INSIGHTS | PERSPECTIVE

Who Pays to Grade AI?

The Independence Problem

By Dr. Ben E. Magee, Founder | October 4, 2026

The credibility of AI evaluation comes down to one design choice, and it is not a technical one. It is this: who pays for the evaluation? Get that answer wrong and no amount of rigor will make the result trustworthy. The markets already learned this lesson once, at great cost. AI is about to learn it again, unless it is built differently from the start.

A lesson the markets already paid for

Before 2008, the major credit rating agencies were paid by the companies whose debt they rated. The issuer chose the rater, the issuer paid the fee, and the issuer could take its business elsewhere if it did not like the grade. The incentives were obvious, and so was the outcome. Complex securities that later proved worthless carried the highest ratings, and the institution whose entire value was independent judgment had quietly sold that judgment to the people it was judging.

The problem was not a lack of talent or method. It was structure. When the party being graded pays for the grade, the grade bends, slowly and then all at once.

The same conflict is arriving in AI

Today, the dominant signal of whether an AI system is safe or capable comes from the company that built it. The builder selects the benchmark, runs the test, reports the result, and decides what to disclose. When independent evaluation does happen, there is growing pressure for the model provider to commission and pay for it.

That is the issuer-pays model, rebuilt for artificial intelligence. It will fail the same way, because the incentive is the same. A provider paying for an evaluation of its own system is not buying the truth. It is buying a result it can use.

The party that bears the risk should pay

Reform Intelligence® is built on the opposite principle. We are paid by the institution that bears the risk, to evaluate that institution's own use of AI. An enterprise deploying a third-party AI system into a regulated process pays us to assess its deployment: the system, the data, the controls, the decisions it touches. We do not take a fee from the model or application developer whose product is under review.

The difference is not cosmetic. The enterprise has every reason to want an honest finding, because the exposure is its own. If the system is not ready, the enterprise needs to know before its customers or its regulators find out. The payer and the party who wants the truth are the same party. That is what makes the evaluation worth anything.

Independence by design, not by promise

Anyone can claim to be independent. Structure is what makes it real. Ours rests on a few commitments that can be checked:

We take no revenue from the model providers whose systems we evaluate, and we never will. Our fee is fixed and does not depend on the outcome, so there is no financial reward for a favorable result. Our methodology is published, so a determination can be examined rather than taken on faith. And we separate evaluating a system from fixing it, so we never profit from finding the problems we are then hired to solve.

Where the conflict can sneak back in

There is one place the old conflict can return, and it is worth naming plainly. If AI developers could pay to earn a Reform mark of approval to help sell their products into enterprises, that would be issuer-pays through the back door, no matter how it was dressed up.

So we hold a hard line. Any certification attaches to an enterprise's deployment, which the enterprise pays to have verified, not to a vendor's product. We would rather give up a revenue line than give up the independence that is the entire reason the mark would mean anything.

Why this is the whole game

An evaluator has exactly one asset: the belief that its judgment is not for sale. Everything else, the methods, the technology, the team, exists to serve that. Lose it and there is nothing left to sell. Keep it and you become the reference the market relies on, the way auditors and rating agencies did for capital markets before the ones who forgot the lesson gave the role a bad name.

Independence is not a feature of trustworthy AI evaluation. It is the precondition. The payer should always be the party that wants the truth. We built Reform Intelligence so that it always is.

Reform Intelligence AI, Inc. is an independent evaluation and governance company, organized as a Delaware public benefit corporation.
reformintelligence.com

© 2026 Reform Intelligence AI, Inc. All rights reserved. Reform Intelligence® is a trademark of Reform Intelligence AI, Inc. This material is provided for general informational purposes only and does not constitute professional, legal, or investment advice, or an offer of services. No part may be reproduced or distributed for commercial purposes without permission.