

Filters First:

USING AI WITHOUT LOSING JUDGMENT

Kabir Narang, Co-Founder of Two Global Firms

February 2026

GFTN

Innovating Finance. For Everyone.



**COMPOUNDING
CAPITAL**


This report has been prepared by Kabir Narang of Compounding Capital and is distributed in collaboration with Global Finance & Technology Network (GFTN). The views expressed in this document are solely those of the authors and do not represent any endorsement by Global Finance & Technology Network (GFTN).

INTRODUCTION

Filters First - Invert always invert

Things to avoid

- Automating the wrong bottleneck and accelerating noise through the system
- Allowing exceptions to become style because AI makes them look rational
- Confusing governance with reporting and losing accountability in the process
- Thinking in relative performance while ignoring absolute resilience and coverage
- Paying for access and narratives instead of incentives and real contribution
- Relying on tool driven adoption rather than process driven adoption
- Forgetting that AI amplifies culture, it does not create it



**“Invert, always invert”
— Carl Gustav Jacobi
19th c. mathematician)**

Most investors will not lose money because they lacked intelligence. They will lose money because they lost discipline.

In many conversations with long term investors, capital allocators and CIOs across Asia, the Middle East, Europe and North America, a consistent pattern keeps repeating. The investors leaning into AI are not necessarily the most technical. They tend to be the ones who already have discipline in their decision making and are now using AI to harden that discipline into something closer to an operating system. The investors struggling with AI adoption are often chasing tools instead of process, and the result is a new form of noise that looks like sophistication.

Most AI conversations in wealth circles are framed around productivity: automating research, generating memos, summarizing meetings, reducing headcount. That is fine, but it is not the real prize. The real prize is decision quality over time. In long duration investing, decision quality is not only about returns. It is also about reputation, governance, and the ability to preserve good judgment across generations. AI can help with that, but only if you are clear about what AI is, what it is not, and where it introduces risks that do not show up in spreadsheets.

The wrong question is how do we use AI. The better question is where does AI meaningfully improve judgment without replacing it. When you frame it this way, AI stops being a shiny capability and starts becoming what it should be: cognitive infrastructure that supports a principled investment process.

THE RIGHT MENTAL MODEL

AI should not be treated as a CIO, an allocator, or a replacement for experience. It should be treated as a system that improves the inputs and consistency of human decision making.

In practice, AI performs best in four roles. It functions as an exceptionally fast analyst that can process more material than any human team. It acts as a pattern recognizer that surfaces inconsistencies across memos, models, and narratives. It serves as a memory system that captures context that would otherwise evaporate across meetings, cycles and staff changes. And it enforces decision frameworks meaning it applies the same evaluation standards repeatedly, without fatigue, politics or drift.

That fourth point matters more than it sounds. Over time, most investment organizations or family offices do not fail because they lacked intelligence. They fail because of inconsistency. They apply tight standards in one cycle and loosen them in the next. They let narratives override discipline. They forget why a prior decision was made. They allow exceptions to become an operating style. AI, used well, does not magically create alpha. It can dramatically reduce unforced errors and narrative driven decision drift.

The governing principle is simple: AI accelerates thinking, but humans retain judgment and accountability. The real work starts once the philosophy is clear. Tools matter less than workflow design. The edge comes from building simple repeatable loops: intake, rubric, adversarial review, verification, and decision capture. The next piece of this report will go deeper into how to structure these loops so AI improves judgment without diluting accountability.

WHERE AI ADDS REAL VALUE TODAY

Deal triage and attention management

The most immediate high value use case is deal triage. Every serious investment organization and family office eventually hits an attention bottleneck: too many inbound deals, too many manager pitches, too many special situations, too many compelling narratives wrapped around unclear fundamentals.

Humans are not built to maintain consistent skepticism at high volume. Even strong teams get anchored by story, status, or social proof. AI can be used as a first pass filter that screens large volumes quickly and flags which opportunities deserve human time. Done properly, it does not approve deals. It narrows the funnel and highlights what to interrogate, especially where risks cluster and where the story does not match the numbers.

A practical way to do this is to treat AI as a structured pre IC layer. Every deal gets the same questions, the same scoring logic, the same set of failure modes. The goal is not to automate conviction. The goal is to allocate human attention to the few situations where judgment is actually required.

Manager diligence and incentive alignment

Most experienced investors intuitively understand that incentives drive behavior but many investment processes still get trapped in surface signals: brand, pedigree, proximity and access. AI can help push diligence down into the layer that matters: incentives, behavior under stress and the real engine of returns.

In private markets especially, there is a recurring mismatch between reported outcomes and realized outcomes. Some managers optimize for asset growth, narrative strength and smooth marks rather than liquidity and cash returns. This is not a moral judgment, it is incentive math. AI can help you pressure test this by systematically extracting signals from decks, letters, and track record disclosures, then forcing the discussion back to the few variables that matter: realized distributions, time to liquidity, concentration risk, loss ratios and how valuations are actually being set.

The point is not to let AI decide who is credible. The point is to make it harder for your own process to be seduced by polish.

Decision framework creation and enforcement

The second underrated use case is decision framework creation and enforcement. Most investors have implicit values around liquidity, governance, downside risk, and time horizon. Those values often are not operationalized. AI can help an investment organization or family office turn values into a practical scoring rubric that forces explicit trade offs.

The value is not the score. The value is the framework. Once the framework exists you can apply it consistently across managers, deals and time. Over years, this consistency compounds. It reduces post hoc rationalization, where an investment committee decides emotionally and then uses analysis to justify the decision after the fact.

A well designed framework also makes governance easier across generations. It becomes a shared language. It prevents the same debate from being re litigated every cycle. It also highlights where differences are real, such as risk tolerance or liquidity preference versus where differences are simply noise.

Knowledge compounding and institutional memory

The third use case is knowledge compounding. Investors accumulate years of tacit learning through meetings, calls, internal debates, and post mortems and then they lose it. People change roles. Memories fade. Context disappears.

AI can turn those fragments into a searchable institutional memory: why you passed on a manager five years ago, what concerns emerged in diligence, what warning signals were visible early, which patterns repeated across cycles, and where you have historically overestimated your ability to influence outcomes. This is not a nice to have. It is one of the most defensible ways to protect judgment across time and generations. Over time, the investors who compound institutional memory will compound capital more intelligently, not because they are smarter, but because they are less forgetful.

A subtle benefit here is cultural. When the institutional memory is accessible, decision making becomes less personality driven. The office becomes less dependent on the loudest voice in the room, and more anchored in evidence and continuity.

WHERE AI QUIETLY BECOMES DANGEROUS

The most important part of adopting AI is understanding how it fails.

The first danger is letting AI slide from analysis into allocation. Allocation decisions embed values, responsibility, and accountability. AI can inform allocation discussions, but it should not be the decision maker. The moment a family office treats an AI output as objective truth, it starts outsourcing the very thing it exists to preserve: judgment under uncertainty.

The second danger is trusting single pass outputs. AI is probabilistic. A single response is a draft, not a conclusion. The families using AI well tend to run multiple passes. They ask for the base case, then the bear case, then the quiet failure scenario. They ask the model to critique its own reasoning, then to argue the opposite side as if it were an adversarial reviewer. This is not overkill. It is the difference between using AI as an intern and treating it like an oracle.

The third danger is citation and reference risk. Models can produce polished answers with references that look credible. If you allow that habit into your process, you will import fragile confidence into high stakes decisions. A serious family office workflow must force verification. If a source matters, it gets checked. If a claim matters, it gets triangulated. AI can help you do the checking faster, but it cannot be allowed to invent certainty. The fourth danger is using chat history as memory. Chat tools feel persistent but they are not a knowledge architecture. Threads fragment, context drifts, old assumptions get lost, and the system becomes a confusing swamp. What works is moving toward durable files and structured repositories, where memos, meeting notes, rubrics, and decisions are stored deliberately and can be queried later. A family office AI system should look more like a well run research library than a long chat log.

The fifth danger is allowing unsupervised agents near sensitive systems, especially financial systems. Even if it works, you are introducing operational and reputational risk that is not worth the efficiency gain. A blunt but accurate mental model is that AI is an intern with a PhD and zero common sense: brilliant in bursts, sometimes dangerously literal, and absolutely requiring supervision and governance.

WHAT TO DO NEXT: THE QUESTIONS WORTH ASKING

If this topic is interesting, the next step is not buying another tool. The next step is asking sharper questions that reveal whether AI will strengthen your process or weaken it. These are not checklist questions. They are questions designed to expose where the operating system is disciplined and where it is not.

Start with the attention bottleneck. Where does time actually break: deal flow, manager diligence, portfolio monitoring, governance, reporting, or family communication for family offices. Most investors say all of the above, but the right answer is usually one or two choke points. If you do not know the bottleneck, you will automate the wrong things and simply accelerate noise.

Then define the non negotiables of your investment culture. For example, you may have a rule around transparency, governance, liquidity path or evidence of realized returns. If you cannot articulate these clearly, AI will amplify inconsistency because it will mirror the ambiguity in your human process. Clarity is a control surface.

Next, decide what decision framework you want to enforce across cycles. A powerful move is to define the criteria you want applied to every deal and every manager, then have AI enforce that decision filters consistently. That is how you turn values into repeatable decision hygiene. Most organizations and families already know their values. They just have not translated them into an instrument panel that actually shapes behavior.

Then audit your institutional memory. How do you store and retrieve learning. If your best insight lives in someone's head or in scattered email threads, you are leaking compounding. The goal is not more notes. The goal is a system where a future investment committee can ask, have we seen this pattern before, and get a high quality answer in minutes.

Finally, set your governance boundary for AI. What can AI do. What can it not do. How are outputs reviewed. Who remains accountable. The investors that get this right move faster with fewer mistakes. The investors that get it wrong move fast until they hit a reputational or operational wall.

CLOSING

AI will not make a disciplined investor reckless. It can make an undisciplined investor feel sophisticated while quietly degrading judgment. Used properly, AI gives experienced investors more time to think, more consistency across decisions, and stronger institutional memory. Those are compounding advantages that matter in long duration capital.

The next stage is execution. The next sequential edition of this report will go deeper into what this looks like in practice with frontier tools like Claude and other models: how to structure workflows around them, how to design prompts that improve investment decisions rather than generate clever text, and how to build lightweight guardrails so the machine augments human judgment without diluting accountability. The goal is practical: prompt patterns, governance checklists, and examples of where these tools are already creating edge for family offices.

ABOUT THE AUTHOR

Kabir Narang is a technology investor, builder and entrepreneur, whose work has been shaped by a simple idea: compounding is not just about capital. It is about relationships, decision quality and the ability to learn faster than your environment changes. Over the past two decades, he has founded and built two global investment firms and invested in technology companies globally. He also works with incredible founders from early stage through growth and pre-IPO phases.

Kabir is building Compounding Capital around a simple philosophy: long-term outcomes come from relationships that deepen with time, trusted LP relationships, and backing exceptional founders and teams through multiple cycles. Prior to this, Kabir was Founding Partner at B Capital and served as the first Chair of B Capital's Global Growth Investment Committee, helping build the organization into a multistage global franchise across Silicon Valley, Singapore, and India.

Kabir is also building HumanTronik, an enterprise AI company building the execution layer that helps large organisations translate leadership intent into measurable business outcomes. HumanTronik is designed as an intent-to-outcome engine: a system where humans and AI work in the same execution loop to help plan, orchestrate, and measure execution across teams, workflows, and systems. He is building it with Monish Darda, co-Founder of Icertis, one the most consequential investments for 4 VC firms and a world-class team & advisors across San Francisco, Seattle, Singapore and India.

Earlier in his career, he worked with Warburg Pincus and spent time in management consulting at McKinsey in India and the Boston Consulting Group in the United States. He holds an MBA from the University of Oxford and a BA in Mathematics from St Stephen's College, Delhi University.

His training in mathematics and long-standing practice of photography have shaped how he thinks about investing and decision making. Mathematics taught him to reduce problems to first principles and reason from structure rather than story. Photography taught him to look for patterns, light, and negative space and to distinguish signal from noise before pressing the shutter. Writing forces clarity. It exposes sloppy thinking and rewards precision. He writes and invests with a first-principles bias, and that shows up in how he evaluates incentives, how he thinks about governance, and how he handles uncertainty across cycles. It also shapes his view on AI. The point is not to be impressed by outputs or speed. The point is to design a system that preserves judgment when the environment becomes noisy, narratives get louder, and decisions matter more.

If you have been experimenting with AI in your investment organization or family office, and you are considering how to do it without creating new risks, the author would enjoy comparing notes. The author has been discussing these questions with investors across twenty countries and the differences in approach and governance maturity are the real story. If any of this resonates, or if your experience has led you somewhere different, reach out to the author at kabir@compoundingcap.com, who is always keen to learn and exchange notes with other likeminded investors.

Author's LinkedIn: <https://www.linkedin.com/in/kabir-narang-b293a9/>



AUTHORS

Kabir Narang

Co-Founder,
HumanTronik along with Founding Partner,
Compounding Capital

CONTRIBUTOR

Kaitlyn Thinn

Head of Strategy and Research,
GFTN

PRODUCTION

Carol Ann Christy

Graphic Designer,
GFTN

GLOBAL FINANCE & TECHNOLOGY NETWORK (GFTN)

6 Battery Road, #28-01, Singapore 049909
gftn.co | hello@gftn.com

This document is published by Global Finance & Technology Network Limited (GFTN) as part of its FutureMatters insights platform. The findings, interpretations, and conclusions presented in GFTN Reports reflect the views of the author(s) and do not necessarily represent those of GFTN, its Board, management, stakeholders, or any individual participant and their respective organisations.

© 2025 Global Finance & Technology Network Limited, All Rights Reserved. Reproduction Prohibited.