
# AI, In Our Own Terms

**Xin Wu, CIO | August 21, 2026**

When I first wrote about AI in September 2025, we were still exploring what AI could do for us. By the time of our second piece in April 2026, both the technology and our use of it had already advanced substantially. Since then, the pace has only accelerated. Our toolkit has expanded, the range of applications across the firm has widened, and AI has moved from a collection of individual experiments toward something more durable: a firm-specific AI ecosystem that brings together powerful models, repeatable workflows, institutional knowledge, and human judgment. Throughout that evolution, one principle has remained constant: we use AI on our own terms, with our philosophy at the center of how different models and tools are selected, shaped, and applied.

Hence, this progress report will be our last dedicated piece on AI, not because the work is finished, but because the subject is evolving too quickly for periodic updates to remain particularly useful. The tools we use today are already different from those we were using six months ago, as is our understanding of how best to use them. Six months from now, that will almost certainly be true again.

Our use of AI began with individuals rather than through an organized, top-down system. Different people on the team experimented with different American and Chinese models for relatively straightforward tasks such as research, analysis, drafting, and translation. At that stage, AI was mainly something people consulted through conversation: one question, one answer, one person at a time. It was useful, but the benefit generally remained with whoever happened to be asking the question.

Our thinking began to change as we drew on our own experience alongside what we were learning from leading firms and subject-matter experts, including a conversation with a professional at a leading quantitative investment firm. What struck us was not simply how extensively such a sophisticated organization was using AI, but how much its practices had changed in only a matter of months. The more important lesson was that there was no durable playbook to copy; the real work was to build an approach that could evolve with the technology.

That shifted our focus from individual experimentation toward a more systematic, firm-wide AI ecosystem that could support different functions, workflows, and roles across the organization.

The distinction we increasingly make is between a task and a capability. Producing a draft is a task. Building a process that can draw on relevant external information and firm context, produce a draft, check it against our own standards, improve it, and route it for human review before it goes any further is a capability. The first saves someone time once. The second becomes part of how the firm operates.

Rather than replacing our existing processes, we have tried to build on them. AI now assists with work across research, investor relations, content, translation, document analysis, and internal knowledge. In some cases, it helps gather and structure information. In others, it can compare sources, work through large amounts of material, retrieve relevant prior context, prepare an initial analysis, or take content through several stages of drafting and review. We are also using it to publish material across languages more efficiently and to make knowledge accumulated in one part of the firm more accessible elsewhere.

The important point is that these activities are increasingly connected through repeatable processes rather than dependent on one person opening a model and asking it a question. That is the difference between a collection of people using AI and an institutional capability. The latter has to be built around the firm's own information, processes, standards, controls, and judgment. Much of that cannot be bought off the shelf, because what gives the system value is precisely what is specific to the organization using it.

For us, the most important questions remain the ones that cannot be automated away. What constitutes a high-quality business? How much margin of safety is enough? Which evidence deserves trust? How should management incentives be interpreted? When does a market price reflect temporary fear, and when is it telling us something we have missed?

The underlying discipline is an old one: understand what a business is worth, pay less than that value, and leave enough room for error that being wrong does not become fatal. That intellectual lineage goes back to Benjamin Graham, but applying it in practice is never mechanical. For Banyan Partners, the framework has also been shaped by years of following Chinese equities closely, evaluating businesses case by case, and learning from both correct and incorrect judgments.

A model cannot supply that experience for us. Nor should it.

AI's role is to help us apply our judgment more effectively and consistently: to process more information and challenge our thinking in research; surface what we may have missed and preserve relevant context in our due diligence packages; and reduce the time spent on work that does not itself require investment judgment, such as preparing our monthly reports. The decisions, and the accountability for them, remain human.

As our use of AI has expanded, we have also learned that reliability matters almost as much as capability. We have worked with both free and subscription models and have encountered usage limits, service interruptions, access restrictions, and other disruptions along the way. The lesson has been straightforward: nothing important should depend on a single model, provider, or login.

Where a process matters, we increasingly build in alternatives. Different models can perform different parts of the work, fallback procedures exist when one is unavailable, and human review remains at the end of any process where the output matters. The objective is not to identify one winning model and build the firm around it. It is to build a way of working that can continue as the models themselves change.

Some of the clearest examples have come from content we now publish simultaneously on LinkedIn and WeChat in different languages. It would be easy to describe such output as "AI-generated," but that would miss what actually happened. Capable individuals, given access to the right tools, firm context, and clear standards, can research, analyze, and communicate at a level and speed that would previously have required considerably more time.

The technology matters, but so do the person using it and the institution surrounding that person. The result comes from all three working together.

That is really the point of this evolution. We did not set out to use more AI. We set out to build something ours: a way of working that reflects how we think, preserves what we have learned, improves how information moves through the firm, and holds up regardless of which model happens to be fashionable this year or what gets announced next quarter.

That is the part that does not expire.

We do not expect the development of this AI ecosystem to stop simply because we will no longer write periodic updates about it. We will keep adapting, testing, discarding, and rebuilding as the technology changes, quietly and continuously, much as we have built everything else at the firm since 2010.

The tools will change. The processes will improve. What stays fixed is the philosophy, the judgment, and the accountability underneath them.