STORM50 2021

Statistical Techniques, Optimization Frameworks and Risk Models

Ranking and Analysis

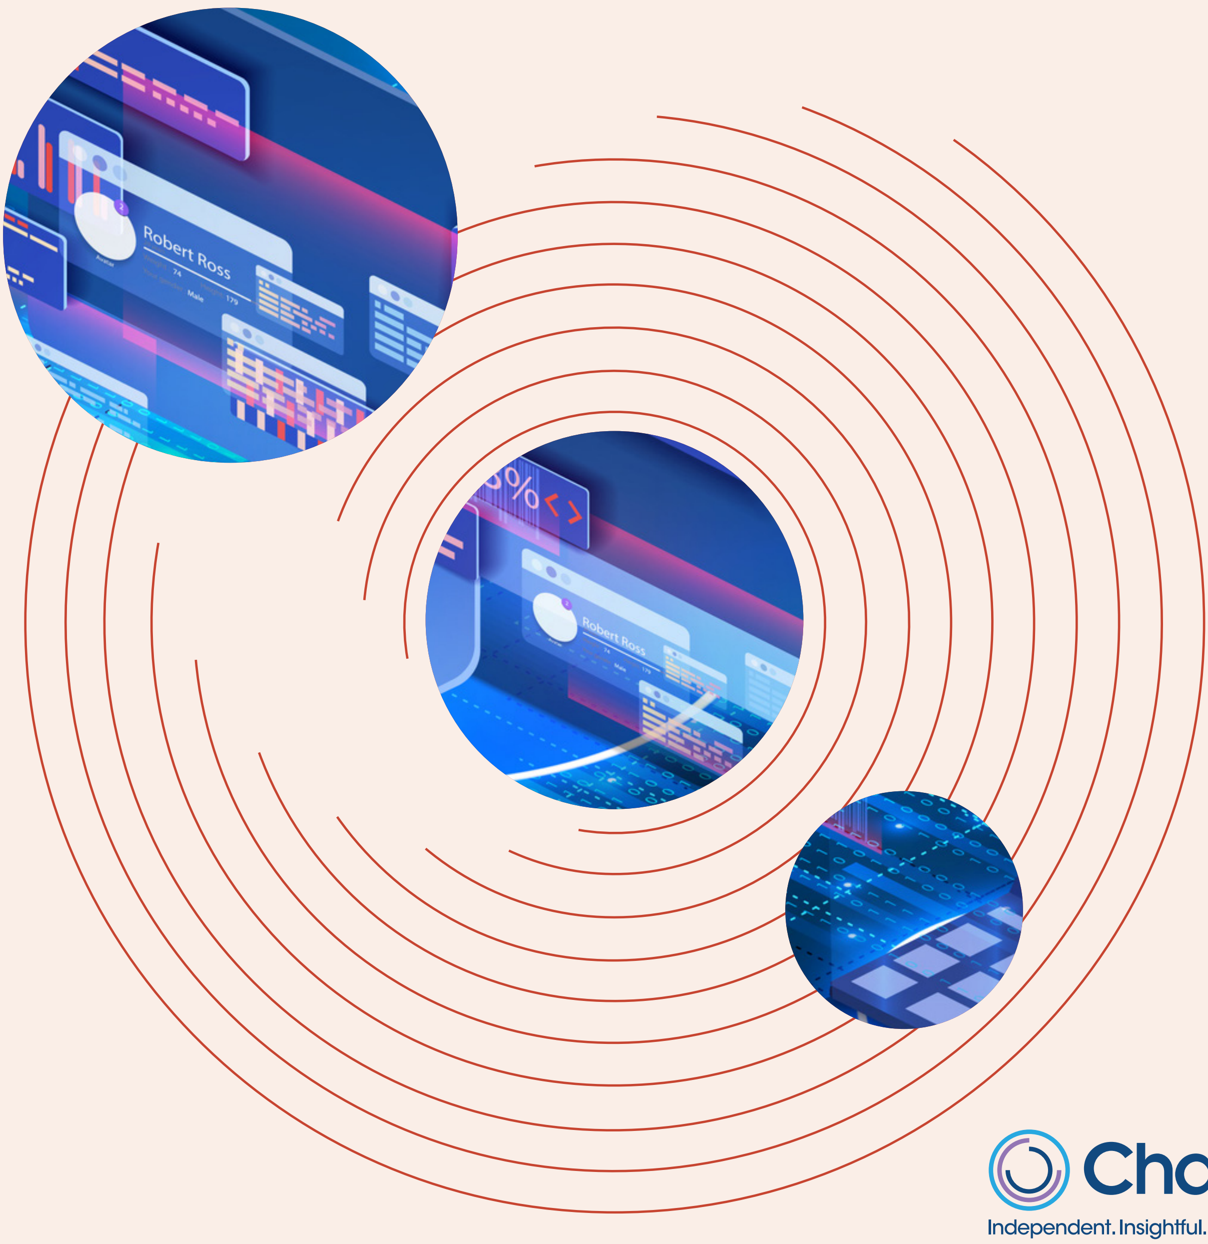

Chartis

Independent. Insightful. Actionable.

About Chartis

Chartis Research is the leading provider of research and analysis on the global market for risk technology. It is part of Infopro Digital, which owns market-leading brands such as Risk and WatersTechnology. Chartis' goal is to support enterprises as they drive business performance through improved risk management, corporate governance and compliance, and to help clients make informed technology and business decisions by providing in-depth analysis and actionable advice on virtually all aspects of risk technology. Areas of expertise include:

- Credit risk.
- Operational risk and governance, risk and compliance (GRC).
- Market risk.
- Asset and liability management (ALM) and liquidity risk.
- Energy and commodity trading risk.
- Financial crime including trader surveillance, anti-fraud and anti-money laundering.
- Cyber risk management.
- Insurance risk.
- Regulatory requirements.
- Wealth advisory.
- Asset management.

Chartis focuses on risk and compliance technology, giving it a significant advantage over generic market analysts.

The firm has brought together a leading team of analysts and advisors from the risk management and financial services industries. This team has hands-on experience of implementing and developing risk management systems and programs for Fortune 500 companies and leading consulting houses.

Visit **www.chartis-research.com** for more information.

Join our global online community at **www.risktech-forum.com**.

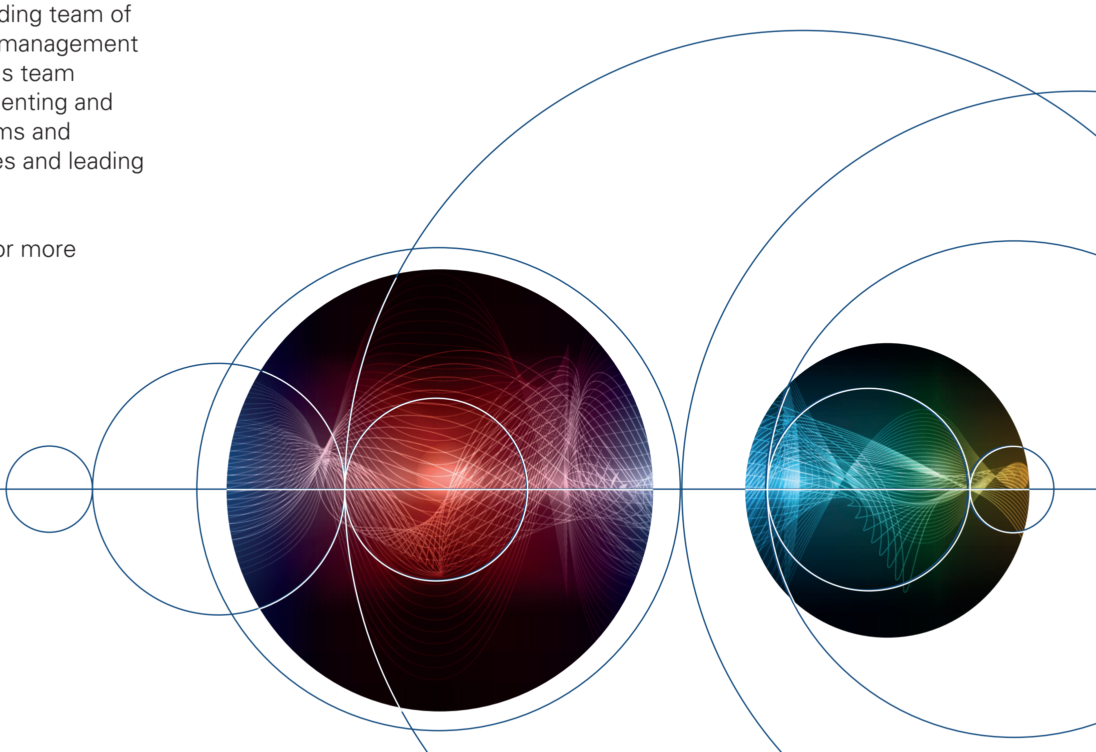

STORM50

Table of contents

List of figures and tables

1. Foreword

I'm delighted to welcome you to the inaugural STORM50 ranking, and its accompanying report. Chartis has been developing industry-leading rankings for many years, employing its detailed and rigorous research and deep industry knowledge and experience. Chartis' flagship report, RiskTech100®, for example, is widely regarded as the authoritative ranking of technology firms across the risk spectrum.

In our research, and our interactions with financial institutions (FIs), vendors and consultants, we have sensed a growing demand for a ranking that focuses on the computational infrastructure and algorithmic efficiency of the vast array of technology tools used across the financial services industry. In addition, structural transformations in the past few years have produced novel statistical modeling approaches (such as machine/deep learning and a range of heuristics) that are now foundational elements of the computing and modeling landscape.

The intersection of process (model governance and validation), technology (programing languages, databases and hardware), new regulations (such as FRTB, IRRBB and PBR[1]) and new modeling methodologies has driven a structural shift toward a more systemic, industrialized modeling sector. In essence, a new discipline is emerging: 'algorithmic engineering in finance'.

In that context, we focus on the following key areas for our new research:

- **S**tatistical **T**echniques.
- **O**ptimization frameworks.
- **R**isk **M**odels (of all types).

…or **STORM** for ease of reference.

The term 'revolution' is often overused in the context of technological change. Nevertheless, we believe that the ongoing process we have identified – the impact of underlying systems software and hardware on firms' choice of models and modeling approaches (and vice versa), and an increasingly structured approach to quantitative techniques in finance – is indeed revolutionary. Indeed, it may even constitute a series of intersecting 'mini-revolutions', because deep structural change is occurring across many different business lines and sub-sectors of the finance industry.

In this report, its updates, and a series of planned reports that will explore specific aspects and dynamics of the STORM landscape, we will document the ongoing revolution, and outline the vendors that are at its leading edge.

With that in mind, I hope you enjoy this report, and will join me in congratulating the vendors and category winners featured in the first STORM50 ranking.

Sid Dash, Research Director

[1] *Fundamental Review of the Trading Book, Interest Rate Risk in the Banking Book, Principles-Based Regulation*

2. A guide to this report

This report combines the STORM50 ranking (and associated category winners) with an overview of Chartis' STORM research – its coverage, conceptual foundation, trend analysis and methodology.

It is arranged as follows:

- *Chapter 3*. **Introduction**.
- *Chapter 4*. **STORM and STORM50: coverage and taxonomy**. A look at the scope of our STORM research, and how we classify the models, analytics and providers we analyze.
- *Chapter 5*. **Revolution and impact: why and how models have evolved**. How and why modern analytics have evolved, and the impact of this evolution on key analytical areas.
- *Chapter 6*. **Building blocks: the technology drivers of change**. The hardware and software developments that have enabled the analytical revolution in financial services.
- *Chapter 7*. **The state of the art: retail analytics in finance**. An examination of the key models and approaches in one of the first areas to make extensive use of artificial intelligence (AI)/machine learning (ML) and other statistical heuristics.
- *Chapter 8*. **Algorithmic engineering and statistical mechanics: the industrialization of STORM**. How the convergence of AI and heuristic tools, and the adoption of DevOps-style approaches, are industrializing statistical methods.

The STORM50 2021 rankings and awards

- *Chapter 9*. **STORM50 2021: methodology**
- *Chapter 10*. **STORM50 2021: ranking**
- *Chapter 11*. **STORM50: category winners**

3. Introduction

Mathematical models have long been the bedrock of finance – the typical modern financial institution can even be viewed as a 'house of models' (see Figure 1). But models are mixed, and vary widely. Historically, some models could be overly simplistic and a poor reflection of the real world – reflecting instead the specific contexts in which they were used[2]. Models developed in one context (actuarial models, say) rarely interacted with or influenced models in other contexts.

This is no longer the case. Over a period of many years, models have started to converge and multiply across different businesses (banking, insurance, the buy- and sell-side, etc.) and approaches (such as the use of optimization for convergence). Not only are the statistical and analytical models that drive modern finance now ubiquitous in the industry, as they evolve and converge they are also causing fundamental changes in the analytical tools most financial institutions now use.

In parallel with these modeling innovations has been an equally radical advance in computing technology. In fact, one cannot think of modern financial systems without considering computers and the evolution of computational methods. A host of functions and capabilities[3] – all of which are fundamental processes in finance – now rely on often sophisticated computational tools, and are often underpinned by implicit or explicit models and analytical frameworks.

This is the context in which the STORM report, and its taxonomy and ranking, have been developed. In this research we track the growing ubiquity of statistical and analytical models driving modern finance, and consider how these are converging and changing. We also examine the interactions between models and the computer technology that underpins them (see Figure 2). Because of rapid advances in computational technology in every relevant area – hardware (graphics processing units [GPUs], AI chips, field-programmable gate arrays [FPGAs], storage and the cloud); databases (open-source, NoSQL, array databases); and languages (domain-specific, MATLAB, Python, Julia, and others) – the process by which models are implemented and executed is now of crucial importance. Within

Figure 1: The typical financial institution is now a 'house of models'

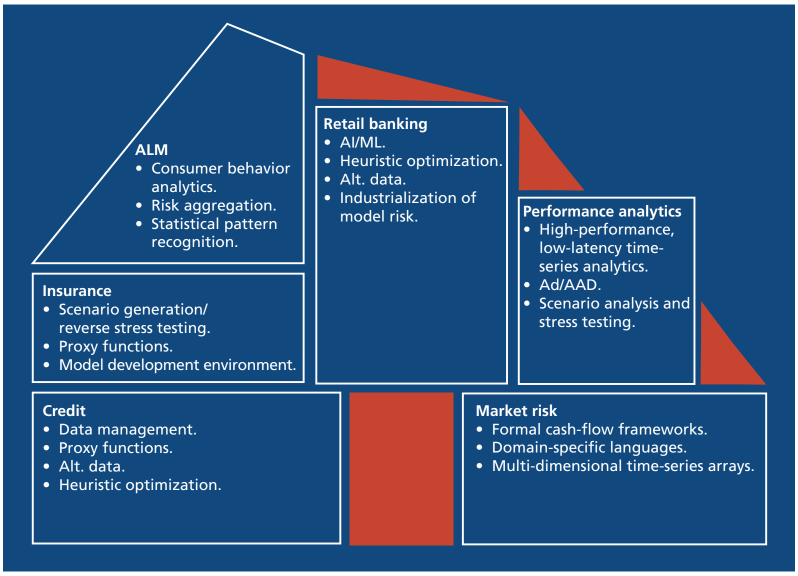

Source: Chartis Research

Figure 2: STORM – examining the interaction between analytics and the underlying computational technology

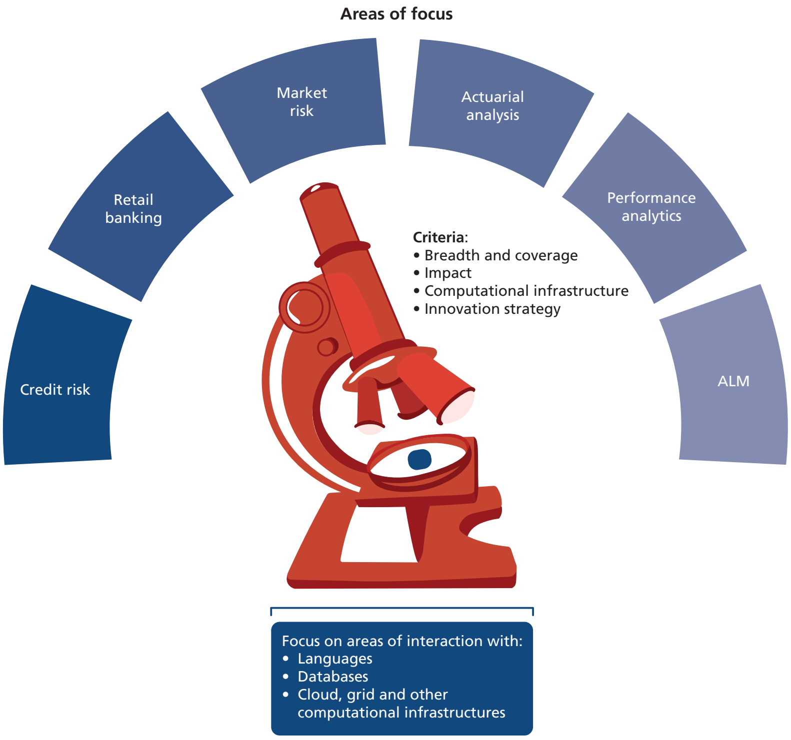

Source: Chartis Research

[2] *Performance attribution models, for example, were developed for asset managers and owners so they could have informed discussions about managerial performance, whereas derivatives P&L explain frameworks were developed for sell-side traders.*
[3] *Including automated trading, efficient bookkeeping, timely clearing and settlements, real-time data feeds, online trading, day trading, large-scale databases, and the tracking and monitoring of market conditions*

this discussion we also emphasize those models that receive relatively little attention elsewhere: the 'intermediate' models (such as cash-flow generators, proxy functions and efficient optimization techniques) that play a vital role in the computational stack.

Finally, we document and analyze the arrival of what we call algorithmic engineering: the increasing industrialization and engineering-style operationalization of the vital models now employed in financial markets, institutions and transactions.

4. STORM and STORM50: coverage and taxonomy

Coverage: what and who

We consider a broad range of providers for our STORM50 ranking and overarching research (see Figure 3), although because there are many we do not attempt to cover all of them. As our focus is on models, model development and the performance of models in the context of 'algorithmic engineering', we place the greatest weight in our scoring[4] on providers of core models, rather than vendors or financial institutions that have repacked or rebadged these libraries, tools or modeling environments.

The Chartis STORM taxonomy: lenses, categories and classifications

Four lenses through which to view models

Whether an institution focuses on capital markets or insurance, every aspect of its business is fundamentally determined by models: how it assesses the state of the market or determines market structure (with yield curves, say, or volatility measures and correlations). All but the simplest situations are now analyzed using models – even apparently robust and simple concepts such as cash flow and net present value (NPV) often have a considerable set of embedded assumptions (see Figure 4). Identifying and analyzing firms' dependence on models (and models' inherent risk) varies in difficulty, depending on where and how the models are being used.

To make sense of the huge variety of analytical and quantitative models, we can examine them through four 'lenses' of analysis, which relate to specific modeling elements:

- Structural mechanics.
- Business line.
- Asset class.
- Time.

These elements have a strong influence on the nature, type, complexity and intensity of the models being used, and the infrastructure and data required to run them.

Structural mechanics

In **mechanical** terms, a model – as many are – can be *statistical, an optimization framework,*

Figure 3: The providers covered by STORM and STORM50

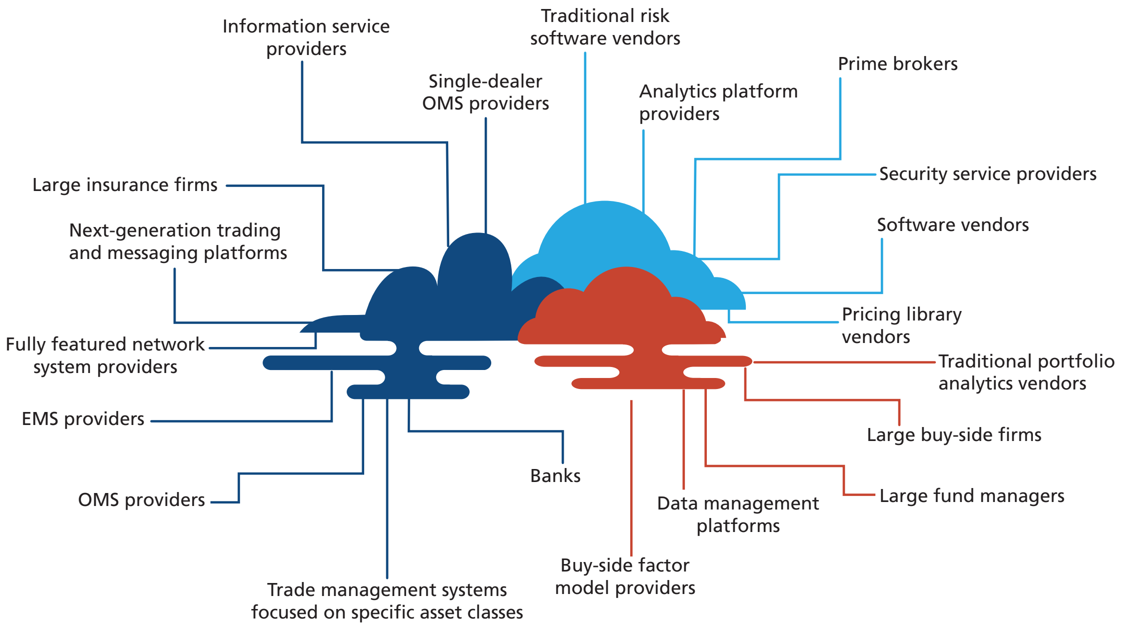

Source: Chartis Research

[4] *For more detail on our methodology, see Chapter 9.*

Figure 4: Everything is a model

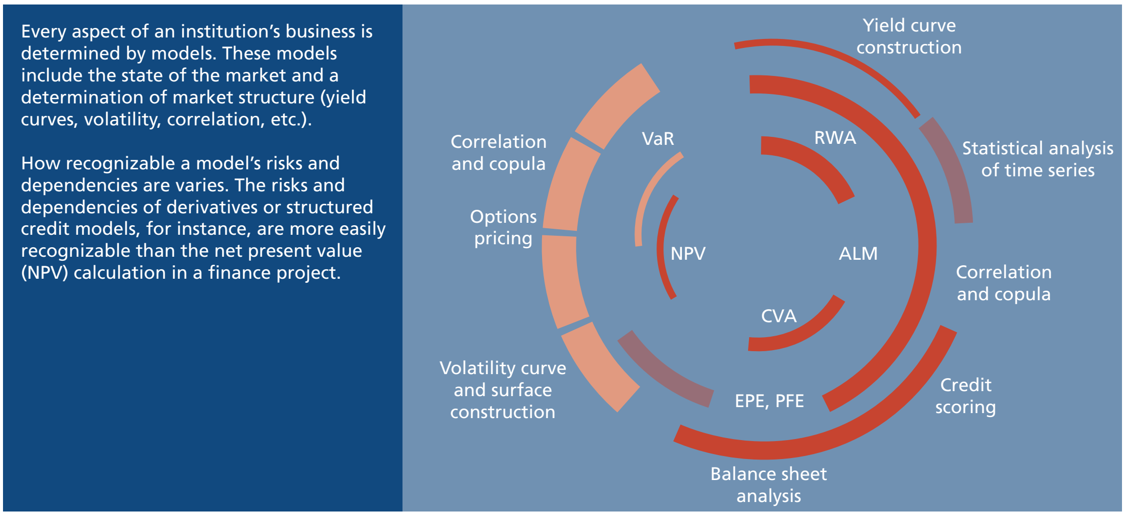

Source: Chartis Research

a partial differential equation (PDE) solver, or a *'state of the world' model* – or even a hybrid of some or all of these. Most derivatives-focused models, for example, have traditionally been some form of PDE solver; portfolio construction involves optimization, while retail banking often takes a highly statistical approach. The method and mechanics a firm chooses to employ will affect its view of the world, and how that in turn is abstracted into models.

Business line

The **business line** in which a model is being used[5] will impose different requirements on the model. Traditional derivatives models, for example, are based on the assumption that all derivatives can be hedged out as much as is operationally and theoretically possible. Retail banks, however, operate under a completely different framework: they interact with many people, so they can extract a large pool of data that can be analyzed statistically. Wholesale banks, conversely (given the relatively 'lumpy' nature of their exposures), adopt a very different approach: the models in that sector are more scenario-based and feature less tightly coupled frameworks.

Asset class

Financial products are wrapped in different regulatory, structural and institutional characteristics, and have fundamentally different pay-off features. Not only do fixed-income products have highly varied pay-off structures, for example, they also have underlying liquidity that differs from that for, say, cash equity products. Equally, highly liquid foreign exchange markets have very different price-formation and institutional dynamics. In essence, this means that derivatives and other products structured on different **asset classes** can have very different modeling requirements.

Time

Time is a vital consideration in model development. How long does a firm expect to have to deal with a transaction or exposure? In insurance-style models, for example, firms know

One among many

The Chartis STORM taxonomy is not the only way to classify risk and analytics, nor is it the only system that Chartis uses. We also employ our pillar classification framework, outlined in our research agenda*, and specific sector and sub-sector classifications**. Our STORM approach, however, is designed to examine common algorithms and execution styles together, focusing on the nature of algorithms and how they are constructed and managed, and how analytics are framed.

* *Chartis Research – Research Agenda: Q1 2021 Update*

** These include credit risk (*Technology Solutions for Credit Risk 2.0: Vendor Landscape, 2019*); market risk (*Sell-Side Enterprise Risk Management Technology, 2019: Market Update and Vendor Landscape*), and several different ways to classify risk delivered to the buy-side (*Risk as a Service for the Buy-Side, 2018*).

[5] *Retail banking, say, or wholesale banking, derivatives trading, fixed-income trading, equities trading or insurance.*

that they will have an exposure for the lifetime of a contract. Conversely, the model in an exchange is based on an expectation that a firm can exit a transaction at any moment.

Surprisingly, few firms think deeply about latency, despite its strong influence on the type of analytics they employ. In this context the use of scenario analysis in particular, which employs a variety of unstructured analytics or less clearly derived structured analytics, is critical.

Six categories of analytics

The STORM classification has six broad categories (see Figure 5) – *retail analytics, credit risk, actuarial risk, market risk, performance analytics* and *asset and liability management (ALM)* – and aligns closely with the overarching business lines and structures in financial firms. While there are overlaps and increased convergence across business lines, in our view these separate segments have structural differences that warrant significant separation and demarcation in our analysis.

Three main classifications

Finally, we classify individual models according to three primary categories: *model type, key features* and *analytical features* (see Figure 6).

Figure 5: STORM covers six categories of analytics

Retail analytics
- Retail credit scoring.
- Leveraging behavioral analytics.
- Structuring linkages between core analytics and downstream advisory and services.

Credit risk
- Revolution in credit risk triggered by shifts in accounting norms and/or the rise of risk-aware accounting (IFRS 9/CECL).
- Integrated credit and interest rate modeling.

Actuarial risk
- New modeling tools (such as extended scenario generation and reverse stress testing).
- Large datasets.
- Hardware-driven acceleration.

Market risk
- Growth of domain-specific languages.
- Computational accelerators having significant impact.
- Proxy generators and functional approximations.

Performance analytics
- Convergence between different styles of performance analytics.
- Increased analytical sophistication of fixed-income attribution, as well as TCA for non-equity asset classes.
- Increased formalization of P&L explain and other risk-based allocation/attribution frameworks.

ALM
- Revolution in the mathematical core.
- Growth in behavioral models.

Source: Chartis Research

Figure 6: Three categories of analytical classification

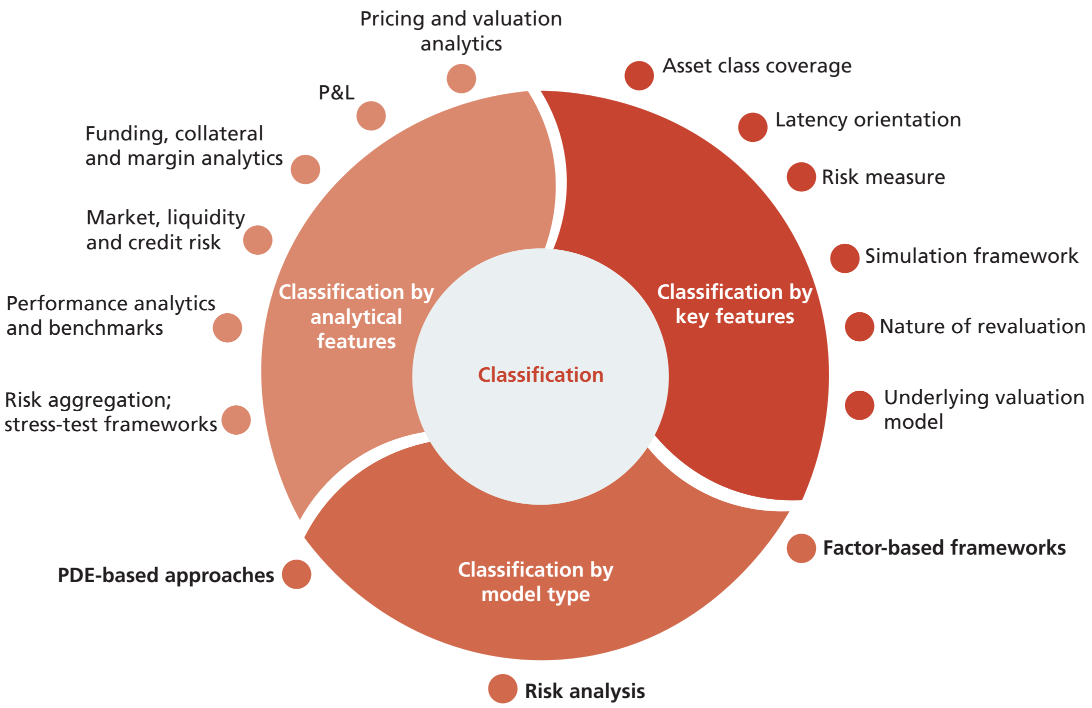

Source: Chartis Research

5. Revolution and impact: why and how models have evolved

Convergence into STORM

As the structural revolution in the statistical techniques and optimization frameworks in financial services has progressed, a variety of overlapping heuristic approaches (variously described in different contexts as AI, ML, deep learning [DL], evolutionary programming [EP], etc.) have revitalized the modeling universe. As we explore in Chapter 6, this has been enabled by rapid evolution in the computational infrastructure for hardware and software. Many statistically driven approaches (such as natural language processing [NLP]) are now ubiquitous in areas such as text analysis and data management.

Equally, this heuristic optimization has also broadened the range of problems that firms can target with analytics. Many problems around data filters, behavioral analytics, customer segmentation, customer risk profiling and data transformation, for example – often traditionally targeted with rules-based methods – can now be driven by statistical frameworks.

There is an element of risk in this, however, as these approaches become more widespread. Small errors and inadvertent assumptions can easily creep in, with a potentially deep structural impact. Underlying errors and assumptions can be buried in modeling structures that can be conventional and difficult – although not impossible – to analyze. As we repeat throughout our research, machine/deep learning (and indeed most other heuristics – even EP) can ultimately be analyzed, and analysis of the surrounding data is central to understanding the problem. However, in some respects, unfortunately, terminological inexactitude and the highly anthropomorphic terms employed (such as 'bias', 'learning' and 'explainabiliy') obscure the data-centric issues at stake.

We believe that novel techniques will ultimately evolve by merging into a broader collection of relatively conventional tools, driven by several market and technological forces (see Table 1). Crucially, this is because, for us, there is no separation between standard statistics, optimization, and risk management and heuristic or AI-driven techniques. Rather, they constitute an overlapping set of statistical techniques, optimization frameworks and risk management models.

Explosive growth in implementation intermediates

Alongside convergence, the use of several types of intermediate models[6] has increased in recent years (or, rather, the importance of these intermediate statistical structures has been increasingly recognized). Of particular note has been the growth in the variety, diversity and power of statistical and optimization capabilities – much of which is directly related to the corresponding growth in computing power.

Intermediate technology stacks (see Figure 7), which sit between the applications themselves and the underlying core compute layer, are important in the development of any set of applications, and must be considered when implementing algorithmic techniques. At the very top are the big chunky intermediates, then comes

Figure 7: Intermediary models are growing in importance

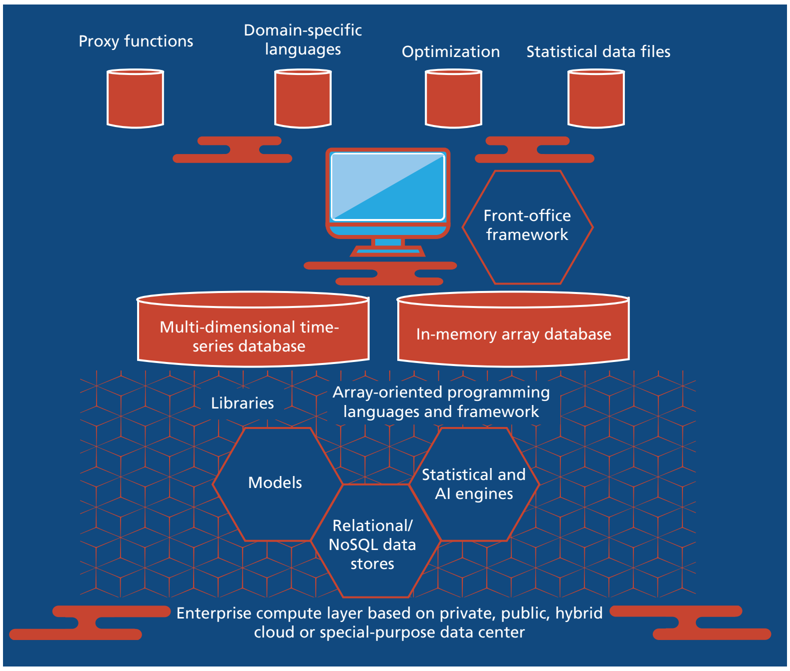

Source: Chartis Research

[6] Models that provide intermediate analytical results – such as time-series statistical frameworks, function approximators and adjoint algorithimic differentiation techniques.

Table 1: The forces driving model convergence across technology stacks

Application/technology stack	Market dynamics/forces	Examples
Models in each of the major verticals	Cross-vertical integration and convergence, such as the increasing convergence of insurance and banking ALM, and the convergence of different credit models and approaches across verticals.	Credit scoring, derivatives pricing, behavioral modeling, funds transfer pricing (FTP) curve construction, ALM modeling
Model intermediates	Universal use of new and non-linear statistical and optimization intermediates, particularly in areas (such as retail banking) where standard frameworks have had mixed results. Even in wholesale finance there is rapid growth in the use of statistical intermediates. Other types of intermediates, such as domain-specific languages and function approximators, are increasingly entering the mainstream in areas such as over-the-counter (OTC) derivatives.	AI/ML and other statistical constructs, heuristic and traditional optimization, domain-specific languages, cash-flow generators
Databases* (and data management infrastructure)	There has been a huge proliferation in the variety and diversity of databases. Ideally, providers of modeling applications should focus on aligning the necessary data structures and databases. Many of the more popular databases in use are open-source, and there is a strong trend toward applications that employ heterogeneous databases.	MongoDB, Cassandra, Hazelcast, Oracle
Languages*	The rapid evolution and sophistication of languages and their available ecosystems has been an underrated element of the trend toward 'algorithmic engineering', as has the ability of programmers to handle a variety of languages in multilingual environments.	Python, Julia, R, MATLAB
Hardware*	The variety of hardware-driven programming options is vast, and includes GPUs, central processing unit (CPU)/GPU combinations, ARM-based platforms and emerging AI-driven chips and FPGAs. This abundance of choice has created something of an 'option paralysis' – or at least lower use of these tools than we might expect.	GPUs and FPGAs provide application- and context-specific acceleration

** See Chapter 4*
Source: Chartis Research

a range of technically oriented intermediates that systems must interact with, with the correct linkages to enable algorithms to perform well and appropriately.

A variety of time-series and optimization techniques have been in use for some time. But following the relatively recent growth in ML/DL tools, and the development of a huge range of heuristic optimization techniques, firms are re-evaluating how they think (and need to think) about the supporting data frameworks. Other mechanisms are entering the mainstream, enabling the broader industrialization of models in finance that we discuss in Chapter 8. Some of the new mechanisms include domain-specific languages that can enable users to develop concise notations and expert vocabulary and abstractions.

While domain-specific languages have had a deep and fundamental impact in the OTC derivatives sector, we believe they can be leveraged in every area of finance, and as such will have a profound impact on the way that vendors develop and structure their applications. Also worth noting is that the importance of appropriate programming languages for different applications is increasingly receiving significant and widespread consideration. The choice of language and ecosystem – such as Python with its vast ecosystem of applications and tools – is becoming increasingly important[7].

The availability of intermediates, as well as new data structure alignments, hardware alignments and language intermediates, and the impact on performance they have, highlight how the mechanics of implementation (i.e., **algorithmic engineering**[8]) is now as important as the model framework itself.

The dominant evolutionary themes

Against a background of increasing model convergence, we have identified four dominant themes in the evolution of models:

The increasing intersection and convergence of models is leading to greater '**hybridization**'. We are seeing many of the approaches developed in insurance (such as scenario generation, for example) increasingly being applied in a banking context. We are also seeing strong use of statistical techniques to preprocess, transform or modify the context of derivatives pricing.

The **success of the 'brute force'** approach to analytics (i.e., computational styles that are either model-free or model-light, including AI/ML/neural networks and other similar statistical mechanisms). The cost of brute-force approaches is lessening (enabled by lower-cost hardware and other system software innovations, such as GPU databases), and we are seeing greater use of statistical and brute force mechanisms across the board, in the final and intermediate stages.

The **growing use of scenario analysis**, and the **development of scenario and stress testing as a science in its own right**. Following the 2008 financial crisis, scenario analysis made a comeback, having been displaced in previous years by models such as value at risk (VaR). A realization that the structure of the market – and by extension the structure of models – was not perfect highlighted the risk inherent in every model, regardless of its structure. Despite fumbling their opportunity to tackle imbalance and structural problems in the market, regulators got one thing right: their increasing mistrust of the precision of models led to a heightened focus on scenario and stress testing. This theme also links to the first two. It encourages both hybridization (a model in one class can be moved into another) and brute force computing (many scenario models are well equipped for mass-scale brute force computational capabilities, given the appropriate scenario structures.

The **growth in reverse stress testing**. In many ways reverse stress testing is the inverse of scenario generation: it seeks to translate broad situations into specific contexts (the range of prices at which a portfolio will drop, for example), by examining a set of findings and devising appropriate scenarios that create that result. Again this can encourage hybrid and brute force computing techniques. While reverse stress testing is not a cure-all, in our view it represents a powerful force for standardization and unification, and for the analysis of contexts that can be difficult to achieve with other approaches.

[7] We will explore this issue later in this report and in subsequent publications, paying careful attention to algorithmic engineering and the physical substrates that can dramatically widen firms' modeling options, often with a very significant economic impact. We will also consider the ecosystems of several quant-friendly languages and programming environments. These include Python, R, Julia, MATLAB and the Java ecosystem (as well as mainstays such as C++ and Fortran). We will also examine how vendors and developers across the board should think about the optimal language for their particular contexts and applications.
[8] Explored in more detail in Chapter 8.

Figure 8: The technical and analytical evolution affects all categories of analytics

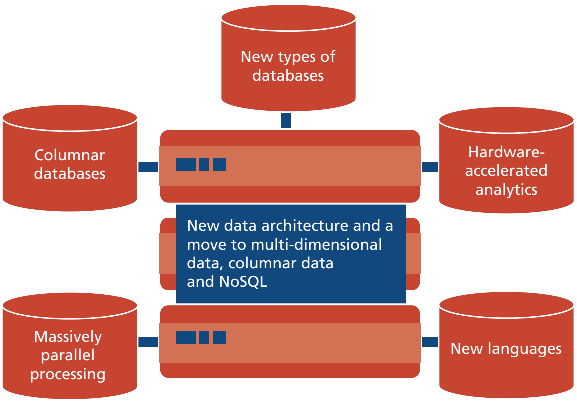

Source: Chartis Research

Figure 9: The changing face of credit in four key areas

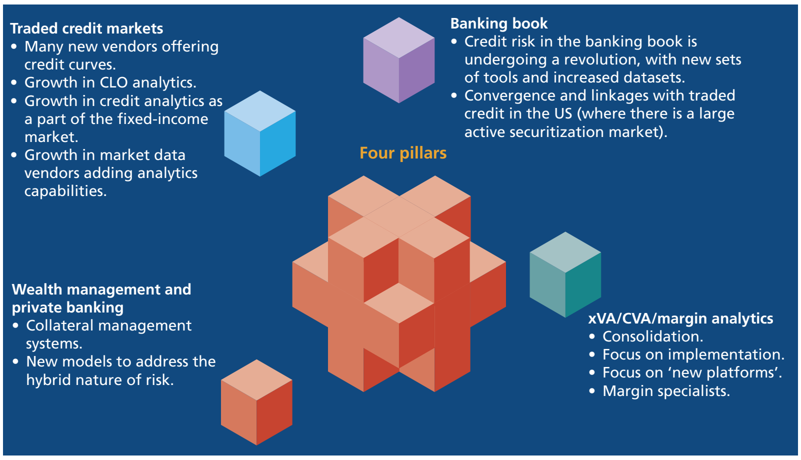

Source: Chartis Research

The most significant challenges with reverse stress testing frameworks usually concern computational technology and algorithmic engineering. Because of the huge difficulty involved, reverse stress tests constitute one area where careful attention to computational substrates can pay off dramatically.

The impact of evolution on the six key analytical areas

The ongoing evolution in models is affecting the main areas of analytics, which interconnect with the technical developments outlined later in the report (see Figure 8). But the impact is especially profound in four areas: *credit, market risk, ALM* and *actuarial modeling.*

Credit: different applications require different analytical components

Alongside existing statistical frameworks, a variety of tools (such as ML and graph analytics) have transformed the environment for credit processes and analytics, especially in banking (see Figures 9 and 10). So while ML-style models have had the most dramatic impact in retail banking, firms have been leveraging traditional tools (such as simulation engines and stress testing) for banking book frameworks. The evolution of algorithms and heuristics has had an especially significant impact on credit portfolio management (see Figure 11).

Figure 10: The re-engineering of credit risk

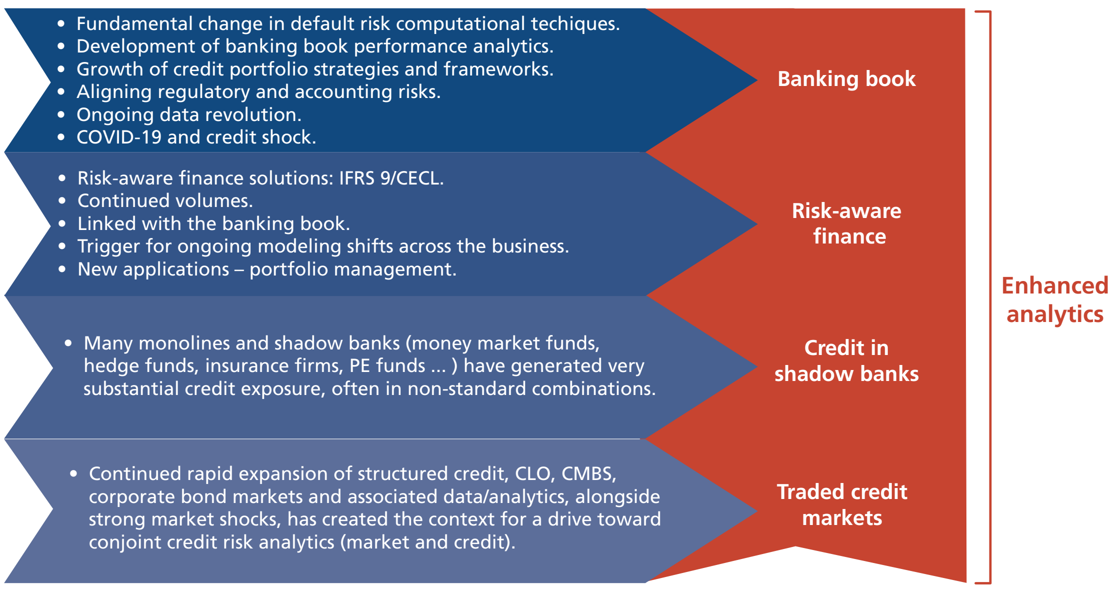

Source: Chartis Research

Figure 11: A diverse range of analytics use cases in credit

Credit analytics/credit flow process	ML	GA	SE	ST	MLB
Margin analytics	●	●	●	●	●
Derivatives counterparty risk	●		●	●	
Credit modeling	●	●	●	●	●
Credit portfolio management	●	●	●	●	●
Risk-aware accounting			●	●	●
Fraud analytics	●	●	●	●	●
Credit trading	●	●	●	●	●
Credit-adjusted credit and balance sheet optimization	●	●	●	●	●
Retail credit scoring	●	●	●	●	●
Behavioral analytics	●		●	●	●
Credit control and limits management	●		●	●	●

ML Machine learning **GA** Graph analytics **SE** Simulation engines
ST Stress testing **MLB** Market-linked behavioral analytics

Source: Chartis Research

Banking: the big story

Banking is evolving and becoming a much broader activity, and conventional banking-like activities have migrated from banks to a broad array of FIs ranging from private equity firms and asset managers to FinTechs. So while particular institutions may behave like banks, certain elements of their operations (such as funding models, for example) will differ, while others (such as credit pricing models) may remain the same. Even in the credit space, there are many opportunities for other types of institutions to structure the nature of credit in new ways.

This ‘shadow banking’ now accounts for a large share of what could be considered ‘banking activities’, and is driving interest in more ‘classical’ banking-style analytics (including ALM, credit, etc.). This is creating a much bigger universe for banking-related algorithms, and driving a structural change in the market that is influencing the way that banks and non-banks use analytics (see Figures 12 and 13). The delivery of banking-style services by a range of non-banks is, in our view, one of the major structural shifts in finance. These ‘new’ styles of institution will, of course, need to leverage banking-style analytics, but their precise requirements and contexts are significantly different.

Figure 12: As banking changes, the universe of banking book-style models is widening

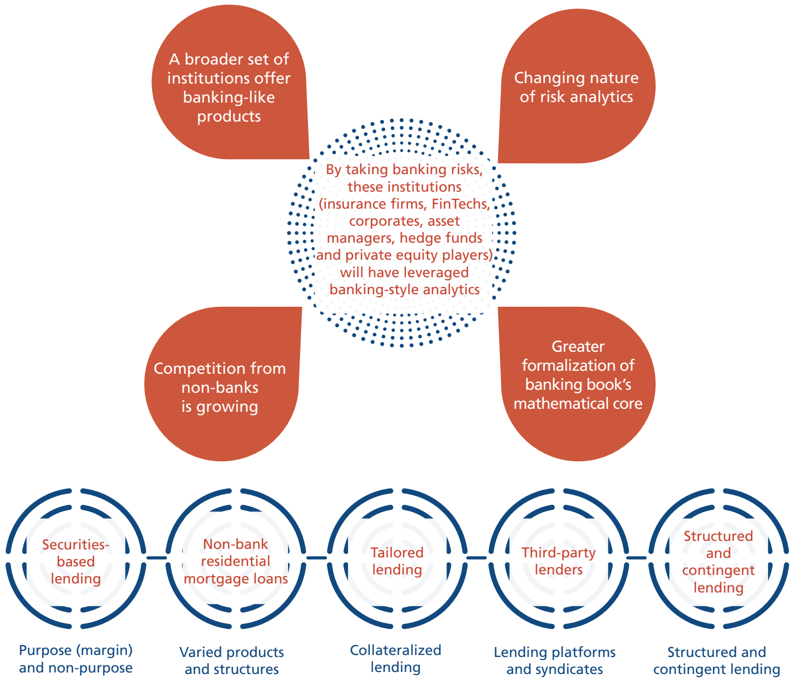

Source: Chartis Research

Figure 13: Banking book-style models are being used in a broader range of institutions

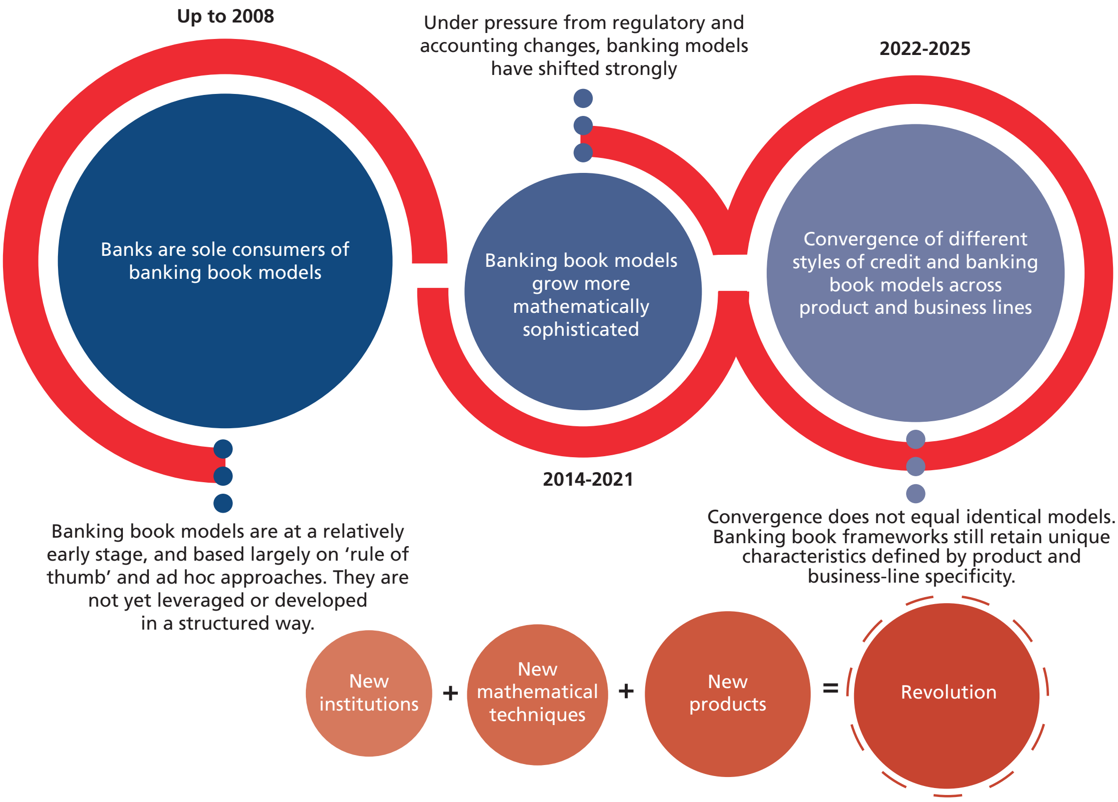

Source: Chartis Research

Market risk: a regulatory-driven structural change

The shift in the market risk modeling environment for banking (see Figure 14) is an area of significant growth for analytics.

This has been driven largely by regulation, especially FRTB, which has caused a slow but steady transformation in market risk practices, and over time has helped to introduce a more industrialized market risk environment. FRTB has focused firms' attention on the data management practices associated with enterprise market risk, including the process of generating and managing risk factors.

FRTB has also triggered a wider revamp of market risk architectures in a broad range of institutions, and enabled broader adoption of a variety of useful technology tools, including in-memory databases, GPU databases and data grids.

ALM: more sophisticated techniques

Like much else, the process of asset and liability management in the banking book is undergoing a structural revolution, with rapid

Figure 14: The forces driving change in market risk

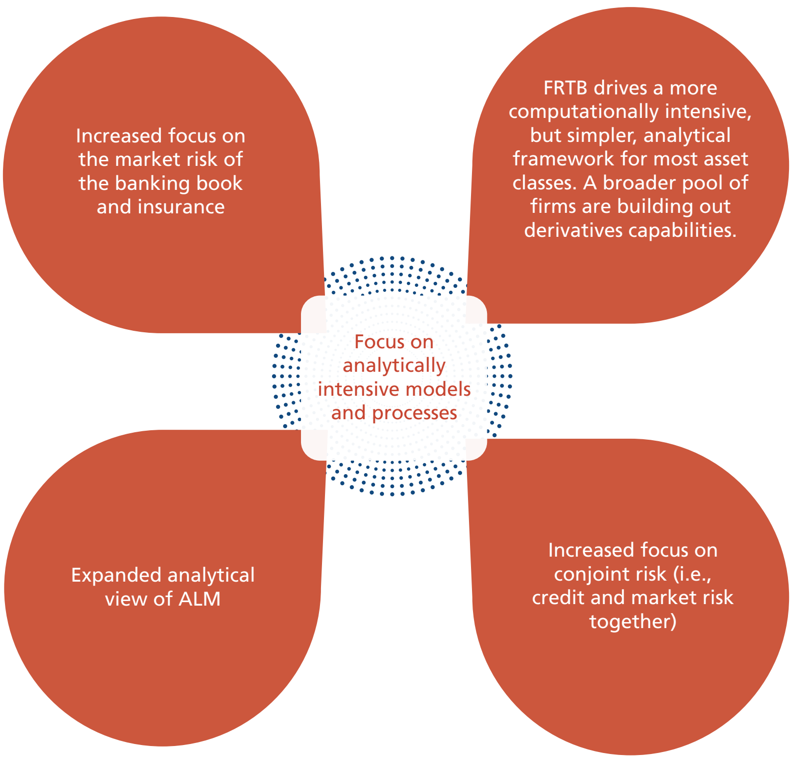

Source: Chartis Research

growth in behavioral modeling, a formalizing of statistical analytics and data management, more sophisticated data aggregation techniques, and the convergence of models across geographies and financial institutions (see Figure 15).

The focus on capital optimization and management has also strengthened, as the cost of regulatory capital has risen. A low (and in some cases negative) interest-rate environment has also made managing interest-rate risk a critical task.

Actuarial modeling: buffeted by regulation

The structural revolution in actuarial modeling is being catalyzed by firms' incorporation of a range of statistical and quantitative models from the banking sector (see Figure 16), as well as an increase in the size of the datasets available. This is happening as actuarial modeling is buffeted by regulatory shifts (such as International Financial Reporting Standard [IFRS] 9 [on the asset side], IFRS 17, Long Duration Targeted Improvements [LDTI] and PBR). Meanwhile, the underlying mathematical frameworks – economic scenario generation, scenario modeling, reverse stress testing, flexible simulation and behavioral modeling – are all evolving into more rigorous formal disciplines.

Figure 15: The changing structure of ALM

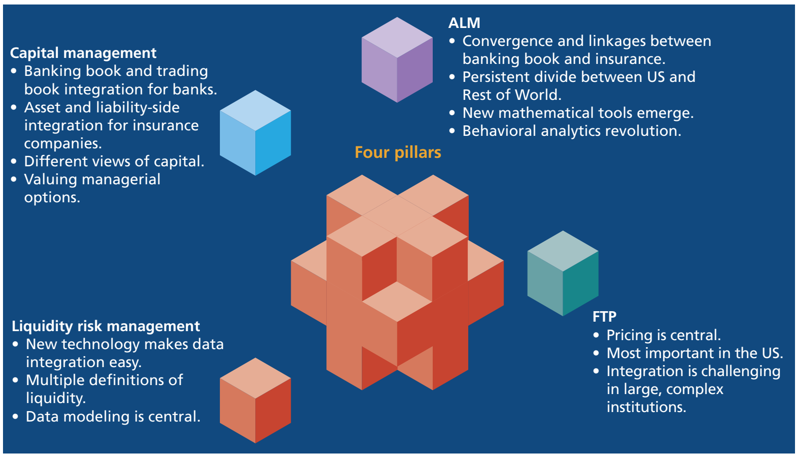

Source: Chartis Research

Figure 16: Actuarial modeling is incorporating models from banking

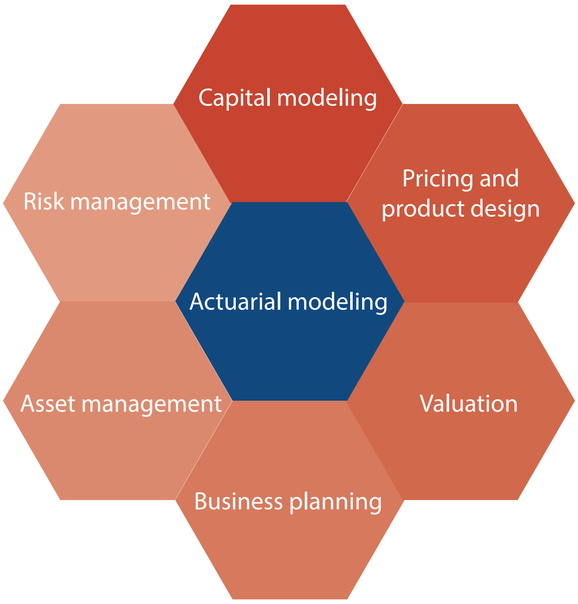

Source: Chartis Research

6. Building blocks: the technology drivers of change

Three main drivers of change

Within the broader trends that are shaping the evolution and development of algorithms (see Figure 17), three main technology drivers stand out:

- Changes in hardware and computation.
- The emergence of mathematically friendly languages with huge open-source ecosystems (notably Python and its array-oriented flavors).
- The proliferation of open-source tools and databases, and the growing options for exploring how these work in particular situations.

Hardware and computation: a daunting choice

For FIs, the diversity and variety of available hardware options can be dazzling and daunting. The enormous range of hardware combinations and strategies they can use to build applications and accelerate the performance of their models (see Figure 18) can create possible choices so diverse as to be potentially debilitating. In our research we have observed that in the vast majority of cases, modelers/quants and IT professionals will build quantitative applications firmly within their comfort zones, rather than truly exploring the alternatives.

The data-parallel approach to hardware in particular, in which the compute element moves to data, rather than the other way around, has been gaining considerable momentum in the past 10-20 years. Much of this momentum is coming from outside financial services – from pure graphics, for example, and the video game and consumer software industries. Given that this is likely to strengthen the data-parallel industry and make the technology more cost-effective, financial firms would do well to familiarize themselves with these approaches and their considerable computing power.

As for the hardware itself, GPUs offer a powerful data-parallel infrastructure, but often benefit from redesigned algorithms. And although FPGAs can enable custom parallelism, simultaneously optimizing the structure of an algorithm and its parallel computational architecture can be a highly complex task.

Figure 17: The broad technology trends shaping the analytical revolution

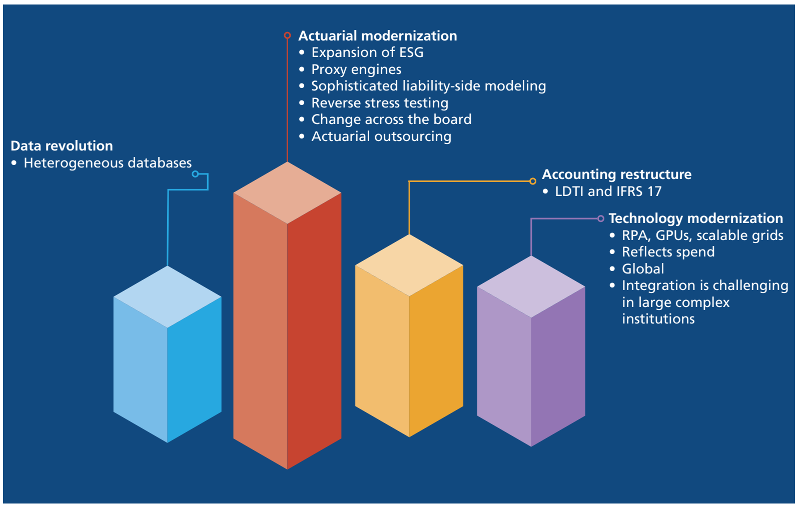

Source: Chartis Research

Figure 18: Hardware – a complex range of options

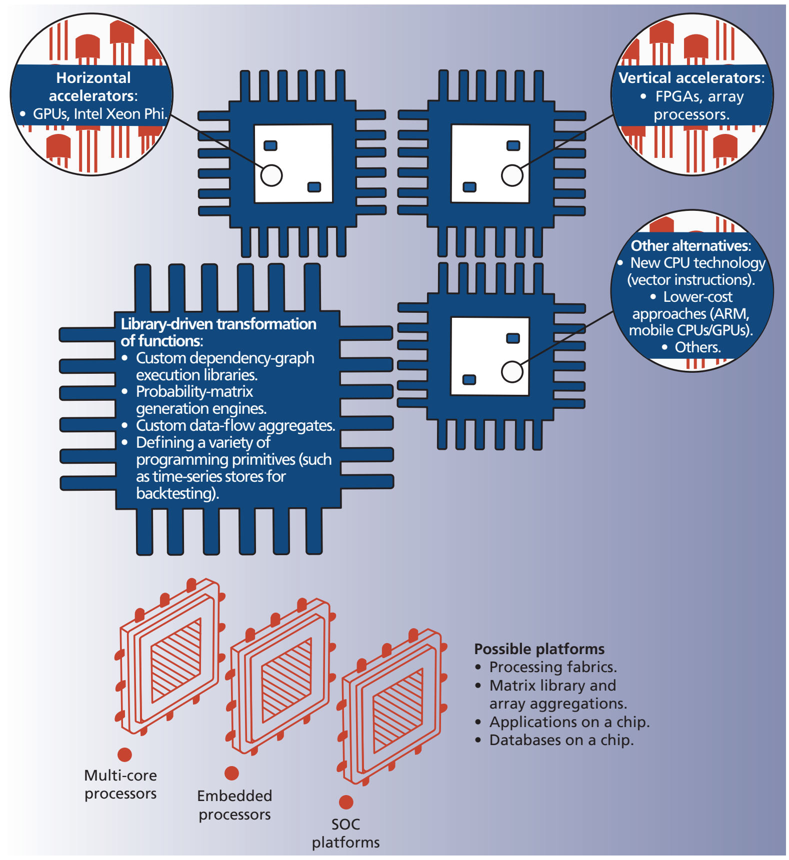

Source: Chartis Research

Figure 19: FIs should view the cloud architecturally

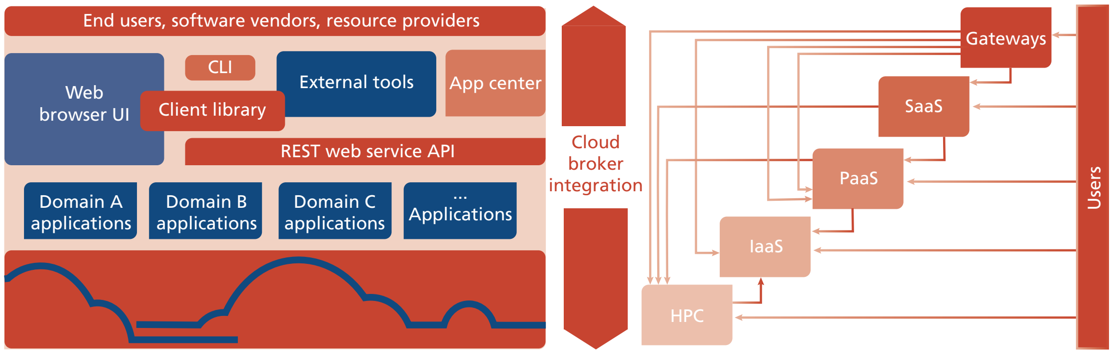

Source: Chartis Research

CPUs, GPUs and FPGAs: strong traction

Our research with vendors suggests that one possible solution to several key functional and methodological challenges in the development of analytics would be to combine CPUs and GPUs, to provide mutual support for different data-parallel styles. Meanwhile, the use of GPU/FPGA databases (as explored later in this section) has been increasing.

Cloud: an architectural view

Creating a robust architecture to support algorithmic development is crucial, and clearly the cloud promises a flexible, cost-effective option. For the vast majority of vendors and FIs, this cost-effectiveness is the lens through which they choose to view the cloud. We believe, however, that firms should view the cloud from an architectural perspective: as a set of architectural principles that can enable algorithms to work from virtually anywhere (see Figure 19). This can even provide a cloud architecture that is not necessarily on the cloud – the system could be hosted internally (or in a custom data center), for example.

At the moment, however, most firms' main focus is on simplifying elements of the cloud architecture. Chartis believes that fewer than 10% of firms on the cloud take full (or anywhere near full) advantage of cloud environments and architectures. A failure to take the architectural view, we believe, could create longer-term issues for firms, because they may find that they can't 'unbuckle' or combine algorithms easily. They could end up with a static implementation that prohibits them from embedding their technology easily into other apps, and vice versa – ultimately constraining how they operate.

Computation: HPC makes a comeback

Models that rely on Big Data-style processing – which has dominated in recent years – don't often work well for many large-scale statistical routines and AI tools: they may operate effectively for ad hoc applications, but significantly less so for massive queries and statistical analysis. Consequently, high-performance computing (HPC[9]) environments have been enjoying a return to popularity: by integrating them with Big Data-style operations, financial institutions can create dynamic systems.

Ideally, an HPC architecture should:

- Support various aspects of a long-term risk architecture.
- Incorporate key functional features, such as:
  - Multi-year storage of market and intermediate data.
  - Editable and dynamic queries.

Novel, but not that novel

When talking about 'novel' hardware approaches (such as AI chips, for example), it's worth noting that they tend to succeed because they are not *completely* novel, but essentially an extension of the GPU framework. By contrast, a completely novel approach – like quantum computing – presents a much greater challenge in terms of widespread adoption.

By adopting techniques that effectively build on what is already in use, firms can develop their systems on an existing base, creating novel approaches that are more likely to catch on quickly. The main concern for financial services firms as far as AI chips are concerned, for example, is how to program them to solve their specific problems.

[9] A loose term for a variety of highly scalable computing architectures with varying parallelization styles. We will outline HPC in more detail in future reports, mapping different computational architectures to specific algorithms and business contexts.

◦ Varied risk measures for all instruments held in their portfolios (some of which are calculated using 'brute force' computational strategies).

◦ An ability to support novel statistical mechanisms at scale.

In some ways, in fact, the ideal configuration is a combination of databases, classical HPC middleware and emergent hardware stitched together by an efficient quantitative framework based on a computationally 'friendly' programming language.

A multi-lingual world: programming languages come of age

Mathematical languages were once complex and hard to use, and few were structured to express mathematical or scientific ideas functionally. This is no longer true: if anything, we are entering a golden age of mathematical languages, with a heightened focus on the strengths of specific languages in the context of programming environments (see Figure 20). It is now possible to construct domain-specific frameworks quickly using sophisticated language creation tools that are increasingly becoming standard, enabling users to rapidly build new analytics.

Newer languages are also more aligned with matrix thinking (they have embedded libraries or language constructs that enable efficient matrix algebra) – and this represents a very powerful step in the evolution of computational infrastructure. And in contrast to earlier, more traditional languages, the latest breed have huge and formidable open-source ecosystems. By learning from other industries (notably science and engineering), financial services firms can devise new ways to adapt these tools to their own particular environments.

In computational terms, Python and Julia are transforming the landscape, enabling a broader range of people to explore different computational styles. So far, Python frameworks are proving the most versatile and popular, serving as all-purpose 'integrationware' in a wide variety of hybrid quantitative applications. Nevertheless, despite Python's dominance, other languages are increasingly making their presence felt, as Scala (and its language ecosystem), OCaml, MATLAB, Erlang, and F# continue to expand – or in some cases dominate – in a variety of different contexts.

Figure 20: Programming – a multilingual world

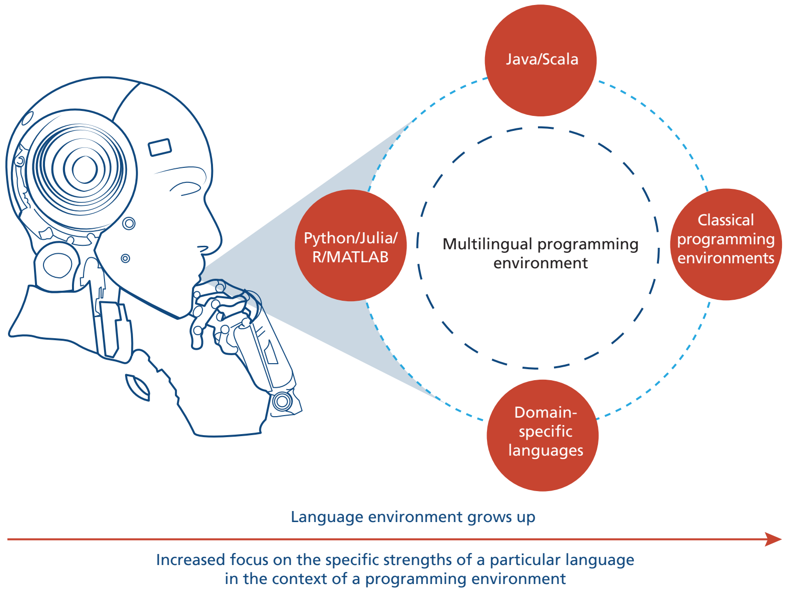

Source: Chartis Research

Focus on Python

Key to Python's growth is the success of not just the language itself, but its ecosystem. Many Java and C/C++ users disparaged Python and other languages, forgetting that the key dynamic wasn't the language itself, but its vast ecosystem of tools and libraries (many of which are written in C/C++). In fact, there are so many options that users are spoiled for choice – and few seldom use anywhere near the full range of options.

Since then, quantitative frameworks built predominantly in Python (or indeed other languages, including Julia, MATLAB, Scala/Kotlin and Erlang) have been gaining traction in a variety of areas in the banking sector, from market frameworks to the management of capital and ALM. These frameworks enable banks (and vendors) to deliver integrated cross-style analytical environments that balance rapid development and control (see Figure 21).

Python (and its ecosystem) is so versatile, in fact, that it can now be used to develop frameworks to address a wide range of technical and operational issues, including:

- Developing a data-quality engine for risk.
- Enabling integration between the front- and middle-office risk engines.
- Developing institutional analytics frameworks.

Figure 21: Python – powering analytical development

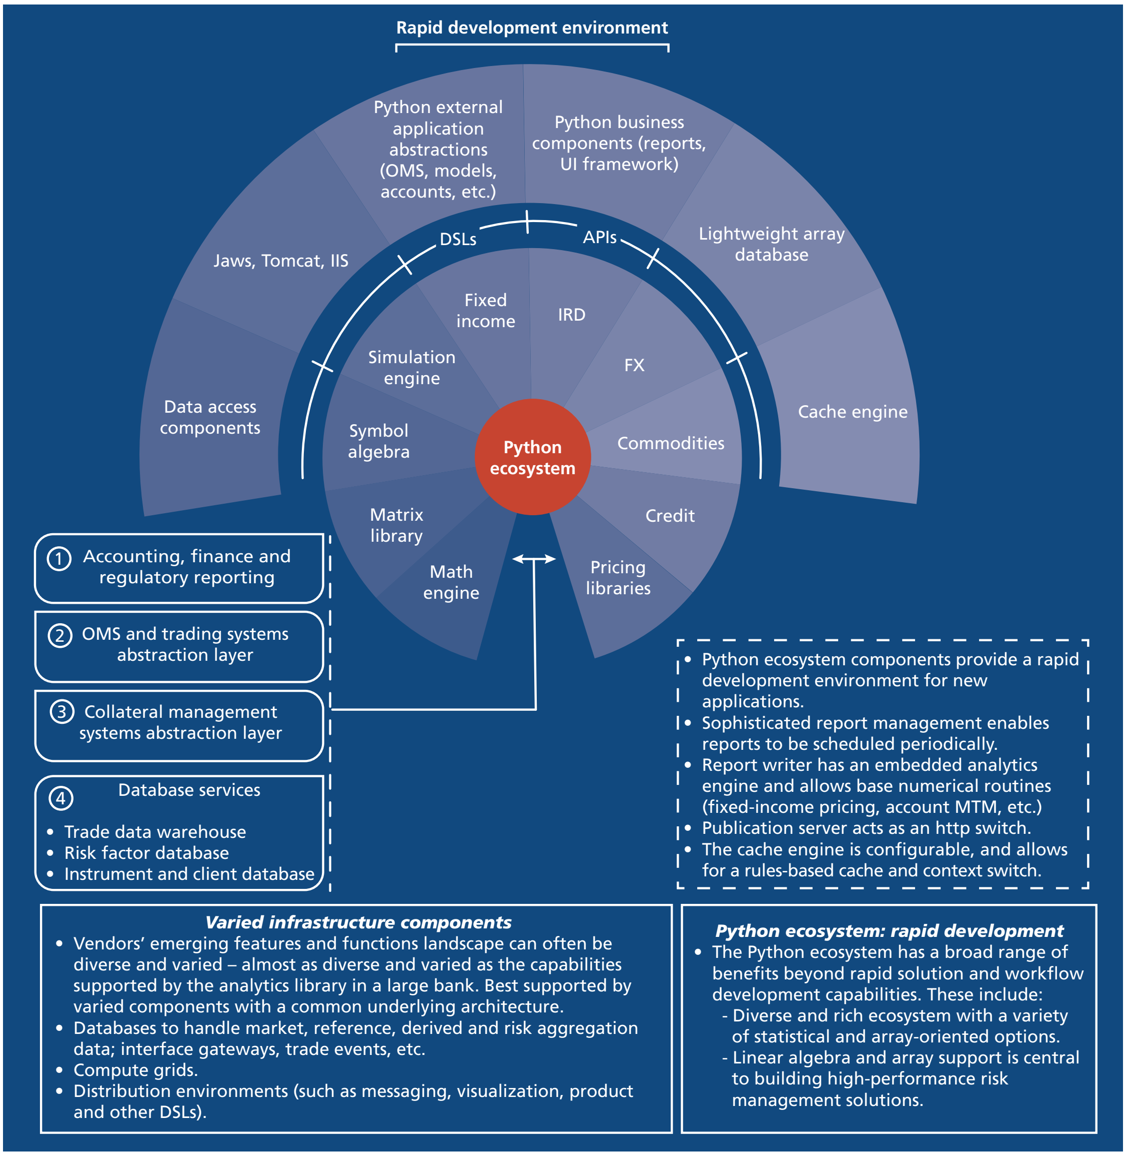

Source: Chartis Research

- Acting as a micro-applications development environment, and a workflow and development manager.
- Developing a virtualized trading platform.
- Constructing and optimizing portfolios.
- Developing credit analytics.

Python-based frameworks have become standard throughout wholesale finance, and are spreading across the financial services universe. However, they are not the last word in languages and programming concepts – other languages (such as Julia and MATLAB) have different ecosystems and advantages, while the Java/Scala/Erlang ecosystems have significant advantages in event-driven applications. Indeed, from a purely quantitative perspective it is possible to see roles for mathematical languages and environments (such as symbolic algebra).

As previously discussed, we are now in a multilingual world in which a knowledge of varied languages will be vital. More importantly, a significant pool of developers now exists that understands, and can develop in, a variety of languages.

Figure 22: All-pervasive – open-source software can be applied across the finance landscape

	Databases	Statistical tools	Open HPC stack	Big Data	Languages
Front-office analytics	High	High	High	Medium	High
Enterprise risk	High	High	High	Medium	High
Finance IT	High	High	High	Medium	High
Collateral management	High	Medium	Medium	High	Low
Operations	High	High	Low	High	Low
Equity trading	High	High	High	Medium	High
FI trading	High	High	High	Medium	High

Applicability High Medium Low

Source: Chartis Research

Figure 23: The GPU-driven database – an increasingly popular option

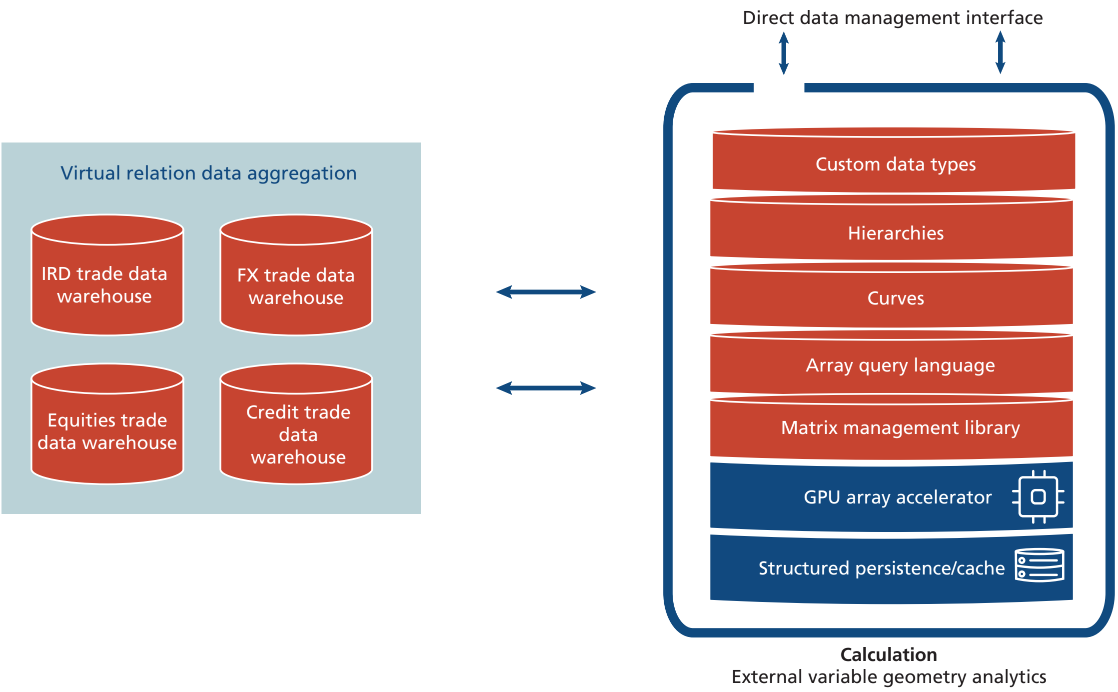

Source: Chartis Research

Open source: pervasive and powerful

Open-source software tools are now pervasive in the finance industry (see Figure 22) – the entire risk stack, for example, can now be built using open-source components. The most successful open-source elements of the STORM landscape are those that have effectively replaced closed-source system software. Not only has this development dramatically reduced costs, it has also provided an environment in which application software developers (vendors and end users) can manage their performance optimization, using open-source tools to design specific algorithms for databases, say, or middleware. Also successful are those tools that have served as the basis of an open standard, such as open-source matrix algebra frameworks and open-source languages.

Open-source databases

Arguably, open-source software components' biggest impact has been on the database element of the stack. Open-source databases have taken over the landscape – Chartis estimates that at least several hunderd open-source databases (and a similar number of closed-source ones) are currently available. Growth in the use of databases driven off new data-parallel hardware (such as GPUs and FPGAs) and programming styles, for example, has been increasing rapidly (see Figure 23).

Figure 24: Matching database to use case

Problem category	NDB	GDB	ADB	HDB	HrDB	DOM	DG
Regulatory compliance	●			●		●	
Fraud analytics		●	●	●	●	●	●
Credit analytics	●	●	●			●	●
Cyber security	●	●	●				
P&L analytics	●	●	●		●		●
T&Cs extraction			●	●			
Equity strategy analysis	●	●	●	●			●
Customer engagement/conduct risk			●	●			
KYC	●	●	●	●	●		
AML/Patriot Act	●	●	●	●	●		
Trade surveillance	●	●	●		●	●	●

NDB Non-traditional SQLDB **GDB** Graph database **ADB** Array database **HDB** Hadoop database
HrDB Hierarchical database **DOM** Document object model **ODM** Object database **DG** Data grid

Source: Chartis Research

This abundance of options has created much more space in which firms can explore how and where analytics work better – which database supports the latency and data structures they require for their algorithms. Firms no longer have to compromise: they can now select a database to suit their particular latency and speed requirements, one that is optimized for their specific analytics use cases (see Figure 24).

The ability to control, structure and examine the underlying code while developing performant system substrates that match the overarching algorithms has been an enormous boon to developers of application software solutions. Not only does this make applications more cost-effective, it also moves them toward heterogeneous database environments, in which firms can focus on employing the appropriate database architecture for specific problems.

7. The state of the art: retail analytics in finance

The retail finance sector is being reshaped by several transformative forces (see Table 2), but four key overarching trends in particular:

- **Digitalization**. The wave of digitalization in retail finance is almost complete, and a new generation of core banking (and credit underwriting and decisioning) platforms are making significant inroads.
- **Lower computational costs**, allowing firms to use a broader set of statistical, risk and optimization tools that can be applied to retail contexts. Previously many of these approaches could only have been applied in wholesale contexts.
- A vast and rapid increase in the availability of **data**.
- More **flexible business practices**, and the possibility of mass customization[10].

An expanded universe

Pre-existing models (and standard statistical techniques) have often struggled against the non-linearity inherent in generating cross-sectional aggregates. They have also had issues with analyzing and consolidating retail behavior, whether in retail banking or other retail finance activities (such as retail insurance or retail brokerage). As a result, the universe of retail modeling has expanded, leveraging alternative and customer-performance data. Retail analytics are now varied and diverse, and an area in which heuristics and advanced statistics have had a significant impact.

Alternative data

A powerful force in retail finance analytics has been financial institutions' use of *alternative* and *transaction data*. While firms have tended to focus on social media, alternative data comes in a variety of formats (see Figure 25). It has been leveraged widely by retail FIs to build and reinforce behavioral analytics, create risk profiles, aggregate and manage retail data, and manage credit portfolios.

In addition to external alternative data, a broad set of client perspectives and needs (such as liquidity, risk appetite and restrictions) can incorporate

Figure 25: Alternative data: many different types

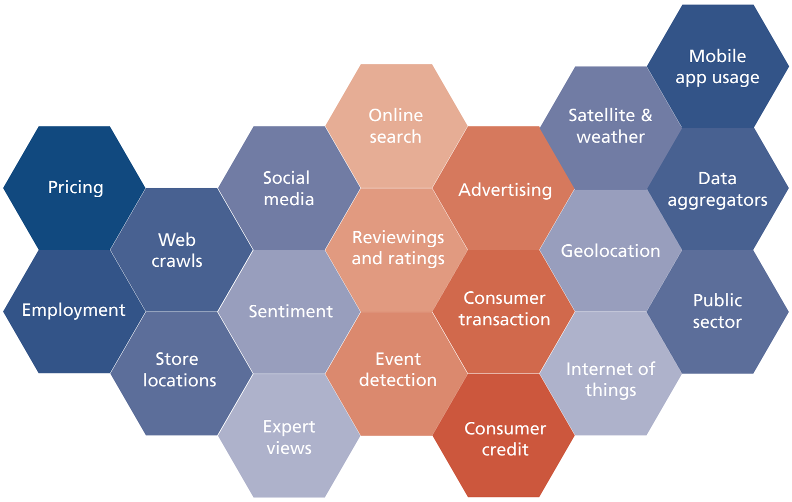

Source: Chartis Research

Figure 26: Customer perspectives and alternative data

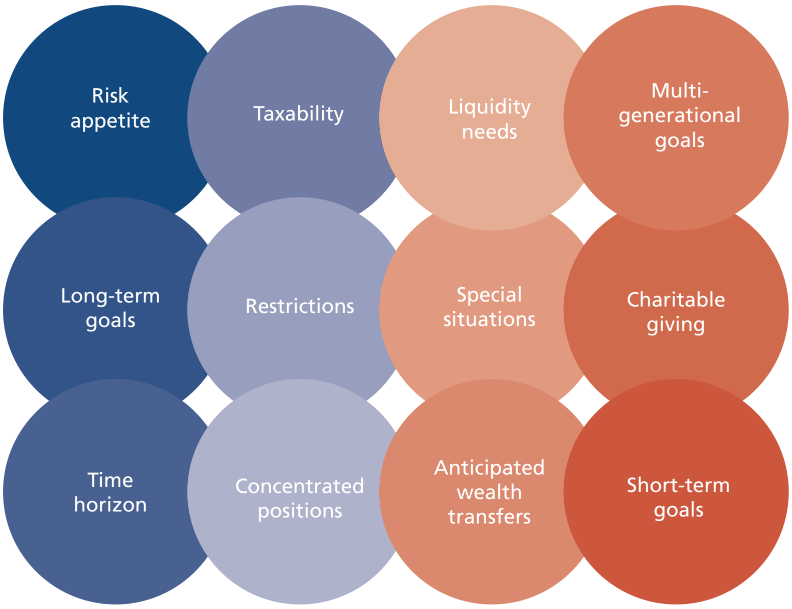

Source: Chartis Research

sophisticated and often institutional-style modifiers for activities such as portfolio management, advice provision, structuring long-term goals and driving charitable giving (see Figure 26).

Many, if not most, of these relationships are non-linear, and managing and controlling alternative data of all types would be almost impossible without the rapid expansion of advanced statistical tools.

[10] *A number of factors – rapidly reducing computational costs, automation and new statistical techniques – imply that 'mass' customization is not far away. In our view, however, mass customization is about the careful design of specific product and analytical structures, and the management of risk-aggregation programs. If, for example, a bank were to issue structured retail products customized for each individual, it would require a highly efficient process in terms of risk management, hedging, regulatory compliance and accounting.*

Table 2: Transformative forces in retail analytics

Force/dynamic	Impact/effect
Challenging times for global retail banks	Retail banking is largely a regional business, regardless of whether it is delivered by FinTechs/shadow banks or traditional institutions. Building and managing the appropriate analytical environments for a wide variety of businesses, regional requirements and cultural norms can prove far more challenging than firms expect.
More sophisticated credit and credit analytics	This is being triggered by IFRS 9/Current Expected Credit Losses (CECL)/Basel reporting requirements, and by the greater availability of data and a broader set of non-linear tools.
The impact of accounting standards on credit risk	This goes beyond mere reporting – the standards have triggered a fundamental restructuring of credit modeling.
The restructuring of banking book credit	This is well underway, and has many moving parts (e.g., performance analytics, data, credit models and curves).
The use of independent credit models	There has been huge growth in the variety, diversity and sophistication of credit analytics and credit data vendors. Increasingly, these models are being used across different business lines and asset classes.
The need for explainable models	These are a critical area of focus for retail banks. We believe, however, that explainability is a matter of data and context analysis rather than just algorithmic design.
The growing importance of pricing optimization	This is becoming an increasingly pertinent feature of the landscape.
The rise of shadow banking	Shadow/non-banks are becoming central to both retail and wholesale banking, as FinTechs, asset managers and other non-banks expand their share of retail finance. Credit is being created through a much more diverse set of channels, and a variety of institutions are required to model credit and fraud risk, and to manage the risk in the models themselves – with a host of regulatory implications.
COVID-19 effects	The pandemic is moving credit operations and operational analytics to the top of the agenda for financial institutions.
Analytics in credit risk management	Analytics and external data continue to have an expanded role in the overall process.
The quantification of ALM	ALM is increasingly becoming a fully quantitative discipline, while behavioral modeling has come of age.
More risk as a service (RaaS)	Highly successful instances of RaaS are plentiful in the capital markets, although they are less common in the banking book.

Source: Chartis Research

Figure 27: Challenges across the retail banking analytics value chain

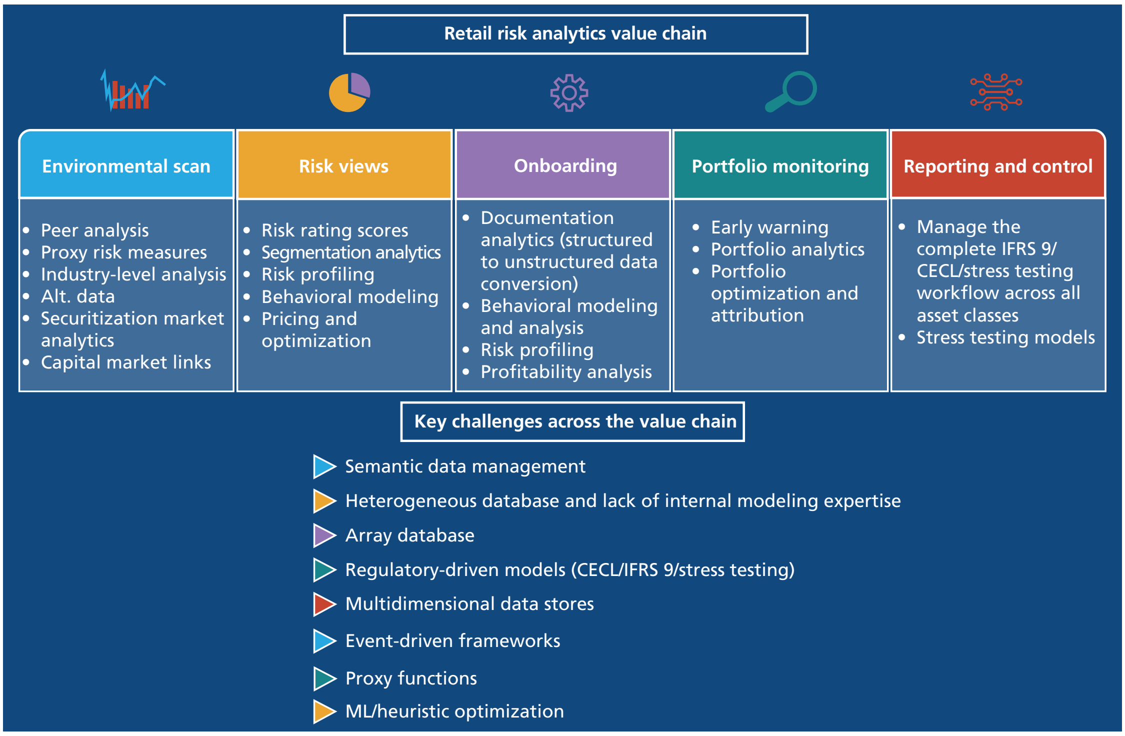

Source: Chartis Research

A complex story

Putting the rapid growth in retail and personal finance analytics aside, the story of their adoption and effectiveness is more complicated, involving multiple challenges (see Figure 27). Some institutions have been highly successful in building sophisticated quantitative and automated approaches. Others, meanwhile, have addressed the growing complexity in credit and credit analytics (see Figure 28) by focusing on credit portfolio management, a critical element of all aspects of credit.

Figure 28: Increasing complexity in credit analytics

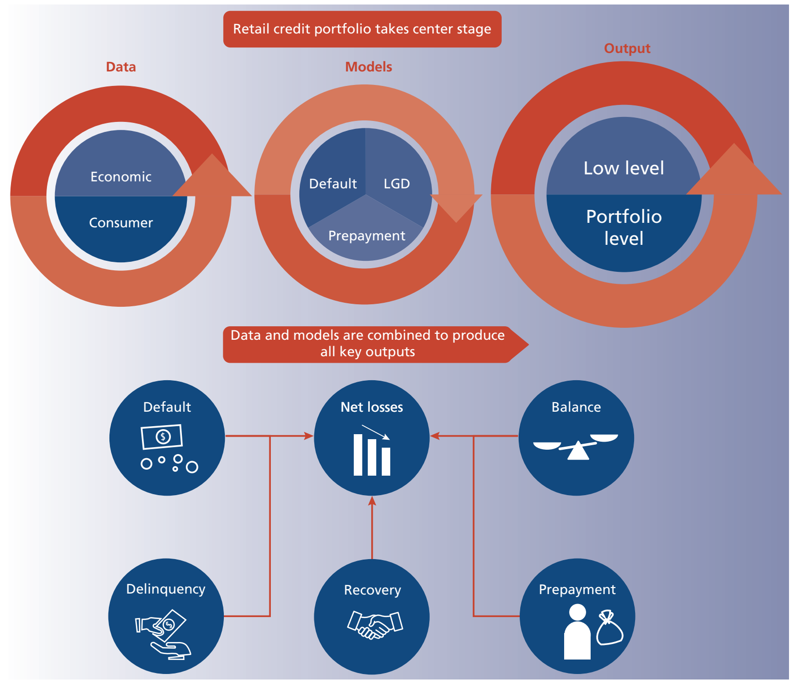

Source: Chartis Research

Three evolutionary themes

Against this background, we have identified three key themes in the evolution of retail analytics:

- Retail behavioral models and aggregation.
- Retail pricing optimization.
- Retail credit scoring.

Behavioral dynamics: increasingly embedded in models

Behavioral dynamics are becoming increasingly embedded within banking credit models and credit portfolio management tools. Behavioral models leverage data gleaned from the implied risks of traded assets and specifically securitized retail products (see Figure 29). This ability to leverage the data is the main reason why several prepayment data suppliers increasingly sell their models to retail and commercial banks.

Behavioral models cover a wide variety of statistical approaches and frameworks:

- Traditional statistical models.
- Neural networks.
- Option-theoretic frameworks.
- A hybrid structure that features some or all of the above.

Many of the emerging behavioral models are ensemble hybrid frameworks that incorporate traditional statistical approaches, neural networks and optional structures.

Retail finance of any form, with its large statistical pools of data around engagements, is obviously the prime candidate for behavioral modeling. At the same time, however, behavioral models and related risk aggregation frameworks are the bridge that links retail and wholesale finance (in whatever market structure). By analyzing customer behavior through the lens of behavioral models, firms can assess deposit and mortgage payment patterns, and create effectively consolidated structures.

The leveraging of prepayment models in US mortgage finance has been a long-running success. Increasingly, however, firms are using novel analytical techniques to build ML/DL models for prepayment frameworks. Development has been advancing rapidly in terms of both academic and operational practice.

Dynamic pricing and price optimization: rapid global growth

Pricing optimization and client strategy management are increasingly sophisticated activities that can incorporate a wide variety of external transactional and alternative data, as well as internal client perspectives and constraints (such as environmental, social and corporate governance [ESG] and portfolio considerations).

Price optimization is a multi-faceted strategy that covers a variety of entities, including product versions, terms and conditions (T&Cs) and coverage levels. These facets are stitched together with AI and ML tools, which factor in the optimal uptake price and maximized profitability (see Figures 30 and 31). Firms can leverage non-standard data, from external and internal sources.

Figure 29: Behavioral models rely on data from several areas

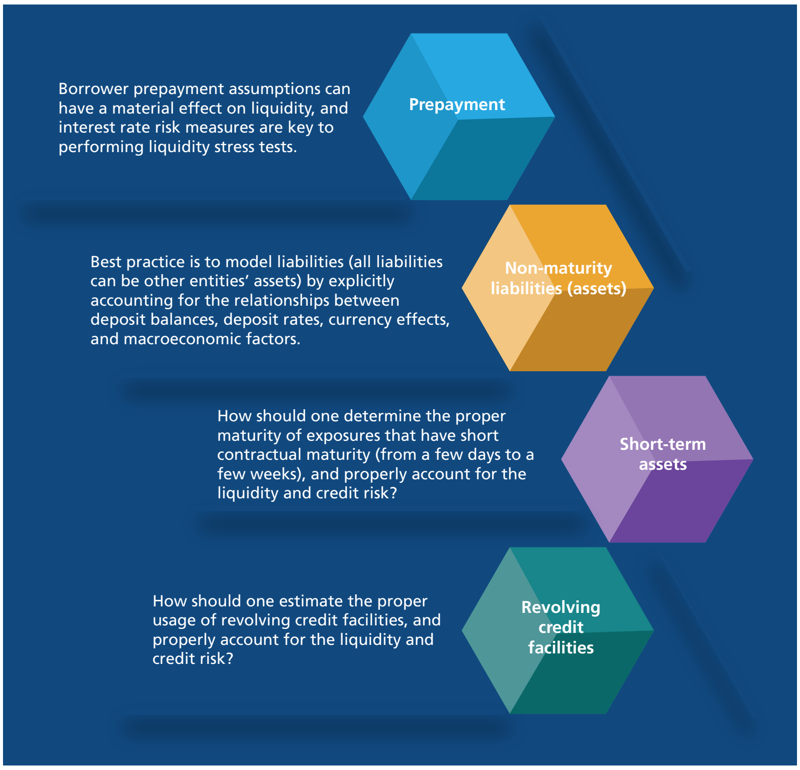

Source: Chartis Research

Figure 30: Dynamic pricing – using AI to stitch together different product facets

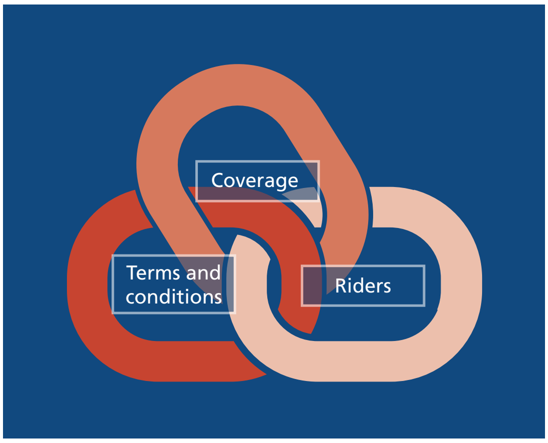

Source: Chartis Research

Traditionally, price optimization strategies have had more salience in North America, although they are spreading worldwide as banks digitalize, leverage broader optimization strategies and embed more sophisticated risk profiles into the process.

We believe that price optimization technologies can not only enable banks (and other retail finance providers) to optimize their portfolios, they are also vital in helping retail financial institutions

Figure 31: ML in pricing optimization

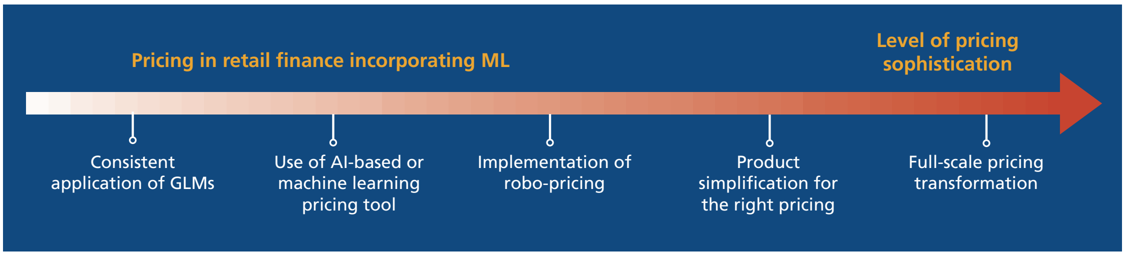

Source: Chartis Research

understand their portfolios (and, by definition, their clients).

Retail credit scoring: old and new

Some of the earliest examples of quantification, the use of alternative data and the use of novel statistical techniques (such as AI and ML) were seen in the retail credit space. Nevertheless, the ongoing structural change and innovation in this segment is as strong as in any other.

Increasingly, firms leverage ML and other heuristics to develop more sophisticated consumer models, which take into account many more variables, as well as the changing structure of the market. And while digitalization is becoming the norm for retail credit processes, different products can have different lifecycles and speeds of approval. Personal loan applications are far more dynamic than mortgages, and subject to fewer regulatory constraints. But digitalization is bringing more and varied data into the retail credit process – notably transaction data. By carefully analyzing transaction and other supporting documentation, firms could generate elements such as risk profiles, appetites and transactional constraints, automatically and in a data-driven way.

With access to a broad range of data about borrowers, firms can now consult a much broader set of retail finance models. The first obvious set of model extensions are those which enable firms to create scoring models based on a borrower's bank transactions. These non-standard models, developed using transaction data, can take a much more behavioral and transactional approach, and can include elements such as the stability of borrowers' expenses, their available savings, and the management of their account balances and overdrafts. The results obtained by the pioneers in this area are excellent, with performance that can outstrip traditional models. Overall, however, non-standard models should be handled with care. As the datasets that firms leverage become broader, the statistical framework gets more complex, and firms will have to manage it carefully, keeping other elements (such as data quality and privacy concerns) in perspective.

Retail finance: not an isolated entity

Figure 32 illustrates different branches of banking – such as wealth strategies and tax planning – highlighting how retail finance has many linkages, especially to the wholesale sector.

Equally, retail analytics (such as credit scoring, behavioral analytics and risk profiling) are increasingly being linked across the enterprise, and more aggregated models are emerging via two distinct channels:

Figure 32: Retail banking is not isolated, but linked to other branches of banking, financial advisory and investment management

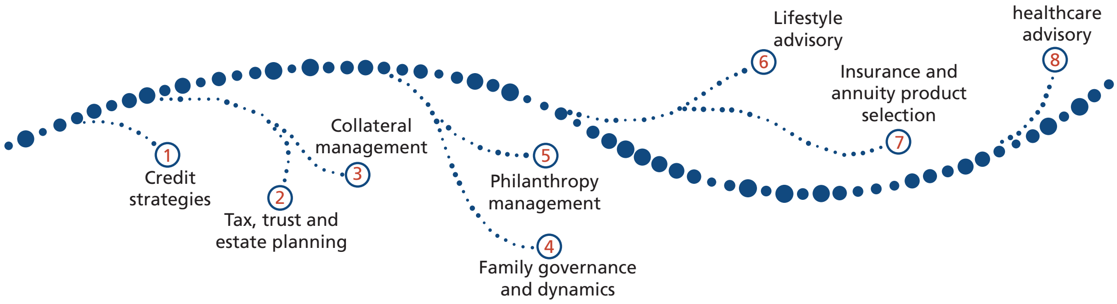

Source: Chartis Research

- **Risk aggregation**, which takes individual risk probabilities and creates aggregates across organizations, business lines, institutions and legal entities.

- **Securitization**. Securitized markets can provide prices and valuations from implied retail risk profiles.

Because both of these relationships are inherently non-linear, AI/ML and heuristics are required to optimally aggregate, imply and analyze the data. We believe that this will continue to be a very significant area of long-term growth[11].

[11] *In the STORM50 ranking we recognize vendors that have innovated in the area of retail analytics, but also those firms that have built (or are building) newer, more efficient ways to expand these links.*

8. Algorithmic engineering and statistical mechanics: the industrialization of STORM

There is growing clarity around how different modeling strategies can be used, validated, managed and controlled efficiently. This is what we refer to as the *industrialization* and *operationalization* of financial modeling. It represents a shift in quantitative methods – no longer do they constitute 'a bit of art and a bit of science'; now they form an engineering discipline in its own right.

The mechanics of convergence

As we have seen, AI and other heuristic styles are converging across the STORM universe, creating a form of 'statistical mechanics' (see Figure 33). This convergence varies according to use case.

Figure 33: The convergence of heuristics and standard analytical techniques is happening across the STORM landscape

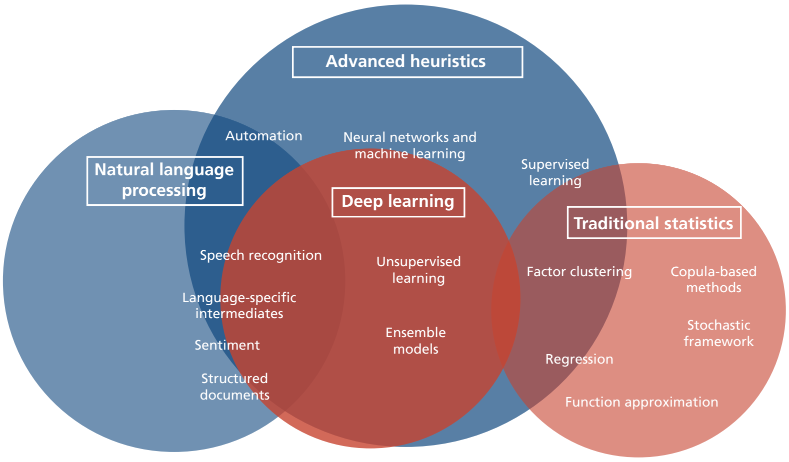

Source: Chartis Research

In assessing the various mathematical functions in which they want to employ AI, firms should consider the input/output of their AI systems to more easily reconcile AI with conventional statistical models (see Figure 34). AI can be added to every element of the risk and analytics workflow, but mostly as an intermediate underpinning statistical capabilities.

The non-linear approach

In achieving this integration, the key challenge for FIs will be enabling AI tools to operate within the context of existing financial market models and theories. The simplest theoretical integration entails using AI as a replacement for linear statistical and optimization requirements. This involves keeping theoretical frameworks constant but adding a layer of non-linearity through various AI styles (such as DL, EP and tabu searches). This approach has the value of preserving a high level of explainability, at the cost of some coherence. In some contexts where inherent non-linearity means that considerable adjustments already have to be made to the inputs (such as a 'smile curve' in options markets), using faster, non-linear interpolators may make models theoretically more consistent. Equally, in most portfolio optimization contexts with non-linear assets (such as derivatives or fixed-income products, or physical commodities

Figure 34: AI's impact on system inputs and outputs

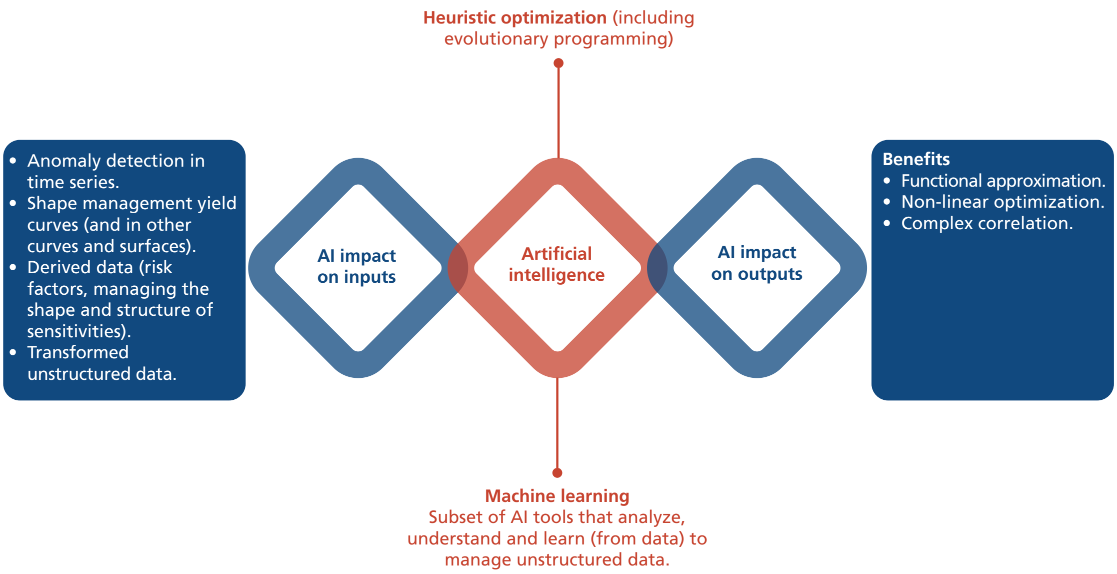

Source: Chartis Research

Figure 35: Assembling an algorithmics framework

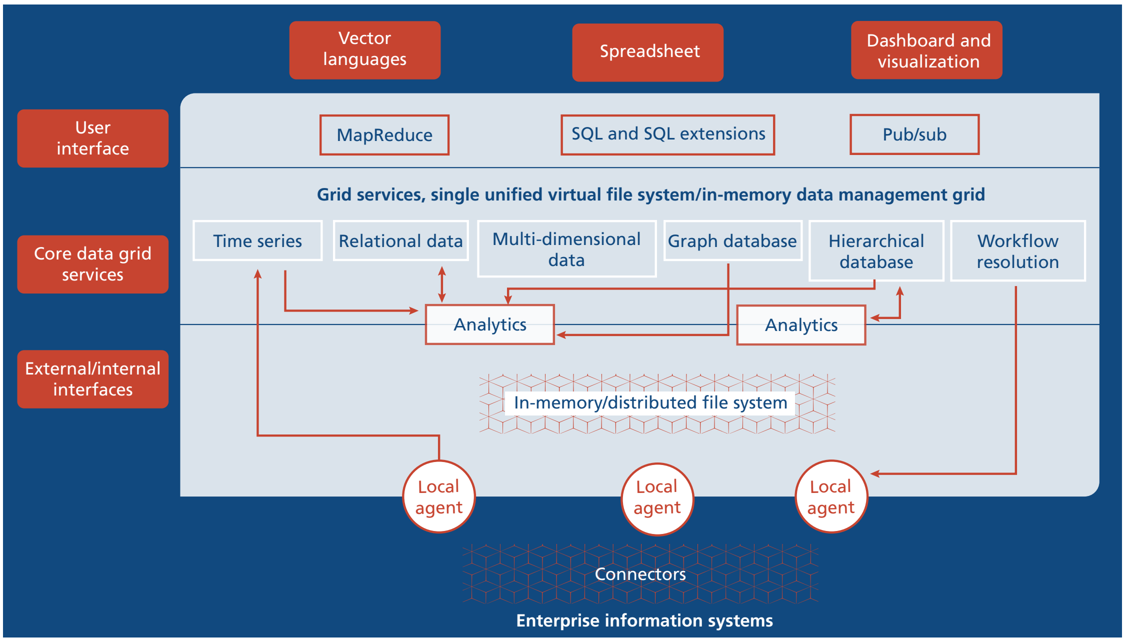

Source: Chartis Research

contracts with embedded options), using non-linear approaches is a substantial improvement, and is likely to produce more accurate results.

Other areas where the impact of advanced non-linear techniques is likely to be seen include:

- Proxy functions for complex non-linear simulations and numerical methods.
- Economic scenario generation.
- Credit default models.
- Complex commodity contract modeling.

Algorithmic engineering: what needs to come together?

Before considering the wider industrialization of analytics development, firms should assemble a universal algorithmic framework, leveraging heterogeneous databases and hardware, and comprising the following elements (see Figure 35):

- A massively parallel processing (MPP) grid, supporting multiple data types.
- Querying, using standard languages.
- Analytics that are available in easily abstractable libraries.
- Data filters, and securities resolution and cross-referencing capabilities.
- Trade data warehouses, integrated into the risk data model.

The key for firms is to recognize that many of the developmental primitives either exist as conceptual artifacts or are available in specific applications. Firms can leverage them to build new applications, rather than continuing to build entire risk application stacks from scratch.

Algorithmic engineering: MLOps and beyond

As modeling and its underlying software tools converge, many of the approaches currently being employed in software development are being embedded into model development practice, including testing, offset release cycles, storing model frameworks in large databases and model development. All of these have undergone the same standardization that has occurred in IT development more generally. As a result, model development is becoming a more industrialized and less conceptual discipline; in addition, this industrialization is now far advanced, and tools are now available to track, test, store and manage the whole cycle (see Figure 36).

In some ways, it is the newest class of models – ML models – that have become most industrialized in their developmental cycle: several firms are focusing on 'MLOps' in much the same way as they have software tools for, say, DevOps (the management and productization of software tools). The rationale behind this operationalization of ML is that at one level it is a process that can be assembled from separate blocks; because there is such a wide variety of styles and techniques, a great deal of data testing is involved. Given the nature of ML models, and how data-dependent their results are, the extent of this data testing means that firms need a systematic program: unless the ML process is industrialized, any reasonably sized organization will struggle in situations that don't involve 'toy' or one-off problems. In addition, many ideas from MLOps approaches can be generalized into other statistical operations and frameworks.

Crucially, operationalizing ML and statistical modeling is a cross-industry phenomenon: firms in other industries are increasingly industrializing and standardizing their AI model management processes, and leveraging ideas that have evolved around DevOps. In this, firms in the finance industry have much to learn from their counterparts elsewhere.

Figure 36: A standard risk workflow has emerged

Source: Chartis Research

Models under scrutiny

Model risk management, including the process of model validation, is an increasingly central aspect of all financial industries. While the challenges and risks of models have been clearly outlined by regulators, senior managers and quants have tended to focus on the risks of OTC derivatives. Regulatory and operational concerns have now expanded to virtually all aspects of the financial services industry: the banking book (credit analytics), insurance (actuarial, risk, business optimization), and financial crime and compliance (anti-money laundering [AML] models). In the face of this expansion one could argue that the model risk management and validation exercise now concerns the whole financial institution: in some ways, every brick in the financial house of models is now under scrutiny.

9. STORM50 2021: methodology

Research process

We employ a multi-step process for our research and analysis for STORM50 (see Figure 37).

Scoring criteria

We use five scoring criteria for the STORM50 ranking (see Table 3). Although there are small overlaps between *impact* and *innovation*, and between *strategy* and the other categories, together they emphasize and stress the primacy of computational techniques, and how the innovation that is occurring in these areas is a key criterion in the STORM (**S**tatistical **T**echniques, **O**ptimization frameworks and **R**isk **M**odels) scoring.

Figure 37: The STORM50 research process

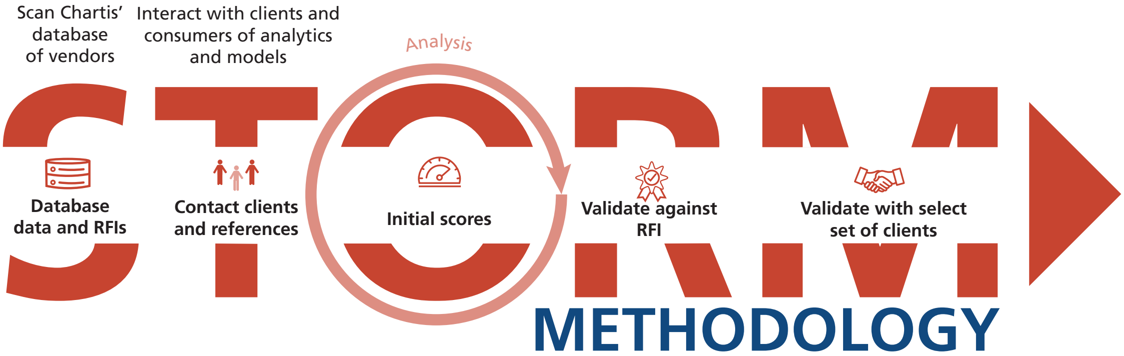

Source: Chartis Research

Table 3: The STORM50 scoring criteria, with definitions

Criteria	Definition
Breadth and coverage	The breadth and coverage of a vendor's analytical techniques, with a focus on different modeling styles (so OTC derivatives modeling, for example, is seen as a distinct style from, say, CAPM-driven portfolio optimization). Vendors receive no additional scores if similar modeling techniques are employed in multiple applications.
Impact	The originality and reuse of an approach. How far have organizations leveraged original approaches from the vendor in question? Some library-driven vendors may, for example, be reused by a broad spectrum of third parties. In such instances they are given positive points for that type of reuse.
Computational infrastructure	A measure of the computational infrastructure's scalability, optimization and alignment with the methodological framework supported by a particular offering/solution.
Strategy	Enabling strategy. This measure evaluates the effectiveness of a particular firm's strategy in enabling its foundational models and computational techniques to be delivered to the market.
Innovation	An overall measure of how innovative the organization is in terms of its modeling approach, analytical methodology, statistical techniques, optimization and risk models.

Source: Chartis Research

10. STORM50 2021: ranking

2021 Rank	Company	Average	Breadth and coverage	Impact	Computational infrastructure	Strategy	Innovation
1	FIS	74.5	78	70	76	77.5	71
2	Moody's Analytics	74.4	80	78	57	79	78
3	Bloomberg	71.2	72	68	74	70	72
4	FICO	67.6	66	74	60	70	68
5	Numerix	67.6	63	68	66.3	70.5	70
6	SS&C	67.5	63	66	70	74	64.5
7	IHS Markit	67.4	68	68	63	70	68
8	MSCI	66.8	68	70	56	74	66
9	MathWorks	66.6	55	74	68	70	66
10	SAS	66.0	64	68	68	70	60
11	Qontigo	65.8	64	68	68	69	60
12	QRM	65.6	64	68	68	68	60
13	Murex	65.0	63	60	70	70	62
14	Quantifi	63.6	53	66	68	68	63
15	Oracle	62.8	56	50	72	74	62
16	Vichara	62.7	53	56.5	68	72	64
17	Numerical Technologies	62.6	58	54	69.5	67.5	64
18	Conning	62.5	55	60	67	67.5	63
19	Kamakura	61.8	65	66	58	50	70
20	Thetica Systems	61.0	55	56	68	68	58
21	RiskSpan	60.8	47	60	64	70	63
22	FactSet	60.6	54	60	58	70	61
23	swissQuant	60.4	60	55	62	64	61
24	LSEG	60.0	54	60	58	68	60
25	Cboe	59.2	49	60	60	66	61

2021 Rank	Company	Average	Breadth and coverage	Impact	Computational infrastructure	Strategy	Innovation
26	Calypso	59.0	51	60	68	62	54
27	Gurobi	58.8	48	60	62	64	60
28	ICE	58.6	48	58	61	66	60
29	FINCAD	57.6	47	60	59	60	62
30	Lacima	57.4	46	60	53	68	60
31	ION	57.0	51	48	58	68	60
32	BlackRock	56.8	58	60	50	60	56
33	Raise Partner	56.4	50	44	60	64	64
34	Andrew Davidson & Co.	56.2	45	62	50	60	64
35	PolySystems	56.0	48	58	54	60	60
36	Beacon	55.8	41	60	54	64	60
37	KDS Global	55.4	48	60	50	55	64
38	Imagine	55.0	51	40	60	64	60
39	CME Group	54.3	47	44	60.5	60	60
40	Confluence	54.2	39	56	58	58	60
41	RPC Tyche	53.9	41	54.5	54	60	60
42	Morningstar	53.8	51	43	56	59	60
43	IDS Gmbh (Allianz)	53.6	45	49	50	64	60
44	Fiserv	53.4	49	44	58	58	58
45	Vector Risk	53.0	45	44	60	56	60
46	Kalotay Analytics	52.8	45	49	50	60	60
47	SecondFloor	51.8	45	44	50	59	61
48	Northstar Risk	51.7	45	41.5	52	62	58
49	RNA Analytics	50.8	42	44	48	60	60
50	ITO 33	50.6	45	42	50	56	60

11. STORM50 2021: category winners

Category award	2021 winner
Chartis categories	
Breadth and coverage	Moody's Analytics
Impact	Moody's Analytics
Computational infrastructure[1]	Oracle
Strategy	Moody's Analytics
Innovation	Moody's Analytics
Solution categories	
AI innovation in capital markets	RiskSpan
AI innovation in retail banking	FICO
AI innovation in wholesale banking	Moody's Analytics
Alt. credit	Moody's Analytics
Application of optimization to analytical problems	FICO
Asset and liability management (ALM)	QRM
Asset-backed securities	Moody's Analytics
Balance sheet optimization	QRM
Climate risk analytics	Moody's Analytics
Collateralized loan obligations (CLOs)	Moody's Analytics
Commercial mortgage-backed securities (CMBS)	Vichara
Convertibles	FIS
Credit portfolio optimization	Moody's Analytics
Credit risk analytics for retail credit	Moody's Analytics
Credit risk analytics for wholesale credit	Moody's Analytics
Domain-specific languages	Numerix

[1] *The scoring for 'computational infrastructure' includes the technological infrastructure available across the entirety of the firms under consideration. In the case of Oracle, this includes its cloud business and capabilities.*

Category award	2021 winner
Economic scenario generation	Conning
Energy commodities analytics	Lacima
Equities performance attribution	MSCI
Equity derivatives	ITO 33
Exchange-traded derivatives	Cboe
Factor modeling	Qontingo
Fixed income: cross-asset	Bloomberg
Fixed income: leveraged loans	Bloomberg
Fixed income: municipals	Bloomberg
Fixed-income performance attribution	Bloomberg
Funds transfer pricing (FTP)	Oracle
High-performance computing (HPC) on the cloud	Oracle
Innovation in computational frameworks	NAG
Innovation in computational languages	Julia Development
Innovation in mathematical environments	MathWorks
Insurance risk analytics	Moody's Analytics
Loans	IHS Markit
Liquidity risk	Bloomberg
Model validation	S&P Global
Mortgage analytics	QRM
Optimization engine	FICO
Over-the-counter (OTC) derivatives	Numerix
Portfolio optimization	Qontingo
Power and networked assets	FIS

Category award	2021 winner
Quant management framework	Beacon
Residential mortgage-backed securities (RMBS)	Vichara
Retail price optimization	FICO
Securitized products	Moody's Analytics
Simulation frameworks	MathWorks
Structured credit	Quantifi
Transaction cost analysis (TCA)	Virtu ITG
xVA	Numerix

12. How to use research and services from Chartis

In addition to our industry reports, Chartis offers customized information and consulting services. Our in-depth knowledge of the risk technology market and best practice allows us to provide high-quality and cost-effective advice to our clients. If you found this report informative and useful, you may be interested in the following services from Chartis.

Advisory services

Advisory services and tailored research provide a powerful way for Chartis clients to leverage our independent thinking to create and enhance their market positioning in critical areas.

Our offering is grounded in our market-leading research, which focuses on the industry and regulatory issues and drivers, critical risk technologies and leading market practices impacting our sector. We use our deep insight and expertise to provide our clients with targeted market and industry analysis, tailoring content to assess the impact and potential of relevant regulatory and business issues, and highlighting potential solutions and approaches.

Chartis' advisory services include:

Market dynamics

The markets that our clients – vendors, institutions and consultants – address are changing at an ever-increasing pace. Understanding the market dynamics is a critical component of success, and Chartis uses its deep industry and technical knowledge to provide customized analysis of the specific issues and concerns our clients are facing.

Market positioning

In today's highly competitive market, it is no longer enough to simply have a leading product or solution. Buyers must be able to appreciate the differentiating capabilities of your brand and solutions, and understand your ability to help them solve their issues.

Working with our clients, we generate compelling, independent co-branded research, targeting critical business issues. This helps our clients to position their solutions effectively, 'own' key issues, and stand out from the crowd.

Collaborating closely with our clients, we develop pragmatic, resonant thought-leadership papers with immediate industry relevance and impact.

Our offering includes:

- **Co-branded research** on key market topics to provide a unique and compelling point of view that addresses a key industry driver and highlights the relevant issues. Reports can be tailored to varying levels of depth and can be powered by quantitative survey fieldwork, qualitative industry interviews, our deep domain expertise, or a blend of all three.

- **Chairing roundtables and/or facilitating events and workshops** to support clients in hosting compelling events that put them at the heart of the discussion.

- **Targeted marketing through our sister brands**, leveraging the power of our parent group – Infopro Digital – to reach across leading brands such as Risk.net, WatersTechnology, FX Week and Central Banking.

Competitor analysis

Our unique focus on risk technology gives us unrivaled knowledge of the institutions and vendors in the sector, as well as those looking to enter it. Through our industry experts, Chartis clients can tap our insights to gain a much deeper understanding of their competitors and the strategies they should pursue to better position themselves for success.

Regulatory impact analysis

The analysis and assessment of regulatory change and implementation is one of Chartis' core strengths. We can apply our insights to assess the impact of change on the market – both as it applies to vendors and the institutions they serve, or on a client's specific product and customer base. We can also provide insights to guide product strategy and associated go-to-market activities, which we can execute for internal use to drive our clients' strategy, or as a co-branded positioning paper to raise market awareness and 'noise' around a particular issue.

13. Further reading

RiskTech100® 2021

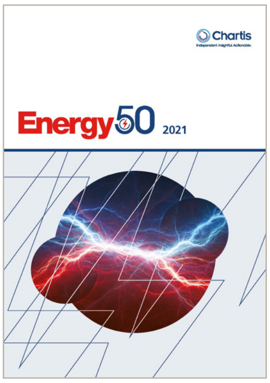

Energy50 2021

Big Bets 2021

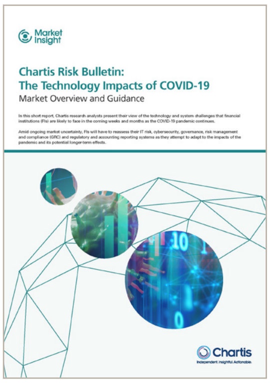

Chartis Risk Bulletin: The Technology Impacts of COVID-19

Artificial Intelligence in Financial Services, 2019: Demand-Side Analysis

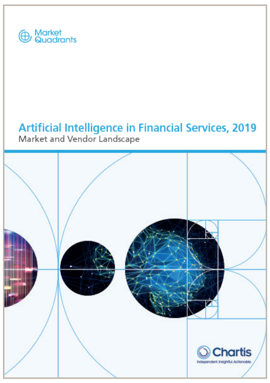

Artificial Intelligence in Financial Services, 2019: Market and Vendor Landscape

For all these reports, see **www.chartis-research.com**