DIGITAL TRUST

Michael Bonaventura *explains why digital fraud is on the increase and what can be done to prevent it*

It's a sobering thought that spending on eCommerce sites in 2021 will be around $5 trillion. More broadly, the Financial Action Task Force (FATF), the intergovernmental organisation tasked with developing policies to combat money laundering, estimates 60 percent of world GDP will be digitised by the end of 2022. Like all markets, these digital markets are built on trust. Without it, merchants and individuals are exposed to a wide range of risks from the obvious financial ones associated with creditworthiness and fraud to the less obvious risks of social and economic exclusion.

In the material world, trust is developed over time. It begins with an introduction, often through a common acquaintance or a recommendation. The strength of the relationship is determined by the way that people subsequently behave towards each other. By contrast, in digital commerce, transaction volumes and the desire for almost instantaneous response have marginalised human involvement; automation has replaced much of what was previously done manually, further accelerating innovation. Technology has changed the nature of our relationships.

Trust has been reduced to an algorithm driven by proof of identity, which currently remains heavily reliant on formal documents. National identity cards and drivers licences provide proof of who we are; payslips and bank statements, proof of what we earn; and utility bills, proof of where we live. As digital identities emerge, these too will be reliant on documents for their validation and verification.

Anyone looking to misrepresent who they are, where they live or what they're paid would need to reflect their fiction in their documentation. And, highly automated workflows are particularly vulnerable to this type of manipulation. Preventing document forgery therefore becomes crucial for maintaining the integrity of the market. Unfortunately, that's impossible.

Blending real and fictitious identity fragments to create synthetic identities is notoriously difficult to detect and constitutes some of the most damaging fraud with losses projected at $48-billion by 2023. Criminals using synthetic identities are often playing a long game, building up a credible and legitimate credit history over months and sometimes years before executing one big ticket fraud and disappearing for good. And, some people who start out honest end up not so, perhaps as a result of changes in their personal circumstances over which they lose control. Moreover, forgery is only one type of document fraud: counterfeits, fraudulently obtained genuine documents and genuine documents misused by an impostor are almost certainly undetectable.

Trust has been reduced to an algorithm driven by proof of identity, which currently remains heavily reliant on formal documents

So, if document forgery cannot be prevented, can the losses be reduced by protecting automated workflows without sacrificing the convenience they offer for the customer experience? Recent advances in AI, especially machine learning [ML], would seem to offer a potential solution.

BESPOKE DEVELOPMENT

Some elements of documents, electronic signatures for example, can be validated against legitimate examples of those signatures using freely available, open source, machine learning libraries. Countering more sophisticated forgeries requires more bespoke development. Detecting manipulation of documents with graphics editors or 'print-manipulate-scan' evasion techniques requires specialist knowledge of the metadata and digital footprints left by scanning and printing devices. These ML techniques work with most document types and deliver around 75 to 80 percent accuracy with minimal training. Good, but hardly sufficient; even more specialised modelling is generally required.

Visual and structural modelling can be used to assess the look and feel of specific types of documents provided by third parties, comparing them against examples of authentic documents provided by document originators – banks, utility companies, government agencies. Taken together, these techniques probably catch most document forgery… but not all.

The consequences are apparent in the 5 percent of loans that are underwritten every year against forged documents where every dollar lost incurs a further $3 in fees, labour costs and recovery expenses; in the 6 percent of global healthcare spending lost annually – around $500-billion in 2020; and in the 3 to 4 percent of fraudulent insurance claims costing carriers between 5 and 10 percent of their annual revenues.

The problem with establishing trust based on documents alone is that it simply provides a snapshot of identity. With the advent of commercial-grade AI and machine learning capable of modelling behaviours, we have an opportunity to film the movie. In the process we can start to develop more holistic models of people (customers), not just identities.

In financial services, retail, gaming and similar oriented services, once identity has been established in the onboarding process, day-to-day activity in an account is subject to on-going monitoring for fraud and laundering. But, AML compliance is imposed by the regulator while fraud prevention is designed to protect the brand and its customers from financial damage. Consequently, these specialities have developed independently of each other: two teams, frequently physically separate, with no common shared data model, still no single view of the customer.

The original fraud prevention and AML systems were rules based. Any alerts generated by those rules were resolved by analysts. While any additional information generated or used in the resolution of those alerts may have been used to improve the rules, the detail was lost once a transaction had been approved (or not). Risk modelling never really evolved, missing the opportunity to create a more holistic picture of the customer that captured both their identity and their behaviour.

FOR HIGH AND VERY HIGH-RISK CUSTOMERS, VERIFICATION IS MORE FREQUENT AND FORENSIC

The next generation of systems using simple machine learning did little better, for the most part only automating the labour-intensive maintenance of the rules themselves. Where they did bring innovation, the models were typically 'black boxes', making it difficult for analysts to understand the rationale for the decisions being made and, therefore, unacceptable to the regulator. The opportunity to radically transform risk management with a systemic approach that modelled risk holistically, remained stubbornly out of reach.

In an ideal world the goal is to establish that a customer is who they say they are and provide them with their requested access to goods and services as quickly as possible. Any decision-making should be transparent and explainable. A more adaptive approach is therefore required to find a balance between accessibility and security, simultaneously minimising the merchant's risk while optimising friction in the customer journey on a case-by-case basis – let's call it adaptive friction.

Adaptive friction is underpinned by 'context aware' machine learning – the third and most recent generation of risk management and monitoring systems.

CONTEXT MATTERS

The models used by these systems make decisions in context, not in isolation. Each new customer interaction is assessed by considering all previous interactions. The more data, the more accurate the assessments become. Different models are used concurrently to detect different types of behaviours: overlapping identities, account takeover, hoarding, basket switching, layering, integration, *etc*.

Like first-line fraud analysts, contextual analysis scores customers across a risk spectrum whose breadth and thresholds can be set to reflect a company's risk appetite. When an account is first opened, following the KYC checks that establish their identity, a customer is assessed as 'medium' risk by default, somewhere towards the middle of the risk spectrum. Their behaviour is then continuously monitored and their risk rating adjusted accordingly based on the models' understanding of the characteristics of different types of risk.

As their risk decreases, an increasing range of services and activities can be made available to them, cross and up-selling without further inspection. If a customer explicitly requests access to services incompatible with their risk profile, additional step-up validation can be applied. If any behaviours increase their risk score, the customer then becomes subject to further inspection. Any criminal behaviours that are identified place them in the highest category and access to their services is suspended.

EACH NEW CUSTOMER INTERACTION IS ASSESSED BY CONSIDERING ALL PREVIOUS INTERACTIONS

The vast majority of customers will quite quickly be scored at the low or very-low end of the risk spectrum and, if they behave reasonably, will never interact directly with fraud or compliance teams. For the medium-risk customers, the process will remain the same as it was initially, with occasional manual checks before high-risk products and services are offered, much like the process we are used to now. For high-risk and very high-risk customers, the verification process will be more frequent and more forensic. Malicious behaviours will be identified sooner, due to both continuous risk assessment at the customer level, or more detailed monitoring, identifying new categories of threats. Contextual reasoning has two immediate benefits. First, the models mirror the behaviour of analysts who evaluate alerts and close most of them immediately – in context-aware systems any alert that is likely to be a false alarm is immediately de-prioritised. Second, for financial crime to be cost effective criminals must be able to replicate transactions at scale, typically involving multiple stolen, invented or synthetic identities. Contextual machine learning models not only make countermeasure evasion difficult, they can also identify behaviours that may not have been seen previously, but which can be recognised as potentially malicious nonetheless.

Ongoing behaviour monitoring contributes to the realisation of continuous KYC. The customer models so created can potentially be used not just by fraud analysts and AML teams – where anomalous behaviours would quickly highlight account takeovers or money mules for example – but also by sales and marketing to improve the customer experience, helping them to make offers that are more likely to convert.

Lack of information is a great leveller. Without any distinguishing information everyone looks the same or at least they are treated that way. This makes for a frustrating experience for legitimate customers and financial crime analysts who spend a disproportionate amount of their time evaluating what turn out to be low-risk applications. Linking identity with behaviour through contextual machine learning promises to transform risk management by delivering a platform for continuous KYC using models that are representative of the many different types of customer that exist in the material world, each with their own individual risk profile ●

Michael Bonaventura
is an analyst at Resistant AI

In an ideal world the goal is to establish that a customer is who they say they are and provide them with their requested access to goods and services as quickly as possible