AI in Food Safety: Useful Tool, Not an Authority

Advanced Food-Tech Solutions · advfood.tech · May 2026

A food process engineer's view on what AI can — and cannot — do in thermal processing and food safety. And why a number like $F_0 = 3$ min is never the full answer, even when it is the right number.

George N. Stoforos, Ph.D. — Food Process Engineer
Advanced Food-Tech Solutions
gnstoforos@advfood.tech

Walk into any food safety conference in 2026 and the same question keeps coming back, phrased five different ways: *Can AI design my thermal process?*

The honest answer is no. The useful answer is that AI cannot design your process but can do a great deal of the work around it — provided it is treated as an engineering tool, not as an authority. That distinction is where most of the current AI conversation in food safety drifts off course.

The F_0 = 3 min problem

Ask any general-purpose AI chatbot: *"What F-value do I need for my low-acid canned product?"*

You will get a number — almost certainly $F_0 = 3$ min at $T_{ref} = 250$ °F (121.1 °C) and $z = 18$ °F (10 °C): the textbook 12D *Clostridium botulinum* cook used as the minimum benchmark for commercial sterility of low-acid canned foods. The answer is literature-based and generally correct. As a process specification, it is meaningless.

An F-value is not a property of "low-acid canned food." It is the integrated lethality of a specific time-temperature history at the cold spot of a specific product in a specific container, calculated against a specific target organism (proteolytic *C. botulinum* spores set the minimum, but spoilage organisms such as *Geobacillus stearothermophilus* often require more), with a specific z-value, validated against specific heat-penetration data, and audited under specific regulations — 21 CFR Part 113, PMO, EU hygiene rules, retailer standards. Change one variable and the appropriate F-value changes, sometimes by an order of magnitude.

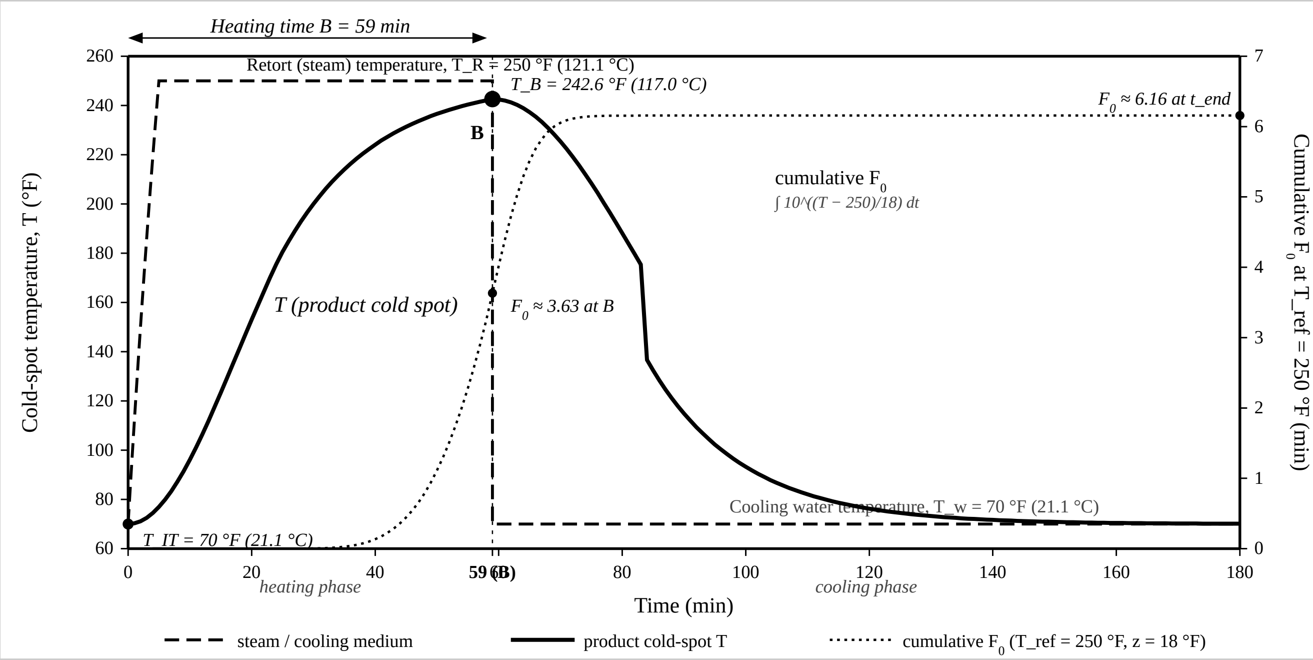

*Heat-penetration profile of a thermally processed product. Steam medium (dashed) reaches the retort set-point, holds to steam cut-off at point **B**, then drops to cooling water. The product cold-spot temperature T (solid) approaches the retort line asymptotically — reaching T_B at B — then decays through cooling. Cumulative F_0 (dotted, T_{ref} = 250 °F, z = 18 °F — the conventional 12D botulinum-cook reference) accumulates almost entirely between the late heating phase and the early cooling phase: F_0 ≈ 3.63 at B and continues to ≈ 6.16 by the end of the process. This is the engineering work AI does not replace: a validated thermal calculation anchored to measured heat-transfer data. AI sits around it — parsing recall feeds, classifying notices, summarizing regulations, flagging anomalies — but the validation itself stays measurement-based.*

The chatbot hands you that number not because it understands your process but because thousands of papers about botulinum cooks end with "F_0 = 3 min." Pattern-matching, not engineering.

What AI is genuinely good at

None of this is an argument against AI. We use AI inside our food safety platform every day. The point is to apply it where it earns its keep:

- **Modeling and simulation.** Statistical fits, kinetic parameters, CFD, and engineering code can be built and iterated far faster than by hand, with broader sensitivity analyses. A heat-transfer model that used to take a week of meshing can be drafted in hours.
- **Parsing unstructured text at scale.** USDA FSIS, FDA, CFIA, EFSA, RASFF, Rappel Conso, MHLW, FSANZ — each publishes recalls in its own format and language. AI turns that mess into clean structured data.
- **Classification under rules.** Given a clear written rule — *"Salmonella is always Tier 1"* — AI applies it consistently and at scale.

- **Spotting outliers.** A model that has seen 10,000 recall notices reliably flags the one that doesn't fit.
- **Summarizing.** Twenty pages of regulatory text into a one-page, audit-ready brief.

These are **language** tasks, not engineering tasks [1, 2].

Where AI is not the right tool

Process-engineering decisions are a different category of problem:

- *"Is this F-value sufficient?"* — Depends on product, validation data, target spore load (*C. botulinum*, *G. stearothermophilus*), z-value, container geometry, cold-spot location, filling temperature, and come-up time. Without those inputs, any answer is generic.
- *"What should the hold-tube length be?"* — Depends on flow rate (gpm), inner diameter (in or mm), density (lb/ft³), viscosity model (Newtonian, Power-law, Bingham, Casson, Herschel–Bulkley), Reynolds regime (`Re < 4000` per FDA's conservative laminar boundary in 21 CFR 113), residence-time distribution, target lethality, and worst-case particle validation. Without a characterized fluid, no model can size the tube.
- *"Is this recall a Class I?"* — Depends on agency classification, pathogen (*Listeria monocytogenes* vs. *Escherichia coli* O157:H7 vs. *Bacillus cereus* cereulide), exposure population, and evidence chain — none inferable from a press release.

AI is most assertive precisely where it should be most cautious [3].

How we actually use AI at AFTS

When we built the AFTS food safety recall-intelligence pipeline, the first design decision was that **no AI agent makes a final call.** AI *enriches* the scraped text — classifying, extracting fields, normalizing terminology, mapping multi-language sources to one schema — and the reviewer prompts are checklists applied verbatim:

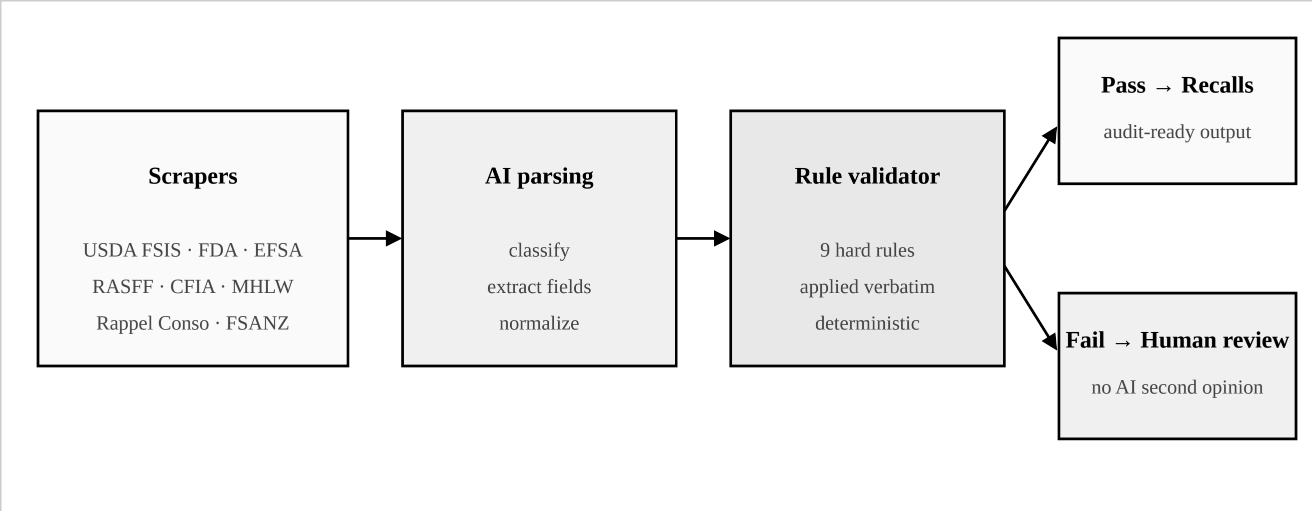

AFTS recall-intelligence pipeline. AI enriches scraped data; the rule validator accepts or rejects; everything is retained for audit and for retraining.

Hard rules — applied verbatim, no improvisation:

- *Salmonella* is Tier 1.
- Cereulide (*Bacillus cereus* emetic toxin) is always Tier 1; classification follows the toxin, not the source's terminology.
- Outbreak = 1 requires actual case counts, not the phrase "linked to."
- Retailer-as-Origine is valid for downstream recalls.
- "Sans marque" maps to "Unbranded."
- "Imposé par arrêté préfectoral" maps to "Mandatory."
- News-mirror sites are blocklisted — primary source only.

The validator accepts or rejects; failed records go to a human, never back to the AI. Every record — passed and rejected alike — is preserved with full provenance: source, AI-enriched fields, rule outcome, reviewer disposition. That archive is both the audit trail and the training set: the labeled history, especially the rejections, feeds prompt refinement and fine-tuning. The hard rules become the curriculum.

> That is how AI earns trust in a regulated environment: by being narrow, deterministic, and auditable.

The engineering principle

Food safety is a discipline of proof. We do not *believe* a product is safe — we have a heat-penetration study, a residence-time distribution, a challenge study, a scheduled process from a recognized process authority, a validated CCP, a documented HACCP plan. AI does not generate proof; it generates plausible text. Confusing the two is one way process-control errors enter the system.

Every thermal-process decision — aseptic, retort, hot-fill, HTST, low-moisture — is anchored to measured data: fluid rheology, flow rate (gpm or L/h), inner diameter (in or mm), hold-tube length (ft or m), density (lb/ft^3 or kg/m^3) at process temperature, viscosity model fit, Reynolds number, worst-case residence time (s), validated F-value (min) at the cold spot. Even the kinetic parameters (D, z) are not constants but distributions with measurable variability [4]. AI surrounds the engineering — surfacing regulation, normalizing data, drafting reports, flagging anomalies — but it does not replace measurement.

Ethics and responsibility

Any serious discussion of AI in food safety arrives at ethics. The relevant questions are concrete, not abstract.

Disclosure. When a recall classification, an F-value report, or a deviation memo has been touched by AI, the documentation should say so. Auditors, regulators, and the public have a legitimate interest in knowing what was assisted, what was generated, and what was independently verified by a named human.

Accountability. AI does not hold a process authority license, does not sign a HACCP plan, and cannot be held legally responsible for a misclassified recall. The named human or organization remains accountable for every output, regardless of how it was produced. Pretending otherwise dilutes the chain of responsibility on which food safety ultimately depends.

Misuse. The same generative capability that drafts a clean audit summary can fabricate a plausible-looking compliance record, a forged challenge study, or a counterfeit lab report. Industry has not yet developed the cryptographic provenance standards that would make such documents distinguishable from genuine ones. Until it does, every regulated artifact warrants additional scrutiny — not less.

Reality check. When AI handles most of the classification work without visible error, human reviewers begin trusting it by default. The very accuracy that makes AI useful is what erodes the vigilance needed to catch the rare misclassification — and food-safety failures are rare events.

Maintaining sharp review discipline, especially on outputs that "look right," is the cost of using AI in any regulated function.

Closing remarks — and where AI will matter most

The practical takeaway: when AI offers a number such as $F_0 = 3$ min, treat it as a starting reference, not a specification. Pull the heat-penetration data, check the z-value, consult a process authority, and run the calculation against your actual product, container, and line. The model is not wrong to recite the textbook value; it just cannot know whether that value is sufficient for your system.

A personal view on where this is heading. AI is genuinely strong at what it does best — processing data at scale, modeling, analyzing complex datasets, completing in seconds work that once took weeks. Three areas of the food sector stand to benefit most.

Food engineering

Modeling, CFD, kinetic-parameter estimation, and engineering code are activities AI accelerates dramatically without changing what counts as a correct answer. Simulations that once needed weeks of meshing are drafted in hours; sensitivity studies that were skipped become routine. The downstream effect is more efficient equipment, lower heat losses, fewer over-processed batches, and less nutritional and sensory damage to the food. AI will not replace the food engineer — it will give every food engineer the bandwidth of a large R&D department.

AI-accelerated design of a counter-flow heat exchanger for an aseptic juice line. The schematic at top shows the optimized geometry with vibrant thermal-field overlay (cold blue → hot red); counter-current arrows indicate the heating medium flowing right-to-left in the annulus while the product flows left-to-right in the inner tube. The lower-left panel shows the optimal temperature profile (bold) emerging from 12,000 Monte Carlo candidates (faint) that span uncertainty in rheology, fouling factor, and geometric tolerances. The lower-right tornado ranks how each design parameter shifts the lethality F-value at the outlet — hot-medium temperature dominates, followed by flow rate and overall U-coefficient. None of this replaces measurement: the AI loop is anchored to laboratory-measured U on a small-scale unit and validated against rheological data for the actual product. AI compresses what was once weeks of CFD meshing, sensitivity analysis, and Pareto optimization into minutes; the food engineer interprets, validates, and decides.

Food safety

Recall reporting, regulatory surveillance, and outbreak attribution all suffer from one problem: fragmented data across agencies, languages, and formats, with gaps that grow whenever a new ingredient, process, or pathogen appears. AI consolidates records, harmonizes terminology, links ingredients through complex supply chains, and flags hazard signals before they become outbreaks. Traceability stands to benefit especially: the technology to track an ingredient from farm to shelf has existed for years; making sense of the data at global scale has not.

Nutrition

And yet the greatest benefit will come in **nutrition**. The data is enormous and heterogeneous — dietary surveys, food composition databases, metabolomic and genomic profiles, biosensor streams, clinical outcomes — and translating that complexity into useful guidance has historically been one of the slowest problems in applied health science. AI changes that pace. I expect real progress in:

- **Personalized nutrition** — models combining genetics, metabolism, microbiome, lifestyle, and intake to produce truly individualized advice rather than population averages.
- **Refined dietary guidelines** — synthesis of fragmented, often contradictory evidence into coherent, data-driven recommendations.
- **Product reformulation** — optimization for nutritional outcomes (reduced sugar, sodium, saturated fat; improved fortification) without sacrificing safety, stability, or sensory quality.
- **Aligning food supply with health** — modeling production, distribution, and consumption so global food systems can be steered toward better health alongside sustainability.

Used responsibly — with the same insistence on validation, transparency, and human oversight that any regulated discipline requires — AI in nutrition could improve public health on a scale few other technologies can match. That, in my view, is where these systems can do the most good.

AFTS Food Safety Intelligence — and regulatory homes

- AFTS Food Safety Intelligence System (AFTS-FSIS) — our curated, rules-validated recall feed across USDA FSIS, FDA, EFSA, RASFF, CFIA, MHLW, Rappel Conso, and FSANZ.
- USDA FSIS — Food Safety and Inspection Service
- FDA Recalls, Market Withdrawals & Safety Alerts
- EU RASFF Window
- 21 CFR Part 113 — Thermally Processed Low-Acid Foods Packaged in Hermetically Sealed Containers

At AFTS, we build software, pipelines, and recall-intelligence systems that take AI seriously — by giving it the work it can do, and relying on validated engineering for the rest.

References

1. Garre, A., et al. (2025). AI's Intelligence for Improving Food Safety: Only as Strong as the Data that Feeds It. *Current Food Science and Technology Reports*. https://doi.org/10.1007/s43555-025-00060-0

2. Kanicheril Ambikalekshmi, A., & LakshmiMohan, M. (2026). Advancing Food Safety Through Artificial Intelligence: A Detailed Review of Monitoring and Detection Technologies. *Food Safety and Health*, 4(1), 82–90. https://doi.org/10.1002/fsh3.70050

3. Zhou, Q., et al. (2022). Artificial Intelligence, Big Data, and Blockchain in Food Safety. *International Journal of Food Engineering*, 18(1), 1–14. https://doi.org/10.1515/ijfe-2021-0299

4. Stoforos, G. N., Kontopanou, M., Giannakourou, M. C., & Stoforos, N. G. (2025). Distribution of Kinetic and Thermal Processing Parameters Based on α-Amylase Inactivation Data. *Journal of Food Science*. https://doi.org/10.1111/1750-3841.70357

About the author — George N. Stoforos, Ph.D., is a food process engineer and **co-founder** of Advanced Food-Tech Solutions. AFTS designs, validates, and troubleshoots thermal processes across aseptic, retort, HTST, hot-fill, and low-moisture lines, and develops audit-ready software including the AdvThermaLogic suite and the AFTS Food Safety Intelligence platform.