or money laundering ring, with involvement of ATMs, or multiple accounts being set up at local branches. That postcode area will be temporarily set to high risk to allow more alerts to be created and analyzed, in order to identify the issues and mitigate the problem. The perpetrators are likely to move on to other areas, if not caught by law enforcement. Either way the risk for that area decreases and likely either the same or a new ring will pop its head up somewhere else, thus making the rule obsolete or dysfunctional without change.

Such manual dynamic rules are not common in the world of transaction monitoring, but those which adjust thresholds automatically, are. And they come in various degrees of complexity.

The *modus operandi* known as smurfing, as described above, is a good example of the need for more dynamic rules, supplementary to the fixed threshold rules, such as the one for cash deposits. To detect smurfing, one needs to zoom out from the micro event of the single transaction and take into consideration a range (or history) of transactions, preferably across a number of accounts potentially associated with a number of account holders. Let us set aside, for now, the

complexity of identifying multiple customers (the “smurfs”) with different accounts – these customers are not usually blue skinned (as the label given to them would suggest) so there is no straightforward way of telling whether they are all colluding as part of a wider smurfing scheme. There are ways of doing so, from a transaction monitoring perspective, but let us look into simpler cases first.

The simplest form of looking at transactions is dynamically. This means to look at transactions dynamically from the perspective of the account, namely the *history* of the account. One can either look for gradual or sudden changes on the account that may arouse suspicion, or for certain patterns *over time*. For the latter a lookback period is defined (usually an  $x$  number of days) and the thresholds need to be exceeded within that time frame. In the smurfing example the scenario could look like this:

*Create an alert when over the past 7 days at least three cash deposits were made with an amount between the fixed threshold value (\$10,000) and 95% of that amount (\$ 9,499).*

Although this rule would look over time, it is not dynamic, as the principal threshold value is fixed and the secondary threshold is derived from that. It would be dynamic if the percentage of the lower threshold were made dependent on the number of cash transaction over that period. For example:

*Create an alert when over the past 7 days at least three cash deposits were made with an amount between the fixed threshold value (\$10,000) and  $100 - CD(n) * 2\%$  of that amount, with a minimum of 75%.  $CD(n)$  is the number of Cash Deposits.*

The more cash deposits are made to the account the lower the alerting threshold for the cash deposit. Such a rule could be applied to a single account, or to all accounts of one customer, or even to all accounts by all customers who are connected to each other by at least one joint account. This would give a much better and more dynamic chance of detecting smurfing activity.

Dynamics of much higher complexity are possible, whereby a range of analytical measurements is applied to an account to detect

a pattern and predict within what range the next transaction will be in case of normal behavior. Transactions that exceed this dynamically calculated threshold will be considered unusual (for a particular account, given that history of the account). Such a rule, for example, would run like this:

Create an alert when at least two cash deposits exceed the amount of two times the standard deviation for *cash deposits on that account for the last 28 days*.

Whilst still looking at cash deposits, this scenario compares the latest cash deposit with the others over the last 28 days and allows for one-off spikes as well as gradual increase. Although money launderers could still avoid alerting by gradually increasing their cash deposits, they would have to do so patiently and be aware of such a rule and, preferably, its threshold parameters.

Current practice is to have such dynamic rules supplementary to the fixed rules, and the fixed cash deposit threshold rule acts as a safety blanket, whilst the dynamic rule deals with those schemes that seek to circumvent the (known) fixed threshold rule.

How this makes sense can be seen when taking into consideration a mule account operation. In a typical mule operation money launderers seek to recruit account holders and use their accounts to channel funds. These accounts are attractive to criminals as they often already have an (untainted) history and account holders are classified as low risk. Typical examples are students, job starters, and low- or middle-income elderly, who are either naively ignorant or deliberately pretend to be to the true nature of the funds channeled through their accounts. For the dynamic rule described above to be effectively circumvented, the money launderers would have to know what normal behavior is for each of the recruited mule accounts. But this information is not easily disclosed without raising suspicion. If the accounts are recruited by means of earn-money-easy-from-your-home types of job advertisement, it would be easier to trick people into receiving and sending specific funds (as part of the job) than it would be for them to obtain and share with their employers the history on their account, especially that history from before they took the job. This would easily

set off alarm bells. On a side note, obviously one of the reasons for the money launderers to recruit mules is because they do not have to worry about alarms bells going off, since certainly they would not have given away their true identity to the mule account holders.

More complex dynamic rules build profiles for accounts and potentially also peer groups (numerous accounts clustered together on the assumption that the behavior on them will be very similar) on a monthly, yearly, or even broader basis, whereby profiles are constructed automatically and expectation for the next days or months are set. Averages and transaction volumes and amounts could be set for incoming and outgoing payments and could be differentiated for payment types or channels.

The dynamic rules described so far do not consider the outcome of the alert analysis or investigation. In a way, the epitome of dynamic rules is where the line between automated parameter settings and rule creation converge. This brings us into the realm of machine learning.

## **Latest Trends: Machine Learning and Artificial Intelligence**

In the last decade the notions of artificial intelligence (AI) and machine learning have gradually found their way to virtually every domain of automation, AML being no exception. The developments can be seen as ultimately replacing the human element in the processes, i.e. completely relying on machines, algorithms, and the self-learning capability of the software. They can also be viewed as augmenting humans involved in the process, whereby data will be processed in such a way that it allows the human worker to concentrate on that part of the work that requires judgment.

When thinking in terms of augmenting or even replacing the human element in AML processes, one must distinguish between (at least) two parts in the overall business process where humans play a role: the scenario design and optimization process, and the triage/investigation process. At face value, both seem to lend themselves to machine learning. But at closer look this seems less evident.

Let us start with the scenario design and optimization part of the process. This part is where the objective of detecting suspicious activity

in the use of financial products and services – suspicious from an ML or TF perspective – is translated into scenarios and analytical models that can calculate the propensity of certain behavior based purely on the data that is fed into the system about the account holder, the account, the financial transactions related to that account, and the non-financial transactions (changes) to both account and account holder.

Transaction monitoring for AML purposes is, in its core, an industrial application of forecasting. Forecasting is a field in analytics whereby data is used to set up and train models to optimize the output of these models in terms of the desired end result. For money laundering this end result and the optimization thereof are, as discussed above, to maximize the hit-rate whilst at the same time minimizing false positives.

Analytical software, such as that developed by SAS, already assists analysts in analyzing data and training the models. This software, in the right hands, can be used to analyze data, identify statistically relevant indicators, and create new rules. Artificial intelligence (AI) is these days usually associated with neural networks as the analytical method that, backed by sufficient computing power, allows machines to become autodidacts. Neural networks as an advanced analytical technique made its introduction into the world of analytics some time ago. It has also entered the realm of applied analytics in AML.

Predominantly analytical software is used to optimize the performance of both new rules, through initial threshold settings, or existing rules, e.g. by doing what-if analysis on slight changes of the threshold settings and what is called below-the-line analysis. We will discuss these in more (technical) detail in following chapters. From an analytical perspective rule optimization is in fact the training of an analytical model. Training of models is either done supervised or unsupervised. Supervised training means that data is added that captures the outcome of the analysis in terms of something being either correct or incorrect. In training AML transaction monitoring models, the correct outcome is a productive alert; the incorrect outcome is an unproductive alert. Instead of alerts, one could also look at cases and the conclusions of the case investigation: submitting a regulatory report would count as the correct outcome; dismissing the case and closing it without follow-up action would be the incorrect outcome. The point

is that these outcomes can tell the system what to look for specifically and the forecasting software can use that to find strong correlations in the data that are indicators of money laundering, or at least productive alerts. Forecasting software, which has been in use for more than a decade in the world of money laundering, is essentially a form of machine learning, whereby the outcome is fed back into system that (re)trains the models and that can suggest new rules or new settings to rules to optimize its productivity (which is the ratio between productive alerts versus non-productive alerts). If the feedback loop would be closed, meaning that alert and case disposition data would be automatically fed back into the system, and if the system would automatically retrain itself periodically and tweak the threshold parameters on the fly, then in fact this would be full-on machine learning.

However, financial institutions are reluctant, for good reasons, to close the feedback loop and automate this part. At the moment, even for those financial institutions that apply advanced analytics to train the models, rule optimization is still a periodic human-driven effort. The (re)training software is not considered sufficiently sophisticated to let the system be in charge of running multiple analyses *and* let the computer decide which set of threshold parameters work best. And probably the regulator who has to assess the robustness of the approach would refuse to sign off on a closed loop mechanism where a computer can tweak the detection scenarios without human interference. Instead, financial institutions choose to periodically engage in an exercise to manually retrain the models. The same analytical software is being used and the same methods applied, but all driven by analytical experts who evaluate the outcome together with domain experts. The decision to change threshold parameters, to deactivate a non-productive rule, or to introduce a completely new scenario is still made by humans.

Even if financial institutions were to move to a closed feedback loop system, transaction monitoring would still not be fully automated to the extent that human judgment would be fully taken out of the equation. The machine would still train itself on the basis of the supervision data, which is a column in the data set that tells it if the outcome is a correct or incorrect one. That final verdict for each alert or case, whether productive or non-productive, is being given by a human. It is the AML analyst, whose job is to look at alerts and the

data contained in them, who has to decide whether to discard it as a false positive or continue investigating until a decision can be reached as to whether to file a regulatory report or not. The keystone holding the metaphorical arc of transaction monitoring together is the human analyst assessing the system generated alert as a true or false positive. And the same applies to below-the-line analysis where periodically thresholds are lowered to allow for the generation of alerts in situations that would not have been generated under the actual settings. These alerts are scrutinized by human analysts, just to see if the actual settings do not miss too many false negatives, i.e. situations that are indicative of money laundering but were not scooped up in the net of the thresholds.

Ultimately, one can foresee the in-house analyst's judgment being replaced by the external feedback from the recipient regulatory body. As long as the recipient FIU does not communicate to the financial institution that it receives too many false positives, or periodic reviews and investigations do not bring to surface situations where reports should have been filed but weren't, the financial institution could continue to trust their self-learning software to do a proper job. This day still seems far off.

Thus, whilst notions of machine learning and AI are not new to AML transaction monitoring and will claim an increasingly prominent position, there will still be a need in the foreseeable future for data analysts developing and training the models and optimizing the rules, and domain analysts investigating and assessing alerts as true or false positives.

## **Latest Trends: Blockchain**

Some say that what the internet and the World Wide Web meant for the online exchange of information, blockchain will mean for the online exchange of value. If true, one can hardly overstate the importance of blockchain. And what will be its impact on money laundering and the software we offer to help our customers combat money laundering?

Blockchain is a technology and a process to decentralize and harden a ledger. Perhaps this is better explained by what blockchain

is a reaction to. Up until blockchain, online or electronic transactions were always mitigated by a central guardian of the ledger, such as the bank holding your account or the clearing house processing batches of payments. What money sits where and goes from whom to whom is kept record of in a secure file or database, providing a single point of vulnerability to malicious actions. Tampering in some way with the one single electronic ledger would suffice to commit payment fraud. And it also requires a central agent who has de facto power over the transactions for all of its customers. Customers have to trust the banks. Blockchain seeks to mitigate the risk of tampering and take out the central agent. At its very core, blockchain seeks to provide a safe way of transferring value and keeping records in a world where trust is not presumed.

Blockchain technology rests on five principles:

1. The ledger is public: verified copies are freely downloadable.
2. It is distributed: many synched copies of the ledger exist across the world.
3. There is consensus: for transactions to be added (hence executed) a majority or threshold consensus must exist amongst the miners who verify the ledger with the new transactions and those miners are competing and not colluding.
4. There is transparency: the ledger will contain a full history of all transactions.
5. There is ownership: one has to own credits before these can be transferred, going short is not an option.

Ownership verification is possible due to the transparency of the full ledger. A batch of transactions to be added to the ledger is presented as a new block to be added to the existing ledger, the chain with the full history of transactions of the same currency. A block (of new transactions) will be presented to be added to the chain and each block contains a cryptographic puzzle based on the contents of the block but also of the entire ledger. This puzzle needs to be solved and the block will only be added to the chain if a predetermined number of checkers, called miners, agree on the outcome. Miners are self-subscribing participants that utilize vast computing power to solve the

puzzle as quickly as their machines allow. The first who comes up with an answer that is similar to the answer found by the others, will get an incentive. For Bitcoin, for example, the first to solve the cryptographic puzzle, upon verification, will unlock a Bitcoin. Hence the mining metaphor. The more complex the cryptographic puzzle, the more computing power (and time) is needed to solve it. Recently it was estimated that the energy required for one Bitcoin transaction could provide electricity for almost 9 US households for a day. At the moment, Bitcoin has set the complexity of the encryption, so that on average every 10 minutes a new block is added to the chain. Once a new version has been confirmed it will discard any variations and all copies are synched to the one new version. With the next block presented this cycle repeats itself.

Blockchain is safe because it requires many miners to confirm a version before the chain (and all of its copies) is updated. This means that any alterations by fraudsters will have to be on the majority of copies, otherwise it will perish in this distributed verification process. Since there are many miners and these are unrelated and unlikely to collude, it is more difficult for fraudsters to tamper with the ledger. Also, computer glitches, creating different versions of the ledger are ironed out this way. Blockchain is also safe because the full ledger is public (anyone can obtain the latest copy online) and unchangeable. Any change to historic blocks will create a deviating version, meaning that the outcome of the cryptographic puzzle will change, and that version will not survive.

Blockchain is believed to revolutionize payments, or even broader, the exchange of value. Blockchain enables peer-to-peer transactions in a world where trust is lacking. It will also allow taking out the middleman, the financial institutions that currently operate as the guardian keepers of their own ledgers (as a global collective) process all transactions. Financial institutions fear that blockchain will do to them what online retail commerce has done to many brick-and-mortar stores. Blockchain, however, may be unstoppable and banks are now exploring ways to incorporate block chain into their ways of doing business. Nasdaq has opened a new trading platform, Linq, specifically for blockchain-based currency. In other words, blockchain is here probably to stay and it will be transforming the financial system;

hence, it will be transforming AML and Compliance. If blockchain is relevant for our own customers, which in the space of AML are mainly banks, then sooner rather than later it will become relevant to those providing software to support them.

The impact of blockchain on AML is threefold. Firstly, on a strictly technical level, it will transform the way (part of) the transactions are fed into the AML system. Rather than an extract from the core the chain itself will be offered. Would this really change the way data is loaded into our system through mapping and ETL? Not necessarily. Does it mean that there will be no limits regarding the history, as the blockchain will be available in its entirety? I am not convinced. For data management and volume and related performance reasons, it would be preferable to keep the flow of data constant and increase will only be permitted if related to organic growth (more transactions on a single day). It would almost make transaction monitoring easier, as blockchain itself has an inherent data quality control aspect to it. From a data integration point of view, blockchain's impact is probably limited to the data pertaining to the actual transaction, and not so much the other data, such as for accounts, customers, or households.

Secondly, there are many who fear blockchain will undermine the current AML effort, mainly because it seems to facilitate anonymous transactions. Whereas an intrinsic part of blockchain and blockchain safety is in its transparency, this is only relative transparency. Only within the chain can one see which transactions involved which accounts. Accounts are coined wallets, and are in fact secured IP addresses. Some cryptocurrencies allow for people and entities to open an account without proper identification or verification of the true identity behind those who own or operate the wallet. In a way this is not different from the, now forbidden, number accounts operated by Swiss and Austrian banks until a decade ago, or from the still existing practice in some tax havens where financial institutions are legally protected from disclosing the identity of the ultimate beneficial owners behind the accounts or corporate entities. In a way this does not seem an issue created by blockchain technology, but by how we use this technology. As a federal prosecutor stated in her Ted Talk: "most successful technology is often first adopted by criminals." Like with the conventional financial system, blockchain providers should

ensure KYC at account opening. Like with conventional banking, the ledger itself may contain customer references that make the customer non-identifiable, as long as they keep these identities and can map them. It is a matter of legal debate and, if necessary, legislation to make sure blockchains used for financial purposes and cryptocurrency are considered or brought under the umbrella of regulatory governance.

Thirdly, naturally following on from the above, the transparency of the blockchain and the inherent security may also benefit AML. Granted, this transparency and the hardness of the public ledger primarily benefits fraud prevention and only to a lesser degree AML. Irregularities in the data loaded into the transaction monitoring system have never been a main concern: the correctness of the transaction data is almost never questioned, at least from an ML concern. Having the full history of the entire ledger freely available obviously will help investigators, both on the side of law enforcement and on the side of the financial service providers to trace back money to their original accounts. From a data visualization perspective, the full ledger seems to take care of data preparation and is probably easier than getting the same data over a similar period of time together from traditional core banking systems. The capacity to process, mine, and visualize Big Data may become a key factor, simply from a sheer data volume perspective. It is widely known that blockchain crypto mining requires vast amounts of computing resources, but that is primarily due to the way the decryption algorithms work. AML data analysis of the ledger is probably less resource intensive, but the sheer size of the ledger may cause issues. One possible solution could be not to analyze the ledger in full, but use it as source data and load the relevant data (or delta) into the core database as is currently common practice.

### *Should the Compliance Domain Prepare for Blockchain?*

The answer is an unreserved Yes. There are too many voices raising blockchain and cryptocurrency as a game changer, for the financial industry to ignore. Customers are already asking what we can do for them in terms of blockchain, also within the subspace of AML and financial crimes. At least we should prepare ourselves by starting to understand blockchain and the impact it will have on existing and

new (types of) financial institutions and the software used. Banks are now starting to explore blockchain technology, and regulators will also get involved in the conversation sooner rather than later. If anything, blockchain will present the Compliance domain with an opportunity to combine transaction monitoring analytics with big data and meaningful data visualization. AML, pushed by blockchain, will move into a space where multiple disciplines will come together . . . and right up the sleeve of data analytics.

### **Risk-Based Granularity and Statistical Relevance**

For statistical analysis in general it applies that the more data, the better. This is particularly true when the data is structured. Financial institutions, even those of smaller size, produce high volumes of data, especially around their core commodity services and products. However, data density is not always equally spread. Large financial institutions stretch their arms widely and may also provide products and services tailored to niche markets, servicing only small and selective groups of customers. As a result, the statistical analysis may become less meaningful and lack statistical relevance.

In addition, the shift from rule- to risk-based system has introduced the notion of profiling: to determine if a transaction or a pattern of transactions is unusual, one must make sure the customer is compared against like customers. This can easily backfire; the more a financial institution knows about the specifics of a customer, the more it knows to what extent this customer can or cannot be meaningfully compared with others. Customer segmentation is necessary to meet the requirements of a risk-based approach, but a profiling that is too granular will come at the cost of data volume, to the extent that statistical relevance will become an issue.

A good example is trade-based money laundering, whereby price levels or volumes of traded goods are deliberately overstated in the contracts and not aligned with the underlying commercial reality, if there is one: the trade can also be wholly made up. This requires enough reference data that is comparable. In a trade of goods, a unit price is affected by type of goods (of course), country, season, expiry date, unit, and also trading contract terms. With all these parameters,

it can hardly get enough data for reference. In one situation we studied two-year trade transactions. The largest few groups of unit price (same type of goods, same country, unit, etc.) had only 200 data points, not satisfactory for a statistical analysis; most other groups had much fewer data points. The move from a rule-based to a risk-based AML framework can also be reconstrued as a move from a strict formal know-your-customer to a genuinely know-the-business-of-your-customer. On the one end of the (rule-based) spectrum, financial institutions merely had to establish the veracity of a customer's identity (or the identity of their registered controlling persons) by way of a relatively superficial check of the ID documents. Since its early days KYC and ID verification as part of that has taken a flight of its own, which we will not discuss in this book. At the other end of the spectrum there is the notion that a financial institution, mainly through its account managers, knows all financial ins and outs of a business or an individual that they provide services for . . . and in case of any suspicion of irregularities the account manager should contact the Compliance or Security department. This turns every customer-facing employee into a part-time AML enforcement officer. This is very much a human-driven approach, effectuated through training, awareness, internal policies, and codes of conduct. Theoretically, there could be a data-driven approach if sufficient data would be gathered. This could be done if AML were to include external data, e.g. on demographics and/or if financial institutions in a country would share data. However, the development to include external data for improved statistical accuracy is not driven by financial institutions as long as the law and regulations do not require them so. Other than in the field of fraud prevention, in AML there is no direct incentive for financial institutions to further increase the complexity of their AML operations beyond what at any given point in time suffices to meet regulatory demand. Even if data would be obtainable or become available, restraint may be required.

From a data-driven transaction monitoring software design and delivery perspective, commercial software vendors have to serve the interest of their commercial customers and optimize a system within the (risk and cost) appetite of the financial institution. Ultimately the financial institution, not the software or its vendor, is held responsible for compliance with law and regulations. It is therefore the financial

institution's decision how and what to implement for the purpose of detection of money laundering in transaction monitoring.

## Summary

With a history of over three decades, AML has made a long journey that has not yet come to an end. This journey has been one of learning and mutual education of the regulator and financial institutions governed by the regulatory framework.

In the childhood years of AML, simple rules based on known money laundering *modus operandi* associated with specific forms of crime led to an extremely high and unworkable number of false positives and Compliance departments and FIUs alike were overburdened and operating neither effectively nor efficiently.

The gradual shift to a risk-based system, which is still ongoing, opened the doors widely for a data-driven approach, where computer-driven analytics provide the answers to ever bigger data analysis and meaningful transaction monitoring and customer risk assessment, whilst at the same time reducing operational costs to an (already very high) minimum.

Whilst human judgment, for the foreseeable future, will remain to be an important factor, financial institutions and software vendors alike are introducing advanced technology in an increasing number of steps on the journey from transaction monitoring and regulatory reporting. Automated computing is not only introduced more widely across this chain, but also embedded more deeply than before.

Models, and how these are trained, are becoming ever more sophisticated, enabled by a parallel growth of computing power (at a relative low-cost increase or even decrease) and the ability to process higher volumes of data (we're talking billions of transactions these days) at higher levels of sophistication.

Financial institutions seek to meet ever more stringent regulatory requirements and avoid the huge regulatory fines and maintain acceptable levels of operational costs and resource demand. And they do so by investing in AML transaction monitoring and case management software that allows them to go bigger, better, faster . . . and still enable them to explain to the regulator that (and, equally important,

*how*) they are doing a proper job as gatekeeper, protecting the financial system from (at least the most blatant) abuses by money launderers.

With the emergence of artificial intelligence and (deep) machine learning and with the introduction of cryptocurrency and blockchain, a new era is beginning for AML Compliance. Rather than invoking a paradigm shift in how we monitor transactions and risk-assess our customers, these developments, from a software manufacturing and implementation perspective, may well be a matter of scaling and absorbing the latest and greatest from the field of data management and analytics.