Features

The daisy chain – Assumption, interconnection and propagation

How AI's hidden web of commitments and financing could amplify risk. **BY AARON P ROFFWARG, PARTNER, BRACEWELL LLP.**

Consider a well-leased office tower. On its upper floors sits a law firm; below it, an accounting practice; and on the lower floors, a management consultancy. To the owner and the bank that financed it, this is a comfortable arrangement, because the tenants do not rise and fall together. There is enough diversification to survive differing economic cycles.

Now consider the infrastructure being built for AI – data centres, chips, power lines, financing and the contracts binding them together. At the level of any single transaction, it looks diversified. A chip sale is not a cloud lease; a cloud lease is not a construction loan.

Appearances can be misleading. The tower is diversified because its tenants depend on different things. The AI buildout may appear diversified, because its participants increasingly depend on the same thing: a shared expectation about the future economics of AI adoption. What appears diversified at the transaction level may be concentrated at the assumption level. This is assumption concentration: a portfolio that is diversified by asset, tenant and counterparty, but not by the premise on which each position depends.

This is not the bubble argument. The concern is not that the shared expectation is wrong, but how the system would respond if that expectation changed.

Following one AI dollar

To see how a shared assumption embeds itself, follow a single dollar of expected AI demand as it moves through the system.

It begins as a forecast inside a model developer. On the strength of that projection, the developer commits to a cloud provider for an enormous quantity of capacity, often years before the revenue to pay for it exists. The cloud provider contracts in turn with a data centre developer to build the physical capacity. The data centre developer borrows, and the lender takes comfort from the chain above it: a signed cloud contract, behind which stands a credible model developer. The loan funds construction, which requires power, and so a utility commits capital to generation and transmission with operating lives measured in decades. The chips that fill the building come from a semiconductor manufacturer, whose own production and investment plans are scaled to the same expected demand. And investors supply equity, underwriting valuations that presume the forecast will be borne out.

Participants may believe they are underwriting different risks while depending on the same economic outcome. The dollar of demand has not been diversified by its passage through the system. It has been counted again at every link. Call it the daisy chain.

A careful reader may object that some of these assumptions are binding obligations. Contractual demand reduces propagation risk; it does not eliminate it. A contract is only as reliable as the counterparty behind it and the end-use economics that make honouring it rational, and those economics rest on the same adoption premise with which the chain began. The daisy chain does not disappear when each link is properly documented. It runs one layer deeper.

Those contractual protections are the ordinary, sound machinery of project finance, refined over decades to allocate risk to the party best able to bear it. What no single instrument is built to reveal is that the risks it has so carefully separated may, this time, be the same risk wearing different paper. The structures are working exactly as designed. The question is whether what they were designed to assume still holds.

From forecasts to infrastructure

The stakes depend on what the assumption has become by the time it is tested. Forecast becomes commitment; commitment becomes financing; financing becomes construction; construction becomes infrastructure, and with each step the original expectation grows less reversible. A forecast can be revised in an afternoon. A transmission upgrade cannot. A half-built data centre cannot be talked back into capital.

The system therefore embeds a revisable assumption into irreversible assets, at enormous scale and considerable speed.

The chain becomes a web

Irreversibility is only half the problem; the other half is how the participants are connected. A chain transmits stress in one direction; a web transmits it in several. The system's connections, which in good times allow capital and capacity to flow with remarkable efficiency, are the same connections along which a revision would travel.

Greater interconnection is not, in itself, greater risk; it is greater interdependence

A web would be no cause for concern if the bets it linked were genuinely independent. That is not the case here. The strands are not independent: followed back, nearly every one arrives at the same place, a common reliance on a future demand that no participant can verify.

Multiple hats

The web is tightest where roles overlap. Increasingly, one firm is at the same moment an investor in a second, a customer of it, a supplier to it, and its strategic partner besides. A semiconductor manufacturer invests in the model developer that is also its customer.

Value moves between them through several channels at once and not all of those channels are cash. A strategic investment stake, for example, carries a mark that depends on the very forecast the commercial contract assumes.

Features

True exposure is no longer captured by the cash that passes between parties. What grows here is not risk but opacity, and opacity has a way of concealing concentration.

To see how opacity conceals concentration, trace one hypothetical linkage through the documents that would govern it. Suppose a chip supplier holds an equity stake in a model developer to which it also sells; the developer has signed a take-or-pay commitment to a cloud provider; and the cloud provider has borrowed against that commitment through a facility whose covenants include a cross-default. Now let the developer's demand forecast soften. The supplier's equity mark falls and its order book thins. The take-or-pay may strain the credit behind it; if that credit trips a covenant, the cross-default reaches lenders who were never parties to the original forecast. No single agreement in that chain discloses the others.

Adoption is not economics

Each of these linkages rests on the same premise, and that premise deserves direct examination. The shared assumption, stated more precisely, is that if AI adoption is vast, the economics must follow. History is not kind to it.

The railroads remade commerce and bankrupted many of the investors who built them, because capacity outran the economics of operating it. The fibre boom produced a glut that ruined firms whose technology had been entirely vindicated. The internet's first commercial wave erased a generation of valuations built on the correct belief that the technology would prevail.

AI adoption may be extraordinary and the economics may still disappoint, because value may flow to users, to a few winners, or to later entrants, rather than to those who financed the first great buildout.

The infrastructure must be paid for from revenue the use of AI is expected to generate. If today's pricing reflects competition for market share rather than present cost, that revenue is itself a forecast: the same forecast embedded in the contracts, the loans, and the power commitments.

Defining propagation

If a shared premise is the source of the exposure, the mechanism that spreads it is propagation: the process by which a change in assumptions affects participants beyond those involved in the original transaction. A revision by a model developer would travel to the cloud provider, the lender, the utility, the chip supplier, and the investors, none of whom was a counterparty to the original revision.

A diversified system propagates little, because its parts are not joined; a concentrated, interconnected one propagates a great deal. The concern is not prediction. How far a change would travel if the assumption did not hold is a matter of structure, and structure is something decision-makers can examine now.

The financial crisis offers a precedent for protection that depends on the very conditions it insures against. In 2008, Goldman Sachs held credit exposures it had partially hedged through AIG, and sought further protection against an AIG default from other market participants. Had AIG failed in the systemic crisis, that protection might have proved inadequate, because those counterparties faced the same conditions. The hedge was purchased from the system it was meant to protect against.

Why this matters

The practical implication is a change in the unit of analysis. Evaluating each project, contract, and loan on its own merits remains necessary but is no longer sufficient, because a portfolio of individually sound exposures may not be a sound portfolio if the exposures rest on a common assumption. The lender's well-collateralised loan, the utility's prudent commitment, and the investor's diversified book of AI positions may prove to be the same bet held many times over.

Questions to ask

Assumption concentration becomes practical when turned into questions a board, investment committee, or credit committee can put to its own book:

1 – Are we diversified by counterparty but concentrated by assumption?

2 – Which forecasts support multiple investments across our organisation at once? A demand projection that appears in three models is one bet, not three.

3 – How many of our partners depend on the same AI-adoption assumptions we do?

4 – What happens if adoption is slower than expected, even if the technology ultimately succeeds? Timing, not magnitude, is what strands long-lived assets.

5 – Where could changing expectations propagate through our own business model?

The load-bearing assumption

Financial systems, like structures, have load-bearing assumptions: premises whose revision is felt not locally but throughout. The question this article has explored is whether the future economics of AI adoption has quietly become one of them: whether a growing share of the infrastructure now being built, across many firms and many contracts that appear pleasingly unrelated, in fact rests upon that single premise, and would feel its revision in concert.

Nothing in this argument requires the assumption that AI adoption will disappoint. The premise may hold; the buildout may be vindicated; the forecasts may prove, in time, to have been conservative. The point is not to predict the load, but to know what is load-bearing – for that is the knowledge that lets a structure be built, and held, and judged with open eyes. An ecosystem can be at once extraordinary and concentrated, diversified in its contracts and singular in its dependence. To see that is not pessimism. It is the ordinary prudence one owes to anything one means to keep standing.

AI may become one of the defining technologies of the century. The infrastructure now being built may prove entirely justified. The forecasts may ultimately prove conservative.

The question explored here is not whether AI succeeds. It is whether the economic architecture forming around AI proves as resilient as the technology itself. ■