10 THE RISE OF REGTECH AND THE DIVERGENCE OF COMPLIANCE AND RISK

Amias Moore Gerety
and Lev Menand

The rise of regulatory technology (RegTech) promises to transform financial-sector risk management, dramatically improving compliance, lowering costs, and reducing counterparty risk. But in the financial sector, RegTech also threatens to increase frictions between banks and their government supervisors. This is because, as banks have become larger and more complex, supervisors have increasingly relied on compliance as a proxy for risk: assuming that a bank that cannot follow its own rules or comply with applicable regulations is engaged in the sort of operational and financial risk taking that might jeopardize the firm's financial health and financial stability more generally. Even when this assumption breaks down, supervisors often couch risk-related judgments in procedural terms to avoid charged disagreements about the likely outcomes of a bank's business decisions.

New automated audit processes will drive a wedge between compliance and risk: allowing financial institutions to engage in universe testing and to prove compliance—evaporating any relationship between the two. Accordingly, government supervisors may be forced to address issues of business risk directly. And internal risk managers and executives will also need to drill down on business-level decision-making as they too will be less able to rely on audit and compliance records to approximate business risk.

This chapter proceeds in three sections. Section 10.1 reviews the traditional approach to audit and compliance. Section 10.2 explores the promise of RegTech and how it is likely to transform audit and compliance over the next ten years. Section 10.3 considers some of the possible perils of "provable compliance" and how the ability of regulated institutions to demonstrate compliance with bright-line rules will increase pressure on supervisors assessing bank safety and soundness.

10.1 TRADITIONAL AUDIT AND COMPLIANCE

Most audit and internal compliance today uses a sampling methodology. Auditors ask for 7 percent or so of the files from a business process. Then people review those files and write a report about any errors they find—covering both deviations from stated policies and procedures and systematic problems that may have contributed to the pattern of errors they see in the sample.

The atomic unit of analysis in any audit is the work of creating a detailed sequence of what happens in a business process: the audit log. With the rise of automation, the work of auditors has become more difficult, not less. Anyone who has created a computer program is familiar with the challenge of re-creating and understanding a precise sequence within an automated process. An unexpected result in the final output will force the programmer to go back and check each input and each intermediate step until she has identified where the program took an unexpected turn.

In systems that are combinations of manual steps and automated steps, often with millions of lines of code, auditors must master the step-through process of what the computer program has done in the business process and try to re-create the mind-set and decision framework of the people who took the manual steps.

This system is inefficient, error prone, and subject to bias in sampling error. Business processes in a modern financial institution can be thought of as decision trees with thousands of branches. Auditors often cannot hope to sample records that reflect all those branches. This system is also subject to error because the auditors' work is itself manual and completed by humans. Human auditors cannot be expected to diagnose mistakes or errors that exist in the branches of those decision trees with 100 percent accuracy. This is one of the reasons redundancy is built in—internal audit, external audit, and supervisory examination.

10.2 THE PROMISE OF REGULATORY TECHNOLOGY

Despite the attention paid to technological innovations such as machine learning and blockchain technologies, many of the advances in regulatory compliance technology are driven simply by the plummeting price of data storage. As data storage has become cheaper over the past decade, systems are now designed to record and store data about all inputs into a system and all outputs (both intermediate and final), as well as to record metadata about human interventions into the business process. In practical terms, this means that auditors have access to such items as the identity of any employee who took an action, the exact time of day the action was taken, and even how long the employee looked at a screen before taking action. In compliance terms, this means that modern workflow systems can be, and now often are, built with detailed audit logs created automatically. The hardest part of any audit—simply understanding what happened when—is now recorded in real time and has become machine readable. Moreover, the drastic reduction in the cost of storage means that financial institutions not only can build, but should be building, these automatic audit logs into all-new enterprise

systems.[1] As the science fiction writer William Gibson said, the future is here, it's just unevenly distributed.

Already today, systems that are built with these automatic audit logs enable compliance professionals to move beyond sampling methodologies to *universe testing*. They can build systems that represent their policies and procedures and analyze and compare every audit log that a system produces against those policies and procedures. Where policies and procedures are deterministic, universe testing can be similarly deterministic. For example, an auditor can in minutes see every file in which an employee issued a mortgage that deviated from the firm's mortgage origination policy in the second quarter of 2017 in the state of California.

But even where policies are not deterministic, universe testing is already taking place. For example, broker-dealers are required to record all communication between their employees and their clients.[2] For many years, auditors have reviewed these communications using sampling, but technology has made these communications alternately easier and harder to track. For example, phone conversations used to be practically impossible to record at scale; but now, broker-dealers have the technological capability to capture and store those communications. And although chat-based platforms make communications capture easier in theory, the proliferation of encrypted personal chat services has made it harder to track and capture this next wave of communications technology. Many start-ups and internal teams are using voice-to-text and natural language processing technology to run 100 percent of these recorded conversations through a set of compliance rules and flags to identify potential risks. Some companies are even experimenting with pushing these compliance flags out to end users, such as software that monitors text as it is being typed and, in effect, asks users, "Are you sure you want to say that?" While these rules are not foolproof—there are, no doubt, both false positives and negatives—they demonstrate

that the possibilities for universe testing are not limited to fully automated or deterministic systems.

From a business manager's perspective, audit and compliance boils down to two hard tasks. The first is to create systems, whether through training, checks and balances, or automation, that seek to ensure the company will do the right thing in the moment of any transaction or business process. The second, historically much harder than the first, is to prove to a regulator or an external auditor at some given point in the future that the company *did* the right thing at any given point in the past.

Even the first is much harder than it appears: modern financial institutions, including ones that we think of as small, will complete thousands of customer transactions each day, and each transaction will have dozens of steps. We're all familiar with the paperwork that accompanies a mortgage, but the business manager must also codify the steps necessary to have his employees reliably and accurately produce that paperwork in a way that follows both the law and the business imperative of good customer service. The second task of proving compliance is where the difficulty and discipline of audit and compliance have developed. It is this second mission that has governed the practices described above of sampling, audit log, and business process re-creation. And it is in this second task that the power of automation begins to transform the discipline of compliance.

Once a company is capable of universe testing against a business process, that company can not only create and analyze the audit log of that process but also reliably represent to any outsider whether the process was compliant. Universe testing enables the digital equivalent of signed checklists with both granularity (any transaction) and scale (all transactions). The state of the art: provable compliance.

Provable compliance is not the end of the story, however. In fact, as we map out the future of risk and compliance, it is

just the beginning. When it is possible to automatically tell an auditor or an internal control system, "Yes, I know what actions were taken by which people, at what time, according to what policies," this introduces the potential for a drastic change in the compliance mind-set. If it is possible to monitor systems in real time, then applying rules in real time can become possible as well—imposing business logic inside of enterprise systems. For example, today many banks approach anti–money laundering (AML) by training staff during onboarding and through written policies and procedures. Training is then combined with traditional sampling-based audits. But these policies do not typically constrain the actions of a bank's AML personnel, who generally conduct free-form manual investigations across a variety of internal and external data systems, documenting their actions in checklists and a written risk report. Today's RegTech solutions allow a bank to directly enforce its policies inside the workstations of its employees—giving them guideposts on expected or required next steps, automating the data sources from which research must be conducted, and storing reports in an audit log of actual actions taken and a structured data format that can be repurposed or analyzed for a variety of purposes inside a bank.

The difference between these two approaches is like the difference between driving while reading a map and driving with the aid of Google Maps. If you take a wrong turn while navigating according to an analog map, no one in the car can be exactly sure where you went wrong or what to do next. But if you use Google Maps, the GPS not only can alert you at the point of each turn to help avoid incorrect navigation but also can automatically reroute you back to your destination. Moreover, the data collection from these systems is now being used to power the possibility of self-driving cars, just as RegTech systems are helping financial institutions collect the data necessary to automate more and more complex decisions about AML risk.

This extension of rules and policies into real-time feedback or constraints can be termed "programmatic compliance." In its strong form, programmatic compliance describes a state where it is not possible for a business system to take a noncompliant action. To date, the complexities of business processes and required decision-making have made programmatic compliance impossible for all but the simplest systems, but the frontiers of programmatic compliance are moving closer all the time. Soon we will be able to use machine learning and other advanced analytics to embed nuance and judgment by capturing human decisions over time and replicating them in more and more cases.

To understand the potential implications of these advances, consider how provable programmatic compliance would affect counterparty risk. For example, in a derivatives trade between two broker-dealers, both dealers would be able to attach an audit log showing how the trade proceeds will be directed and the status of any collateral. The ability to accept risk from a counterparty starts to change drastically if systems can immediately measure that party's compliance, reducing the uncertainty premium embedded in risk calculations. Consider another example: money laundering compliance. If a counterparty can use an audit log to prove that they did not take significant risk originating a transaction or if the data history is strong enough that a bank can trace the history of any transaction in an automated way, both parties' AML burden drops because they are no longer facing these unknowns. This is not something that is possible today, nor is it something that will happen tomorrow. But it is something that is fast approaching as IT departments get better and better at testing and proving compliance in high-volume, scalable ways.

10.3 THE PERILS OF PROVABLE COMPLIANCE

Like any technological transformation, provable and programmatic compliance will bring its own challenges and create new

risks. One of these challenges is likely to involve the way universe testing will alter the relationship between heads of business lines and risk managers (on the level of the firm) and between financial institutions and their regulators (on the level of the system overall).

Risk managers, boards of directors, and government regulators are all charged with monitoring and managing risk taking at and across financial institutions. Risk management involves the identification and evaluation of risks—measurable probabilities that certain negative outcomes will occur in the future—and the application of resources to minimize and control the likelihood of these negative outcomes. Some risks are easy to measure and address: interest-rate risk, for example, can be managed using simple derivatives. Other risks are amorphous and difficult to control—for example, counterparty risk, the likelihood that one of the parties fails or is otherwise unable to perform its contractual obligations as promised. Counterparty risk today necessarily involves an estimation of other firms' solvency risk management capabilities, as well as systems to monitor and aggregate those estimates over time. It is particularly difficult for third-party monitors, such as board committees and outside examiners, to assess these sorts of risks and determine whether the business is managing them appropriately. It is well known in banking that managers have incentives to take more risk than would be preferred by society at large.[3] Risk management and oversight requires a firm and its overseers to mediate between optimistic and pessimistic views of an uncertain future state. It is also true that regulators, board members, and risk managers will have less detailed and less direct knowledge about the risks embedded in any business transaction.

Compliance, by contrast, is the practice of monitoring business actions for consistency with rules—including laws and regulations, and policies and procedures. Unlike risk, which involves judgments about the future (based on understandings

gleaned from the past), compliance involves judgments about the past (based on information gathered about the past). While assessing compliance can be difficult, it is more susceptible to objective evaluation than risk management. Perhaps as a result, supervisors and other officials charged with managing risk taking across the firm tend to focus on a business's compliance with bright-line rules. Examiners often assume, either implicitly or explicitly, that poor compliance means inappropriate levels of business risk. Indeed, for the past two decades, bank supervisors in particular have focused on compliance. Bank supervisors likely rely on compliance checks because they are easier to conduct and more difficult for banks to dispute. They are also more objective: supervisors can point to definitive evidence that a bank did not follow the rules. There is never definitive evidence—until it is too late—that a bank has taken on too much risk.

Take the "London Whale" scandal, for example, which involved a $6 billion loss in a single quarter on a twelve-figure bet on exotic derivatives at J.P. Morgan's commercial bank.[4] These losses precipitated the largest safety and soundness penalty ever levied. Yet, despite the fact that the traders involved had been empowered to make a levered bet about the direction of the global economy and had increased the bank's balance sheet exposure to the risk as the bet began to move against them,[5] the rationale provided by supervisors was procedural: the consent orders faulted the bank merely for failing to adequately supervise its traders, properly value its investments, "implement adequate controls," and "ensure significant information . . . was provided in a timely and appropriate manner to examiners."[6]

By couching supervision in such technical, procedural terms, the banking agencies rendered their orders easily defensible, especially to constituencies with different perspectives on risk and regulation. No one could argue with the Fed that J.P. Morgan was noncompliant. But, had the Fed faulted the

bank for its risk taking and not for its compliance violations, commentators and bank executives might have argued that the bank was justified in taking such large risks.

This technocratic approach is also borne out in recent research that exposes new data about the supervisory process and the pressures that have come to bear on both risk managers at banks and regulators. For example, data released by the Federal Reserve Bank of New York shows the different topics supervisors focus on, called "Matters Requiring Attention" (MRAs), in their confidential letters to banks.[7] Figure 10.1 breaks out MRAs by issue type, with the most procedural matters in dark gray at the bottom of the stack and the most substantive ones lightly shaded at the top. Procedural supervision is more common for the larger banks than for the smaller ones. Only 8 percent of MRAs at large banks were related to their loan portfolios, compared with 27 percent of MRAs at state member banks. Nearly 80 percent of supervisory activity at the larger banks fell into the first three, largely procedural, categories, whereas 55 percent and 56 percent fell into these categories for state member banks and smaller bank holding companies, respectively.

Part of the significant focus on risk modeling among the large banks is that this data covers the period when the largest banks became subject to the Comprehensive Capital Assessment and Review, the Fed's supervisory stress tests. The data suggests that despite the popular conception of the stress test as a quantitative assessment of bank balance sheets, it was the qualitative component with a focus on risk modeling that drove the MRAs. The Fed assumed that if a bank was not good at managing its capital planning process, it probably should not be allowed to pay out capital even if the Fed's quantitative analysis suggested that the bank's balance sheet was strong. Whether a bank has enough capital to sustain its business through a sharp macroeconomic contraction is difficult to determine because it involves projections about the future. It

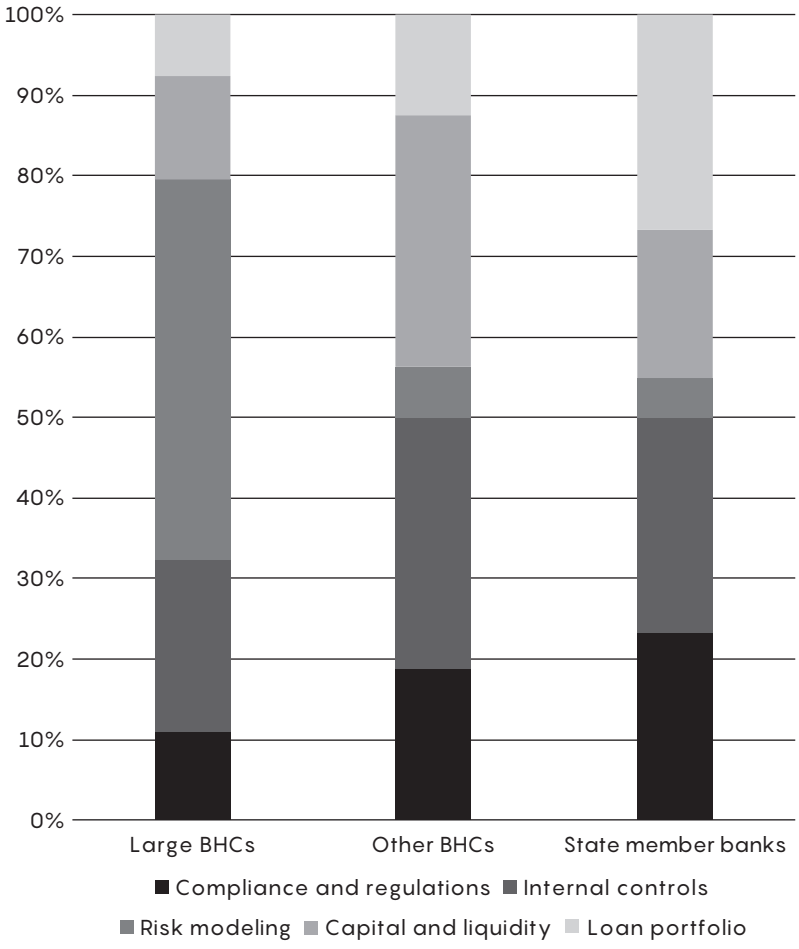

FIGURE 10.1
MRAs by issue type.
Source: Paul Goldsmith-Pinkham, Beverly Hirtle, and David Lucca, *Parsing the Content of Bank Supervision*, Federal Reserve Bank of New York Staff Reports 770 (March 2016), https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr770.pdf?la=en.

is far easier to assess whether a bank followed a careful capital planning process to decide how much capital to pay out to its shareholders because this involves merely collecting information about the past.

This sort of *proceduralism*—the focus on compliance instead of risk and couching judgments about excessive business risk in terms of compliance—reduces conflict: conflict between supervisors and bank executives, conflict between supervisors and senior officials in Washington (who are often lobbied by large financial institutions to rein in overeager junior examiners), conflict between supervisors and Congress (which is populated by many legislators who are uncomfortable with broad supervisory discretion), and conflict between supervisors and lawyers, who, themselves, prefer the legal certainty of bright-line rules. As one might expect, the greater the imbalance of power between supervisors and the firms they oversee, the more supervisors will fear conflict and seek refuge in the neutral discourse of process and the bright-line clarity of regulations.[8]

Even if this supervisory emphasis on compliance as a risk indicator has been helpful, the rise of automated compliance and programmatic compliance will decouple the relationship. Both compliance and risk have traditionally been disciplines that marry process discipline with human expertise and experience, and therefore weaknesses in one were presumed to indicate weaknesses in the other. The ability to impose compliance discipline through implementation of automated systems means that this may soon cease to be the case. If so, supervisors will no longer be able to predict who is going to be well positioned for macroeconomic shocks or operational risks tomorrow, based on who was following the rules yesterday. This decoupling is likely to place new pressures on supervisors and disrupt an equilibrium where supervisors can use the neutral language of compliance failures to support judgments about excessive risk taking. The judgments that banks make about risk are much closer to business judgments, such

as about the future performance of the financial system and whether it is appropriate to undertake a business transaction. These judgments will force regulators directly into the zone of conflicts, where they have to second-guess business executives without being able to claim superior expertise or concrete evidence of mistakes.

As we look forward, we should think not only about the benefits of RegTech but also about the pressure it will place on the policy apparatus and the existing political equilibrium between the regulator and the regulated.

10.4 CONCLUSION

The rise of RegTech offers enormous benefits. RegTech will help banks improve compliance with their own policies and outside rules and regulations. This adherence should lower the cost of capital, reduce counterparty risk, and further public policy aims. But it may also begin to change, at least in the banking sector, the relationship between the regulators and the regulated. As the compliance gap shrinks, it will become harder for bank supervisors to use compliance as a proxy for business risk. Bank risk managers and regulators will have to move their dialog into the zone of uncertainty about future financial and economic performance. Regulators represent the public's interest in restricting financial firms' risk taking. But doing so without the comforting language of procedural weakness will require more trust and more explicit conversations about society's appetite for risk in the banking system.

Both financial firms and regulators will need to take steps to prepare for the reality of compliance automation and to benefit from its potential. Financial firms should take the following steps:

- *Focus on high-value experiments*. Take advantage of technology and data improvements by choosing focused experiments

in areas where compliance is more data driven and, in particular, where it includes high volumes of third-party data sources. The difficulty of working with third-party data will be lower, and the value of processing it well will still be material. Focus on scaling the learnings from each experiment, not just the specific programs that succeed.

- *Map business processes from end to end*. Track the full journey of a business process from the provenance of each data point through to the transactional record, controls, compliance, and audit. Firms do not need to rebuild their processes all at once, but they need to be aware of their needs at each step to make sure that the control functions get the right data inputs and can attach early enough to benefit from the potential advances in automation.
- *Focus on the frontier of the easy*. Many of the possibilities discussed above do not require technological breakthroughs; they require only careful application of existing, often open-source, technology. Too often, innovation groups focus on breakthrough tech like quantum computing, which is usable only in a lab, and lose sight of the very real gains that could come from building cloud native applications or automating routine steps. Especially for innovation in compliance, financial firms don't have to chase the frontier of what's possible.

Financial regulators should take the following steps:

- *Practice jujitsu, not tug-of-war*. Given how quickly technology is spreading, regulators (mainly lawyers and economists) may worry that they don't have the technological expertise to keep up. As described above, technology should make it easier to create high-quality audit logs and manage large volumes of data. Regulators don't have to compete with financial firms in a tug-of-war over who has better tech capabilities; they can act more like a jujitsu master: using the greater power of technology to ask for and receive better, cleaner data from financial institutions.

- *Don't lose focus on risk*. The end goal for financial regulation is the creation and maintenance of a stable financial system that provides services to the economy in ways that are consistent with a society's values. Regulators and, in particular, policy makers should resist the temptation to accept the equilibrium where rote compliance is the focus of policy, supervision, and enforcement. Automation will make it easier to achieve rote compliance and will increasingly reveal that compliance systems cannot substitute for business judgments about risk.

NOTES

1. Fifteen years ago, no IT department could have created an automatic audit log of every action that was taken in every system in a bank. The technical challenge of recording every action was too hard (most computer programs stored intermediate steps only long enough to calculate the next action) and data storage was far too expensive.

2. FINRA Rule 3110, https://www.finra.org/rules-guidance/rulebooks/finra-rules/3110, accessed July 2021.

3. See, for example, L. Bebchuk and H. Spamann, "Regulating Bankers' Pay," *Georgetown Law Journal* 98, no. 2 (2010): 247–287, also Harvard Law and Economics Discussion Paper No. 641. See also A. Admati and M. Hellwig, *The Banker's New Clothes* (Princeton, NJ: Princeton University Press, 2013).

4. To put the loss in perspective, J.P. Morgan at the time typically made a profit of approximately $5 billion per quarter across its entire business. It is also worth noting that the Whale losses happened in a benign credit market when interest rates were stable and interbank lending conditions were normal. See *JP Morgan Annual Report*, 2014, https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/investor-relations/documents/JPMC-2014-AnnualReport.pdf.

5. "Those holdings were created, in part, by an enormous series of trades in March, in which the CIO bought $40 billion in notional long positions which the OCC later characterized as 'doubling down'

on a failed trading strategy." US Senate, Permanent Subcommittee on Investigations, *JPMorgan Chase Whale Trades: A Case History of Derivatives Risks & Abuses*, 2013, 4, https://www.hsgac.senate.gov/imo/media/doc/REPORT%20-%20JPMorgan%20Chase%20Whale%20Trades%20(4-12-13).pdf.

6. JP Morgan Chase Bank, N.A., Order No. AA-EC-2013-75, O.C.C. (September 2013), at 3 ("the Bank's oversight and governance . . . were inadequate"); JP Morgan Chase & Co., Docket No. 13-031-CMP-HC, F. RES. (September 2013), at 4 ("JPMC exercised inadequate oversight"); OCC Order, supra note XX at 3 ("the Bank's valuation control processes and procedures . . . were insufficient to provide rigorous and effective assessment of valuation").

7. P. Goldsmith-Pinkham, B. Hirtle, and D. Lucca, *Parsing the Content of Bank Supervision*, Federal Reserve Bank of New York Staff Reports 770 (March 2016), 44, https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr770.pdf?la=en.

8. L. Menand, "Too Big to Supervise: The Rise of Financial Conglomerates and the Decline of Discretionary Oversight in Banking," *Cornell Law Review* 103 (2018): 1527.