FXE Technologies

Digital Insights

Rethinking commercial credit assessment

Why agentic AI is changing how lenders assess risk, context and complexity

PERSPECTIVES FOR
FORWARD-THINKING LENDERS

Agentic AI and the future of credit assessment in commercial finance:

Is nuanced assessment transforming complex lending?

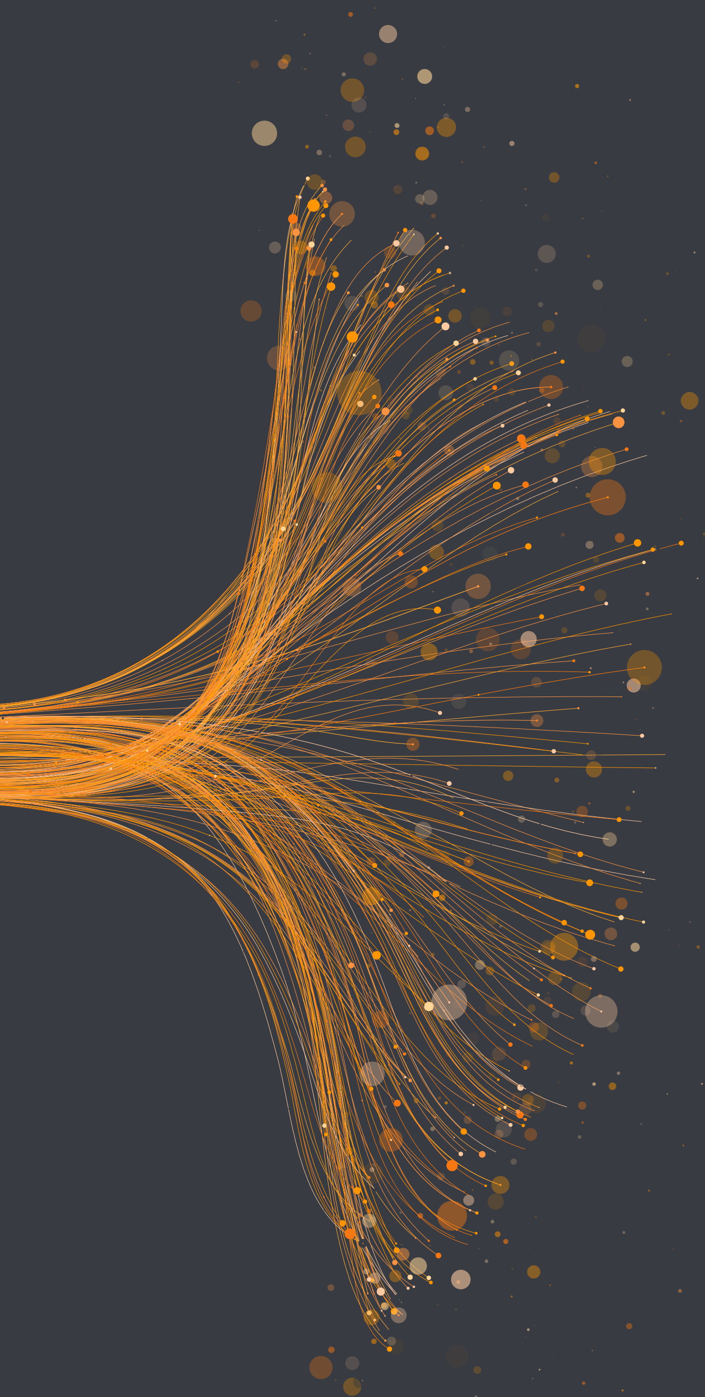

For more than two decades, credit scoring has been the dominant paradigm in credit decisioning. It transformed consumer lending by distilling complex financial behaviour into a single predictive metric that lenders could use to make rapid decisions at scale. Yet commercial finance has never experienced the same transformation.

The reason is simple: commercial lending is fundamentally different. Businesses are not homogeneous entities like consumers. They range from sole traders to multinational groups, operate in vastly different industries, and seek financing structures that vary widely in purpose, scale and complexity.

The next transformation in commercial credit assessment will not come from better scores. It will come from **agentic AI**—systems capable of reasoning across multiple forms of data and contextualising risk signals in ways that resemble how experienced underwriters think.

> Rather than replacing human judgement, agentic AI will augment it. It will provide underwriters with the analytical depth of large-scale data processing combined with contextual reasoning that avoids the blunt simplifications of traditional scoring.

This Insights Report provides a roadmap to how agentic AI can be adopted – not sometime in the future but today.

Why scoring has always struggled in commercial finance.

Credit scoring is, by design, a simplification mechanism. It compresses a wide range of behavioural and financial signals into a single numerical value representing predicted risk.

This abstraction works well in consumer credit because the underlying entities are relatively uniform: individuals or households applying for standardised products such as credit cards, personal loans or mortgages.

Why scoring works in consumer credit

Consumer scoring models operate effectively because they are built on:

- Large volumes of comparable borrowers
- Relatively simple financial structures
- Consistent product types
- Standardised repayment behaviour

When combined with affordability checks, scoring models helped modernise consumer credit more than twenty years ago.

Why scoring struggles in commercial finance

Commercial finance operates in a very different environment.

- Businesses at different stages of development
- Entities ranging from sole traders to multi-entity groups
- Sector-specific revenue cycles and irregular cashflow patterns
- Complex funding structures
- Management capability and strategy

Reducing this complexity to a single score inevitably strips away critical context.

This limitation is visible in practice. FXE data shows that even in **lower-end commercial lending for micro-companies**, traditional commercial credit scores are not performing particularly well. In fact, relying solely on **consumer scores of business owners often produces better outcomes[1].**

Many lenders already recognise this implicitly. Commercial credit scores are frequently ignored or treated as a minor input rather than a primary decision variable.

> The industry has long understood that scoring is an imperfect tool in commercial finance. Until recently, however, there was no viable alternative capable of operating at scale.

Source:

1 See FXE Lending Monitor Q3 2021 page 7 for analysis of performance of director scores vs. commercial scores: https://fxe-technologies.co.uk/app/uploads/2021/07/FXE-Lending-Monitor-Q3-2021_Final-1.pdf

The core idea behind agentic AI

Agentic AI represents a fundamentally different approach to analysis.

Instead of compressing many variables into a single predictive score, **agentic systems reason across multiple data points and maintain their context** throughout the assessment.

At its core, agentic AI operates through a set of specialised AI "agents". Each agent performs a specific analytical task, such as:

- Extracting financial signals from bank statements
- Analysing management accounts
- Identifying sector-specific cashflow patterns
- Reviewing credit bureau data
- Interpreting contractual documents

> These agents interact with one another, combining their findings to form a broader analytical picture.

Unlike traditional predictive models, which output a probability or score, agentic systems generate **structured reasoning.** They can identify:

- Which risk signals exist
- How severe those signals are
- What contextual factors may explain them
- Whether mitigating evidence exists

> This ability to connect signals across multiple structured and unstructured data sources is what makes agentic AI particularly powerful in commercial finance.

Context is everything in commercial credit

A simple example illustrates the difference between scoring models and contextual reasoning.

Consider a business showing the following indicators:

- Rejected payments due to insufficient funds
- Instances of exceeding its overdraft limit
- A late loan repayment

In a traditional scoring framework, these events would likely be treated as negative signals that increase the risk score.

But context matters.
The same pattern might represent three very different situations:

Scenario:
Liquidity distress

The business is struggling financially and may be running out of working capital.

Scenario:
Seasonal volatility

The business operates in a highly seasonal industry and historically experiences temporary cashflow pressure before entering a period of strong seasonal earnings.

Scenario:
Seasonal volatility but additional factors matter

While the business operates in a highly seasonal industry and historically recovers as strong seasonal earnings kick-in, the business has significantly increased use of debt in the last 12 months and expanded into foreign markets.

A scoring model recognises the events but does typically not understand the narrative around them.

Agentic AI, by contrast, can incorporate contextual information such as:

- Historic cashflow patterns
- Changes in the financial profile of the business
- The magnitude of fluctuations relative to previous years
- Expected seasonal earnings cycles
- Shift in market / customer focus

This allows the system to distinguish between **structural financial weakness and temporary volatility.**

This type of contextual interpretation is precisely what experienced underwriters do manually today.

> Agentic AI allows that reasoning process to occur at scale.

Rules-based decisioning remains the foundation

Despite its advanced reasoning capabilities, agentic AI does not replace rules-based credit frameworks. In fact, **robust rule systems remain essential.**

Commercial lending decisions often involve exposures of hundreds of thousands or millions of pounds.

Lenders must be able to demonstrate:

- How decisions were made
- Which policies were applied
- How risk appetite was enforced

Rules-based engines provide this governance layer.

These systems typically evaluate:

- Lending policy compliance
- Credit bureau profiles
- Affordability metrics
- Debt service capacity
- Sector restrictions

Each rule produces a clear outcome:

- **Pass** – within risk appetite – proceed
- **Fail** – outside appetite and cannot be remedied – decline
- **Flag** – requires further assessment – proceed to agentic investigation

> Agentic AI operates primarily within the flagged cases.

The grey area in commercial lending

In commercial finance, the majority of applications do not fall neatly into “good” or “bad” categories.

Even in lower-value lending below £100k, only around 30% of cases are straightforward approvals or declines.In larger commercial transactions, this can drop to 1–10%.

The vast majority of applications sit in a **grey area:**

- The business may not fully meet every lending criterion
- There may be risk signals present
- But those signals may have legitimate mitigating factors

This grey area is where credit underwriting expertise is most valuable—and where most underwriting time is spent.

However, not all of this time is productive: those involved in underwriting often spend **90% or more of their time** gathering information to build the full context required for a decision. Only 10% of the time is spent on making decisions.

For smaller loans, the economics often do not justify this effort. As a result, many grey-area cases are simply declined.

For larger transactions, a significant amount of time is spent on requesting and analysing documents such as:

- Bank statements
- Property Schedules & Rental agreements
- Supplier contracts
- Management accounts

The administrative overhead of assembling this information is substantial.

> Agentic AI can dramatically reduce this burden by orchestrating the collection of relevant information for a specific lending case – ensuring that relevant information is available to assess a case fully.

Agentic AI as the underwriter's analytical apprentice

Agentic AI should not be viewed as a replacement for credit underwriting. Instead, it functions as an analytical apprentice.

The system performs tasks that currently consume the majority of underwriting time:

- Collecting data from multiple sources
- Extracting relevant signals
- Proposing potential mitigating evidence to flags identified by rules-based assessment
- Summarising findings in structured reasoning

The result is effectively a **first draft of a credit paper.**

Underwriters remain responsible for:

- Reviewing the reasoning
- Challenging assumptions
- Applying judgement
- Making the final decision

> This preserves the principle of **human-led decisioning**, while dramatically improving the speed and depth of analysis available to underwriters.

CASE STUDY:

Agentic AI in commercial mortgage assessment

FXE has explored the use of agentic AI in the assessment of commercial mortgage applications.

Commercial mortgage underwriting is among the more complex commercial credit decisions that banks underwrite ‘at scale’ with the value of contracts typically ranging from £250k-£5M in lending.

The analyst must consider:

- The financial performance of the trading business
- The financial position of the business owners
- The structure of the property transaction
- Sector-specific trading dynamics
- Balance sheet strength
- Cashflow management

To support this process, FXE uses a rules-based decision engine that evaluates several hundred criteria across areas such as:

- Lending policy fit
- Affordability
- Debt service capacity
- Credit appetite

Each criterion produces one of three outcomes:

- **Pass** – within appetite
- **Fail** – outside appetite and cannot be remedied
- **Flag** – further investigation required

> The agentic system is designed to investigate the flagged criteria – only after the rule-based assessment has been completed.

Designing a multi-agent architecture for 'flags' in commercial credit assessment

Agentic AI is applied not only to investigate individual risk signals that have been flagged, but to structure the entire credit assessment process in a way that is proportionate, auditable and aligned with how experienced underwriters think.

FXE worked with Cloud Combinator to design and build a multi-agent underwriting system for commercial finance applications using Amazon Bedrock, Amazon Bedrock AgentCore, AWS Lambda and Amazon DynamoDB.

The objective was not simply to automate discrete analytical tasks. It was to create an architecture that mirrors the stages of real credit reasoning: forming an initial view of risk, investigating relevant areas in depth, challenging ambiguous cases, and producing a structured recommendation for human review.

The resulting system is organised into four layers, each performing a distinct role in the overall decision process.

Layer		What happens	Benefits
0	Establishing risk posture	Intake borrower classification, data collection, and initial risk view	Proportionality Governance Reduces unnecessary analysis
1	Specialist domain assessment	Parallel domain agents assess people/conduct and business transaction factors	Depth Specialisation Comprehensive coverage
2	Structured challenge in the grey area	Adversarial agents present opposing interpretations (one argues for progression, one for decline)	Surfaces bias Tests resiliance Stregthens judgement
3	Recommendation synthesis and credit paper generation	Sythesis agent produces recommendation, credit paper, risk register and proposed conditions	Clear output Auditability Decision readiness

0

Establishing risk posture

Every application enters through an initial risk profiling stage.

The system classifies the case across two dimensions:

- Business risk
- Proposal risk

Together, these define four broad categories ranging from lower-risk applications through to cases requiring full forensic analysis.

This initial classification determines how the system behaves downstream. Lower-risk cases are processed with minimal analytical overhead. Higher-risk cases trigger deeper investigation and more extensive evidence gathering.

The design is intentional. Rather than applying uniform analysis to every application, the system first determines what type of case it is dealing with.

This creates three important benefits.

First, it introduces proportionality. Analytical effort, cost and latency scale with complexity rather than being fixed.

Second, it strengthens governance. The reasoning path is shaped by an explicit and explainable assessment of risk posture at the outset.

Third, it reduces unnecessary analysis. Straightforward cases are not subjected to the same level of scrutiny as complex or ambiguous ones.

Notably, the risk profiling stage does not access mitigation data. It classifies the application based only on the information available at intake. This prevents the system from prematurely softening risk signals before the appropriate level of scrutiny has been established.

1

Specialist domain assessement

Once classified, the application is assessed by specialist agents operating in parallel.

One agent focuses on people and conduct-related factors such as personal credit profiles, commercial bureau data and account behaviour. Another focuses on business and transaction-related factors such as company performance, deal structure, funding requirements and property characteristics.

This separation reflects how commercial underwriting is performed in practice. Different categories of information require different analytical approaches and draw on different data sources.

The benefit of this structure is analytical clarity.

Rather than a single model attempting to interpret all signals at once, each agent reasons within a coherent domain:

- Credit and conduct signals are assessed together
- Financial and structural signals are assessed together
- Related evidence is evaluated in context

This improves both precision and interpretability.

At this stage, flagged issues are investigated using a structured mitigation framework. For each flag, the system:

- Identifies relevant mitigating factors
- Locates supporting data
- Extracts specific evidence
- Assesses that evidence against defined criteria
- Documents any remaining uncertainty

This transforms binary flags into contextual assessments.

Instead of simply identifying risk signals, the system evaluates whether those signals are genuinely concerning or have reasonable explanations supported by evidence.

This is a key distinction. The system is not generating narrative explanations after the fact. It is following a disciplined, repeatable mitigation process that mirrors how experienced underwriters analyse risk.

2

Structured challenge in the grey area

For cases where evidence points in more than one direction, the system introduces an additional layer of analysis: structured adversarial challenge.

This stage is only triggered for genuinely ambiguous applications.

Two agents are assigned opposing roles:

- One argues for progression
- One argues for decline

The purpose is to surface competing interpretations of the same evidence.

In manual underwriting, these are the cases that prompt discussion between analysts or credit managers. Different perspectives help determine which risks are material and which are acceptable.

Embedding this logic into the system improves decision quality in three ways.

It ensures that ambiguous evidence is actively challenged rather than passively accepted.

It makes competing interpretations explicit, rather than embedding them in a single line of reasoning.

It provides a clearer basis for human review, as underwriters can see both sides of the argument before forming a final view.

Importantly, different model families are used for this stage. This introduces genuine diversity in reasoning patterns and reduces the risk of uniform bias in how evidence is interpreted.

3

Recommendation synthesis and credit paper generation

The final stage synthesises the analysis into a structured recommendation.

A single arbiter agent produces the outcome and generates a formal credit paper.

This separation between analysis and recommendations is deliberate. It ensures that no single assessment component becomes the implicit driver of a recommendation and preserves a clear chain of reasoning.

The output includes:

- A clear recommendation
- A structured summary of the case
- Analysis of borrower, financials and security
- A documented risk register
- Proposed conditions where applicable

This is not simply a recommendation. It is a complete first draft of the credit paper.

For underwriters, this shifts effort away from assembling information and toward reviewing and applying judgement. For institutions, it improves consistency in how decisions are documented and communicated.

Why the layered structure matters

The effectiveness of this approach lies in its structure. Commercial credit assessment is not a single-step prediction problem. It is a staged reasoning process.

The architecture reflects this by separating:

- Classification of risk posture
- Investigation of relevant domains
- Challenge of ambiguous evidence
- Synthesis into a recommendation

This delivers several advantages for lenders:

It enables proportional analysis, where effort scales with complexity.

It creates clearer governance, with explicit reasoning paths and separation of responsibilities.

It improves auditability, as each stage of the process is structured and traceable.

It aligns closely with how experienced underwriters think and work.

And it provides a foundation that can evolve as new data sources, products and policies are introduced.

It is about structuring judgement so that it can operate at scale without losing context, discipline or transparency.

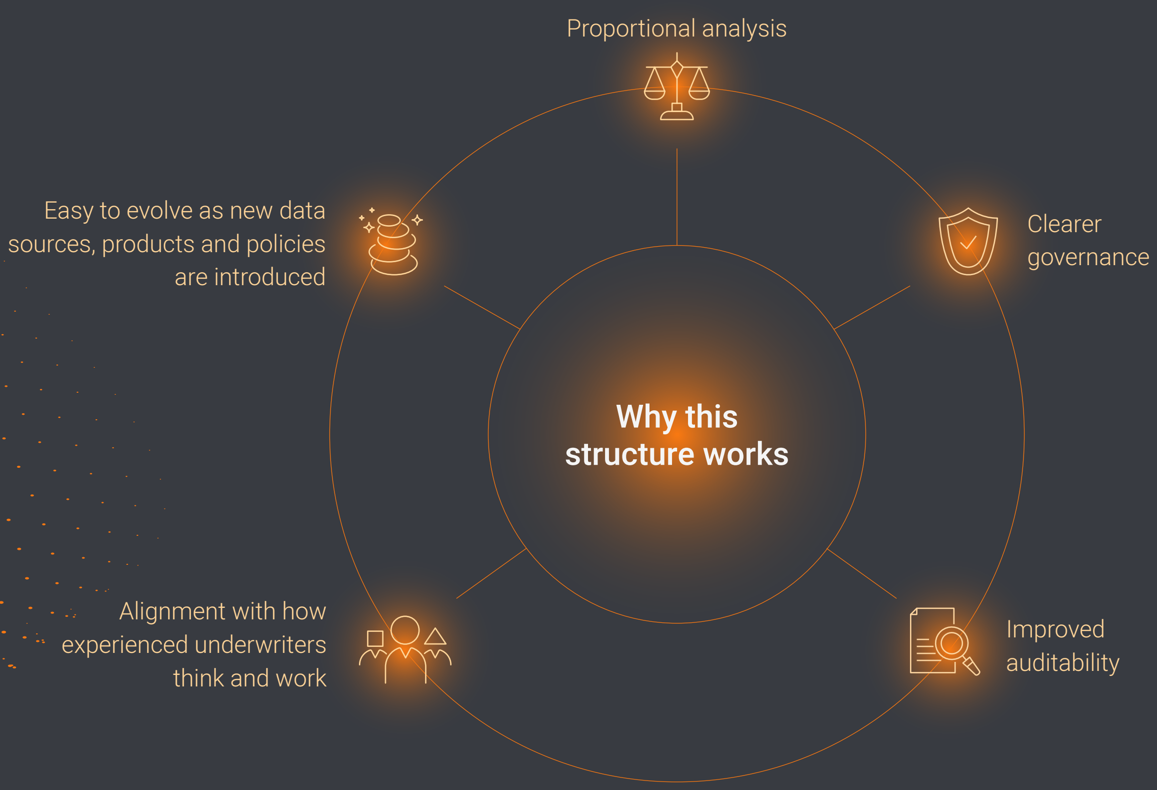

> Most importantly, it demonstrates that the future of commercial credit assessment is not about replacing judgement with automation.

Assisted decisioning in the age of agentic AI

Commercial lending decisions are inherently multi-dimensional. Agentic AI allows lenders to move toward a more sophisticated model of assisted decision that puts the human at the centre of where to deploy funding but supported by nuanced analytics:

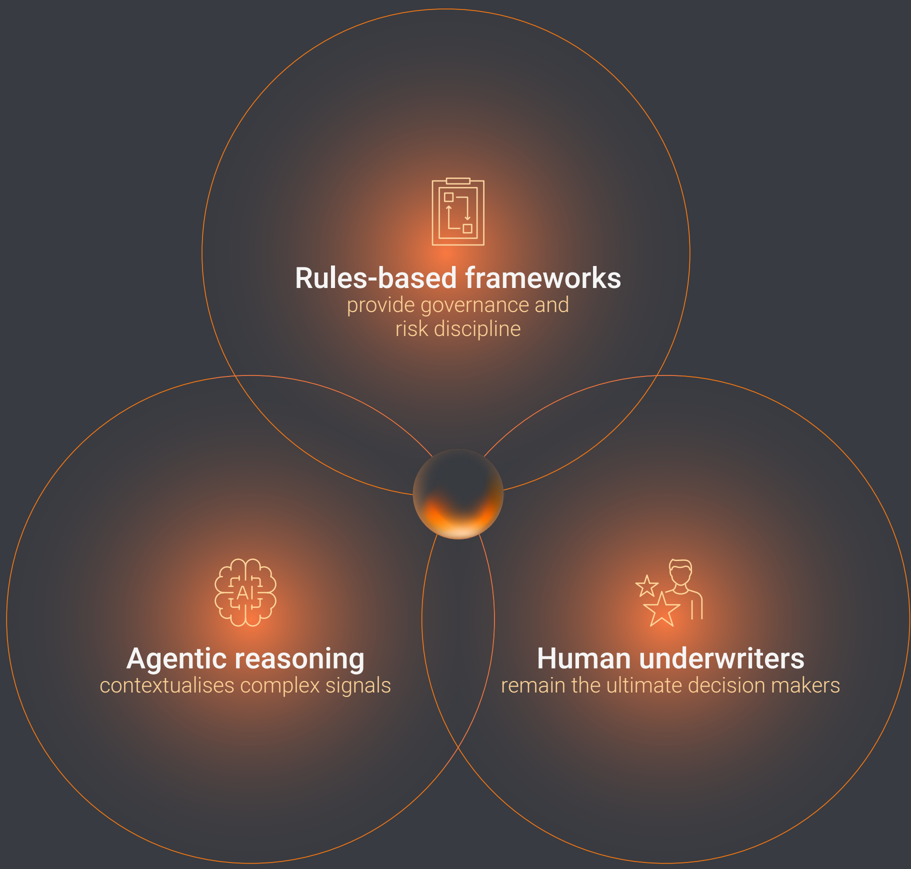

> This combination offers something the industry has long sought but never fully achieved: **Scalable underwriting that preserves human judgement rather than replacing it.**

In commercial finance, where context often matters more than data volume, this shift may prove transformative. The opportunity is to reduce decision-making time by 50–80% by leveraging agentic AI across data ingestion, risk analytics and assessment, and orchestration. This approach not only reshapes credit decisioning itself but also the coordination of activities between customer, business, and risk—shifting away from traditional workflow-driven loan origination towards a model focused on core components, with agents orchestrating the process dynamically. Alongside more efficient development of credit papers, assisted processes also offer greater consistency, underpinned by feedback loops that track outcomes and continuously refine future recommendations.

> In this model, competitive advantage will increasingly belong to lenders who can combine speed, depth, and judgement— not as trade-offs, but as standard practice.

FXE Technologies

FXE Technologies offers a suite of digital SME lending solutions that enable banks, brokers and lenders to instantly triage customers against underwriting models while transforming customer conversion and engagement. FXE Technologies' solutions are used by a range of customers including Tier 1 banks and boutique lenders.

View the FXE Technologies solution suite at **fundingxchange.co.uk/fxe-technologies**

Funding Xchange

Funding Xchange has been a leading provider in the digital assessment of SME lending applications since 2014. The Funding Xchange MarketPlace puts businesses in control of their funding, providing access to 70+ lenders from one simple funding request, enabling them to easily compare terms and apply with confidence and not impact their credit score.

Funding Xchange SME Lending Monitor

We believe that collaboration between banks, alternative lenders, digital technology providers and policy makers is vital to ensure businesses have access to the critical lifeline that funding often represents. This collaboration brings together different capabilities, providing business owners with the ease of access to business finance that the consumer finance market has enjoyed for more than a decade.

Through our marketplace, which is used by over 30,000 businesses across the UK every quarter, we have a front row seat to observe any changes in funding needs – and the funding solutions available to them from more than 40 providers.

For further information on our capabilities and to learn how we help small businesses, please visit: **fundingxchange.co.uk**

FUNDING XCHANGE Rise London | 41 Luke Street | London EC2A 4DP

FUNDING XCHANGE Rise London | 41 Luke Street | London EC2A 4DP