The Lot That Could Not Be Found

Food traceability, uncertain inventory and the evidentiary limits of artificial intelligence after the FDA's November 2026 consultation

José M. López, BSc, MD, PhD, LLD (C)

Abstract

The United States Food and Drug Administration's November 2026 consultation on lot-level traceability exposes a fundamental problem in contemporary food logistics: an inventory system may know how many cases remain without knowing which production lots they represent. The distinction matters whenever products are commingled, repacked, transformed or moved through facilities that record transactions more reliably than physical identity. This article analyses the legal architecture of the US Food Traceability Rule, its deferred enforcement, and the different obligations operating in the European Union, Spain and Great Britain. It then develops a technical framework for distinguishing observed lot identity from inferred provenance, quantifying uncertainty, validating algorithms and testing recall performance. Particular attention is given to mixed pallets, continuous production, ingredient transformation and microbiological incident response. The central proposition is that digital traceability should be evaluated as an evidentiary system, not merely as a database or software procurement exercise. A proposed research agenda examines whether probabilistic reconstruction can reduce investigation time without creating unjustified certainty or narrowing recalls unsafely.

Keywords: *food traceability; FSMA 204; lot genealogy; uncertainty; artificial intelligence; recall simulation; evidentiary integrity; food law*.

1. A consultation that reveals a structural weakness

On 6 November 2026, the FDA convened a public meeting on challenges and solutions in lot-level food traceability. The meeting was not an amendment to the Food Traceability Rule. Its significance is more diagnostic: Congress had directed the agency to engage regularly with regulated businesses on possible flexibilities, and the FDA invited discussion about the practical difficulty of preserving lot identity through complex distribution systems. The October Federal Register notice also provided a route for written comments through 6 December 2026 [1–3].

The rule's original compliance date was 20 January 2026. FDA proposed moving it by thirty months to 20 July 2028; subsequently, federal legislation directed the agency not to enforce the rule before that latter date. These are distinct legal events. A proposed amendment to a compliance date is not identical to a statutory non-enforcement direction, and neither should be represented as repeal of the substantive requirements [1,2].

The central practical difficulty is not that businesses lack barcodes. It is that business events and physical product movements are not always recorded at the same granularity. A warehouse-management system may track item, quantity, location and shipping date while retaining only a

partial connection between each outgoing case and its original lot. When a recall begins, the system can produce a plausible inventory narrative, but plausibility is not evidence of identity.

This distinction has become especially consequential as machine-learning systems are marketed for supply-chain reconstruction. A model can infer which lot probably moved through a warehouse; it cannot retroactively observe an unrecorded event. The scientific and legal question is therefore how an operator should represent, validate and use uncertainty rather than concealing it beneath a single asserted answer.

2. The physical meaning of a lot

A production lot is an operational and documentary construct. Its boundaries may reflect a manufacturing run, a defined time interval, a raw-material receipt, a packing period or a regulatory traceability lot code. None of these boundaries necessarily coincides with the microbiological distribution of a hazard. A contamination episode may span several nominal lots; conversely, a single lot may contain products with very different exposure histories.

Food safety depends on the relation between hazard-generating events and product populations. Consider a line producing a chilled ready-to-eat food. If a contamination source emerges after a sanitation failure, the relevant population may begin at the first affected unit and end only after validated restoration of control. A daily lot code could be too broad or too narrow depending on the process. A traceability system that preserves codes perfectly can still misrepresent the hazard boundary.

Conversely, an ingredient supplier may ship a dry powder in several deliveries from one production batch. If each delivery receives a new warehouse identifier, an enterprise-resource-planning system can inadvertently fragment the genealogy. The record appears precise because each identifier is unique, yet the causal relationship among deliveries has been lost. Identifier proliferation is not traceability.

The defensible design principle is to maintain several linked identities: manufacturing batch, regulatory traceability lot, logistics handling unit, ingredient input and finished-product output. These identities should be connected by explicit transformation events rather than forced into a single universal code. This is an engineering interpretation, not a newly imposed legal requirement.

3. What US law requires, and what it does not

The US Food Traceability Rule, promulgated under section 204(d) of the Food Safety Modernization Act, establishes additional records for covered foods on the FDA Food Traceability List. The structure uses critical tracking events and associated key data elements, with a traceability lot code providing continuity across relevant transactions. Its reach includes qualifying domestic and foreign operators handling food for the US market. Coverage depends on the listed food, its form, the activity and applicable exemptions; it is not a universal requirement for every ingredient or food business [2,4].

The FDA framework requires records to be available within twenty-four hours, or another reasonable time agreed with the agency, in specified circumstances. A business therefore needs not only retained information but the ability to retrieve and assemble it quickly. The difference between archival retention and operational recall readiness is frequently underestimated.

The November consultation concerns potential flexibility in achieving lot-level tracking, particularly where real-world handling breaks straightforward identity chains. Any inference that the FDA has

already approved probabilistic substitution for mandatory records would be premature. Discussion papers, public comments and meeting statements are inputs to policy development, not binding permissions [1,3,5].

Operators should also distinguish non-enforcement of this particular rule before July 2028 from continuing obligations under other federal food-safety, recordkeeping, recall and adulteration provisions. Delayed enforcement does not suspend the need to identify affected products or cooperate in investigations. In practical governance terms, the period is an opportunity for testing, not a justification for postponing data architecture.

4. Europe, Spain and Great Britain: similar objectives, different legal architecture

Article 18 of Regulation (EC) No 178/2002 establishes traceability throughout production, processing and distribution and requires operators to identify their suppliers and business customers. The familiar one-step-back/one-step-forward model is a legal baseline, not a complete description of a modern manufacturing genealogy. Its breadth covers food, feed, food-producing animals and substances intended or expected to be incorporated into food or feed [6].

Commission Implementing Regulation (EU) No 931/2011 provides additional information requirements for food of animal origin. These include, as applicable, accurate descriptions, quantities, supplier and recipient details, lot/batch/consignment references and dispatch dates. For an EU poultry or meat processor, the interplay of these records with hygiene and microbiological controls is more important than a generic statement that the enterprise has “full traceability” [7].

In Spain, the directly applicable EU instruments operate alongside national competent-authority structures and official-control procedures. A software certificate or ISO certification cannot replace the operator’s obligation to produce accurate records. The Spanish implementation question is frequently evidentiary: can the establishment show that the actual lots entering and leaving a process correspond to its documented genealogy?

Great Britain retains its own food-law framework derived in part from the former EU system, administered through domestic institutions and legislation as amended. Northern Ireland requires separate treatment because its relevant food-law arrangements remain linked to EU rules. Businesses supplying both Great Britain and the EU should verify current consolidated territorial texts rather than treating “UK compliance” as a single undifferentiated category.

The legal comparison is asymmetrical. The US framework prescribes additional event-based information for a defined high-risk list; EU general food law imposes broad traceability duties, supplemented by sectoral requirements. Neither architecture automatically validates a particular AI inference. Where the original transaction is unknown, the operator must distinguish recorded facts from reconstructions and avoid representing the latter as contemporaneous observations.

5. From transaction records to an auditable genealogy

A useful digital representation is a directed event graph. Nodes identify physical or logical product populations, and edges represent observed operations: receiving, splitting, mixing, transforming, packing, storing and shipping. Each edge should carry time, facility, operator or device, quantity, unit of measure, source record and validation status. The graph is not necessarily a blockchain. An

ordinary relational database with robust event history, access controls and immutable audit logs can be more suitable.

The most difficult edges occur where materials merge or divide. If two ingredient lots are added to a blender and the resulting bulk material is divided into four finished-product lots, the genealogy is many-to-many. A simplistic database that assigns one input lot to one output lot destroys information. A rigorous model records all plausible contributors and the relevant mass-balance or residence-time assumptions.

Continuous processing adds another complication. Input material introduced at 09:00 does not necessarily emerge at 09:00, and residence-time distributions can create overlap between nominal batches. Flow rate, mixing volume, dead zones, rework and start-up transients all affect the mapping. For a risk investigation, the relevant product window may need to include uncertainty bands derived from engineering data.

Event provenance must also be preserved. A scanned case code is an observation. A code entered manually after a shift is a retrospective declaration. A lot selected by an algorithm is an inference. These are different evidence classes even if all three are stored as the same character string. Systems should never erase the distinction when data are exported to an authority or customer.

6. The mathematics of uncertain identity

Suppose a distribution centre has three candidate incoming lots for a shipment whose original case-level association was not captured. The facility can infer a probability distribution from receipt times, pick sequences and inventory movements. That distribution may support internal prioritisation, but it does not establish that any one lot was actually shipped.

A probabilistic model should express its assumptions. For example, first-in-first-out behaviour may be a reasonable prior in a controlled warehouse but invalid during stock rotation, replenishment, partial pallets or manual overrides. If the operational process violates the assumption, an apparently confident probability is an artefact of the model rather than a property of the evidence.

Calibration is indispensable. Among cases assigned 80% probability to a particular lot, approximately 80% should truly belong to that lot in a representative validation set. Discrimination alone, -ranking the right lot above alternatives-, is insufficient. In a recall, the decision threshold depends on asymmetric consequences: excluding a contaminated lot may cause severe harm, while including additional lots imposes economic and waste costs.

Bayesian methods, constrained optimisation and graph algorithms are all plausible tools, but each has distinct failure modes. A mass-balance constraint can eliminate physically impossible mappings; it cannot determine which of two equally feasible pallets was picked. A machine-learning model can identify anomalies; it cannot create ground truth absent from the underlying process. An honest system may need to return “identity unresolved”.

Where public-health decisions are involved, uncertainty should widen the candidate product population unless there is validated evidence to narrow it. This is a proposed risk-governance principle rather than a numerical legal threshold specified by FDA or EU law. Any operational threshold must be set through documented hazard analysis and competent-authority expectations.

7. Case study: a mixed-pallet distribution failure

Consider a hypothetical manufacturer producing six lots of a chilled ready-to-eat product during one day. A logistics provider receives 480 cases, builds mixed pallets for several retailers and records the total quantity by stock-keeping unit, but not the original lot of each case. Two days later, a positive environmental investigation identifies a possible contamination window affecting two manufacturing lots.

If the distributor can reconstruct only the destination of each SKU, the manufacturer may need to recall all shipments that could contain either affected lot. The cost is driven not merely by the contaminated production volume but by the loss of information downstream. If pallet-to-case scans had preserved the original codes, the potentially affected customer set could be smaller.

An inference engine might estimate that retailer A received mostly earlier lots because its order was prepared first. That inference may help investigators focus verification efforts. It should not, without validation and authority acceptance, be used to exclude retailer B from a protective action when physical mixing makes the affected lot plausible there.

The case demonstrates a counterintuitive economic point: traceability expenditure can create value by preventing unnecessary recall expansion, even when it has no effect on the underlying contamination probability. A credible business case should therefore model both preventive controls and the conditional costs of investigation, withdrawal, recall, reputation and waste.

The scenario is illustrative and contains no real customer data. Its purpose is to expose the analytical difference between knowing product quantity and knowing product provenance.

8. Microbiology and hazard boundaries: where traceability meets process science

Traceability becomes meaningful only when coupled to a scientifically justified hazard model. For Salmonella in a dry ingredient, the risk may involve persistent contamination across equipment or lots. For Listeria monocytogenes in chilled ready-to-eat foods, environmental persistence, post-lethality exposure and shelf-life growth may determine the affected population. For chemical contaminants, the relevant boundary may be a raw-material source or a mixing event rather than a sanitation interval.

Consequently, a recall engine should not assume that lot code alone determines exposure. It should accept incident-specific evidence: environmental monitoring, challenge-test results, process deviations, temperature histories, supplier alerts and analytical uncertainty. These inputs must be interpreted by qualified personnel. An AI model that merges them without documented rules risks converting weak correlations into asserted causation.

The strongest architecture links three layers: the physical flow of material, the microbiological or chemical hazard mechanism, and the legal record required for market action. Failure in any layer can undermine the other two. Excellent microbiological testing cannot reconstruct unrecorded distribution; perfect distribution records cannot determine the contamination boundary without process evidence.

This is particularly relevant to novel preservation technologies. A treatment’s demonstrated efficacy under defined conditions is a separate question from whether treated lots can be identified later. Industrial validation should preserve application parameters, product identifiers, sampling times

and subsequent movement records. The evidence chain must survive both routine quality release and an adverse-event investigation.

9. AI governance: explainability is necessary but insufficient

Food companies increasingly discuss artificial intelligence as a solution to traceability fragmentation. The more useful question is what decision the algorithm is permitted to support. An anomaly detector that flags inconsistent inventory balances presents a different risk from a model that recommends excluding a customer from a recall. Governance should be proportional to the consequence of error.

For every algorithmic reconstruction, the operator should retain model version, training-data provenance, input snapshot, assumptions, confidence measures, validation outcomes and human approvals. Reproducibility requires more than a screenshot of a dashboard. A regulator or auditor may need to reconstruct why a particular set of lots was included or excluded on the date of a decision.

Generative language models can assist with document retrieval, multilingual supplier communication and drafting investigation narratives, but should not be treated as primary evidence of lot identity. Retrieval-augmented systems may reduce transcription burden while still producing incorrect associations when documents are ambiguous or identifiers overlap. Deterministic checks and human verification remain essential.

Blockchain deserves equally careful treatment. Tamper-resistant logging may improve confidence that recorded events have not been altered, but it does not prove that an event was true when entered. An immutable false record is still false. The bottleneck is often physical data capture, identity resolution and disciplined exception handling rather than the cryptographic properties of storage.

10. Economic and operational design

A traceability investment should be assessed against the total cost of maintaining reliable lot identity. That includes scanners, labels, system integration, training, exception handling, warehouse redesign, supplier onboarding and the time required to correct bad records. A low-cost application that relies on continuous manual reconciliation may be more expensive than a better-integrated system over its life.

An illustrative scenario clarifies the stakes. If a company experiences one major recall every ten years and better lot-level capture reduces the expected affected volume by 20,000 cases, the avoided direct product cost can be material even before accounting for retailer penalties and disruption. Yet these benefits depend on incident type and on whether the narrowed population can be defended. No universal return-on-investment multiplier is justified.

The key performance indicators should therefore include: proportion of movements with observed lot identity; unresolved genealogy edges; time to assemble a complete traceback; discrepancy rates between physical inventory and records; percentage of mass balanced through transformation; false exclusions in simulations; and the size of precautionary recall populations under uncertainty.

For ingredient suppliers, especially those serving several jurisdictions, the most useful near-term strategy may be a standard event-export schema rather than a proprietary consumer-facing

dashboard. The ability to transmit trustworthy lot data to processors, distributors and authorities is a commercial capability in its own right.

11. Research & Innovation Opportunity: uncertainty-aware recall reconstruction

The unresolved research problem is whether a traceability system can use incomplete logistics evidence to accelerate investigations while explicitly preserving the uncertainty that protects consumers. Existing work in event graphs, Bayesian inference, optimisation and supply-chain analytics provides components of a solution. What remains insufficiently established is how these methods perform under realistic food-industry missing-data patterns and legally consequential recall decisions.

A testable hypothesis is that a hybrid system combining deterministic event constraints, physical mass balance, calibrated probabilistic inference and explicit human decision gates will outperform both purely manual reconstruction and unqualified automated lot assignment. The expected advantage would be faster identification of plausible affected shipments without an unacceptable rise in false exclusions. This is a proposed research agenda, not an experimentally established result.

The validation programme should begin with instrumented pilot facilities in which ground-truth case and lot movements are captured independently. Investigators would then deliberately remove records to simulate mixed pallets, damaged labels, unrecorded transfers and production overlap. Competing algorithms should be assessed blind against the retained ground truth.

Primary endpoints should include traceback time, recall sensitivity, false-negative lot exclusion, precision, calibration error, audit completeness and staff workload. Stress tests should vary missingness mechanisms: random scanner failures, systematic omission during peak shifts, and adversarial or erroneous data entry. A model validated only on random missing values is unlikely to represent actual industrial failures.

Legal review must occur before field deployment. In the United States, experimental inference cannot be assumed to satisfy FSMA 204 records. In the EU, algorithmic reconstruction does not discharge Article 18 obligations unless the underlying operator identification and other applicable requirements are met. Any automated narrowing of a withdrawal or recall population must remain compatible with food-safety duties and authority oversight.

The industrial opportunity is broader than software licensing. A validated framework could support traceability audits, quality-system design, packaging and warehouse process optimisation, and simulation-based training. For industry, the relevant potential contribution lies in integrating process knowledge, food microbiology and evidence governance and not in asserting that an untested AI platform already exists.

12. A practical governance framework for management

The board-level decision is whether traceability is treated as an IT purchase or as a safety-critical operational capability. The latter requires ownership across quality, production, logistics, procurement, legal and digital systems. Each department holds a different part of the evidence; none can certify the whole genealogy independently.

Before approving a system, management should require a demonstrated traceback exercise using real historical transactions. The exercise should include a mixed pallet, a split lot, a transformed ingredient, a returned product and a deliberately missing scan. The provider should show the complete chain of records, identify unresolved edges and explain how the system prevents unsupported certainty.

Supplier contracts should define the identifiers and data fields that must accompany deliveries, the format of electronic exchange, change-control duties, retention periods and procedures for correcting erroneous records. Contracts cannot replace statutory duties, but they can reduce the operational gaps that make those duties difficult to discharge.

Finally, traceability performance should be tested against incidents rather than only against routine dispatch. A system that works on ordinary days but collapses when an outbreak investigation demands answers within hours is not operationally mature.

13. Conclusions

The November 2026 FDA consultation should be understood as a signal about the limits of present-day food data infrastructure, not as a retreat from the public-health rationale for traceability. The central problem is the preservation of physical identity through commercial operations that routinely split, combine and transform products.

The most defensible path is neither blind technological optimism nor rejection of probabilistic methods. It is a disciplined separation of observed facts, documented declarations, model-derived inferences and unresolved uncertainty. When uncertainty affects a safety decision, systems should be designed to reveal it, preserve it and support proportionate protective action.

For food biotechnology and ingredient companies, traceability is increasingly a scientific, legal and commercial competence. The objective is not merely to retrieve a lot number. It is to reconstruct, within a defensible time, which products could have been affected, why that conclusion follows from the evidence and which assumptions remain open to challenge.

Regulatory comparison at a glance

Jurisdiction	Binding framework	Practical interpretation
United States	21 CFR Part 1 Subpart S; FSMA 204	FTL-specific CTE/KDE records; no FDA enforcement before 20 July 2028 under statutory direction
European Union / Spain	Regulation (EC) 178/2002 Art. 18; sectoral rules	Broad traceability and supplier/customer identification; additional animal-origin information
Great Britain	Domestic food law as amended	Verify territorial consolidated rules and authority guidance
Northern Ireland	Relevant EU-linked arrangements	Separate territorial assessment required

References and primary sources

[1] FDA. Public Meeting: Challenges and Solutions in Lot-Level Food Traceability, 6 November 2026. https://www.fda.gov/food/workshops-meetings-webinars-food-and-dietary-supplements/fda-public-meeting-challenges-and-solutions-lot-level-food-traceability-11062026

[2] FDA. FSMA Final Rule on Requirements for Additional Traceability Records for Certain Foods. https://www.fda.gov/food/food-safety-modernization-act-fsma/fsma-final-rule-requirements-additional-traceability-records-certain-foods

[3] Federal Register. 91 FR 64656–64658, 9 October 2026, Docket FDA-2014-N-0053. https://www.govinfo.gov/content/pkg/FR-2026-10-09/html/2026-20720.htm

[4] 21 CFR Part 1, Subpart S: Additional Traceability Records for Certain Foods. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-1/subpart-S

[5] FDA. Discussion paper on flexibilities for lot-level tracking, 27 May 2026. https://www.fda.gov/food/hfp-constituent-updates/fda-releases-discussion-paper-ahead-traceability-public-meeting-june

[6] Regulation (EC) No 178/2002, especially Articles 18 and 19. https://eur-lex.europa.eu/eli/reg/2002/178/oj

[7] Commission Implementing Regulation (EU) No 931/2011. https://eur-lex.europa.eu/eli/reg_impl/2011/931/oj

[8] FDA. Public meeting, 15 June 2026, including transcript and traceability readiness materials. https://www.fda.gov/food/workshops-meetings-webinars-food-and-dietary-supplements/fda-public-meeting-challenges-and-solutions-lot-level-food-traceability-06152026

[9] FDA. Proposed compliance-date extension, regulatory impact analysis, 2025. https://www.fda.gov/about-fda/economic-impact-analyses-fda-regulations/requirements-additional-traceability-records-certain-foods-compliance-date-extension-ria-proposed

[10] GS1. EPCIS and CBV standards for event-based supply-chain visibility (technical reference; implementation is not itself a legal safe harbour). https://www.gs1.org/standards/epcis

[11] UK Food Standards Agency. Food traceability guidance and food law resources. https://www.food.gov.uk/business-guidance/food-traceability

[12] Regulation (EU) 2017/625 on official controls. https://eur-lex.europa.eu/eli/reg/2017/625/oj

Related INTABIOTECH Knowledge Areas: Food Safety · AI & Food · Regulatory Affairs · Processing Technologies · Biotechnology · Emerging Research

© 2026 José M. López. All rights reserved.