Driving the acceptance of AI for financial crime prevention

As financial institutions continue to leverage AI models for the detection of financial crime, they will need to be prepared for the new regulatory and operational challenges that await.

Model Edge

A PwC Product

Contents

Introduction

The Financial Services industry has long used machine-learning for the purposes of financial crime prevention (such as for the detection of fraud), with more recent adoption being observed to bolster anti-money laundering (AML) programs. At larger institutions, the number of models used as part of Financial Crime Compliance have been increasing at an annual rate of 10-25%[1] with parallel increases in model complexity. Much of these trends can be attributed to decreasing costs associated with computational power and data storage which have made artificial intelligence (AI) technologies more accessible than ever before. Innovations in AI are also changing how financial institutions approach and tackle Financial Crime Compliance, such as through the use of non-traditional data like images, scanned documents, and biometric signals, to name a few. This has helped enable models to bolster financial crime decisioning throughout the customer and account life-cycle rather than just at a transactional or product level. In a recent PwC survey[2], 85% of CEOs stated that they believe that AI will change the way they do business in the next five years. Taking into account the high costs associated with Financial Crime Compliance - up to $1.2B annually for large institutions[1] - and the opportunity to improve on the estimated 75-90% of Financial Crime Compliance model-related false positives, it is clear to see why financial institutions are increasingly turning to innovative methods that incorporate machine-learning in order to help improve model performance and efficiency.

However, the transition to AI-enabled financial crime prevention can be a challenging process. In a recent PwC survey, 84% of CEOs indicated that AI-based decisions need to be explainable in order to be trusted. In the highly-scrutinized regulatory environment of Financial Crime Compliance this is especially true. Institutions beginning the process of implementing machine-learning and AI approaches for financial crime prevention face many new regulatory challenges with interpretability, responsible use, model validation, and ongoing performance monitoring, among others.

85%

of CEOs stated that they believe that **AI will change the way they do business** in the next five years

[1] The evolution of model risk management
[2] PwC 23rd Annual Global CEO Survey

Introduction

The sections that follow are aimed at helping institutions to think through the common challenges and considerations associated with AI adoption in the quest to fight crime.

- **Enable Success by Generating Buy-in from the Start -** the more successful programs utilizing AI did so by building and cultivating relationships between business and technology stakeholders early on in the process and by continuously promoting engagement - doing so can result in effective and efficient models and an accelerated and transparent model design to deployment process

- **Promote Model Explainability Through Informed Modeling Decisions -** model outputs need to be interpretable and explainable to be able to satisfy regulatory examination and model risk committee review. In order to succeed, modeling teams should document decisions made throughout the model development and testing process while also providing supporting evidence - for example, when evaluating and selecting model features

- **Up-to-date Model Validation is Essential -** programs beginning to consider AI for Financial Crime Compliance for the first time should confirm alignment across the model development and validation functions. Without alignment it will be difficult for programs to maintain consistent and explainable model documentation and for model validation to provide a true effective challenge

- **Ongoing Monitoring Practices Need to Keep Up -** traditional model performance monitoring processes still have a role to play, but they no longer paint a complete picture. Ongoing monitoring of AI models should consider more than just basic alert volumes and escalation rates. Depending on the type of model being employed, institutions should also consider monitoring model output distributions based on assigned scores or probability estimations, among other potential metrics

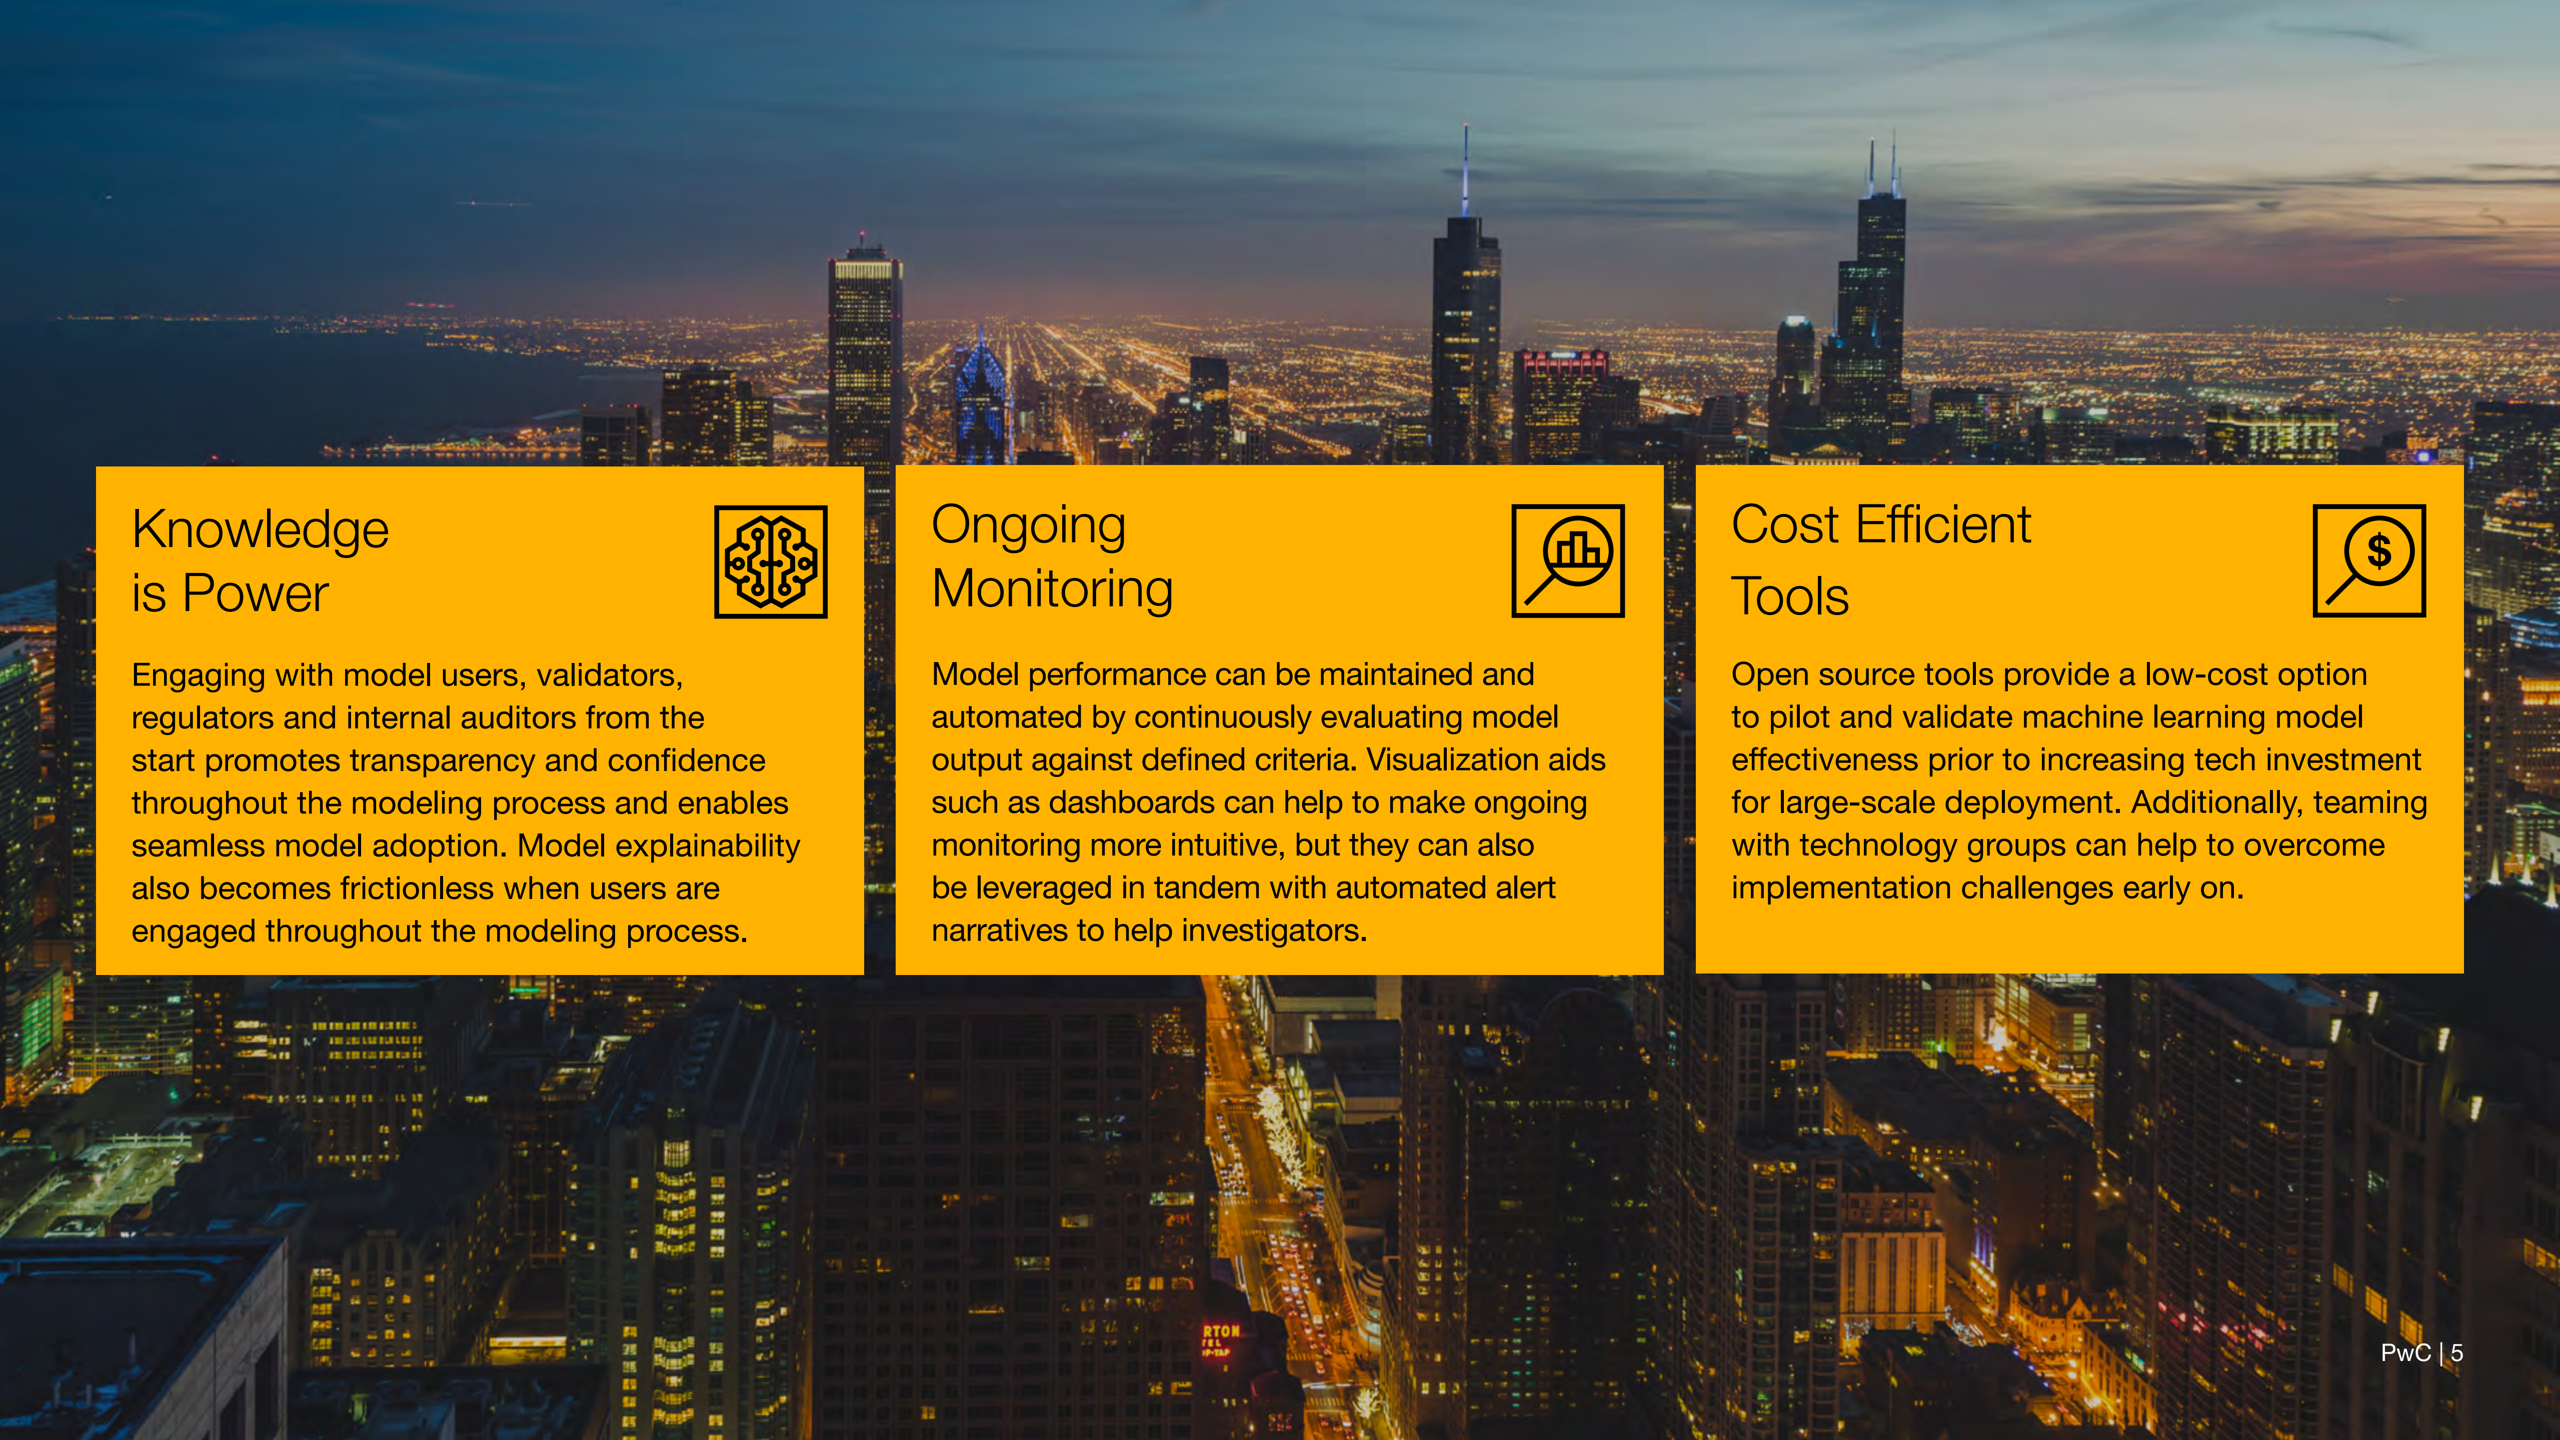

Knowledge is Power

Engaging with model users, validators, regulators and internal auditors from the start promotes transparency and confidence throughout the modeling process and enables seamless model adoption. Model explainability also becomes frictionless when users are engaged throughout the modeling process.

Ongoing Monitoring

Model performance can be maintained and automated by continuously evaluating model output against defined criteria. Visualization aids such as dashboards can help to make ongoing monitoring more intuitive, but they can also be leveraged in tandem with automated alert narratives to help investigators.

Cost Efficient Tools

Open source tools provide a low-cost option to pilot and validate machine learning model effectiveness prior to increasing tech investment for large-scale deployment. Additionally, teaming with technology groups can help to overcome implementation challenges early on.

Generating Buy-in from the Start

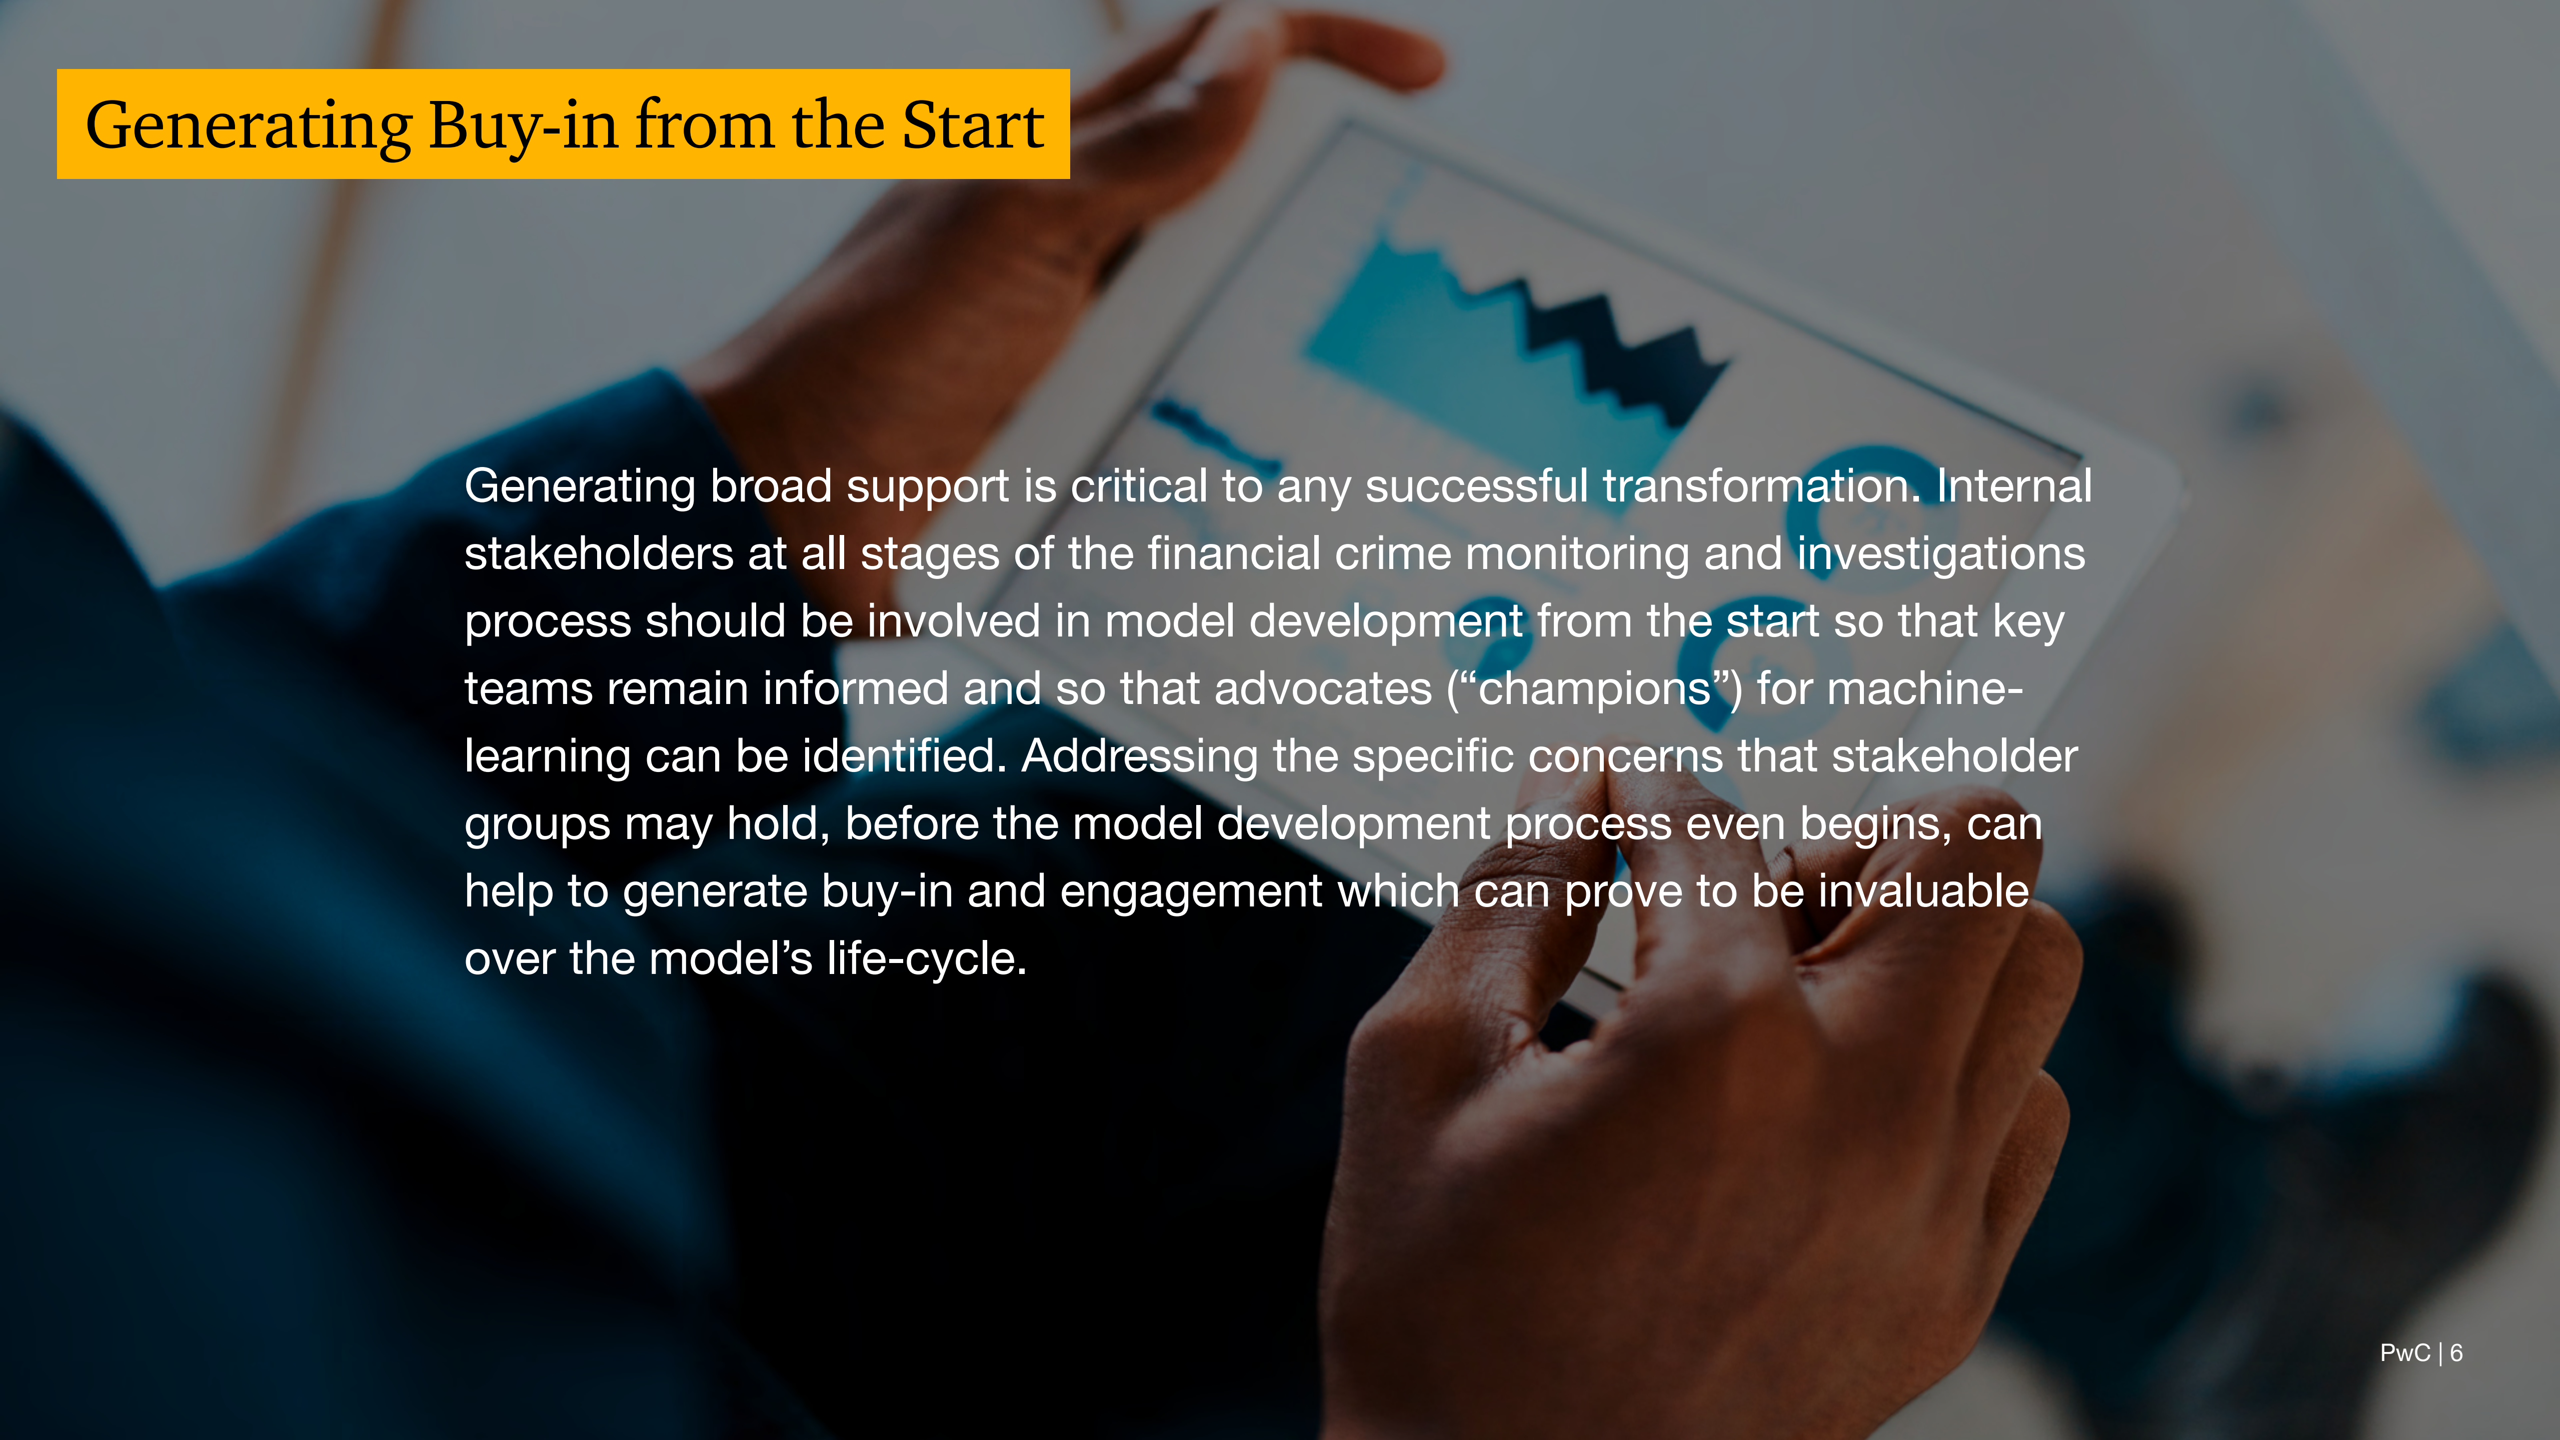

Generating broad support is critical to any successful transformation. Internal stakeholders at all stages of the financial crime monitoring and investigations process should be involved in model development from the start so that key teams remain informed and so that advocates ("champions") for machine-learning can be identified. Addressing the specific concerns that stakeholder groups may hold, before the model development process even begins, can help to generate buy-in and engagement which can prove to be invaluable over the model's life-cycle.

Model Risk Committee

Model Risk Committee and AML Committee will be primarily concerned with the potential for new gaps in risk coverage.

A holistic assessment of risk coverage provided by existing models should be performed to address these concerns. Machine learning models can be trained on historical data to confirm that existing coverage is maintained. Once existing coverage is confirmed, conversation can shift to benefits, such as the expansion of coverage that machine learning provides.

Investigators and FIU

Investigators and the FIU may be concerned about the quality of generated alerts and effort required to investigate alerts.

Institutions implementing ML have seen large reductions in false-positive alert volumes. Since ML models tend to consider more data and context when interpreting risk, holistic justifications of alerts can be automatically provided to investigators. Natural Language Processing can be used to generate alert narratives, making alerts more interpretable and to help reduce investigation times.

Organizational Leadership

Organizational leaders may have concerns about technical capabilities and investment costs.

Teaming with technology groups can help to overcome technology barriers, since many organizations will be able to leverage existing investments in data science platforms made for other business units, such as credit risk. Open source tools can also help reduce the cost of pilot programs, helping to quickly demonstrate the value of machine learning without significant up-front investment.

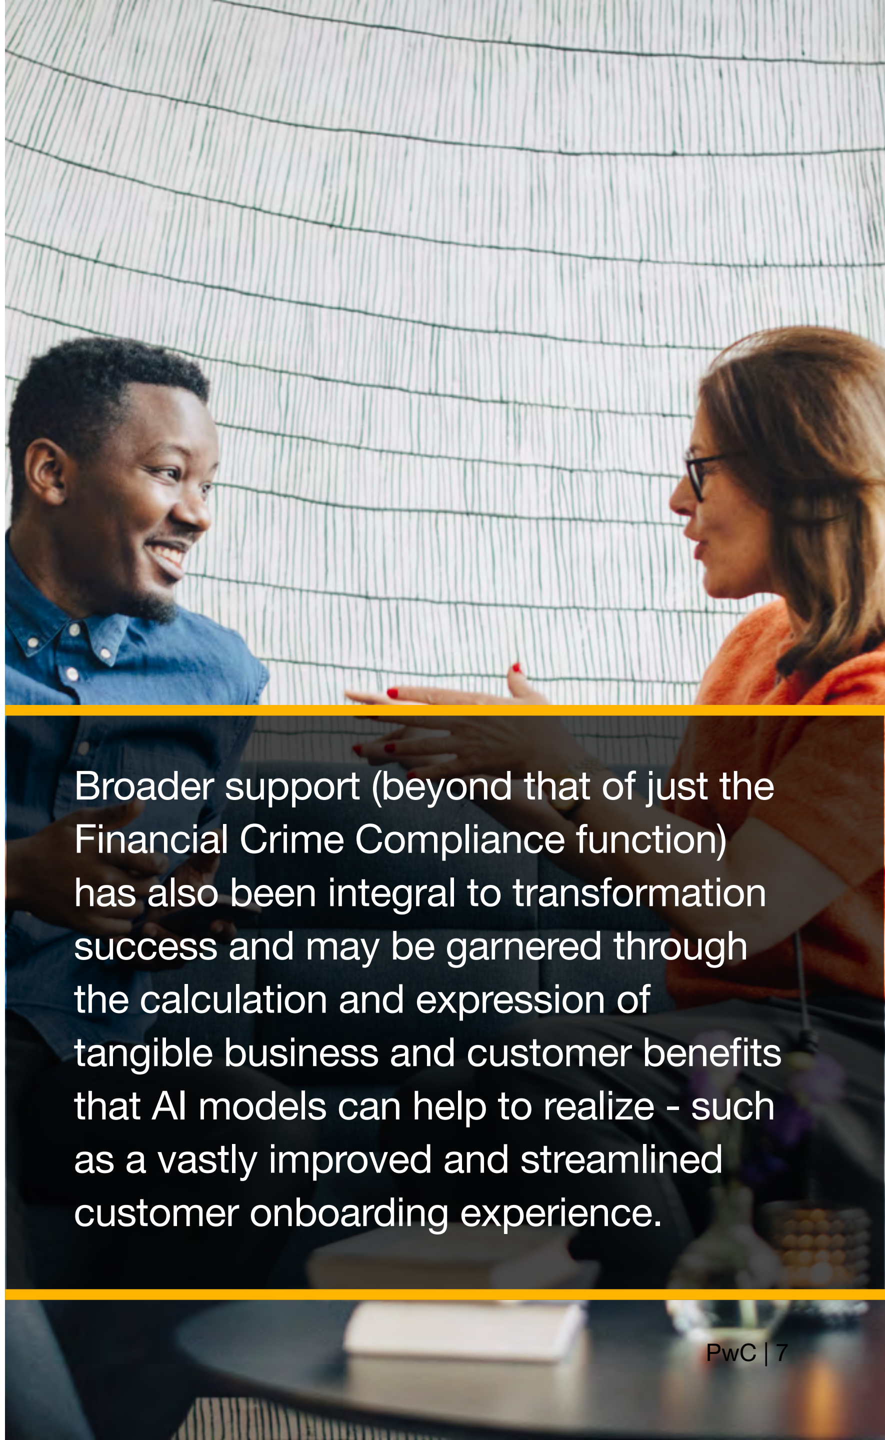

Broader support (beyond that of just the Financial Crime Compliance function) has also been integral to transformation success and may be garnered through the calculation and expression of tangible business and customer benefits that AI models can help to realize - such as a vastly improved and streamlined customer onboarding experience.

As machine-learning and AI have been used in fraud prevention and detection for more than 25 years, there are increasing opportunities to integrate fraud models that leverage AI into broader business processes that help drive better customer experiences, operational efficiencies, and increased business revenue. An example of this can be demonstrated by improvements to the fraud screening process that takes place during new account onboarding. Through the use of biometrics and image analysis, coupled with traditional fraud risk scoring, a potential customer can open a new banking account by submitting a selfie along with a picture of a government ID to fully verify their identity - all within a few minutes. This type of transformation alone has helped to provide clients with more account pull-through that has resulted in potential increases in product revenue to the portfolio by 20% to 30%*.

Open-source vs. Vendor Solutions

Teaming with internal technology groups can help to overcome implementation challenges early on. Many functions within a financial institution have already made investments in AI, and they are always on the lookout for additional use cases to help justify the costs to-date.

In lieu of internal technology group partners or sponsors, open source tools can provide a low-cost option to pilot AI models in order to prove effectiveness and to demonstrate that the approach is valid. Once broad support is established, the path to larger tech investments can become easier. Even so, most vendor solutions (such as analytics platforms) may help to facilitate productionalization and data aggregation, but they will typically not offer any algorithms or packages that are not likely to be found for free in the public domain.

AI Models Stealing the Show

Institutions transitioning away from the use of rule-based systems in favor of AI models for AML surveillance of non-traditional products, such as correspondent banking or trade finance, have observed reductions in false-positive alert volumes by as much as 95%* while also observing equally significant increases in risk coverage and relevant SAR filings. Some institutions have even observed alert to SAR escalation rates improve from as low as 0.5%* to over 15%* after implementing machine-learning models.

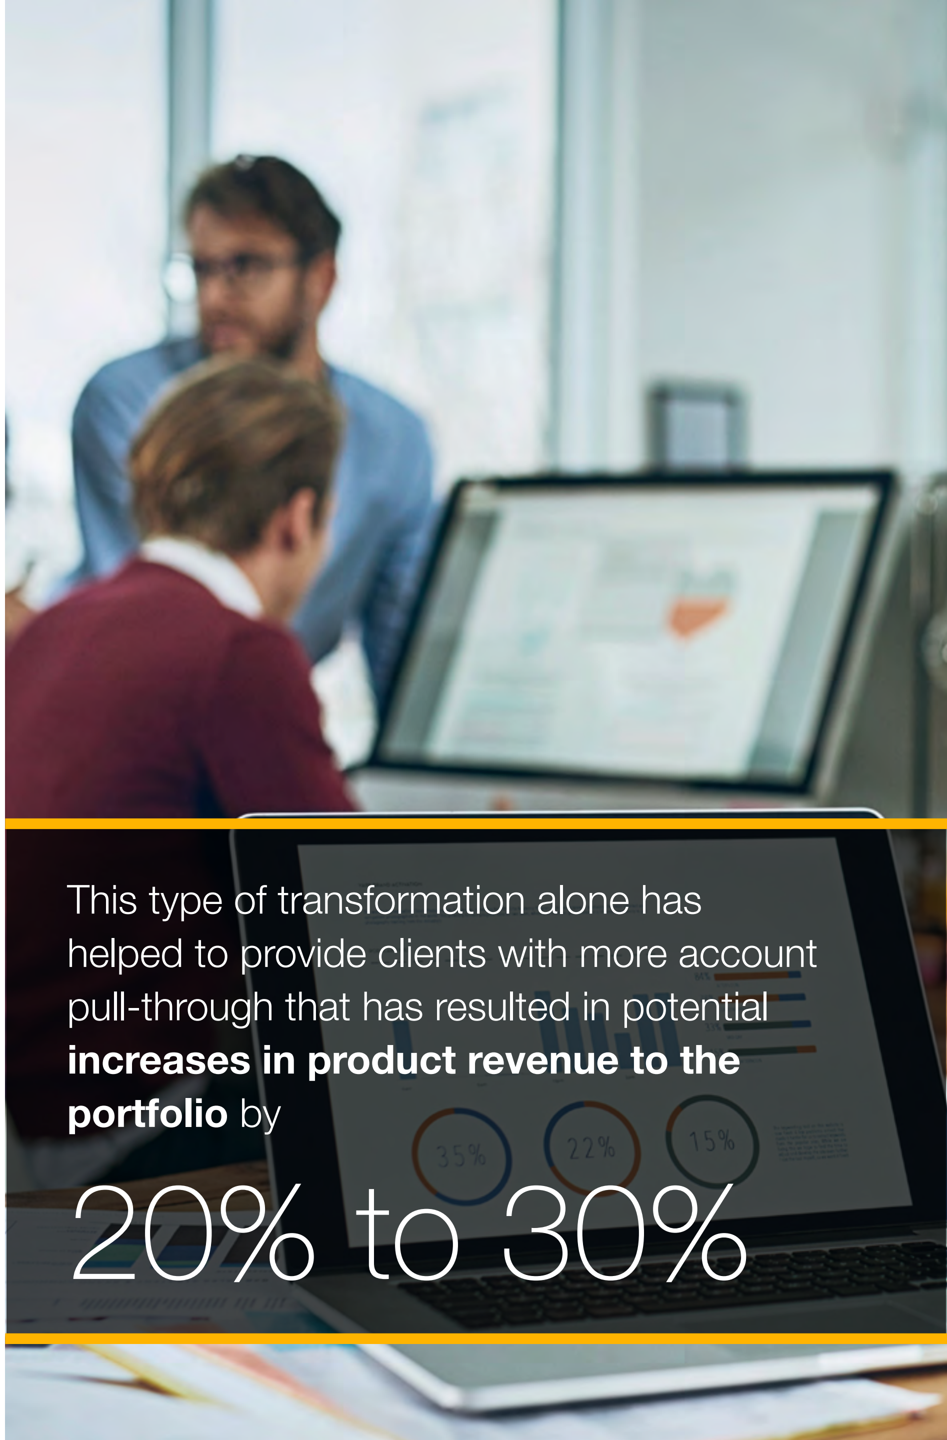

This type of transformation alone has helped to provide clients with more account pull-through that has resulted in potential **increases in product revenue to the portfolio** by

20% to 30%

Making Informed Modeling Decisions

A wide array of machine-learning and AI techniques can be applied for use as part of Financial Crime Compliance. The optimal model or approach will vary by use case, customer or product segment, and respective risk typology, along with many other factors embedded across financial crime prevention and detection processes. As such, a deep understanding of both the business requirements and the applicable machine-learning techniques and methodologies is essential. This can be achieved through partnerships between the modeling and data science team, the risk and investigations teams, technology, business, and data, among others. For example, investigators are often overlooked during the initial stages of model development but they can provide insight into existing threats and help to identify data points that may be invaluable to incorporate for enhanced risk coverage. This insight helps to arm data scientists and modelers with the business and subject matter expertise needed to develop impactful features and proxies to consider during model development.

Explainability is a keystone consideration for any Financial Crime Compliance program utilizing or considering the use of AI. Partnerships across the organization can help address this area by informing the documentation of the rationale for decisions made during the model development process. Proper accounting of the factors that contribute to the choice of data type, features, data sources, frequency of screening, ongoing monitoring processes, among others, are all crucial to confirming that models are defensible and explainable. Given the varying degree of technical knowledge that will be found across potential audiences reviewing each model, the primary goal for model documentation should be approachability and procedural repeatability - audiences should be able to follow in the footsteps of the model development process and be able to logically understand why specific decisions were made before then proceeding to review and understand the underlying data science. This does not mean that proper references or testing evidence are not needed (they definitely are), but it does suggest that in order to promote explainability the audience should be able to understand the business rationale behind modeling decisions and be able to get to the same conclusions before examining the underlying technical processes.

As such, a deep understanding of both the business requirements and the applicable **machine-learning techniques and methodologies is essential**

Setting the Stage for Model Validation

Machine learning models leverage large volumes of data which allow them to incorporate and assess additional information far beyond that which is considered by traditional rule-based systems. Traditional rule-based models are typically constrained by the data that can be incorporated due to a fixed data structure that is predominantly determined by third-parties and that does not allow for customization at the institution level. Machine learning and artificial intelligence models go beyond the conditional reasoning of traditional rules-based systems and allow for Financial Crime Compliance functions to expand into the domain of contextual monitoring. Contextual monitoring enables a more holistic modeling ecosystem where Financial Crime Compliance programs are able to assess interactions between entities, transactions, and other characteristics found in an institution's data to paint a more nuanced classification of risk. However, these advances in modeling necessitate new approaches to model validation.

Model validation of machine learning and AI models can take advantage of many existing validation methodologies, but additional accommodations for model development and testing processes specific to AI models should be adopted. Although the model development and validation functions are independent of each other, there are significant benefits to both functions working together to collate and formalize AI model development and documentation best practices from the onset. These may include defining the processes for algorithm hyperparameter optimization or for the formation of training data sets. The partnership between developers and validators can facilitate the creation of testing plans and standards well in advance of the formal validation process. This approach confirms that both of the functions are continuously aligned. It will also provide the opportunity to explore alternative validation methodologies and for both functions to align on benchmarks and the expected scope of evidentiary materials that will need to be retained. As is the case during the development stage, proper documentation of model validation is critical in generating trust in model efficacy and outcomes. Additionally, the importance of model documentation to be interpretable by a lay audience cannot be overstated.

Additionally, the **importance of model documentation to be interpretable** by a lay audience cannot be overstated.

Getting Ahead of Common Model Validation Challenges

	Common Model Validation Requirements or Expectations	Best Practices During Development to Avoid Potential Issues
Model Data	Model validation or model risk committee teams may necessitate the same governance and sign-off procedures to be imposed when evaluating non-critical data elements or features of a model - even if they do not ultimately get incorporated into the final model itself.	Establish formal roles and responsibilities related to the use of potential or existing model data, including the associated processes for data usage, when having initial conversations with the model validation or model risk committee.
Industry Standards	Model validation & model risk committee teams expect formal documentation whenever there is a potential for departure from 'industry standards'. For less mature programs this may even include documenting pilot modeling efforts or documenting the decision to use machine-learning for the first time within a specific function (such as AML, Fraud, Sanctions, or otherwise).	Since "industry standards" are subjective it's imperative to understand the perspective of each model validation team or model risk committee before beginning a development effort. Establishing formal model development and validation methodologies that are vetted and that have the support of each function will help to reduce friction in the future.
Model Benchmarking	In most situations model validation or model risk committee teams expect AI modeling methodologies to incorporate formal evaluative criteria for comparing models to "challenger" or alternative models	Use regular check-ins with the model validation and model risk committee teams to discuss decisions related to established methodologies where the evaluation of "challenger" models may or may not be necessary - while documenting justifications.
Implementation	Model risk committees typically require that AI model implementations follow established internal and external IT standards to confirm process repeatability and that fail safes have been incorporated. This is due to the escalated model risks associated with Financial Crime Compliance.	Collaborate with IT and the model validation / model risk committee teams early on in the development process to understand the implementation standards that will need to be followed and considered in order to carry out a successful implementation in production.

Explaining Machine Learning Decisions

In the heavily regulated Financial Crime Compliance environment the ability to explain how a model functions and why a model made a particular decision is paramount. Machine learning models are often perceived as black boxes, which can jeopardize trust in the decisions that they make. The opacity of machine learning models and the associated risk of increased regulatory scrutiny has historically led firms to prefer rules-based or less sophisticated modeling approaches, often at the expense of performance.

However, a cause for optimism is that the accelerating adoption of machine learning models has also led to innovations in model interpretability. Model interpretability techniques help to provide increased insight into the relationship between the variables (features) utilized by a model and the model's outcomes. These innovations have not only made it easier to debug potential model issues, but they have also provided stakeholders with more insight into the inner workings of a model's decision making process than ever before. This increased insight has led to increased confidence when discussing advanced models with regulators.

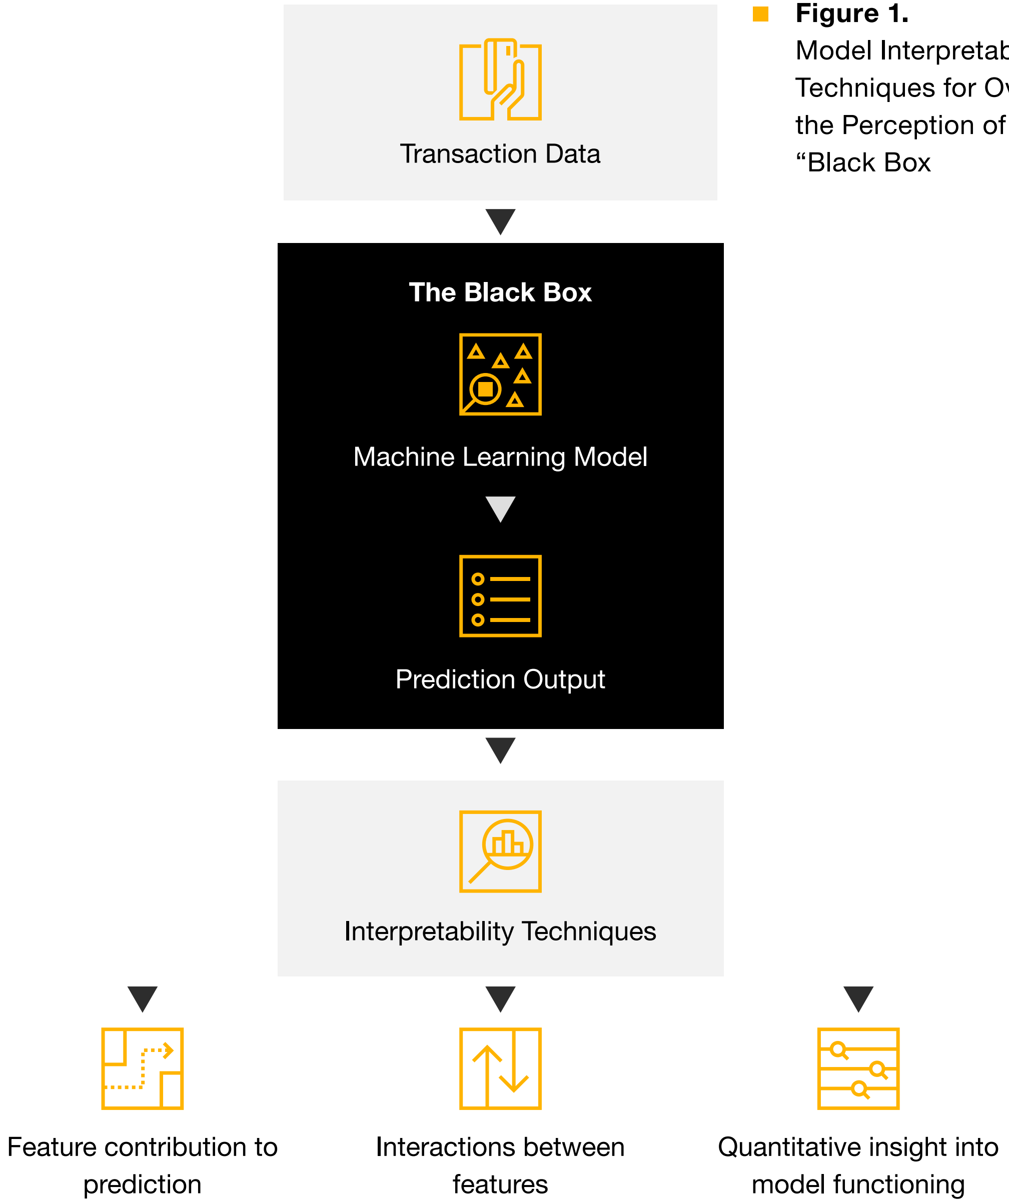
Transaction Data
The Black Box
Machine Learning Model
Prediction Output
Interpretability Techniques
Feature contribution to prediction
Interactions between features
Quantitative insight into model functioning

Figure 1. Model Interpretability Techniques for Overcoming the Perception of the "Black Box

Rethinking Ongoing Monitoring

Maintaining a sustainable and current machine learning process requires persistent review and enhancement of the base model. Customer behaviors inevitably change over time and bad actors continuously find new ways to evade detection. Models and modeling processes that do not incorporate or consider ongoing monitoring and maintenance are destined to become outdated and to experience potentially significant drops in performance. Model monitoring is paramount to a financial institution's ability to identify and adapt to changing threats.

Approaches for maintaining model relevance and performance include, but are not limited to, the following:

1. Probing the unknown - auxiliary unsupervised learning models (anomaly detection algorithms that identify outliers based on the data provided to them) can be run in parallel with existing trained supervised learning models to potentially identify risks never identified before by the institution. Once identified and confirmed, these new risks can be investigated and incorporated into the model training process to bolster and improve the existing model's performance and risk coverage.

2. Defining "normal" output criteria - formalizing an objective definition of what constitutes "normal" model output is a critical step in establishing an ongoing monitoring process that looks to identify deviations in model output metrics. Deviations should be identified and investigated promptly to determine and address the root cause.

3. Automated learning and improvement - AI models trained to flag transactions that resemble historical suspicious activity can be automated to be retrained and to incorporate new escalations and feedback identified by an institutions' investigations team. As investigators find additional productive alerts, those same alerts are fed back into the model to continuously improve the model's future performance and risk coverage.

Models and modeling processes that do not incorporate or consider ongoing monitoring and maintenance **are destined to become outdated and to experience potentially significant drops in performance.**

Rethinking Ongoing Monitoring

Model monitoring dashboards are an important tool for tracking model performance and for identifying potential issues. Monitoring dashboards should be designed to track a variety of model-specific performance indicators and to clearly flag deviations from pre-defined acceptable parameters. Model performance monitoring dashboards are intended to be intuitive and to be easily interpretable by a wide spectrum of stakeholders that may or may not be familiar with machine learning and AI.

Every aspect of the monitoring process should be documented for regulatory compliance. Detailed documentation of model versions, continuous monitoring processes employed, the reasoning behind choice of metrics to monitor, and how metrics are calculated, can confirm that there is always data to support the organization's view of model health.

The shift in Financial Crime Compliance programs towards machine learning and AI makes knowledge of these advanced modeling techniques increasingly valuable. PwC has helped numerous clients implement machine learning models to bolster their financial crime monitoring capabilities, while simultaneously helping those same clients to successfully navigate the corresponding regulatory examinations. Our combination of technical competencies and financial services experience makes PwC a trusted advisor in the financial industry's push to adopt machine learning and to usher in the next generation of enhanced surveillance models.

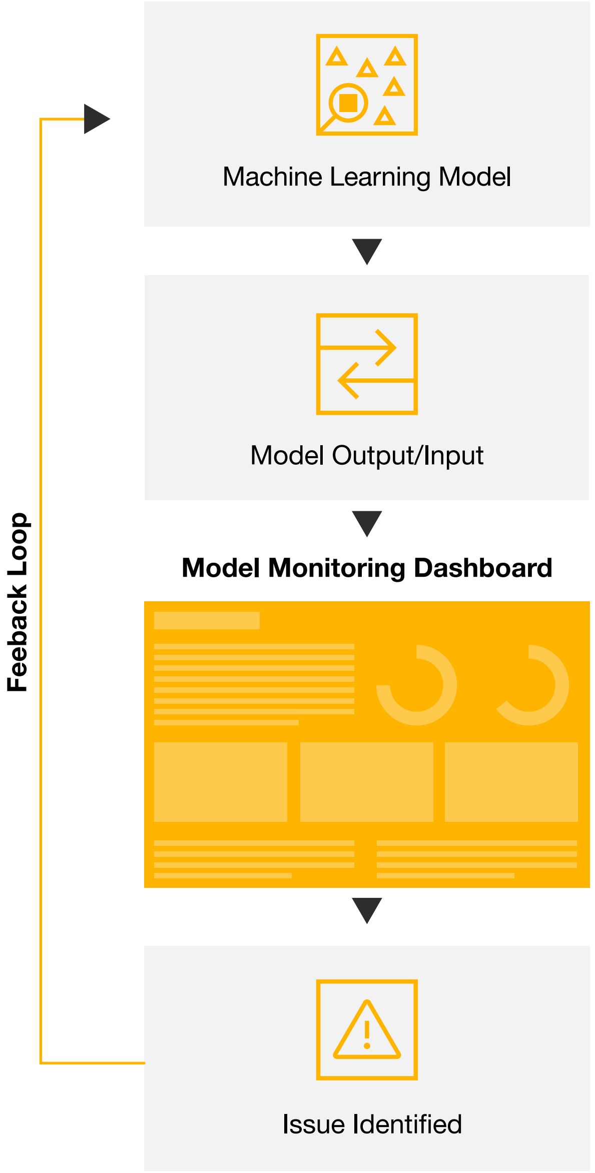

Figure 2. Monitoring Model Performance with Dashboards for Continuous Enhancement

Contacts

Vikas Agarwal

Principal, PwC US
+1 216-789-0314
vikas.k.agarwal@pwc.com

Vikas Agarwal, leads PwC's Financial Crime Unit practice with more than 16 years of financial services, regulatory, and analytics experience. He works with financial institutions and financial technology companies to implement analytics and advanced technology solutions related to fraud, anti-money laundering, watch list screening, trade compliance, and regulatory reporting.

Nirupama Suryanarayan

Principal, PwC US
+1 551-482-0004
nirupama.suryanarayanan@pwc.com

Nirupama is a Principal in PwC's Financial Crime Unit practice with over 14 years of experience in compliance, risk management and analytics. As the Data & Insights Leader, she has worked with financial institutions and fintech assisting them through the regulatory review and remediations of their financial crime programs and transformation of their financial crime programs by leveraging machine learning and advanced analytics.

Andrew Bonslater

Director, PwC US
+1 312-270-3668
andrew.bonslater@pwc.com

Andrew Bonslater is a Director in the Financial Crimes Unit Technology practice at PwC. Since joining PwC, Andrew has applied his background in AML and software development to deliver a perspective beyond the traditional approach by combining innovation and technology to improve quality and increase productivity in financial crimes programs for financial institutions.

Model Edge
A PwC Product

www.pwc.com/us/en/products/model-edge.html