SpencerFane®

The New Distressed Asset: Navigating the Unique Complexities of AI Company Bankruptcies

As the initial euphoria surrounding generative artificial intelligence (AI) begins to rationalize, the market is approaching a critical inflection point. The "growth at all costs" era for AI startups, fueled by massive capital injections and staggering compute spends, is transitioning into a phase of rigorous fiscal scrutiny. For many high-valued companies with unsustainable burn rates, the path forward may ultimately lead to bankruptcy.

Liquidating or restructuring an AI company, however, is fundamentally different from the "SaaS 1.0" era. When a traditional software company fails, the value resides in its codebase and customer lists. In the AI domain, the most valuable assets are often intangible, illiquid, and legally radioactive: vast repositories of training data and the weighted parameters of fine-tuned Large Language Models (LLMs).

Therefore, as interest rates remained elevated and limited partners grew impatient, the inevitable question is no longer hypothetical: what happens when an AI company goes bankrupt?

The answer is far more complicated than most restructuring professionals, investors, and even general counsels currently appreciate. The bankruptcy of a traditional software company is a well-worn process. The estate holds a codebase, a customer list, perhaps some patents, and a handful of enterprise contracts that a buyer can assume. The playbook is established. The AI company bankruptcy is something else entirely – a collision of cutting-edge machine learning assets, labyrinthine data rights, privacy law obligations, and unresolved questions about intellectual property that neither Congress nor the courts have yet answered.

This article examines the four principal legal fault lines that will define AI company insolvency proceedings in the years ahead and offers strategic guidance for the executives, investors, and legal counsel who need to be thinking about these issues today, not after a default notice arrives.

The Hidden Goldmine: Unstructured Data as a Distressed Asset

When a traditional company fails, the inventory of valuable assets is largely predictable: intellectual property, physical equipment, accounts receivable, and enterprise contracts. When an AI company fails, the most valuable assets in the estate may be ones that never appear on the balance sheet and that a conventional restructuring advisor would overlook entirely.[1]

We are referring to what practitioners have begun to call "unstructured data" – the accumulated digital exhaust of a company's operations. Slack archives containing years of nuanced technical discussions among domain experts. Jira and GitHub repositories document how an engineering team reasoned through complex product decisions. Internal email threads, customer service interaction logs, proprietary annotation datasets, and human-feedback records generated through the company's own model evaluation processes are all in play. To a lay observer, this is corporate detritus. To an AI developer seeking high-quality, human-generated text with which to train the next generation of LLMs, it can be extraordinarily valuable.

> The most valuable assets in an AI company bankruptcy may be ones that never appear on the balance sheet – and that a conventional restructuring advisor would overlook entirely.

The emergence of specialized distressed-asset funds and strategic acquirers actively scouting for this type of training data has created a genuine buyer market that did not exist five years ago. Trustees, debtors-in-possession, and their advisors must understand this dynamic. Failing to surface and properly value a distressed AI company's data repositories risks leaving substantial value on the table – value that could meaningfully affect the recovery available to creditors.

The challenge, of course, is that recognizing value is only the first step. Realizing it is another matter altogether – one that implicates some of the most complex legal

questions in contemporary restructuring practice.

The Legal Minefield of Section 363 Data Sales

The sale of assets free and clear of liens and interests under Section 363 of the Bankruptcy Code is one of the most powerful tools available to a debtor or trustee. It is also, in the context of AI data assets, one of the most legally treacherous.

The Privacy Law Collision

Selling a company's accumulated data repositories to a buyer who intends to ingest that data into a large language model immediately creates friction with some of the most demanding privacy regimes in the world. The California Consumer Privacy Act (CCPA) and its successor, the California Privacy Rights Act, impose strict limitations on the sale of personal information and require, among other things, that the purpose for which data is used be consistent with the purposes disclosed at the time of collection. The European General Data Protection Regulation (GDPR) is even more demanding, requiring a lawful basis for processing personal data and imposing restrictions on secondary use that can be difficult to satisfy in an insolvency context.

The problem is structural. When a company collected user data, it made specific representations in its privacy policy about how that data would be used. Training a competitor's language model was almost certainly not among those disclosed purposes. Preemptively addressing this potential future in the company's terms and conditions is something companies should seriously consider.

But not all contingencies can be agreed between the individual and the company. For example, a Section 363 sale may not override GDPR obligations or CCPA requirements and the buyer who acquires and uses data in a bankruptcy sale may not be insulated from enforcement action if the underlying collection or transfer violated applicable law.

The Consumer Privacy Ombudsman

The Bankruptcy Code contains a partial answer to this tension, though an incomplete one. Under 11 U.S.C. § 363(b)(1), when a debtor seeks to sell personally

identifiable information in a manner inconsistent with its privacy policy, the court may appoint a consumer privacy ombudsman – an independent officer whose role is to evaluate and report on the privacy implications of the proposed sale. The ombudsman does not have veto power over the sale, but the court must consider the ombudsman's report before approving any transaction that implicates consumer privacy rights.

In practice, this mechanism was designed for far simpler scenarios: a retail company selling its customer email list to a competitor, for example. Applying it to the sale of complex, interwoven AI training datasets raises novel questions that bankruptcy courts will be ill-equipped to resolve quickly. This will almost certainly lead to a select group of expert ombudsmen able to address these nuances. How does the ombudsman assess the privacy risk of data that will be used not to directly contact consumers, but to shape the probabilistic outputs of a language model? What anonymization or de-identification standards are sufficient? These questions do not have settled answers, and the litigation risk for buyers who proceed without carefully negotiated indemnification protections is substantial.

The Privilege Problem

Perhaps the most underappreciated risk in a bulk AI data sale is the inadvertent transfer of attorney-client privileged communications. A company's Slack archive, email repository, or internal document management system will almost inevitably contain communications that are covered by the attorney-client privilege, work-product doctrine, or both. It may also contain third-party trade secrets – information shared by vendors, customers, or partners under nondisclosure agreements that impose obligations on the company.

Dumping an unreviewed corporate archive into a distressed asset sale – as the time pressure of a bankruptcy proceeding might otherwise encourage – creates severe risk. The privilege holder (or former holder) may have grounds to seek to claw back the privileged materials, disrupt the sale, or pursue sanctions. Third parties whose trade secrets were included may have claims against both the seller and the buyer. Any buyer of AI training data from a distressed estate should insist on robust representations and warranties, a credible privilege review process, and indemnification provisions that survive the closing.

The "Poisoned Model" Dilemma: Selling a Trained LLM

The data asset sale is, in many respects, the simpler case. The more legally novel, and potentially more economically consequential, situation arises when the distressed estate includes a trained large language model itself. A heavily fine-tuned, domain-specific LLM can represent tens of millions of dollars in computer and human capital investment. It may be the most valuable single asset in the estate. It is also a legal liability of uncertain dimensions.

License Revocation and the Severable Model Question

Many AI models are trained on data that the company licensed from third-party sources – news archives, academic databases, proprietary datasets, or content licensed from creators and publishers. Those licensing agreements typically contain terms that govern what happens upon the licensee's insolvency. Many include change-of-control or assignment restrictions. Some grant the licensor the right to terminate upon a bankruptcy filing.[2]

If the licenses to the underlying training data are revoked upon insolvency, a deeply consequential question arises: does the trained model itself become legally unusable? The model weights encode, in mathematical form, patterns derived from that licensed data. Is the model a derivative work that rises or falls with the underlying license? Or are the weights sufficiently transformed that they constitute an independent asset unencumbered by the original data licenses?

Courts have not answered this question. The existing copyright doctrine on machine learning models is thin and developing rapidly. Counsel advising buyers of distressed AI model assets must be prepared to operate in a zone of genuine legal uncertainty – and to structure transactions accordingly, with appropriate representations, price adjustments, and termination rights if the legal landscape shifts unfavorably after closing.

Inherited Liability for Copyright Infringement

The copyright infringement question presents a distinct and, in some ways, more acute risk. Many AI companies trained their models on data scraped from the internet without licenses or permissions from the copyright holders. A substantial

body of litigation is already testing whether that practice constitutes copyright infringement – and early decisions have been more favorable to plaintiffs than the AI industry initially anticipated.

The distressed asset sale context creates a particularly dangerous scenario for buyers. If the debtor company scraped copyrighted content without authorization, and the buyer acquires and continues to use the model trained on that content, does the buyer unwittingly open itself to infringement liability once it uses the purchased copyrighted content? A Section 363 sale order can transfer assets free and clear of many claims – but the "free and clear" protection applies to successor liability for the debtor's past acts, not necessarily to the buyer's own independent act of using the model going forward.

> A Section 363 sale order can transfer assets free and clear of many claims – but it cannot insulate a buyer from infringement liability arising from its own continued use of a model trained on unlicensed data.

Put differently: the moment the buyer uses the acquired model in a commercial product, it is not relying on a historical act of the debtor; it is independently deploying a system that may itself generate outputs derived from infringed works. The scope of Section 363's protective effect in this context is genuinely unsettled, and buyers should not assume that bankruptcy court approval of the sale provides a litigation shield against copyright claims brought by third-party rights holders.

Due diligence on a distressed LLM acquisition should therefore include a detailed reconstruction of the model's training data provenance, an assessment of the debtor's data sourcing practices, and a realistic evaluation of the litigation risk associated with any pending or threatened copyright infringement claims against the estate. Representations and warranties insurance in this context, while expensive, may be worth the premium.

Strategic Takeaways

The legal questions addressed in this article are not theoretical. The volume of AI company distress will increase as the market matures, capital becomes more selective, and the distance between valuation and commercial reality becomes

impossible to paper over. The counsel, investors, and executives who are thinking carefully about these issues now will be far better positioned than those who encounter them for the first time during an insolvency proceeding.

For Portfolio Companies and Their General Counsels

The single most important proactive measure is to structure data licensing agreements with insolvency in mind. License agreements governing access to training data should specify what happens upon an insolvency event, whether the license survives and on what terms, and how the model weights will be treated in the event of a wind-down. The company's privacy policies should be drafted with sufficient flexibility to accommodate the possibility of a sale in a distressed context, which is not the same as drafting them to permit every possible secondary use but does mean avoiding language that forecloses a legitimate sale process. Internal data governance protocols that track what data the company holds, from whom it was obtained, and on what legal basis, are an operational necessity, not a compliance luxury.

For Venture Capital and Private Equity Investors

Investors in AI companies should treat data provenance and licensing structure as first-order diligence items, not afterthoughts. A portfolio company whose value is largely embedded in a trained model built on unlicensed or legally questionable data is carrying hidden liability that may materially impair recovery in a distressed scenario. Board-level oversight of data sourcing practices and licensing architecture should be standard for any portfolio company whose core asset is an AI model. When a portfolio company does face distress, early engagement with restructuring counsel who understand both the technical nature of AI assets and the applicable legal framework will be essential to maximizing the value available to investors.

For Distressed Asset Buyers

The opportunity in distressed AI assets is real; so is the risk. Buyers who understand the layered legal complexity – and who conduct genuinely rigorous diligence rather than relying on the protection of a sale order they assume to be broader than it is – will be the ones who can transact efficiently and confidently. Those who do not will

discover that the Bankruptcy Code, for all its power, was not written with large language models in mind.

Key Diligence Considerations for Data Asset Acquisitions from Distressed AI Estates

- Privacy compliance audit: Map all personal data categories against the debtor's historical privacy policy disclosures and applicable jurisdictional requirements (CCPA, GDPR, HIPAA, where applicable) before signing.
- Privilege review protocol: Require a structured privilege review process as a closing condition; negotiate escrow or holdback arrangements tied to the completion of that review.
- NDA and trade secret sweep: Identify all third-party confidentiality obligations in the debtor's data repositories; ensure the sale order includes appropriate limitation-of-liability provisions.
- Consumer privacy ombudsman strategy: Engage proactively with the ombudsman process rather than treating it as an obstacle; a cooperative approach can substantially reduce the risk of a court-imposed sale modification.
- Cross-border transfer mechanisms: If the data includes EU or UK personal data, ensure the transaction includes appropriate transfer mechanisms (standard contractual clauses, adequacy decisions) that survive the sale.

This blog was drafted by Gregory Ewing and McKay Holley, attorneys in the Washington, D.C. and Las Vegas offices, respectively, of Spencer Fane. For more information, visit spencerfane.com.

——————————

1 A close second in value is a well-known asset – unexpired leases and executory contracts for server storage and compute. These assets are out of scope of this article but certainly will need to be considered by the debtor.

2 11 U.S.C. 363(l) limits the enforceability of insolvency contingency provisions, but 11 U.S.C. 365(n), which governs rejection or assumption of licenses in intellectual property, is often subject to litigation. In the new realm of AI bankruptcies, these clauses will likely be the subject of significant interpretation and litigation.

*Click **here** to subscribe to Spencer Fane communications to ensure you receive timely updates like this directly in your inbox.*