JFI MASTER LIBRARY • WEEK 2 CASE FILE

The Ghost Diagnosis & The Phantom Consult

Liability, supervision, and privilege when a clinical AI flags a condition no human ever confirmed

Format: Litigation-Grade Briefing (PDF) **Compliance Focus:** ABA Model Rules 5.1 & 5.3; standard-of-care doctrine

Core Tech: Autonomous diagnostic inference · closed-loop model logging

Hypothetical training file. This is a fictional case engineered for partner debate and predictive-liability simulation. All parties, facts, holdings, and any case or rule references are invented for instructional purposes and do not describe real persons, events, or binding precedent. Verify all authorities against primary sources before relying on them in practice.

THE FACT PATTERN

A diagnosis that arrived before anyone asked for it

A regional hospital network deploys a Level 4 autonomous diagnostic platform—marketed as a "clinical co-pilot"—that continuously ingests intake notes, lab feeds, and imaging across the emergency department. The system is configured to surface high-acuity findings automatically, without a clinician first ordering the analysis.

A 41-year-old patient presents with vague chest discomfort. Before any physician reviews the chart, the platform autonomously flags a rare cardiac anomaly and pushes a "critical" alert into the EHR, populating a provisional diagnosis field. The on-shift physician, mid-surge and trusting the alert, orders an aggressive intervention keyed to that finding. The anomaly was a false positive generated from a known sensor-artifact pattern. The patient suffers a serious complication from the unnecessary procedure.

Post-incident review reveals two complications. First, no human clinician ever independently confirmed the flagged condition; the chart shows a diagnosis with no attributable author. Second, the platform's closed-loop reasoning log—the only record of *why* the model surfaced the alert—is retained by the vendor under a proprietary-data clause and is not directly accessible to the hospital.

THE PARTIES

Plaintiff: The injured patient, alleging negligence and lack of informed consent for a procedure driven by an unverified machine-generated finding.

Defendant A: The hospital network, which deployed the platform in autonomous-alert mode and set the supervision protocol.

Defendant B: The platform vendor, asserting that the tool is a decision-support aide and that the proprietary reasoning log is trade secret.

Defendant C: The treating physician, asserting reasonable reliance on a "critical" system alert under emergency conditions.

ISSUES FOR DEBATE

Where does responsibility land when the machine acted first?

1. **The unsupervised aide problem.** If the platform produced a diagnosis with no human author in the loop, has the hospital effectively let an unsupervised non-clinician practice medicine? How do supervision duties translate from the law-firm analogy (ABA Model Rules 5.1 and 5.3 on supervising lawyers and nonlawyer assistants) to a clinical setting where a model, not a person, generated the work?
2. **Reasonable reliance vs. abdication.** Can the physician's reliance on a "critical" alert be reasonable when the standard of care presumes independent clinical judgment? At what point does deference to an autonomous system become an abdication of the professional's own duty?
3. **The phantom reasoning log.** The only record of the model's reasoning sits behind a vendor proprietary-data clause. Is that log discoverable? Does labeling it a trade secret shield it, or does the hospital's inability to produce it cut against the defense on spoliation and recordkeeping grounds?
4. **Product vs. service framing.** Is the vendor selling a product (inviting product-liability and design-defect theories) or a service (inviting professional-style negligence)? The "decision-support aide" label is doing heavy lifting—who benefits if a court accepts it?

DEFENSE POSTURES

Two competing strategies

HOSPITAL — CHANNEL TO THE VENDOR

Position the platform as a defective product operating outside its validated parameters. The autonomous-alert configuration was vendor-recommended; the sensor-artifact false positive is a known design issue. Argue the hospital's supervision protocol was reasonable for a tool sold as validated, shifting exposure upstream to Defendant B.

VENDOR — THE AIDE DEFENSE

Hold the "decision-support aide, not decision-maker" line: the tool only surfaces findings; a licensed physician remains responsible for verification and the treatment decision. Frame the harm as a supervision and standard-of-care failure by the humans, not a product defect—and resist disclosure of the reasoning log as trade secret.

THE IOLEBA THROUGHLINE

Why this connects to your sovereign-counsel practice

This file rhymes with a principle from the Guild's recent AI Legal Update: an AI system that touches privileged or high-stakes work is treated as a third party, not a confidant—and outputs you can't inspect are outputs you can't defend. The phantom reasoning log here is the clinical mirror of the public-AI privilege problem: when the record of *why* the machine acted lives outside your control, your ability to verify, disclose, and defend collapses.

The recurring lesson across the JFI series: closed-loop, vendor-controlled AI shifts risk onto the deploying professional precisely when transparency matters most. Sovereign, inspectable infrastructure is not just a compliance preference—it is an evidentiary one.

Discussion prompt for partner debate: If you represented the hospital, would you fight to compel the vendor's reasoning log even knowing it might reveal the false positive was foreseeable? Weigh the discovery value against the spoliation and recordkeeping exposure it could surface against your own client.