

Insights

AI Policies and the Pirate's Code: When Rules Become a Liability

Article
Lowndes
10.20.2026

By Alex Dobrev, Jackie Bozzuto, and Jyllian Bradshaw

Captain Barbossa was not exactly a model of corporate governance. But in *Pirates of the Caribbean*, he offered a memorable interpretation of the Pirate's Code: it was "more what you'd call guidelines than actual rules."

Some companies may be approaching artificial intelligence the same way. Worse, they may have it backward: employees treat important restrictions as optional guidelines, while management drafts aspirational commitments as ironclad rules.

Both approaches can lead to trouble. And unlike Captain Barbossa, corporate executives cannot simply sail away when things go wrong.

The Ship Has Already Sailed

While management debates whether the company needs an AI policy, employees may already be using AI to summarize confidential contracts, draft customer communications, analyze financial information, or screen job applicants.

Some uses are low risk. Others may expose confidential or personal information, jeopardize trade secret protection, produce inaccurate results, or introduce legal risks into consequential business decisions such as hiring, promotion, discipline, and termination.

This is the reality of *shadow AI*: employees adopting AI tools without the organization's knowledge or approval.

The obvious response is to establish some rules. But that introduces a less obvious risk: a company can create problems not only through what its employees do with AI, but also through what its policies promise about AI.

Related Attorneys

[Alexander Dobrev](#)
[Jacqueline Bozzuto](#)
[Jyllian R. Bradshaw](#)

Related Expertise

[Business Litigation](#)
[Corporate](#)
[Labor & Employment Law](#)

When the Policy Becomes Exhibit A

Consider a company that proudly declares:

- “All AI-generated content is independently verified.”
- “Every AI system is continuously monitored for accuracy and bias.”
- “Customer information is never retained or used by AI providers.”

Impressive commitments. But are they true?

Who verifies every output? What does continuous monitoring actually involve? Has anyone examined the AI vendors' contractual terms and data practices?

And does continuous monitoring include weekends?

A policy requiring comprehensive oversight may look reassuring in the boardroom. It can look considerably less reassuring when an opposing attorney asks for evidence that the promised oversight actually occurred.

To be clear, violating an internal policy does not automatically establish legal liability. In *Discount Tire Co. v. Bradford*, a Florida appellate court reaffirmed that internal policies do not independently create a negligence duty. But those policies may still be offered as evidence of what the company knew or considered reasonable, while contractual promises and public representations can create separate exposure. In this regard, employee handbooks deserve particular attention.

The Federal Trade Commission's enforcement action against DoNotPay illustrates another danger: making claims about AI capabilities that the business cannot substantiate.

The lesson is not to avoid AI policies. It is to avoid writing checks that the company's compliance practices cannot cash.

Some Rules Should Actually Be Rules

Employees should not upload confidential, privileged, trade secret, or personal information into unapproved AI platforms. Material AI-generated assertions should be verified before employee rely on it. AI applications affecting employment, housing, lending, insurance, or other consequential decisions warrant appropriate approval and safeguards.

Those are meaningful guardrails.

But not every AI-assisted email requires formal review. Not every productivity application needs an independent bias audit. And not every business needs an AI governance committee with its own organizational chart.

A practical two-page policy that employees understand and follow may provide better protection than a 40-page manual promising audits, certifications, and monitoring that nobody performs.

After all, a written requirement to conduct an audit does not actually conduct one.

Before Writing a New Policy, Check the Old Ones

There is another potential trap. The company may already have an AI policy without realizing it.

Existing confidentiality agreements, privacy notices, cybersecurity policies, employee handbooks, collective bargaining agreements, customer contracts, and vendor questionnaires may contain promises that AI use now implicates.

Some were written before today's AI tools existed. Others may promise safeguards that the company has never evaluated in the context of AI.

Adding another policy without reviewing those commitments may simply add another layer of obligations, inconsistencies, and potential evidence of noncompliance.

The better approach begins with three questions:

1. Do we know how our employees are actually using AI?
2. Which activities require firm rules, and which call for reasonable guidelines?
3. Can we demonstrate that we are doing what our existing policies already promise?

Captain Barbossa could afford a flexible interpretation of his code. Businesses have considerably less room for improvisation when contracts, regulators, and legal obligations are involved.

The goal is not to have the most impressive AI policy. It is to have one that actually works.

Because when trouble arrives, **"we had a policy" is a much weaker defense than "we followed it."**

For questions about developing or updating your company's AI policies, contact Alex Dobrev, Jackie Bozzuto, or Jyllian Bradshaw.