Why Now: The Only ML Curves You Own

Every technology faces its share of skepticism, and with Machine Learning (ML), whose erroneous predictions can cause significant harm, the skepticism is earned. That skepticism is why deferring has probably felt like a prudent call. Unfortunately, deferring isn't free; it delays the compounding of two assets that belong exclusively to your business: the organization's trust in predictions and the data that makes them better. Meanwhile, teams are already experimenting on their own, accumulating regulatory exposure and redundant effort.

Data and Trust Compound

That same caution explains why most businesses begin their ML journey on auxiliary functions, or confine it to a smaller region. Furthermore, operational databases built to support business and reporting rarely include the features and variation needed to learn relationships with outcomes. Therefore, the financial return of the first project often appears modest. However, its return is exceptionally large because it paves the way for organization-wide ML success.

The first success turns skepticism into confidence. Trust cannot be declared in a strategy document or bought with a platform; it is earned each time a prediction is acted on and holds. Trust is what determines whether a prediction gets acted on or circumvented. Trust is also what makes the second asset reachable, because the investments data curation demands, from updating data ingestion to deliberately spending margin to generate evidence, are authorized on the strength of predictions that have proved valuable.

That confidence transforms data curation on two fronts: collection and generation. A model that underperforms because of missing variables tells you exactly what to start recording. Those features get captured from then on, and there is no backfilling a measurement nobody took. *Every quarter of delay is a quarter of data that will never exist.*

The second is harder and more valuable, and it comes later, once trust has been earned. Your operational history records what you did, and not what would have happened otherwise. If pricing has remained steady for years or moved in lockstep, no volume of that history will teach a model what a price change does. The variation isn't in the data because it was never in the business. The same holds for routing, staffing, marketing campaign timing, credit thresholds, and every lever managed for consistency. *Consistency is usually good for operations, but it always yields poor evidence.* Generating evidence means deliberately varying a lever across comparable units and accepting a known short-term loss to purchase information you cannot otherwise obtain.

This is the step most organizations never reach, because it requires the trust that starts with the first success. It is also the step that separates a company with an advantage from just another company with models.

Enriched data raises model accuracy and makes a wider range of models feasible. The confidence drives wider adoption and targeting of high-value problems. The combination of accuracy and high-value problems yields higher returns, strengthening confidence and creating a virtuous cycle. Trust to act on predictions and the data required to generate them begin to compound from the day you start.

The Decision Is Being Made Without You

Teams experiencing acute pain and individual enthusiasts begin exploring while the organization deliberates. The result is shadow AI with the same dynamic as shadow IT, but a materially worse risk profile. These uncoordinated, siloed initiatives create redundant effort across teams and generate regulatory exposure. The absence of a governance structure produces models:

- running without auditable logs,
- built without essential documentation,
- trained on user data beyond consent terms,
- built by feeding confidential data to unaudited SaaS tools.

The regulatory exposure is accumulating now, unrecorded, against a decision that has not been made. Moreover, these efforts often stall because of a lack of adequate tools, wasting invested effort and eroding trust in ML, which creates roadblocks for future initiatives.

The remedy is not prohibition, which does not work and has rarely worked for any prior technology. It is a sanctioned first project, the cheapest way to establish the review, documentation, and data-handling standards that uncoordinated efforts currently lack.

The Barrier Has Already Fallen

ML initiatives no longer require a large commitment, such as hiring a team of data scientists, MLOps engineers, and software engineers, or paying large consultancy fees. An analyst, data engineer, or data scientist can build and productize a robust, compliant model in weeks without expensive GPU clusters. The cost of finding out has collapsed, whereas the cost of not knowing compounds.

The natural objection is that the technology improves every year, so delay is cheap and probably prudent. It would be, if the technology were the scarce asset, which it isn't. Everyone gets next year's models, on roughly the same terms, at roughly the same time. The technology curve is shared; the data and trust curves are yours alone, and they don’t start until you do.

The First Project Sets the Trajectory

The virtuous cycle runs in reverse just as readily. A first project that stalls hardens skepticism into policy. The data-curation budget quietly goes unrenewed. The next three proposals arrive pre-discredited, and the organization loses not one project but years of compounding.

Which is why the choice of first problem carries far more weight than its direct return suggests. You are not selecting a use case. You are selecting the evidence on which every subsequent decision will rest. The first problem is chosen against a different standard than a business case:

- **Would a win here be noticed?** Not the largest return, but the most legible one.
- **Can it be answered with data already in hand?** Or does it need a measurement nobody has been taking?
- **Is there someone positioned to act on the prediction inside their existing workflow?** A prediction nobody acts on changes nothing.
- **Does it sit comfortably within what governs your sector's use of AI?** Or is it subject to multiple regulatory requirements?

The first project sets the precedent, whether you intend it to or not.

Start With What You Already Hold

So the first question is not which vendor or which platform. It is narrower and more answerable: does the data you already hold support a problem worth solving?

That answer is sitting in your system today, and establishing it takes days, not months. Waiting will not improve it; it only delays the first cycle of compounding while the shadow efforts keep growing and the only curves you own stay flat.