Three Lines of Defence Case Study: Grand River Valley Bank

You work at Grand River Valley Bank (headquartered in Waterloo ON). GRVB has recently onboarded a new team that focusses on financing technology companies. The target market includes software/SaaS and AI firms. These firms have been through a number of financing rounds but would like some debt to avoid further dilution of their equity.

The team has been clear that most of the loans would not be collateralized and borrower's operations would not be generating sufficient cash flow to confidently meet timely interest and principal repayment. The equity investors in these firms are founders with a track record along with strong VC partners.

The team will be called the "Tech Banking Team" or TBT and they will bring with them origination bankers and loan specialists, as well as middle and back office support groups.

TBT's value proposition is that, while the loans on their own would be unprofitable, the firm would put its operating deposits with your bank and your wealth group would be invited to speak to their executive team about on-boarding as clients. Their previous bank also had deep relationships with the VC firms and their executives which they hope to bring to GRVB.

GRVB's CEO was personally involved in the team's hire and has publicly and privately stated her excitement and optimism about what the team will bring to the bank. During the last quarterly call with analysts, the CEO and CRO both discussed the new team and stated that GRVB was not changing its risk appetite with this new initiative and that the team's strong risk culture would be an excellent fit with that of the bank.

The second-line risk group (IRM) has formed a team to work with TBT and other 1LOD leaders to start to work through the issues that may arise as they start to do business.

Some background information:

- GRVB's Board has just approved its strategic plan and associated risk appetite.
- The new team would be part of GRVB's Commercial Bank division. The Commercial Bank has run a traditional book of business and has a Moodys equivalent WAR of Ba1.
- Current Board-approved risk appetite metrics were calibrated off of the existing book of business and did not anticipate the acquisition of the new team. The metrics were calibrated to be well inside the risk capacity of the GRVB.

TBT and the head of the Commercial Bank have drafted a risk appetite for the new business. Key highlights include:

- Authorized TBT loans will be no greater than 15% of the authorized commercial portfolio.
- TBT's relationship ROE will meet or exceed the Commercial Bank target of 13%.
- TBT's Chief Credit Officer will be able to approve any loans smaller than $100MM up to a borrower limit of $250MM. Above this amount, the head of TBT and the head of the Commercial Bank will need to jointly approve. Note that credit approvals at GRVB are done within the 2LOD.
- TBT will be exempt from GRVB's credit policies and procedures and subject to their own until such time as they can agree with GRVB on new frameworks.

Your manager in the 2LOD has asked your team to draft a response to this in advance of the meeting. While your manager will be there, this will be yours to present.

The head of TBT and his direct reports will be there, as well as the head of the Commercial Bank. The head of TBT is known as charismatic and charming but highly driven and passionate about his team and their clients. TBT's track record is impressive having delivered ROEs above 25% for the past 3 years and revenue CAGR of 15% over the same period.

Technology Disruption and Risk: Agentic AI, Quantum, and Emerging Industry Insights Case Study:

WhitePine Financial Holdings: AI as a Cross-Cutting Risk

Fictional illustration prepared for discussion purposes.

WhitePine Financial Holdings is a Canadian financial institution with operations spanning retail banking, commercial lending, wealth management, and insurance. Over the past two years, WhitePine has embedded AI systems across its core business functions. Three events in 2026 have raised questions the organization has struggled to answer cleanly.

The Credit Adjudication Platform | Machine Learning

In 2025, WhitePine deployed a machine learning model to accelerate small business and consumer lending decisions. The model integrated traditional credit variables with behavioural transaction data, real-time fraud screening, and dynamic pricing logic. Unlike a rules-based system, the model incorporated ongoing retraining and periodic parameter updates informed by new data, and its credit assessments evolved as those updates took effect. Approval volumes increased, processing times fell, and management considered the deployment a success.

During a recent period of regional economic volatility, the model began declining applications that WhitePine's own credit policies would otherwise have supported, and pricing new credit facilities at spreads wider than WhitePine's standard risk-based pricing would have produced. Applicants quoted higher rates either went elsewhere or accepted terms more expensive than warranted. Those who were declined received a notice of the decision, but branch staff could not explain the basis for it if asked. Certain industries and communities appeared disproportionately affected. Complaints began circulating on social media, and local media began questioning whether automated decision-making was disadvantaging particular borrowers.

Responsibility for monitoring model performance metrics sat with Model Risk; metrics remained within approved tolerances, so no escalation was triggered. Legal and Compliance assessed whether complaint handling continued to meet FCAC standards during the surge. The business interpreted the model's response as appropriate given elevated economic risk. Customer Relations logged the surge in complaints as a service quality issue and routed them to the relevant business lines for response.

Each function observed indicators moving in concerning directions, but no single metric crossed the specific trigger that required escalation. The framework had been designed to respond to discrete breaches within risk categories, not to aggregate signals accumulating across them simultaneously.

By the time the issue reached executive management, OSFI and FCAC had already sent information requests.

The Call Centre Assistant | Generative AI

WhitePine's retail call centre had introduced a generative AI assistant that consulted the bank's actual policy documents in real time to help agents respond to customer inquiries. But in edge cases such as complex mortgage modifications and insurance claim eligibility, it produced responses that were confident, plausible, and subtly incorrect. Call centre agents, many of whom lacked deep product expertise, had little basis to question outputs that appeared authoritative. Several customers received inaccurate information about their loan terms

and entitlements before the pattern was identified. Because interactions were summarized rather than fully retained, WhitePine struggled to determine which customers had received inaccurate information or whether affected customers could be identified systematically.

The Cash Management Tool | Agentic AI

WhitePine's treasury division had deployed an agentic AI system to manage the bank's intraday cash positions, deciding how much money to hold in different parts of the business and moving funds between business units as market conditions changed, within parameters that had been pre-approved.

The system was authorized to execute liquidity reallocations automatically unless interrupted by authorized personnel. During a sudden market disruption, the system began acting faster than treasury staff could follow. As conditions deteriorated, several staff concluded independently that the pre-approved parameters were no longer appropriate for the situation unfolding. The treasury analyst monitoring the system attempted to pause specific transactions. Their manager, unaware the analyst was already acting, issued a different instruction. The head of a business unit whose funds were being reallocated tried to block those movements. The risk function, receiving escalation calls from multiple directions simultaneously, issued its own directive. Simultaneous attempts created conflicting instructions rather than effective control.

Escalation to senior leadership became congested. The organization's escalation framework had been designed for decisions made at human pace, and the system continued to operate throughout.

One More Thing

WhitePine's insurance broker also advised that carriers are introducing AI exclusions across D&O, professional liability, and general liability policies at renewal. Some categories of AI-related exposure WhitePine may have assumed were transferable are no longer insurable at economically viable levels. The board has begun asking whether the organization is accumulating forms of operational and governance risk that markets themselves are becoming unwilling to absorb.

Full Group Discussion: Your Assignment

The board Risk Committee has received separate briefings on all three events. Each business line has described its piece of the problem. The board is not satisfied. No one has connected these events into a coherent account of what WhitePine is actually facing.

As Special Advisor to the Chief Risk Officer, you have been asked to do exactly that. Your CRO has raised the following questions:

Do these events represent weaknesses in individual risk functions, or do they expose something more fundamental about how WhitePine is architected to govern risk in an AI-enabled business?

Credit Adjudication Platform

What risk categories were active in that scenario?

Where was the governance weakness? Was it Model Risk? Compliance? Business leadership?

Model Risk confirmed the model's performance metrics remained within approved tolerances throughout. Yet the model was producing outcomes that WhitePine's own credit policies would not support. How is it possible for a model to be performing correctly on every metric that Model Risk was tracking and still produce outcomes the bank did not intend? What does that tell us about what those metrics were measuring?

Call Centre Assistant

The credit platform produced unexpected behavior under conditions that might have been anticipated and tested for. The call centre system was operating exactly as designed and still produced wrong outputs. What does that difference mean for how you would approach governing each?

Cash Management Tool

Multiple people attempted to override the system simultaneously. None of them succeeded, and their attempts created conflicting instructions. What would have needed to be in place for any one of those override attempts to have been effective?

* * *

Which assumptions embedded in WhitePine's existing ERM architecture appear no longer reliable under conditions of large-scale AI deployment?

Breakout Session Discussion: Questions

You remain in the role of Special Advisor to the CRO of WhitePine. *Each group has been assigned one of the six questions below. Be prepared to share your key observations with the full group.*

Q1 - Credit adjudication: Monitoring and detection

Model Risk confirmed that performance metrics remained within approved tolerances throughout the credit platform event. Yet the model was declining applications and overpricing new facilities beyond what WhitePine's own credit policies would support. What monitoring of model outcomes rather than model performance metrics would have detected this earlier? At WhitePine, who might

be best placed to take responsibility for that monitoring, and what expertise would they have needed to assess what they were seeing in real time? What outcomes, behaviours, or customer impacts should have been treated as risk indicators even though traditional model metrics remained stable?

Q2 - Credit adjudication: Remediation

WhitePine now has a backlog of potentially affected lending decisions: wrongful declines and overpriced new originations. What would a credible remediation program require in terms of reviewer capacity, access to application files and model decision records, and credit judgment exercised independently of the model under review? What does the FCAC 56-day complaint resolution requirement mean for how quickly that capacity must be mobilized, and what would WhitePine, in hindsight, have been well advised to build into the model deployment to facilitate complaint resolution?

Q3 - Call centre: Detection and remediation

The generative AI assistant was operating as designed and yet customers received subtly incorrect information about their loan terms and entitlements. No threshold was breached and no metric flagged a problem. Unlike the credit adjudication scenario, there is no defined backlog of identifiable decisions to review. What does this reveal about the limits of traditional monitoring for this type of AI system, and what would identifying and remediating affected customers require? What design-stage decisions would have made it possible to detect this problem earlier and identify which customers were affected? What level of logging, retention, and interaction traceability would have been necessary to investigate the issue effectively after the fact?

Q4 - Cash management: Governance of agentic systems

The agentic cash management system operated within its pre-approved parameters throughout the market disruption. Yet treasury staff could not determine what it was doing or why, override attempts created conflicting instructions, and escalation to senior leadership became congested. What does this reveal about the limits of parameter-setting as a governance mechanism for agentic AI systems? What authorization architecture and human oversight design would WhitePine need before deploying this kind of system at operational scale?

Q5 - Cross-cutting: Accountability

Across all three events, each function operated within its defined lane and WhitePine was still blindsided. Consider both the first line business functions that deployed and operated these systems and the second line oversight functions responsible for the risk framework. Where should accountability for end-to-end AI system outcomes sit in the first line, and what coordination and integration role should the second line play to ensure that what each function sees is assembled into a coherent picture?

Q6 - Cross-cutting: Governance designed to cope with problems in production

In all three events, problems reached customers and regulators before WhitePine's internal governance detected them as enterprise risk events. WhitePine's governance architecture appears to have been designed around the assumption that controls would prevent problems from reaching production. What is different about governance architecture designed to cope with AI-related problems in production from the architecture WhitePine had in place, and what would need to change?

Risk Reporting and Communication Case Study: Undrawn Commitments

Undrawn (or unfunded) commitments are the difference between the authorized loan amount and the amount drawn, where the lender has a legal obligation to fund the undrawn portion at the borrower's request, as long as certain conditions are met.

Common banking and wealth products that have this feature include:

- Revolving credit lines to consumers – these can be secured or unsecured
- Securities-based margin facilities in wealth management
- Revolving credit facilities for commercial borrowers
- Delayed draw term loans
- Letters of credit or letters of guarantee

Regulatory reporting and capital treatment usually broadly tracks the above.

You've been asked to create a report on undrawn commitments. The primary audience for the report will be risk teams in 1LOD and 2LOD, as well as Corporate Treasury, who want to understand the funding and liquidity implications of unfunded commitments.

For now, a weekly report is adequate but management would like the option to produce it on demand or more frequently in times of stress.

Assume that your firm has 3 main business segments: Retail, Commercial and Wealth. Retail is a Canada-only business but Commercial and Wealth have U.S. and Canadian operations. Each segment has its own loan system but your firm has a common set of reference data and a common credit risk system.

- Loan systems data for each segment include authorized and drawn amount by facility for each borrower and by currency. They also have the SIC code for each facility.
- The common credit risk system has credit facilities with ratings and assigned loss-given-default codes. SIC codes are also included in credit data and can be different from the loan data since credit may over-ride the business input.
- Reference data have single name hierarchy, sector classifications, ratings mapping and LGD codes. They also house legal entity information that lists which of your firm's operating entities holds the loan.

- Your Risk Technology team has a legacy risk and aggregation engine that is fed by these data sources and can be modified to calculate the undrawn exposure and aggregate the data according to your specifications.
- A reporting layer will draw from the engine and produce the report.

A simple illustration of the architecture is below.

Propose a risk reporting structure given the information above.

1. What are some of the key questions we need to address with respect to:
    - Scope of the report
    - Granularity vis-à-vis intended audience
    - Cadence of reporting
2. What are some of the BCBS 239 issues that might arise in the development of the report, especially in terms of:
    - Governance and infrastructure
    - Risk data aggregation
    - Risk reporting practices

Undrawn Commitments Risk Reporting Architecture

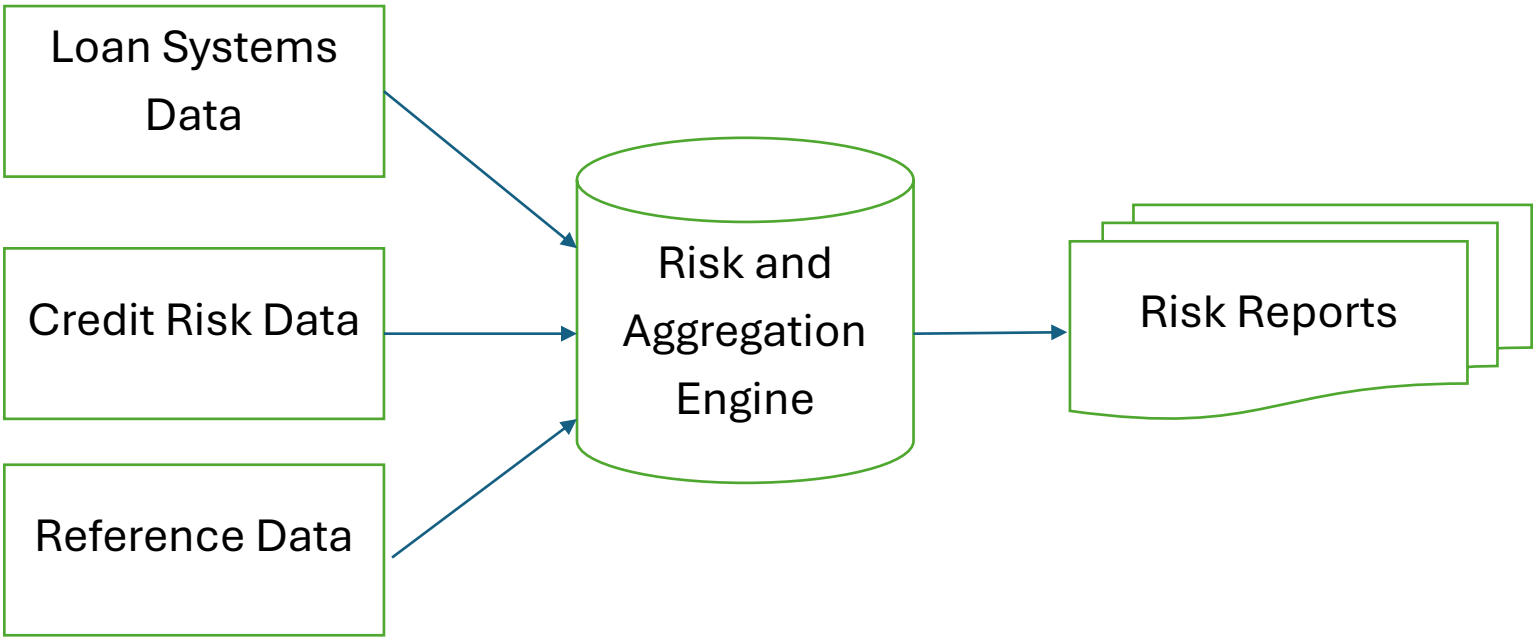

Risk Appetite Case Study: Maplestone Financial Group

Raza Hasan, Executive in Residence

Key Learning Objectives

- To establish familiarity with the types of core principles that might underpin a financial institution's risk appetite framework
- To gain insights into how these core principles can be applied in practice
- To explore the role of the risk appetite framework to help navigate through potential risk/reward trade-offs and establish a clear "tone from the top" for disciplined risk-taking behaviour
- To examine how operational risks (third-party, technology and cyber, model, and compliance risk) should be weighed alongside financial risks when assessing a new opportunity

Background

Maplestone Financial Group ("Maplestone") is a publicly listed (TSX: MSF) Canadian bank with total assets of approximately $55 billion. It has two primary business units (BUs) – the Personal Banking BU and the Commercial Banking BU. The Personal Banking BU offers residential mortgages, home equity lines of credit, credit cards and deposit products to Canadian families in the broad middle market, using a digital-first model supported by a modest branch network in Ontario and Quebec and a national mortgage broker channel. The Commercial Banking BU serves the lending, deposit and cash management needs of Canadian entrepreneurs and small and medium-sized businesses, with a particular focus on the commercial real estate, technology and professional services sectors, through a team of relationship managers. While Maplestone operates nationally, it only recently expanded into Western Canada, which currently accounts for 10% of its total revenue.

Maplestone's success has been supported by its long-standing tradition of continuous innovation, data-driven decision making and high levels of organizational agility in pursuing emerging opportunities. The bank is currently midway through a three-year core banking modernization program, the largest technology investment in its history, which is migrating its deposit and lending platforms to a new cloud-based core system. The program is running approximately six months behind schedule and 20% over budget. A failed system cutover last spring caused a 19-hour outage of online and mobile banking, which attracted significant media coverage and a supervisory letter from OSFI requesting a remediation plan. That remediation work is still in progress.

Over the last three years Maplestone's ROEs have been in the range of 9-10%, below its stated long-term objective of 13-15% (a level it last achieved before rising funding costs compressed its net interest margin and credit provisions returned to more normal levels). The recently tabled draft business plan shows projected profitability largely in line with current ranges, with relatively modest core income and revenue growth (low single digits). The board has expressed its dissatisfaction with the draft plan and has challenged management to *"go back and find some creative ways to bring profitability more in line with our long-term objectives."*

Maplestone's capital position is strong. Its Common Equity Tier 1 (CET1) ratio stands at 13.8%, well above regulatory minimums and its internal target of 11.5%. The ROE "drag" associated with the current "excess" capital position accounts for approximately half of the prevailing ROE shortfall.

Maplestone's stock has traditionally been favoured by income-oriented investors, primarily due to its reliable dividend performance. There are currently seven equity analysts rating the stock, with the following ratings distribution: 2 "sell", 4 "hold", 1 "buy". The two "sell" ratings are recent developments, citing margin pressure, concerns about future growth prospects and execution risk on the technology program following last spring's outage. One credit rating agency recently affirmed Maplestone's rating but flagged operational execution as a key risk to monitor.

Maplestone's articulated corporate vision, mission and values are as follows:

- *Vision:* To be the most trusted bank for Canadian families and small businesses.
- *Mission:* To help our customers build financial confidence and security.
- *Corporate Values:* Innovation; Customer focus; Integrity; Agility; Data-driven decision making

Scenario

LedgerLeaf Inc., a fast-growing, venture-backed fintech based in Calgary, recently approached Maplestone with a partnership opportunity. LedgerLeaf offers point-of-sale instalment loans to consumers and working capital loans to small merchants, embedded directly within merchant checkout and accounting software across Western Canada. As LedgerLeaf is not a regulated lender or deposit-taking institution, it needs a bank partner to originate and fund its loans. Under the proposal, Maplestone would originate the loans under its own name and hold them on its balance sheet, growing to a portfolio of up to $2.0 billion over three years. LedgerLeaf would operate the end-to-end customer journey, including digital onboarding and identity verification, credit decisioning using its proprietary machine-learning model, servicing and collections. Its platform is hosted with a single public cloud provider and would connect to Maplestone's new core banking system through APIs. Given strong demand and limited competition in this segment, LedgerLeaf's financial advisors project pro-forma ROEs in the mid-20% range on allocated capital. This indicative return has been validated by an external consulting firm retained by the sponsoring BUs to assist in the assessment of this opportunity.

Preliminary due diligence has surfaced several further facts. LedgerLeaf's credit model has been trained on three years of data, all from a relatively benign credit environment, and makes extensive use of alternative data. LedgerLeaf has provided a point-in-time (Type I) controls assurance report, rather than one testing controls over a period of time. It experienced a data breach 18 months ago affecting approximately 40,000 customers. It is not yet profitable and has roughly 15 months of cash runway. Its previous bank partner exited the relationship last year for "strategic reasons". In a memo to the CRO, the Head of Operational Risk noted that third-party due diligence is not yet complete, no exit plan has been developed, and the integration work would require changes to the core banking platform while outage remediation is still under way. LedgerLeaf has asked for a signed term sheet within 30 days, as it is also in discussions with another bank.

While acknowledging that this opportunity may be *"a little off of our core strategy"*, the CEO and the heads of both BUs (who would jointly sponsor the partnership) support advancing it. They cite the attractive indicative returns, its ability to quickly deploy excess capital and bolster top line growth, its

being a *"great diversification play"* across new products, customer segments and geography, and *"a chance to show the market our innovation credentials after a difficult year."* The CRO has indicated that based on the analysis provided, she could *"get comfortable"* with the credit risk, provided that the portfolio is maintained within a prescribed limit and monitored closely, and that the operational risk questions *"can be worked through during implementation."* The New Initiatives Policy requires that any new business venture, material outsourcing arrangement or capital allocation to a new product be presented to the Risk Committee of the Board for review and approval, hence the reason it is being tabled on today's agenda. Selected risk appetite metrics prepared for the committee are shown in Exhibit 1.

You are a member of the Board Risk Committee being asked to consider this proposed partnership. This discussion has already exceeded the allotted time, with many strong views being expressed both for and against this proposal. Feeling that the committee is struggling to reach a clear consensus, the Chair pronounces *"This is an interesting opportunity, and we've heard a number of good points being expressed on both sides, but we seem to be spinning our wheels. I don't feel that we've organized our thoughts sufficiently yet to take this to a vote. You will recall that we recently spent a lot of time and effort refreshing our Risk Appetite Principles, including the operational resilience tolerances we added last year, and I know we were all pretty happy with the results. It seems to me that this type of situation is precisely why we need to have a good articulation of our risk appetite. I'd therefore like to suggest that we organize our discussion around those core principles and then take this to a vote once we have the benefit of that holistic perspective in front of us."*

Exhibit 1: Selected Risk Appetite Metrics

Risk Appetite Metric	Appetite / Limit	Current	Pro Forma (full ramp-up)
CET1 ratio	Appetite ≥ 11.5% (early warning 12.0%); tolerance 10.0%	13.8%	12.9%
CET1 ratio under enterprise severe-but-plausible stress	≥ 10.0%	10.8%	10.1%
Unsecured lending as % of total loans	≤ 8.0%	5.1%	9.1%
Critical third parties	Due diligence complete and tested exit plan in place before go-live	Compliant	Due diligence incomplete; no exit plan
Tolerance for disruption: critical customer-facing services	≤ 4 hours	Breached last year (19-hour outage); remediation ongoing	LedgerLeaf contractual recovery time of 12 hours
Model risk	All high-materiality models independently validated before use	Compliant	LedgerLeaf model not validated; vendor restricts access to model code
AML/KYC	No appetite for systemic KYC/AML control failures	No open regulatory findings	Customer onboarding fully outsourced; controls untested
Conduct and customer outcomes	Complaint rates within 1.5× portfolio average	Within appetite	LedgerLeaf complaint rate ~3× Maplestone average

Pro forma figures assume the portfolio reaches $2.0 billion at full ramp-up. Stress results reflect the bank's most recent enterprise stress test, with LedgerLeaf portfolio losses of 13% under the severe-but-plausible scenario, versus expected annual losses of approximately 4.5%.

Risk Appetite Case Study Worksheet

Principle	Opportunity Alignment	Notes
Return for Risk Assumed Company is appropriately compensated for level of risk assumed (financial and non-financial)		
Strategic Alignment Risk enables strategic goals Risk aligns with vision, mission This risk taking is critical for delivery of core customer value proposition		
Total Portfolio Perspective Marginal contribution to existing risk profile Diversification/concentration effects are favourable		
Stakeholder Interests Risk appropriately reflects and balances stakeholders' reasonable expectations (shareholders, customers, regulators, rating agencies) "No surprises"		
Culture/Values Alignment Risk aligns with core corporate values Acceptable reputation/brand risk profile		
Capability/Capacity Risk taking is appropriately supported by available risk-taking capacity (financial and other, including management bandwidth and change capacity)		
Operational Resilience & Control Environment Third-party, technology/cyber, model and compliance risks are understood and controlled Critical operations can remain within tolerances for disruption The bank retains accountability for outsourced activities		
Other Principles?		