

A large, faint background graphic of a human brain, where the internal structures are replaced by a complex network of circuit lines and nodes, symbolizing artificial intelligence.

GenAI in Credit Decisioning

WHITE PAPER

HOW FINANCIAL INSTITUTIONS
CAN EXTRACT ROI FROM AI

Dominique Zosso, Ph.D.
Data Science, Zoot Enterprises Inc.



zootsolutions.com



EXECUTIVE SUMMARY



We have all seen the headlines: Generative AI (GenAI) is all the rage, but most GenAI pilots thus far have struggled to produce measurable ROI. The gap often stems from looking for impact in the wrong parts of the process.

The financial services industry, too, is not immune to this pattern: GenAI has rapidly emerged as one of the most influential forces reshaping expectations in lending. Financial institutions face increasing pressure to demonstrate forward-thinking leadership—all while navigating regulatory uncertainty, internal risk aversion, limited in-house expertise, or legacy operational constraints.

Although the excitement surrounding AI is substantial, the practical pathways to safe, effective, and compliant adoption therefore remain underdeveloped.

Modernization succeeds when automation handles the repeatable, while human expertise strengthens everything else.

This is a lesson exemplified by General Motor's (GM) rushed automation efforts in the 1980s, and it unequivocally applies to the AI arms race, today. In credit, the most meaningful early value comes from pointed applications that target the repeatable operational steps where human reviewers still encounter unnecessary friction. By focusing GenAI on these specific areas, financial institutions can begin to overcome the early ROI barrier and uncover value in places that often go overlooked.

Through this lens, this whitepaper provides a grounded, realistic, and strategically constructive approach to applying GenAI in credit decisioning. GenAI can deliver immediate and meaningful value when used to enhance supportive processes such as document intelligence, validation workflows, structured data transformation, and operational productivity, rather than as a replacement for deterministic decision engines.

The goal is to provide lenders with a responsible and achievable path toward becoming pragmatic adopters of GenAI: a path that begins with focused operational improvements and builds the institutional learning needed to take advantage of more advanced use cases as they mature.



Dominique Zosso

Dominique Zosso, Ph.D.
Data Science, Zoot Enterprises Inc.

WHY AI PRESSURE IS INTENSIFYING

Artificial intelligence is rapidly reshaping U.S. banking and financial services, and the industry is undergoing a period of heightened technological expectation. From automating customer service to enhancing risk models, AI's promised influence now spans nearly every layer of the modern financial institution. In credit underwriting specifically, generative AI is expected to make for faster decisions, improved efficiency, and even expanded financial inclusion. One such promise is that by analyzing vast datasets (including non-traditional data like utilities or cash flow), GenAI can identify patterns that help reveal opportunities where assessing lending risk might otherwise be challenging.

At the same time, consumer expectations have shifted toward near-instantaneous decisioning, especially in areas such as point-of-sale credit interactions, embedded finance offerings, automated lending marketplaces, or automotive lending. These markets reward lenders with low pend rates and penalize organizations whose workflows still heavily rely on manual underwriting. Yet, automation beyond a certain point becomes technically or operationally prohibitive without the assistance of more flexible, adaptive tools. This is precisely where GenAI's actual strengths matter. It is not suited to replace credit judgment, but it can support the processes around it by interpreting documents, preparing data, clarifying ambiguous applicant inputs, and reducing the friction that slows underwriting teams.

For U.S. lenders, the question seemingly is no longer whether to use AI, but how to harness it within a framework that safeguards compliance, transparency, and public trust.

This pressure is complicated by discussions of GenAI often being shaped by extremes:

Many have expected too much of GenAI (overhype), such as assuming that it should be able to automate more of underwriting than it can, while others have feared catastrophic implications (overfear) and fall into the opposite belief that GenAI is fundamentally incompatible with lending. Both narratives continue to heavily influence how institutions approach GenAI today.

According to McKinsey's December 2023 analysis of generative AI in banking (McKinsey & Company, Dec 2023¹) around two thirds of senior digital and analytics leaders believed the technology would fundamentally change the way they do business. McKinsey's subsequent March 2024 review (McKinsey & Company, March 2024²) of 16 large financial institutions found that roughly 70 percent of banks with centralized GenAI operating models had already moved use cases into production, compared with about 30 percent of those taking a fully decentralized approach. This accelerating shift in expectation and execution is creating real strategic pressure, especially for growing institutions that are expected to keep pace without the scale or internal AI engineering depth of large national banks.

Simultaneously, the regulatory environment has become more vocal, even as formal rules remain incomplete. The Consumer Financial Protection Bureau (CFPB) has issued multiple cautionary communications about the use of "black-box" credit models and the need to maintain specific, accurate adverse action notices even when using advanced analytics (CFPB Circular 2022-03, May 2022³).

Other prudential regulators have reiterated, such as the Federal Reserve through SR 11-7, that institutions remain fully accountable for the transparency, governance, and documentation of any models used in credit decisioning. In a highly regulated arena like consumer credit, opacity will not excuse compliance failures. Lenders must ensure that algorithmic decisions meet the same legal standards as human ones, from avoiding discrimination to explaining adverse decisions, or face regulatory and reputational consequences.

Indeed, U.S. regulators have made it clear that existing laws fully apply to AI-driven credit processes. These regulatory expectations can unintentionally feed into an overfear narrative: While the practical limitation really only applies to using opaque models as credit decision-makers, concerns about explainability easily get generalized to all types of GenAI use cases. Understanding the nuance in use cases and application is crucial in discerning risk from noise.

As a result of all these influences, lenders find themselves in a complicated dilemma: pressured to innovate, wary of regulatory missteps, constrained by limited internal ML/AI capacity, and still reliant on a patchwork of legacy tools and vendor ecosystems.

Success with AI in credit will require marrying innovation with responsibility.

Part of that responsibility involves mastering the balancing act between the opposing extremes of overhype and overfear: GenAI cannot—and currently should not—autonomously determine credit outcomes, but it is also not barred from lending workflows, *a priori*. Its value is in the decision-adjacent processes, where it can improve speed, consistency, and operational clarity without replacing the core decisioning models that remain subject to strict governance standards.



UNLOCKING VALUE OF GENERATIVE AI IN CREDIT: A TALE OF CAUTION

GenAI is beginning to show where it can offer practical value in lending, and one of the clearest lessons from early experimentation is that meaningful ROI comes from applying it in the right parts of the workflow. The lessons emerging from GenAI business experimentation mirror those from GM's 1980s automation struggles:

- Throwing technology at unoptimized processes.** GM invested heavily in robots before truly working out the kinks in its existing production processes. The result was automating confusion. Today, companies can jump on the "AI fad" by deploying large language models into workflows that are still poorly structured or rely on tribal knowledge. There are few ways to lose money faster than automating a process that is poorly understood.
- Automation for automation's sake.** GM leadership at the time was obsessed with futuristic technology as a cure-all for Japanese competition, rather than addressing deep-seated organizational and structural issues. The current "AI arms race" encourages a fear of missing out, leading companies to invest billions in AI pilots that yield zero return, because they prioritize the technology itself over a sound business strategy or blended integration.
- Undervaluation of human workers.** GM's strategy implicitly assumed workers were the problem that could be replaced by machines. A similar risk today involves replacing staff with AI agents, overlooking the human expertise still needed to guide, oversee, and refine GenAI-driven systems. Toyota's success with "autonomation" demonstrated that technology succeeds when it complements human expertise and organizational learning, not when it replaces it entirely.

- **Going too big, too fast.** GM attempted to implement entirely new, complex systems overnight and across numerous plants. The lesson for GenAI, today, is to proceed incrementally, starting with small, targeted experiments on predictable tasks, and scaling up as institutional learning gradually progresses.

Clearly, the opportunity today is not in GenAI automating entire intricate credit decisioning processes. It is in making the parts of lending that still rely too heavily on human intervention more seamless and efficient for both underwriters and applicants. When institutions focus first on understanding of, and reducing friction in, these human-centered steps, they create the conditions for early wins and build the internal capability needed to broaden GenAI's role over time.

We believe that GenAI represents a timely and accessible opportunity for financial institutions to enhance their lending operations while preserving the structures that support responsible credit practices. The most reliable value in the near term comes from focusing on the points in the workflow where human reviewers still experience unnecessary friction. By directing GenAI toward these specific areas and adopting an incremental approach, institutions can avoid the ROI challenges that have surfaced in many early GenAI pilots. Early wins in these targeted steps create clarity, strengthen internal capability, and help organizations prepare for the more advanced credit-specific applications that will become possible as the technology matures.

This progression is not only practical. It deepens institutional learning and builds familiarity with GenAI in a way that strengthens governance and supports responsible expansion over time. Financial institutions that understand where to look for value today, and that adopt technology in a thoughtful, staged manner, will be well positioned for the next wave of transformative use cases. These institutions will have the foundation, oversight, and readiness needed to leverage more diverse data ecosystems and more flexible decision support models as the landscape evolves.

From this foundation, adoption naturally expands into decision-adjacent analysis, structured decision support, and eventually more advanced credit specific applications as readiness and governance mature. Much of the most compelling value emerging across the industry aligns with these layers, including customer engagement through AI assistants, document intelligence for file preparation, synthetic data creation for system testing, pend rate reduction in instant-decisioning environments, and new ways of leveraging external data ecosystems.

IMMEDIATE OPPORTUNITIES FOR GEN AI IN CREDIT

In the near term, the most accessible value sits inside the operational steps that introduce unnecessary friction in lending. Many credit processes, particularly in small business, commercial, and other complex segments, still require a human reviewer to interpret information, resolve inconsistencies, or request clarifications from the applicant. These moments are often where pend rates rise and turn times stretch. Financial institutions that want to capture early GenAI value should begin by examining these points of manual effort. Typical examples include the first pass review of submitted materials, the coordination required to gather missing information, and the early triage steps that determine whether a file is ready for underwriting.

GenAI can help by preparing the information in a clearer, more organized form, identifying gaps before a human reviewer ever touches the file, and simplifying the initial exchanges that often delay progress. Banks are already applying GenAI to automate document ingestion, extract income and cash flow attributes from tax forms or bank statements, and generate concise summaries of unstructured documents.

Several institutions in 2024 surveys reported using language models to perform quick Q&A on file contents, identify missing documentation, and flag anomalies within minutes rather than hours. Underwriters may be faced with long histories of financial behavior, multiple sources that must be reconciled, or narrative explanations that require careful reading to understand context. GenAI can help by structuring this information into clearer formats, highlighting areas that warrant deeper review, and surfacing trends or discrepancies earlier in the process.

Beyond, AI tools can synthesize multiyear business bank statements into comparative summaries, highlight abnormal cash movements, reconcile conflicting income data across sources, or provide context from borrower narratives that might otherwise require several minutes of manual reading. This improves the efficiency and consistency of the review, while keeping the underwriter in full control of the final interpretation. The goal at this stage is to help teams handle higher volumes or more varied forms of information without adding burden or risk.

By improving the clarity of what underwriters see, GenAI makes it easier for them to focus attention on the parts of the file that genuinely benefit from professional judgment rather than on the administrative effort required to find those parts.

Another area that fits this early operational category is customer engagement. AI powered chatbots on lender websites or mobile applications can guide applicants through required documentation, answer routine questions, help fill out forms, and provide real time status updates. This reduces back and forth delays, lowers abandonment in application funnels, and prevents situations where underwriters receive incomplete submissions.

GenAI can also support personalization of the borrower experience in ways that remain outside of regulated credit decisioning. For instance, it can tailor the timing or channel of communications, structure product explanations differently depending on borrower needs, or adjust the order in which already approved options are displayed. These optimizations enhance borrower clarity and reduce confusion without affecting eligibility or pricing.

GenAI can also begin to reduce pend rates in instant decisioning channels by triaging the ambiguous cases that traditionally require human review. In automotive lending, POS, embedded finance, or BNPL flows, AI can reconcile mismatches between consumer entered data and third-party feeds, interpret scanned paystubs or identity documents, and turn a large portion of borderline cases into clean files ready for deterministic decisioning that completes timely and presents an opportunity for conversion.



These opportunities do not alter the credit decision. They simply reduce avoidable slowdowns, making it easier for underwriters to focus on evaluating risk and for applicants to move through the process without repeated interruptions. Other industries have already made significant headway using GenAI to streamline similar operational workflows, and with thoughtful compliance preparation, financial institutions can adopt many of these same advancements in ways that align with their regulatory and governance requirements.

NEXT GENERATION CREDIT DECISIONING

From there, institutions can extend GenAI into structured decision support. At this point, GenAI's role shifts from preparing information for human review to preparing information for system review. It helps convert unstructured or semi-structured inputs into organized, machine-readable elements that a deterministic decision engine can evaluate. The core policy logic and credit determination remain unchanged. The engine continues to execute decisions with the same transparency, replayability, and regulatory alignment that are essential for responsible lending.

GenAI simply enhances the quality, consistency, and completeness of the information that flows into that system. This includes mapping raw data from open banking feeds, alternative credit bureaus, employment verification sources, collateral databases, or vehicle ecosystems into the structured attributes a policy requires. GenAI can also reconcile missing or contradictory fields and produce a single, audit-ready representation of borrower information.


For lenders, this stage represents a bridge between operational efficiency and deeper transformation. It enables more flexible use of new data sources and more dynamic ways of assembling information, while maintaining the safeguards that define compliant credit processes.

Synthetic data generation also becomes valuable at this point. Institutions can create realistic but entirely artificial loan files to test decision engines, validate policy changes, or explore edge cases—without handling sensitive customer data.

The final stage is the exploration of advanced structured artifact generation. This approach allows GenAI to assemble pre-approved building blocks that describe how a deterministic engine should evaluate the available information in an application. The engine then executes the decision using its established framework.

This model combines the adaptability of GenAI with the controls of traditional credit decisioning, offering a path toward more sophisticated use of diverse data ecosystems without sacrificing clarity or explainability. While this stage remains forward-looking for most institutions, it represents a natural evolution for organizations that have laid the groundwork through earlier phases of adoption.

Across each of these stages, one principle remains constant: Progress builds on itself.



Every step enhances institutional learning, strengthens governance, and increases confidence in how GenAI behaves within established credit practices. In the early phase, the ROI opportunity resembles what organizations see when adopting GenAI more broadly: improved efficiency, clearer information, and smoother operational throughput.

As adoption matures, more credit-specific opportunities begin to appear, and institutions that have built their foundation through pointed, incremental steps are best positioned to take advantage of them.

This progression reflects a responsible path for lenders seeking to bridge the ROI gap and prepare for the more transformative applications that may follow.

CONCLUSION: A STRATEGIC OPPORTUNITY FOR FOCUSED INNOVATORS

GenAI is beginning to reshape many parts of the lending process, and its value comes from strengthening the proven foundations that already define responsible credit practices. Lenders have spent decades building processes that are explainable, auditable, and aligned with regulatory expectations, and those principles remain the anchor. GenAI adds something new. It brings the ability to interpret unstructured information, surface patterns that are difficult to see at scale, and streamline the operational steps that slow teams down.

The themes across the earlier stages show how GenAI will cascade into lending: AI-enabled chat experiences reduce applicant confusion, document intelligence accelerates file preparation, synthetic data unlocks faster system testing, decision-adjacent assistance elevates underwriter throughput, and structured artifact generation increases flexibility without compromising regulatory expectations. When paired with sound governance, it becomes a set of tools that accelerate and elevate established workflows rather than replacing them.

Successful use of GenAI will require institutions to think creatively about how these capabilities fit within their culture, their talent model, and their existing credit frameworks.

Meaningful adoption will not begin with autonomous decision making. It will start with organizations becoming comfortable applying GenAI inside their operations, understanding its behavior, and building the confidence and oversight structures that allow it to support judgment rather than attempt to stand in for it.

This is not a retreat from innovation. It is a forward step that extends the strengths of responsible lending into a new technological era.

GenAI will emerge through a cascading path. It may not be ready to make independent credit determinations today, but it is already well suited to deliver real lift in areas where lenders face operational bottlenecks or data complexity. Early experimentation and controlled deployment unlock the insight needed to pursue the higher value opportunities that sit just ahead.

As capabilities mature, institutions that have taken the time to crawl and walk will find it far easier to accelerate when the industry reaches the point where GenAI's full potential becomes accessible. In contrast, institutions that wait for a perfect moment are more likely to face the abrupt and costly leap that often accompanies late adoption in technological cycles.

Ultimately, these principles align with the values that have guided innovation in credit decisioning at Zoot and for our customers. GenAI extends this tradition rather than disrupting it. While technology is transformative, it reinforces the fundamentals that matter most in lending: clarity, consistency, strong governance, and a commitment to responsible, transparent decisioning.



ABOUT ZOOT ENTERPRISES



Zoot Enterprises is a global provider of advanced origination, acquisition and decision management solutions for some of the largest organizations in the world. We offer comprehensive and flexible platforms for specific business needs that include loan origination, fraud detection and prevention, data acquisition, and more.

Zoot is a leader in the technology industry. We enable clients to access hundreds of cutting-edge data sources in real time, and provide business user control that empowers our clients to adapt to their evolving strategies. Our cloud-based, secure processing environment delivers decisions in milliseconds and has the capacity to deliver billions of real time decisions annually.

For more than 35 years, we have launched solutions to market faster than the competition. We have partnered with influential U.S. and international organizations including major financial institutions, retailers, healthcare organizations and payment providers to foster excellence in multiple industries.

At Zoot, our promise is simple: To make our clients successful. We work closely with the top financial institutions in the world to deliver state-of-the-art solutions that satisfy the most demanding decisioning needs.

Zoot Enterprises can help you execute data agnostic, one-to-one decisions informed from all available sources. If you're ready to improve your origination process and keep pace with the changing times, please contact Zoot today.



(406) 556-7555



zootsolutions.com

From Data to Decision®



ENDNOTES

1. Vishnu Kamalnath, Larry Lerner, Jared Moon, Gökhan Sari, Vik Sohoni, and Shuo Zhang. Capturing the full value of generative AI in banking (McKinsey & Company, Dec 2023) <https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking>
2. Kevin Buehler, Alison Corsi, Brian Weintraub, Mina Jurisic, Andrea Siani, and Larry Lerner. Scaling gen AI in banking: Choosing the best operating model (McKinsey & Company, March 2024) <https://www.mckinsey.com/industries/financial-services/our-insights/scaling-gen-ai-in-banking-choosing-the-best-operating-model>
3. CFPB Acts to Protect the Public from Black-Box Credit Models Using Complex Algorithms (CFPB Circular 2022-03, May 2022) <https://www.consumerfinance.gov/about-us/newsroom/cfpb-acts-to-protect-the-public-from-black-box-credit-models-using-complex-algorithms>



HEADQUARTERS

555 Zoot Enterprises Lane
Bozeman, MT 59718

PHONE & FAX

Phone: (406) 586-5050
Fax: (406) 587-8414

SALES

Phone: (406) 556-7555
Email: sales@zootweb.com