DATA NUTRITION LABELS: SUPPORTING ARTIFICIAL INTELLIGENCE EXPLAINABILITY IN FINANCIAL SERVICES

March 2026

Financial Services Sector Coordinating Council
for Critical Infrastructure Protection and Homeland Security

Table of Contents

Executive Summary

This whitepaper promotes the safe and responsible adoption of Artificial Intelligence (AI) capabilities within the financial sector, by focusing on the importance of sound data management practices when leveraging AI. Deriving inspiration from the food industry and the widely recognized nutrition labels, Data Nutrition Labels (DNL) aim to help standardize and enhance the transparency and explainability of datasets used for AI systems within the financial sector. This paper proposes a two-tiered approach, Foundational and Enhanced, that supports a more structured approach to increase the transparency and explainability of data quality evaluations in the financial sector. This approach also acknowledges the varying levels of data management maturity observed across the financial sector. In addition, this paper outlines the importance of establishing clear roles and responsibilities, while recognizing some of the challenges that may arise throughout the adoption and implementation journey for each organization. DNLs offer financial institutions an additional mechanism to support AI explainability, enabling them to integrate these tools in ways that align with their business objectives, data management practices, and AI strategies.

Purpose

The purpose of this paper is to recommend a more structured approach for the evaluation of data quality as it relates to artificial intelligence solutions (AI) solutions in the financial sector, to support increased transparency and trust in the use of AI. The rapid integration of AI and Generative AI[1] (GenAI) across all industries further emphasizes the importance of having transparency and standards for the quality of data used for AI solutions. More specifically, GenAI's effectiveness and trustworthiness are significantly affected by the quality of the data, both structured and unstructured, used by the Large Language Model (LLM) or Small Language Model (SLM) powering the solution. Lack of standards, transparency, and quality surrounding the use of data for GenAI systems can negatively impact the outcome, resulting in ineffective and potentially inaccurate GenAI systems.

The pace at which AI capabilities are advancing provides an opportunity to collaborate and drive industry standards that can support enhanced explainability, transparency and results from AI systems. Increasing the explainability of AI systems can help financial institutions accelerate adoption by removing ambiguity and identifying risks while ensuring the prudent use of the technology and alignment with federal and state, and international regulatory standards and guidance.

The integration of data nutrition labels (DNLs) for financial institutions can support three main objectives:

1. **Internally Focused**: Ensuring the appropriate level of observability in data supply chains to support confidence and trust in leveraging AI and GenAI capabilities (near-term).
2. **Industry Collaboration**: Driving standards across the sector to support enhanced transparency and explainability of AI and GenAI systems (medium-term).
3. **Increase Customer Confidence**: Increase transparency and explainability in the use of AI and GenAI capabilities to enhance consumer confidence in Financial Institution's products and services (long-term).

This paper is a deliverable of the Financial Services Sector Coordinating Council's (FSSCC) *Artificial Intelligence Explainability and Data Nutrition Label* workstream.

[1] "Generative Artificial Intelligence (AI) is a system of algorithms or computer processes that can create novel output in text, images or other media based on user prompts."

Introduction to DNLs

The concept of a DNL is not novel, as it mirrors the well-established framework of food nutrition labels. For example, in 2018 the DNL project was established and aimed to replicate the success of food nutrition labeling by providing transparency and standardization into the "health" of datasets used for AI systems (Data Nutrition Project, n.d.). Similar to the manner in which food nutrition labels revolutionized transparency into the food industry, DNLs have the opportunity to drive a similar outcome for financial institutions' governance and approach for the use of AI in a prudent and safe manner.

The adoption of food nutrition labels in the United States parallels the potential for use of DNLs within the financial sector. The United States Department of Agriculture (USDA), established in 1862 to protect and improve America's food supply (Food Safety and Inspection Service [FSIS], n.d.) has provided oversight and supported the evolution of food industry standards and safety protocols that continue. Examples include the Pure Food and Drug Act of 1906 ("The Wiley Act"), which prohibited the manipulation or misbranding of food and drugs for human or animal consumption. The law also required that food and drug labels be accurate and not mask inferiority or impurities that would cause potential harm upon ingestion (U.S. Food and Drug Administration [FDA], n.d.).

Improvements to the food and drug industry continued throughout the twentieth century, which saw the establishment of the FDA in 1931 (Food Safety and Inspection Service [FSIS], n.d.). Increased attention to food safety resulted in subsequent Acts and laws being passed or amended to ensure the safety of the food supply, which continues today.

The evolution of food nutrition transparency and safety illustrates the ability of transparency to support a shift within an entire industry from reactivity to proactive transparency, trust, and safety. The DNL concept applies the same logic and construct to the financial industry's use of AI by codifying dataset "ingredients" (provenance, scope, and sampling, akin to serving size), flagging "allergens" (sensitive attributes and downstream use constraints), quantifying "daily values" (quality, completeness, bias, drift, and stability metrics), and stating "freshness dates" (currency, lineage, and refresh cadence). DNLs operationalize governance at the point of data use. For financial institutions, this standardization can augment current model risk management frameworks, auditability, and consumer protection by making dataset health and transparency legible across organizations, vendors, and consumers. Just as food label became a common language for safety and quality, the DNL can become a sector-wide practice for prudent, sound, and fair use of AI and GenAI, which will serve to accelerate adoption while anchoring accountability and trust.

Given the pervasive nature of data and the increasing use of AI in society, coupled with existing regulatory oversight that financial institutions are subject to, the potential impact (positive) of tools like the DNL concept cannot be overstated. New legislation proposals continue to be brought forward, such as the "Artificial Intelligence Research, Innovation, and Accountability Act of 2024" which aims to establish regulatory requirements for increasing the transparency into GenAI systems and specifically highlights the characteristics of the data used for training and input into the system (U.S. Congress, 2024).

Specific to the financial sector's position of handling sensitive personal and financial data of customers, establishing standards that elevate the transparency of AI is paramount for ensuring continued trust and confidence in the products and services provided by financial institutions.

DNLs remain a relatively new concept in financial services, particularly when applied to AI systems. Frameworks, standards, and best practices are still evolving and will mature as firms gain experience and regulatory guidance becomes formalized. As the sector continues to experiment and learn, the DNL approach must remain flexible, adapting to new technologies, evaluation methods, emerging risks, and the practical realities of enterprise implementation and consistency.

Real World Application of DNLs

Beyond the financial sector, DNLs are gaining traction with various organizations and governing bodies. As previously discussed, the rapid pace and breadth at which AI capabilities are being deployed have created an increased need for transparency, accountability, and standardization. This accelerated pace of development and integration creates an opportunity for DNLs to be a practical mechanism to enhance trust and mitigate risk, in a standardized manner.

Initiatives such as the aforementioned Data Nutrition Project and organizations such as Google, NVIDIA, IBM, and Twilio have embarked on similar journeys to support increasing the trust, transparency and explainability of AI solutions.

Specifically, NVIDIA has been developing "Model Cards" that provide transparency into key characteristics of machine learning and other models for both technical and non-technical consumers. Beginning in 2019, the publication of "Model Cards for Model Reporting" introduced a structured approach to enhancing transparency with AI and machine learning development, particularly given the pervasive use of these models across critical industries (Mitchell et al., 2018).

In a similar vein, Twilio also introduced the concept of "AI Nutrition Facts" to promote greater transparency in the development and deployment of AI models and solutions. The goal was to help address the proverbial "black box" that seemingly plagues AI solutions, specifically aiming to provide visibility into how data is used to train LLMs (Nutrition Facts AI, n.d.). Twilio's efforts

have also been referenced in a 2024 report from the National Telecommunications and Information Administration on Artificial Intelligence Accountability, which provides research on responsible, safe, and transparent use of AI (National Telecommunications and Information Administration [NTIA], n.d.).

With continued advancements in AI capabilities, increased transparency and accountability surrounding the use of data for AI solutions is needed for the prudent and responsible use of AI across all industries.

Benefits of DNLs and Supporting AI Explainability (XAI)

DNLs help people understand how AI works by making the data behind them more visible. By highlighting and defining dataset characteristics and behavior, DNLs reduce the sense of a "black box" and build trust for customers, colleagues, and supervisory teams in financial institutions that use these technologies (Zodage, Harianawala, Shaikh, & Kharodia, 2024).

Just as consumers interpret food labels without specific nutrition education, DNLs support business leaders, risk partners, and frontline staff to identify important characteristics about the data that underpins an AI system without needing to be a data scientist or data expert.

As financial institutions expand their use of AI, questions continue surrounding how outputs are produced and how they impact customers and internal teams. Explainability helps to drive confidence (internal and consumer), regulatory alignment, and responsible decision making in high-stakes environments. The DNL framework offers a pragmatic way to support explainability by providing structured, transparent insights into the datasets that underpin model development and monitoring. When institutions align DNL foundational and enhanced ingredients with recognized explainability principles, they create an opportunity to narrow the gap between technical complexity and human comprehension. This manifests by breaking down, what is at times, opaque processes and language surrounding data quality, movement, and integrity and providing a more digestible view into the characteristics of a dataset. The result can lead to a more consistent approach to dataset transparency that further aids internal data management practices and can strengthen external trust.

How DNLs Help People Understand and Govern

- **Improve dataset transparency**: DNLs document key attributes, collection processes, and intended uses so that stakeholders can understand the data that powers AI systems.
- **Reveal limitations and risks**: Labels can surface gaps, biases, and quality issues supporting teams to make informed decisions and plan mitigations.

- **Support audits and reviews**: Clear, structured documentation streamlines and supports oversight by compliance, risk, and internal audit teams.
- **Enable traceability**: DNLs provide enhanced documentation and insights into data sources, supporting the ability to identify how specific inputs may have influenced final results.
- **Make comparisons easier**: Standardized disclosures help teams compare datasets and identify divergences, strengths, and exposures.
- **Strengthen accountability**: Documented provenance, assumptions, and constraints reinforce accountability for AI system behavior and customer impact.

Why Explainability Matters to People

- **Trust and compliance**: Explainability supports the confidence of customers, colleagues, executives, and supervisors, and it underpins regulatory compliance.
- **Managing opacity:** Many AI systems, more specifically GenAI, are not directly interpretable, so explainability can act as a practical control to support the safe use and responsible outcomes.

In addition to supporting explainability, DNLs can help enhance transparency by documenting how, when, and where data may be legally, contractually, and ethically used. Through structured and consistent dataset documentation, DNLs can provide compliance teams and business leaders with a practical, plain language approach to assess whether data aligns with organizational policies and potentially regulatory requirements. This transparency complements existing governance frameworks and reinforcement of responsible data practices.

Well-governed, well-curated data protects both customers and the institution, and DNLs can assist by surfacing critical details such as provenance, collection methods, and usage boundaries. Given the heightened stakes in financial services, institutions can benefit from tools that support bias identification, issue mitigation, and confirmation that a dataset is fit for purpose before it powers an AI solution.

By adopting DNLs, organizations can standardize how they increase visibility into dataset characteristics and behavior. Teams can identify potential risks before development begins, reduce rework from inconsistent or unfit data, and deliver results that better serve customers and internal stakeholders (Holland et al., 2018).

Core Components of DNLs

As referenced, DNLs offer an approach that provides insight into the construct of a dataset prior to the development of AI systems. Research and development efforts for DNLs continue to evolve, as demonstrated by the incremental advancements made by the Data Nutrition Project. With these enhancements informed and supported by ongoing research and feedback from industry experts and practitioners. They offer valuable insight into the challenges and complexities associated with improving the explainability of unique datasets, particularly in the context of AI systems. Feedback received within the Data Nutrition Project can be summarized into three main points:

A. A single dataset cannot be relevant to all consumers
B. Time constraints and lack of expertise to address underlying issues in a dataset
C. Use Case Bias: Practitioners often have a narrow focus when evaluating a dataset for use

This feedback further highlights the importance of the use of DNLs in successful integration and use of AI capabilities, agnostic of industry. (Chmielinski, et al, 2022).

When considering adoption, it is essential to consider the varying levels of maturity that financial institutions may exhibit as they begin their journey with AI.

For that reason, a two-tiered Data Nutrition Label provides a foundational level of transparency to support identifying initial risk, bias, and dataset health, along with an enhanced tier that provides further insight into datasets with more complex characteristics, such as unstructured data.

A core element of the two-tiered approach is that it allows financial institutions, regardless of their data management maturity or staff expertise, to adopt and benefit from the DNL.

Foundational Tier

The foundational tier was designed to support traditional machine learning and deterministic AI models, drawing on core data management practices that should be consistent across most financial institutions. These ingredients provide a practical approach to assess a dataset's key characteristics, supporting AI developers, data owners, and business sponsors' confidence that the data meets baseline standards before it's ingested into a model or AI solution.

To reinforce this consistency and support broader interoperability, institutions should also consider aligning foundational requirements with recognized data publishing standards. Standards such as common metadata formats, interoperability protocols, and documentation requirements help establish that dataset characteristics are internally transparent and acceptable across external organizations and vendors. Incorporating these practices can

strengthen regulatory compliance and make the DNL framework more practical by aligning well-known standards and guidance that integrate into existing governance processes.

The Basel Committee on Banking Supervision highlights that sound data quality is the foundation for digitalization projects and warns that poor source data is a root cause of downstream issues. Common taxonomies and tools that improve consistency and accuracy, along with scalable and adaptable data architectures such as well-designed data lakes, are cited as critical enablers for meeting the needs of internal and external stakeholders (Bank for International Settlements [BIS], 2021).

The table below depicts the foundational ingredients, along with a description, and approach to a proposed calculation, scoring range, and weighting. It's important to note that the scores and potential weightings are intended as a complementary input for data owners or the AI developer to use while developing a use case. The intent is not to have the scores be punitive for the data owner, developer, or business owner of the AI System. These attributes facilitate a baseline assessment of the underlying ingredients, supporting a standardized approach for determining whether a dataset is fit for a given AI use case. Financial institutions implementing the foundational tier may assess ingredient weightings and scoring based on their existing data management practices.

Data Ingredient	Description	Proposed Calculations	Scoring Range*	Potential Weighting*
Data Types	Diversity of data types identified in the dataset (structured, unstructured, etc.)	Score = 0 if structured, 10 if unstructured	0 to 10	.25
Data Volume	Total volume of contained in the dataset (based on minimum established volume threshold)	Score = 1 if ⩾ threshold; 3 if data is 50–100% of the threshold; 5 if <50%	1,3,5	.10
Data Quality Score	Population Stability Index is useful for tabular / structured data, undefined for unstructured data	Population Stability Index Value * 4	0 to Infinity or Undefined	.25
Data Update Frequency	Frequency (recency) by which the data is updated (indication of data staleness)	Categorical scoring: • Real-time / Hourly = 1 • Daily = 2 • Weekly = 3 • Monthly = 4 • Occasional/Undefined = 5	1 to 5	.15
Data Collection Method	Transparency into the available documentation & Controls associated with the data collection methods (automated, manual, streaming, etc.)	• Fully documented with robust controls and audits = 1 • Partially documented/controlled = 3 • Ad hoc and poorly documented = 5	1,3,5	.15
Confidential Data	PII or other confidential fields, especially in open text fields when the potential exists for off-premises use	• 1-5, low for fewer confidential fields, high for larger numbers of confidential fields	1 to 5	.10

**Illustrative, subject to evaluation and modification by the implementing financial institution*

As noted above, research and development efforts for DNLs continue to evolve, and implementation of frameworks like the DNL will continue to evolve as firms continue with their adoption of AI.

Enhanced Tier

While the foundational tier offers a traditional lens grounded in core data management and machine learning practices, the enhanced tier shifts focus to characteristics that are more relevant to the development of AI systems. This tier emphasizes the importance between managing structured and unstructured data and introduces critical concepts such as the presence of data labels, the data sampling strategy, and the application of train, test, and validation splits. These elements provide financial institutions with deeper insight as they work to strengthen AI governance and navigate the added complexities associated with GenAI systems.

The enhanced tier supports explainability and enables a consistent, scalable assessment framework that accommodates varying levels of data maturity and AI use case sophistication. It is structured around five ingredient categories that span data collection and validation: Motivation, Composition, Collection, Preprocessing, and Maintenance. Each category includes targeted questions designed to surface dataset characteristics that are not captured in the foundational tier.

Although the enhanced tier incorporates expert judgment, its value lies in the standardized and repeatable approach it offers for evaluating datasets. This approach aligns with practices commonly used by financial institutions, which balance structured documentation with professional review. By making assumptions, limitations, and lifecycle controls explicit, the enhanced tier reinforces AI explainability, transparency, and standardization.

Archetype / Ingredient	Question	Potential Values
Motivation	Why was the training dataset created?	Open text
	Who created / funded the dataset's creation?	Open text, should be individuals / groups
Composition	What does each record in the dataset represent?	Open text
	How many records are there in total?	Integer
	Does this dataset contain labels?	True / False
	How was the train / test / validation split applied, if used?	Open text
	Are there any errors, data problems or issues, sources of noise, or redundancies in the dataset?	Open text
	Does the dataset contain PII?	True / False
Collection	How was the data acquired?	Open text
	If the dataset is a sample, what was the sampling strategy?	Open text, or N/a
	Over what timeframe was the data collected? Is the data refreshed on a cadence?	Date Range
Preprocessing	Was any preprocessing / cleaning / labeling applied to the dataset? If so, please describe	Open text, or N/a
Maintenance	Are there continuing maintenance requirements? If so, please describe	Open text, or N/a

Benefits of Leveraging Enhanced vs Foundational Ingredients

Adopting the foundational tier independently or by implementing both the foundational and enhanced tiers together will further support explainability and its associated benefits. The enhanced tier, however, offers distinct advantages. It supports increased transparency, strengthens explainability, and introduces a higher level of rigor in documentation by capturing context around motivation, collection methods, preprocessing decisions, and maintenance practices.

In the same manner that dietary guidance varies depending on age, organizations with different levels of data management maturity and risk tolerance require different degrees of structure and assessment. Financial institutions with established governance, infrastructure, and processes (more mature) are better equipped to meet the more comprehensive requirements of the enhanced tier. However, the application of a risk-based approach should continue to underpin the application of enhanced risk management requirements. For example, a financial institution would consider the risk level of the use case that the AI system would support when determining which DNL tier would be appropriate.

Institutions earlier in their journey (less mature) or delivering lower risk use cases can derive immediate value from the straightforward disclosures provided by the foundational tier. This maturity and risk aligned pathway allows firms to begin with the foundational tier, demonstrate value, and scale to the enhanced tier as capabilities and use cases evolve to support more advanced dataset assessments.

The foundational tier supports the delivery of transparency and direct insight by offering a consistent view of basic data types, volume, refresh frequency, and general data quality. This level of transparency is well suited for evaluating and deploying less complex AI use cases, where the associated risks and potential harm are typically lower.

Larger financial institutions and global systemically important banks (GSIBs) may require more detailed insights and advanced controls for AI use cases, depending on the function or service the AI system is intended to support. For example, these institutions often manage more complex projects, such as anti-money laundering, fraud detection, and direct customer support. As with the food nutrition analogy, these organizations require more advanced “nutrition” information to achieve optimal, risk mitigated outcomes.

Current guidance from the Office of the Comptroller of the Currency (OCC) via the Model Risk Management Handbook establishes the importance of including a rigorous review of data quality, data relevance, and fit for purpose assessment of all data used for model development (OCC, 2021). While this guidance remains relevant for the adoption of AI capabilities, the

evolution of these capabilities and the accelerated pace of integration require continued engagement with regulatory agencies regarding best practices in this area.

Acknowledging the variety of use cases and data management maturity levels across the financial sector is critical when evaluating the viability of the DNL framework. The implementation and maintenance of the enhanced tier could potentially create an economic and/or employee skillset constraint that may limit the adoption to larger financial institutions. Recognizing this potential barrier led to the creation of the two-tiered approach. Smaller financial institutions can leverage the foundational tier, while more mature organizations, with greater resources and higher risk use cases, can adopt both the foundational and enhanced tiers.

Roles and Responsibilities

According to The National Institute of Standards and Technology (NIST) AI Risk Management Framework, "*Trustworthy AI depends on accountability. Accountability presupposes transparency.*" (NIST, 2023).

To drive accountability and sustainability of DNLs within the financial sector, clear roles and responsibilities are critical for success. In alignment with the two-tiered approach (foundational and enhanced), these roles and responsibilities should reflect the understanding that accountabilities will vary across organizations, depending on how they deploy, govern, and maintain the DNL framework.

Clear roles and responsibilities are essential for financial institutions across all operational activities, however when ensuring the prudent and responsible use of AI, this becomes critical.

There are two central roles to support the successful adoption of the DNL framework, data owner and AI developer. However, it is also important to acknowledge the role that the business owner (e.g., AI system business sponsor/owner) and Model Risk Management (MRM) will play, depending on the organization.

Data Owners:

Data owners should be accountable for the identification of authoritative data sources, and the determination of which data source is the authoritative source for each ingredient. According to a McKinsey survey, 24% of respondents cited the lack of available data as a barrier to the adoption of AI capabilities (McKinsey & Company, 2023). This highlights why the role of a data owner is critical to ensure that data is discoverable, reliable, and accessible. To ensure this, data owners should:

- **Ensure Governance Over the Dataset**: Overall quality, integrity, and security of their datasets. This includes enforcing and adhering to internal data governance policies as well as and existing regulatory expectations, such as the General Data Protection Regulation (GDPR) for personal data handling. Which states in Chapter 4, Article 24, that the controller shall implement appropriate measures to ensure processing is compliant with the regulation (GDPR Info, n.d.a).
- **Make Available Ingredients and Analysis Results for Each Ingredient**: Detailed information for each ingredient within the dataset, including collection method, known issues or defects, and data refresh frequency.
- **Aggregation of the Ingredients into Nutrition Label**: Ensuring all ingredients are accurate and available for inclusion into the DNL.

AI Developers:

The data scientists, or AI practitioners, who will be the primary consumer of the DNLs are accountable to ensure the selection of the dataset meets the requirements of the AI use case.

AI developers should:

- **Ensure the Data is Fit for Use**: Evaluate the contents of the DNL against the proposed use case to ensure it aligns with the intended use.
- **Proactive Risk and Bias Assessment**: Assess the output from the DNL and ingredients to ensure the dataset is free from bias and/or does not contain documented issues or inconsistencies that would materially impact the performance of the AI system.
- **Informed Model Selection**: Complete holistic assessment of the dataset to ensure comprehension of the dataset construct, aggregation, and origin.

Designating the data owner and the AI developer as the central roles within the DNL framework does not exclude other stakeholders, such as the business owner of the AI system or MRM. Rather, this designation emphasizes that the DNL is grounded in the hands-on development and assessment activities required for prudent AI system development. The goal is for the DNL to become a practical, supportive artifact that captures dataset health, scope, and quality, and helps control risk within the development process, while still allowing for oversight from business and risk functions.

With any development effort, the business owner must be explicitly informed of the benefits, risks, and limitations, associated with the data used. The DNL can support this by providing a more standardized approach to transparency on dataset health, including population coverage, quality checks, known biases, and usage constraints. This transparency will help business owners make decisions based on tangible evidence, while also maintaining oversight through their organization's change management framework.

Within financial institutions, the presence of MRM is codified, although the specific role MRM plays in validating AI will vary by organizational structure and the continued advancement of GenAI capabilities. Regardless of how responsibilities are allocated between MRM and other functions, organizations should evaluate how to integrate the DNL into existing governance mechanisms, including model inventories, validation protocols, and documentation standards. Where MRM review is required prior to deployment, aligning DNL artifacts to support existing MRM requirements can accelerate deployment readiness, and improve transparency by making data related assumptions and validations more consistent.

Clear delineation of roles and responsibilities between data owners and AI developers is essential for the successful adoption of the DNL framework. Data owners should be accountable for the designation, completeness, and timeliness of authoritative data sources. AI developers should be responsible for selecting data that is fit for use, given the specific use case. By working together, and in partnership with business owners and MRM, these roles help support enhanced explainability, improve reliability, and foster greater organizational trust in AI.

Implementation Considerations

The implementation of DNLs in the financial sector for AI capabilities must navigate a complex set of challenges, including but not limited to strict regulatory guidance and expectations (state, federal, and international) as well as varying degrees of organizational maturity and readiness to adopt AI. Continued adherence to established regulations and guidance as well as identifying opportunities to adopt and influence new requirements is critical for the responsible adoption of AI.

The development, implementation, and use of the DNL will largely complement existing guidance such as the OCC's Model Risk Management Handbook and the Federal Reserve Board's (FRB) Supervisory Guidance on Model Risk Management (Board of Governors of the Federal Reserve System, 2011), by enhancing the processes, oversight, and standardization for model development.

When considering the implementation of DNLs there are four areas of focus that should be addressed to support successful adoption of the DNL with the use of AI.

- **Sustainability:** Ensure the process to define, generate, and use the DNL integrates into existing development and governance frameworks to avoid confusion or redundancy.
- **Clear Expectations:** Document and codify expectations within the firm's standard governance protocols for how DNLs will be applied, including required actions when results do not meet standards.
- **Roles and Responsibilities:** Define roles for AI developers, data owners, and governance functions to support proper oversight and execution.

    - **Additional Consideration**: Define required actions if the DNL is not created or the results of the DNL do not meet standards.

- **Accountability and Performance Monitoring:** Establish an accountability framework aligned with firm policies and procedures, ensuring roles and responsibilities are documented and enforceable.
    - **Additional Considerations**:
        - Incorporate measurable indicators, such as the percentage of datasets with a DNL or quality scores over time, to monitor adherence and drive intended behaviors.
        - Define escalation protocols for lack of adherence and/or poor-quality scores, ensuring consequences are consistent with the firm's standards while maintaining a non-punitive, improvement-focused approach.

Within the FSSCC's AI DNL workstream, a consistent theme emerged around ensuring that DNLs do not become punitive or administratively onerous, unintentionally stifling innovation or adoption.

To support this, financial institutions are encouraged to evaluate how the DNL can be integrated as an enhancement to their existing AI governance frameworks. This integration will support increased transparency in the development of AI capabilities and should also accelerate adoption by leveraging established governance and organizational structures. Financial institutions adopting the DNL should assess their current data management, MRM, and AI governance frameworks to determine the most effective way to realize the potential value of DNLs within their organization and governance structures.

Financial institutions should avoid over-engineering the integration of the DNL into the MRM framework if it does not align with organizational policy and structure for AI evaluation. Specifically, organizations that do not incorporate AI system evaluation within their MRM framework may find it useful to assess the DNL during the initial intake or review of an AI use case. This flexibility further supports flexibility for financial institutions to implement their evaluation in a way that best aligns with their organizational structure.

Building on this, guidance from the NIST Artificial Intelligence Risk Management Framework ("AI RMF") playbook highlights the need to identify risks associated with a lack of control over dataset sourcing, testing, and evaluation. Measure 2.2 within the AI RMF provides guidance such as evaluating data representativeness, establishing thresholds and alert procedures for data within its context of use, and documenting procedures and processes that effectively mitigate data bias (NIST, n.d.). The AI RMF does not dictate where or how this evaluation should occur; instead, it offers organizations suggested actions and guidance for the prudent adoption of AI capabilities.

Successful adoption of DNLs for financial institutions is dependent upon positioning them as enhancements to existing practices and frameworks, versus as a standalone construct. Akin to OCC and Federal Reserve guidance on MRM, DNLs should be viewed as an augmentation, and where necessary, an evolution to current data quality, governance, and documentation controls to keep pace with advancing AI capabilities. To the extent possible, the DNL should integrate into these structures and support the clarification of accountability, decision rights, and handoffs across development, review, and deployment processes. By embedding DNL artifacts into established governance frameworks, such as model evaluations and standards, documentation templates, and change management, institutions can accelerate adoption, increase transparency, and drive sustained risk mitigation, ultimately improving trust in AI systems.

Ultimately, each financial institution is best positioned to determine whether and how the DNL framework should be integrated into its governance structures to realize potential benefits.

Tracking Adoption and Benefits of the DNL

The ability to objectively and consistently track how well an organization adopts a framework is essential to the success of any program. If a financial institution chooses to integrate the DNL framework into its existing governance structures, it will be important to monitor both the adoption rate and the benefits realized.

To support effective implementation, organizations should establish clear metrics to measure both adoption and impact. Below are five suggested metrics designed to track the progress and benefits of the DNL framework:

Metric	Definition	Owner	Notes
Time to Deployment	Median days from data selection to AI solution deployment	AI / GenAI Governance Team	Ownership to be determined based on organizational alignment
AI Solution Rollback Rate	% of AI solutions that required a rollback due to data issues	AI / GenAI Governance Team	Organizations will need to align to existing release & deployment triggers that initiate rollback activities.
Adoption & Coverage	% of AI use cases with Data Nutrition Label	AI / GenAI Governance Team	Long-term goal for adoption / coverage may vary based on an organization's model classification.
Audit Findings	% increase (or decrease) for Internal Audit findings related to data inconsistencies	AI / GenAI Governance Team	Will need to baseline against current Internal Audit finding rates
Data Quality Scores	% increase in average data quality scores for datasets that have an associated Data Nutrition Label compared to their baseline score prior to DNL implementation.	Data Governance Team	Will need to baseline against current data quality scores.

Expectations for Third Party Vendors that Provide AI Capabilities

While many financial institutions may rely on third-party vendors who provide AI and GenAI capabilities to support operational needs and strategic goals, they must evaluate how these service providers will meet the transparency and explainability requirements related to AI capabilities.

Navigating standards for regulated and unregulated parties (third party vendors) require significant investment and coordination from all involved (financial sector, technology/third-party vendors, and regulatory agencies, etc.). However, there is precedent for achieving this. The healthcare industry underwent a rapid technology transformation with the introduction of electronic medical records ("EMR") near the turn of the twenty-first century. Pointing to the ratification of the Health Information Technology for Economic and Clinical Health Act of 2009 ("HITECH Act"), where authority was given to the Department of Health and Human Services to establish programs and standards for the prudent and safe use and adoption of technology across the healthcare industry. The HITECH Act also included standardization for how protected health information should be exchanged via electronic platforms. This regulation eventually supported the development of the Fast Healthcare Interoperability Resources (FHIR) framework, which enhanced the Health Level 7 ("HL7") framework to facilitate the ease of exchanging information between provides, patients, payers, and third parties within the healthcare industry. (Office of the National Coordinator for Health Information Technology [ONC], n.d.a/b/c).

Similar to the healthcare's adoption of interoperability standards, the evolution of AI service FactSheets ("FactSheet") offers a framework for standardization and increasing transparency with third-party vendors. FactSheets are not entirely novel and are informed by Supplier's Declarations of Conformity ("SDoC"), widely used in industries such as telecommunications and consumer products to provide written assurance that a product or service conforms to specified requirements (Arnold et al., 2019). Just as the HITECH Act established a pathway for transparency and standard data exchange in healthcare, FactSheets aim to create a standardized mechanism for communicating the attributes of AI services, including purpose, performance, safety, security, and lineage (Arnold et al., 2019). While initially proposed to be voluntary, FactSheets are envisioned to mature into industry requirements, beginning with a period of collaboration between AI service providers and consumers to refine their structure and content, and ensuring balance between consumer needs and the flexibility for innovation by AI service providers (Arnold et al., 2019).

As financial institutions rapidly pursue the adoption of AI, they can draw lessons from other sectors' technological transformations, which demonstrate the value of collaboration, transparency, and trust. Highlighting these successes is not meant to understate the effort

required for industry-wide coordination; rather, they serve as exemplars for advancing technology responsibly. By establishing standards and committing to collective governance, the financial sector can scale AI responsibly. In this context, frameworks such as DNLs complement FactSheets by providing transparency at the dataset level. Together, these tools support consistent and transparent disclosures that drive trust and partnership between vendors and financial institutions.

Regulatory Guidance

With the rapid advancement and integration of AI into everyday life, there are targeted efforts to ensure and promote the responsible and safe use of these capabilities. The need to adequately control and mitigate risks that are implicit with the use of AI extends well beyond the financial sector. As such, regulations have been established globally to help control and mitigate for anticipated risks, while enabling industries to leverage the advancements in technology.

GDPR: The standard for data protection in Europe defines seven principles that must be addressed when handling the personal data of European citizens (GDPR Info, n.d.b). Specifically, referencing Article 5 and the requirements for processing personal data:

a) Lawfulness, Transparency, and Fairness
b) Purposeful Limitation (only for legitimate purposes; nothing further)
c) Data Minimization
d) Accuracy and Staleness
e) Storage and Retention of Personal Data
f) Security of Personal Data
g) Accountability (of Data Controller/Owner)

The concept of DNLs can assist in interpretating how the GDPR established requirements for transparency, fairness, and accountability when handling personal data, inclusive of the use of AI.

Health Insurance Portability and Accountability Act (HIPAA): The healthcare industry has strict requirements on the accessibility and sharing of personal health information (PHI) and personally identifiable information (PII). Pointedly, when the Health Insurance Portability and Accountability Act of 1996 became law in August of 1996, it provided standards for the exchange of health information. Targeting privacy and security HIPAA has become an industry standard, accelerated by the proliferation of EMRs. The applicability of HIPAA standards and requirements extends beyond healthcare, as the financial sector also handles customers' PII and PHI and must adhere to HIPAA standards. (U.S. Department of Health & Human Services [HHS], n.d.a) and (HHS, n.d.b).

National Artificial Intelligence Initiative Act (NAIIA) of 2020: The United States Government signed into law the NAIIA in 2020 with the goal to advance collaboration, establish guidelines, develop a risk mitigation framework for the deployment of AI systems and support overall trustworthiness of AI systems. Specifically, the law directs the NIST to play a critical role in developing standards and guidelines for the prudent, safe, and effective use of AI (NIST, n.d.b). Through this effort, the NIST has developed the NIST AI RMF, to support the increase of transparency and trustworthiness in the development of AI systems (NIST, n.d.c) and (U.S. Congress, 2020).

White House AI Action Plan and Executive Orders: The signing of Executive Orders in support of the White House's AI Action Plan ("Action Plan") provide additional guidance for the development, implementation, and use of AI in the United States. Specifically, the Action Plan calls for building high-quality datasets as a "national strategic asset" and directs the National Science and Technology Council (NSTC) Machine Learning and AI Subcommittee to recommend data quality standards. The Action Plan further recommends policies to address AI interpretability and system control. Specifically, it directs NIST, along with the Department of Commerce, Department of Energy, and the National Science Foundation, to publish guidelines and develop "the science of measuring and evaluating AI models…" to support the assessment of AI system performance and reliability. The recent publication of the Action Plan allows for insight into how novel guidance and policy recommendations will impact the financial sector. Even without explicit guidance for financial institutions, we can infer correlating efforts such as the need for data quality standards and increased explainability and interpretability of AI systems (The White House, 2025).

Regulatory guidance and expectations will continue to be introduced and codified through legislative actions, both in the United States and abroad. These regulations and guidance are not comprehensive, rather, they highlight the emerging trend of regulatory expectations being established to promote the safe, responsible, and accelerated adoption of AI. As demand for AI increases, we can expect additional, more granular state-level, federal, and global regulatory requirements to follow, which will have direct implications for the financial sector.

Outlook and Considerations

The successful adoption of DNLs across the financial sector will depend on several key components. Continued research, industry collaboration, and regulatory engagement. As AI capabilities continue to evolve, firms must remain nimble in adapting governance frameworks while maintaining a focus on transparency, accountability, and interoperability. Sector wide standards and collaboration, clear regulatory guidance, and active engagement with third-party vendors will be imperative to scaling AI responsibly. Collaboration between financial institutions, technology providers, and regulators should focus on alignment with codified governance

frameworks and refining DNL standards. Furthermore, DNLs must remain flexible to accommodate new capabilities, including synthetic data.

As synthetic data becomes more prevalent with AI development, future iterations of DNLs should incorporate standardized components to support transparency and explainability. Synthetic data is increasingly used to supplement or replace actual datasets in development and testing, but its quality and impact must be assessed to maintain confidence. The evaluation of synthetic data quality cannot be a one time or static activity. It requires continuous testing to ensure fidelity, utility, and alignment with intended use cases (Financial Conduct Authority [FCA], 2025).

Future DNL frameworks should consider the inclusion of standards for documenting generation methods, performance benchmarking results, and Train-Synthetic-Test-Real (TSTR) evaluations. The TSTR evaluation can help AI developers understand the quality of the synthetic data and how well it aligns with the performance of the model using real data (Financial Conduct Authority [FCA], 2025).

Beyond the continued research for the use of synthetic data, there are several other components that will support the proliferation of DNLs across the sector.

- **Continued Research and Development**: Ongoing refinement of the DNL framework and adaptation to emerging technologies and capabilities is essential for addressing the complexities of the financial sector.
- **Industry Standards and Collaboration**: Drive interoperability and partnerships across the public and private sectors, leveraging lessons from other industries such as healthcare. Ensuring the inclusion of practitioners, legal experts, regulatory agencies, and financial sector leaders.
- **Third-Party Vendor Engagement**: Maintain collaboration with vendors providing AI capabilities to promote transparency and responsible partnerships.
- **Regulatory Backing and Incentives**: Regulators and legislators will play a critical role in establishing clear guidance for the responsible use of AI. Securing regulatory support can significantly accelerate adoption by financial institutions.

By continuing to research, collaborate, and iterate on the DNLs, the framework can support the financial sector's goal of leveraging new and emerging AI capabilities in a prudent and responsible manner.

Conclusion

The implementation of DNLs in the financial sector represents an ongoing evolution toward greater transparency, accountability, and trustworthiness in AI. Drawing inspiration from the

familiar food nutrition label, DNLs provide a practical approach to improving the understanding of the data underpinning AI systems. DNLs can also help reduce concerns surrounding the "black box" perception that has sometimes hindered stakeholder adoption of AI. The proposed two-tiered approach recognizes the diverse levels of data management maturity and varying use case risk across the financial sector. This structure promotes the broader adoption of AI while supporting the development of more advanced governance and risk mitigation capabilities over time.

Ultimately, the successful adoption of DNLs will require sector buy-in, clearly defined roles and responsibilities, consistent standards, and careful consideration of consequences. This approach can support a firm's risk mitigation and assessment of data, ultimately helping to safeguard consumers and support the continued AI adoption journey in the financial sector.

References:

Arnold, M., Bellamy, R. K. E., Hind, M., Houde, S., Mehta, S., Mojsilović, A., Nair, R., Natesan Ramamurthy, K., Reimer, D., Olteanu, A., Piorkowski, D., Tsay, J., & Varshney, K. R. (2019). FactSheets: Increasing trust in AI services through supplier's declarations of conformity. Retrieved December 30, 2025, from *arXiv preprint arXiv:1808.07261*. https://doi.org/10.48550/arXiv.1808.0726

Board of Governors of the Federal Reserve System. (2011). Supervisory guidance on model risk management (SR 11-7) [PDF]. Retrieved July 12, 2025, from https://www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf

Bank for International Settlements. (2021). *Progress in adopting the Principles for effective risk data aggregation and risk reporting* (BCBS d559), Section 2.3.5: Banks should ensure sound data quality as the foundation for digitalisation projects. Retrieved December 29, 2025, from https://www.bis.org/bcbs/publ/d559.pdf

Chmielinski, K. S., Newman, S., Taylor, M., Joseph, J., Thomas, K., Yurkofsky, J., & Qiu, Y. C. (2022). The Dataset Nutrition Label (2nd Gen): Leveraging context to mitigate harms in artificial intelligence. Retrieved June 21, 2025, from https://arxiv.org/pdf/2201.03954v1

Data Nutrition Project. (n.d.). Data Nutrition Project. Retrieved May 30, 2025, from https://datanutrition.org/

Financial Conduct Authority. (2025). *Generating and using synthetic data for models in financial services: Governance considerations*. Retrieved January 2, 2026, from https://www.fca.org.uk/publication/corporate/synthetic-data-models-financial-services-governance.pdf

Food Safety and Inspection Service. (n.d.). Our history. U.S. Department of Agriculture. Retrieved June 19, 2025, from https://www.fsis.usda.gov/about-fsis/history

GDPR Info. (n.d.-a). General Data Protection Regulation (GDPR). Retrieved July 12, 2025, from https://gdpr-info.eu/

GDPR Info. (n.d.-b). Article 5 GDPR – Principles relating to processing of personal data. Retrieved July 12, 2025, from https://gdpr-info.eu/art-5-gdpr/

Holland, S., Hosny, A., Newman, S., Joseph, J., & Chmielinski, K. (2018). The dataset nutrition label: A framework to drive higher data quality standards. Retrieved June 21, 2025, from arXiv preprint arXiv:1805.03677. https://arxiv.org/abs/1805.03677

McKinsey & Company. (2023). AI adoption advances, but foundational barriers remain. Retrieved August 23, 2025, from https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2018). Model cards for model reporting. Retrieved June 21, 2025, from https://arxiv.org/pdf/1810.03993

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) [PDF]. U.S. Department of Commerce. Retrieved June 21, 2025, from https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

National Institute of Standards and Technology. (n.d.-a). AI RMF playbook: Measure 2.2. Retrieved June 21, 2025, from https://airc.nist.gov/airmf-resources/playbook/measure/#Measure%202.2

National Institute of Standards and Technology. (n.d.-b). AI congressional mandates & executive orders. Retrieved July 14, 2025, from https://www.nist.gov/artificial-intelligence/ai-congressional-mandates-executive-orders

National Telecommunications and Information Administration. (n.d.). AI accountability policy report [PDF]. U.S. Department of Commerce. Retrieved August 23, 2025, from https://www.ntia.gov/sites/default/files/publications/ntia-ai-report-final.pdf

Network of the National Library of Medicine. (n.d.). Generative artificial intelligence. National Library of Medicine. https://www.nnlm.gov/guides/data-thesaurus/generative-artificial-intelligence

Nutrition Facts AI. (n.d.). Nutrition Facts AI. Retrieved August 23, 2025, from https://nutrition-facts.ai/

Office of the Comptroller of the Currency. (2021). Model risk management. Comptroller's Handbook. Retrieved June 21, 2025, from https://www.occ.gov/publications-and-resources/publications/comptrollers-handbook/files/model-risk-management/index-model-risk-management.html

Office of the National Coordinator for Health Information Technology. (n.d.-a). Health IT legislation. Retrieved July 12, 2025, from https://www.healthit.gov/topic/laws-regulation-and-policy/health-it-legislation#:~:text=The%20Health%20Information%20Technology%20for,that%20relate%20to%20ONC's%20work

Office of the National Coordinator for Health Information Technology. (n.d.-b). HITECH Act excerpt from ARRA with index [PDF]. Retrieved July 12, 2025, from https://www.healthit.gov/sites/default/files/hitech_act_excerpt_from_arra_with_index.pdf

Office of the National Coordinator for Health Information Technology. (n.d.-c). Health Level 7 (HL7) Fast Healthcare Interoperability Resources (FHIR). U.S. Department of Health and Human Services. Retrieved July 12, 2025, from https://www.healthit.gov/topic/standards-technology/standards/fhir#:~:text=TEFCA-,Health%20Level%207%20(HL7)%20Fast%20Healthcare%20Interoperability%20Resources%20(FHIR,and%20promotes%20improved%20health%20outcomes

The White House. (2025). America's AI action plan [PDF]. Retrieved August 5, 2025, from https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf

U.S. Congress. (2020). Conference report to accompany H.R. 6395, National Defense Authorization Act for Fiscal Year 2021 (National Artificial Intelligence Initiative Act of 2020) [PDF]. Retrieved July 12, 2025, from https://www.congress.gov/116/crpt/hrpt617/CRPT-116hrpt617.pdf#page=1210

U.S. Congress. (2024). S.3312 — 118th Congress: To provide a framework for artificial intelligence innovation and accountability, and for other purposes [Bill text]. Congress.gov. Retrieved August 23, 2025, from https://www.congress.gov/bill/118th-congress/senate-bill/3312/text#toc-id18e3ef9e44ae4ca2a683b5aaad2c1ecc

U.S. Department of Health & Human Services. (n.d.-a). HIPAA privacy rule: Laws and regulations. Retrieved July 19, 2025, from https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html

U.S. Department of Health & Human Services. (n.d.-b). Summary of the HIPAA privacy rule [PDF]. Retrieved July 19, 2025, from https://www.hhs.gov/sites/default/files/privacysummary.pdf

U.S. Food and Drug Administration. (n.d.). How chemists pushed for consumer protection: The Food and Drugs Act of 1906 [PDF]. Retrieved August 8, 2025, from https://www.fda.gov/files/about%20fda/published/How-Chemists-Pushed-for-Consumer-Protection--The-Food-and-Drugs-Act-of-1906.pdf

Zodage, P., Harianawala, H., Shaikh, H., & Kharodia, A. (2024). Explainable AI (XAI): History, basic ideas and methods. *International Journal of Advanced Research in Science,

Communication and Technology*, 4(1), 123–130. . Retrieved July 19, 2025, from https://www.ijarsct.co.in/Paper16988.pdf