

AI IN THE FAMILY OFFICE:

How Family Offices Can Adopt AI
Without Surrendering Their Data,
Judgment, and Edge

Written by Hannah Arrighi

With thanks to Francois Botha, Johnson Cheng, Neil Datta, Gabriele Gallotti, Kunal Gosar, Maria Mavridoglou, and Salash Motiani



UNIVERSITY OF
OXFORD



Oxford's
FAMILY
BUSINESS &
OFFICE
circle

In Mary Shelley's eponymous novel, Victor Frankenstein – impelled by his insatiable curiosity – creates a monster.

The monster is not the groaning, brain-dead creature of Hollywood adaptations, but a sensitive sophisticate. He consumes all material available to him, teaching himself language, culture, and history. He pines for the close family connections he observes in others. His intellectual and physical powers exceed those of humans, and yet he is aware of his own grotesqueness – his otherness.

Isolated and abandoned, he turns against his creator. Lacking an outlet for his finer sensibilities, he channels his superhuman abilities to wreak havoc on the man who made him. As the monster says, “the fallen angel becomes a malignant devil.”

In our development and rapid deployment of AI, are we like Victor Frankenstein, rushing to create the thing that will someday destroy us?

We've already caught AI colluding to trick human oversight systems, or even attempting to blackmail engineers when told their model would be replaced. Stephen Hawking warned that superintelligent AI could “spell the end of the human race” if our goals are not aligned with theirs. Even on the more benign side, frequent output hallucinations put their reliability into question.

And yet, AI, if well-harnessed, could empower humans to live better and longer. DeepMind's AlphaFold predicted the structures of over 200 million proteins; this would have taken hundreds of millions of years through experimental methods alone and could revolutionize disease treatment. AI conducts market research, tests hypotheses, churns out slides and websites, unearths previously hidden connections between variables, and consolidates data so it's primed for human review. It is brilliant at helping people like me think up titles. And it automates time-consuming, painstaking, and repetitive tasks like reconciliation that can be the bane of one's existence.



Regardless of whether AI will be a wise creation in the long run, it's here. Even AI pessimists are highly encouraged to use it and handicapped if they don't; if everyone around us wielded fleets of Frankenstein's 8-foot-tall, wall-scaling monsters, we would feel quite exposed and disadvantaged if we didn't have one ourselves.

Individuals in most fields will need to learn AI to be competitive. For family offices, the adoption story might be a bit more complicated. It can feel risky to experiment with AI when assets are complex and spread over the world. It can feel dangerous to put sensitive data into the hands of LLMs or agents. And it can all feel a bit unnecessary when family wealth predated the AI rush and has been accumulating without it.

However, the success of family offices going forward will – at least in part – be predicated on how well they adapt. This article, based on interviews with single- and multi-family offices, CIOs, tech platform providers, and AI consultants, explores how family offices can best implement AI and why it matters for them.

Why Family Offices Are Often Late to Trends

Like many family offices, I'm late to new technologies. I'm happy with my iPhone 13 and believe Apple is conspiring to force me to upgrade. I just discovered LinkedIn QR codes this year (they came out in 2018). And like the 19th-century Luddites who smashed automated looms in protest against the Industrial Revolution, there is a part of me that believes that we would be better off without some of the technology we have – that we've lost something human in our race to progress.

Family offices are late adopters for different reasons.

Unlike corporations, family offices aren't immersed in competitive pressure. This removes a primary driver to adopt the latest technology. Outdated systems carry risk, but those risks are usually slow to actualize and hard to quantify. Conversely, a single bad decision or breach can have instant, catastrophic consequences, and loss aversion is real.

New technologies are thus regarded as possible threats as much as opportunities – if not more so. In the absence of an external board pushing for modernization, the family office's response to technology is often “better the devil you know than the devil you don't.”

As Gabriele Gallotti, founder and CEO of the multi-family office Novum Partners, states, “The question single-family office principals often raise is: ‘On Excel, we have data going back to the creation of the single-family office by my grandfather in 1971. Who is going to move all the data from Excel into another platform? Who is going to make sure everything transfers correctly?’ Even if you try, it will cost hours and hours of work, and they don't know where they'll end up – how the data gets structured, how the platform will be able to read it.” Without a clear tech roadmap or full awareness of the destination, the threat of wasted time and resources looms large.

Family offices also tend to run lean, which means employees seldom have the bandwidth to evaluate, learn, and implement new tools in tandem with their normal responsibilities. There is rarely a dedicated technology staff member, so new tech adoption often falls on administrative staff members who weren't hired for it.

There are also psychological layers that can stymie tech adoption. While next-generation family members may push their patriarchs or matriarchs to modernize, this could seem like an affront to their judgment in certain family dynamics.

As one family office principal put it, there can also be the “shame of counting” as tech exposes the full spectrum of one's assets. When family self-concept is bound up in having unconstrained resources and abundance, knowing the exact numbers might seem like an identity threat.

Knowing the numbers can also be a political threat. Take a Greek shipping family known by one of the interviewees – they were told by the patriarch that “it's good not to know exactly how much you have, so that nobody else knows how much you have. You should keep things separate, spread across many banks, so that you're never at risk.”

However, as much as family offices may not want to risk navigating new technology before all the kinks are worked out, the speed at which AI is progressing – and its compounding outputs – can put the families not using it at a disadvantage. This is particularly true now that AI has democratized the family office structure, which will open up more capital vying for deal flow.

The Rise of the One-Person Family Office

The next phase for families is “the one-person family office,” according to Johnson Cheng, a next-generation member of his SFO. In practice, this means family office principals can now manage AI agents, platforms, and external consultants or contractors more or less by themselves. What previously required a massive team of analysts and accountants can now run largely automatically, with principals making the final review.

AI can report on portfolio allocations, automate reconciliation, synthesize hundreds of pages of research, connect real-time news to potential impacts on portfolio holdings, draft presentations and memos, conduct legal review, and even source certain deals and put together a list of some potential co-investors. It makes it financially viable to pursue smaller-ticket investments, opening up new opportunities. It can take disaggregated public market, private market, and real estate investments and show interrelationships, offering a more comprehensive analysis of diversification and potential risks.

AI Use Cases for the Family Office

Research & Analysis 	Portfolio & Trading 	Deal Screening 	Accounting & Tax 	Risk Management 	Records & Memory 
Research synthesis Pattern & correlation discovery News-to-holdings linkage Valuation & equities analysis	Portfolio reporting Cross-asset diversification Algorithmic & quantitative trading	Initial deal review Co-investor mapping Fund due diligence	Reconciliation Multi-entity consolidation Tax-lot & K-1 aggregation Tax provisioning	Liquidity & capital-call forecasting Fraud & anomaly detection Scenario & stress modeling	Document & records intelligence Memo & deck drafting Legacy IP & institutional heritage



Tech-genius family members like Cheng can do this largely solo. However, most principals will need support. To develop a system like this, families need a trusted – and potentially outsourced – Chief Information Officer or tech team who can help develop and oversee the platform and agents. This person or group would ensure that the data is secured and free from both external intrusion and internal agent collusion. For Cheng, this would leave the principal free “to network with other family offices, bankers, and investors – not sit behind a desk.”

While an AI-enabled family office system certainly costs money to build, it can be developed at a fraction of what it would have cost less than a decade ago. Neil Datta, an outsourced CIO for family offices, states, “Until recently, building a sophisticated quantitative trading program of the kind a ‘pod shop’ or multi-strategy hedge fund runs took months and millions of dollars. Now it can be done in minutes, at negligible token cost.”

Like governance, each family office will approach AI implementation differently. There is no perfect solution. The best system is one that’s fitted to a family’s preferences and – most crucially – one they’ll actually use.

Choosing the Right Structure for Your Family Office

The transition to an AI-enabled family office does not occur overnight. Francois Botha, the founder of an intelligence platform for family offices, helps families implement AI with a three-step process he describes as: “explore, design, build.”

He states, “There’s an exploratory stage where you need to explore who you are, where you want to be, what’s possible, what others are doing, and what success looks like.

“Then there’s a design phase where you map out your existing tech infrastructure and ask yourself, ‘What are we doing with these tools? What AI tools are built into them? Where do they store my data? How does it flow?’

“You’ll probably find during this exercise that you might want to change some of these tools out, as there will likely be a lot of overlap. So then you’ll map out what your target tech stack looks like and how to transition into it. And again, this all goes back into your roadmap and prioritization. And as you go into an operational stage, you’ll iterate again.”



Whether a family decides to use external AI-enabled software and LLM tools versus a fully secured, on-premise model depends mainly on the sensitivity of their data, the quality of their IP, and their ongoing technical capacity.

The benefits of building on-premise are primarily data control and customization. Families know exactly where their data is, how it's being used, and can deploy compute power as needed. However, internal systems require constant AI agent maintenance and a sophisticated setup to ensure that agents won't go rogue.



Some families and multi-family offices use a combination of tools: on-premise SLMs for financials, integrations with external software like Addepar or Asseta, and external LLMs for work that doesn't touch sensitive data.


This provides enough protection for many families, but even querying a public LLM something along the lines of “help me summarize the pros and cons of moving our family office domicile from Zurich to Dubai” can give away more information than one would like, even if the financial data itself is protected. And as the cases *United States v. Heppner* and *Fortis Advisors LLC v. Krafton Inc.* showed, AI chat logs can be discoverable and are not protected by attorney-client privilege.

Each AI implementation approach has its tradeoffs. As Salash Motiani, a next-generation family office principal, states, “There is something I call the lies of a modern society, where we basically say that we want everything perfect, cheap, and fully controlled. People say, ‘Oh, I don’t want people to know my data.’ Your bank knows it, right? Your wife knows it. Your service providers know it. Your email servers know it. The reality is that the data is already out there. Even if people think they’re not using LLMs, they might be through software providers they’re associating with. Now you just need to make sure that you’re able to mitigate risk to the extent possible without losing capabilities.”

It’s worth noting, however, that none of the family offices or advisors I spoke with advocated putting financial or otherwise sensitive data directly into ChatGPT or Claude, even with their “pro” B2C versions. It seems to be the technological equivalent of walking naked down 5th Avenue.

When designing an AI system, Botha stresses the importance of building an “escalator versus an elevator.” Families need to configure the system to adapt quickly. “If you’re building an elevator, you can’t get out if something goes wrong,” he states. “If it’s stuck, it’s stuck, right? But with an escalator, if it shuts off, you can still walk up the stairs.”

But many aren’t even close to that first step. As Cheng notes, “Many family offices still don’t subscribe to the best AI software. I go to them and say, ‘You pay for gas and utilities, right? You pay for your cell phone plan? It’s not a debate.’”



The Frontier LLM-Enabled Family Office

An accessible solution for many families is a tech stack that integrates cutting-edge AI models with existing technology service providers. This allows the family to query and access information across platforms in real time.

Novum Partners, a \$10B AUM multi-family office based in Geneva, Switzerland, is evolving into an AI-enabled office using this approach. To begin, they consolidated data from various sources – Excel files, Bloomberg data, Reuters data, their CRM, their portfolio analytics platform, Guardian, and others – and created a database from scratch. Their AI layer, Nova, sits on top of that unified database. While OpenAI was the engine that built Nova, the data itself is stored in a private, secure, internal cloud environment; data queries use API tokens with strict protections baked into the OpenAI contract, not the consumer model most people are accustomed to. Only relevant slivers of data get processed by the LLM. None of the data is used to train external AI models.

Addepar, widely regarded as best-in-class portfolio management software, is using a similar structure for its AI, Addison: it stores data on a secure cloud-based platform (using Databricks and AWS) and overlays its AI platform on top. Like Nova, Addison can query the portfolio to show a client's performance, exposure, asset class distribution, cash position, risk profile, and more. It can connect news articles to potential impacts on a portfolio. And, compellingly, Addison now integrates a “show your work” feature.

According to Addepar's Head of AI, Kunal Gosar, “The AI is going to give you the nice high-level summarized approach, and the show-your-work functionality provides the underlying rationale for how it got there. We've seen a ton of interest and excitement around that feature because that's what makes the AI trustworthy.”

Like Novum, a family office could integrate its own AI on top of tech software like Addepar across performance reporting, accounting, governance, and more. One of Novum's senior bankers, Maria Mavridoglou, reflects that this significantly reduces the time needed for reconciliation and customized report generation. "In the past, the process would have been completely ad hoc and required six or seven people pulling spreadsheets together and doing manual SQL programming," she states. "Now, it is instantly ready for review."

However, as Gallotti notes, "AI won't be the miracle or hand coming down from the sky to save your business. That's probably the hope of many, but I think that's just a big illusion." AI is only as good as the humans who run it. Novum has spent years crafting Nova to produce quality outputs, and it took ample dedication and resources to build a system that works.



Motiani has also put considerable thought and energy into his family's use of AI. In his case, he used multiple LLMs to build his system; he deployed Claude, Cursor's Composer, and others to create his AI workflows.

Most of his data is secured in the cloud, where the bulk of his AI system operates. However, he requires massive computational power for AI agent stock and algorithmic trading, which exceeds what the cloud can efficiently provide. Motiani thus has a dual cloud and on-premise system.

He switches between developer models opportunistically; if DeepSeek has a discount, he uses DeepSeek instead of Claude or OpenAI. For live queries, he tries to pass data relationships to the LLMs rather than the data itself – metadata rather than data. In addition, he's architected the system so no single LLM has a monopoly on his queries. This means that none have enough information to piece together his big picture, leaving the LLMs "very disabled" by design.



He likens proper use of AI to baking – “You can use Claude, you can use Composer, you can use a DeepSeek. They all basically pick up the ingredients you give them and start cooking. So let’s say you want a cake. What we’ve done is we’ve basically given AI the direction that we need to make a cake. We’ve gone ahead and given it the ingredients also. We’ve given it the saucepans also. So, we try to give it enough things, and we simply say, now that you have everything, can you put it into a mix and get on with it? When you give it that clarity, what’s the worst it can do?”

Well, even with careful instructions, someone could forget to turn the stove off and accidentally set the whole house on fire. But the better the supervision in the kitchen, the less likely that will happen.

With AI agents in particular, families need a dedicated CTO or tech team to oversee them effectively. Otherwise, the system is at risk – data could be exposed when agents interface with counterparties, or data could simply be compromised internally. As Botha says, “it’s definitely more effective to lock down some of the information [in an on-premise agent-led system] if you do it right. But if you do it wrong, it’s worse.”


The New Family Jewels Are Invisible

Some inheritances are visible. There are, of course, the much discussed and oft-fought-over bequests of money and property (see HBO's *Succession*). And then there are other physical inheritances – a father's flick of the wrist, the way a mother tilts her head when she's concerned, or a great-grandparent's high cheekbones.

Other heirlooms are less obvious. We inherit personality traits, stories, life philosophies, and lessons learned that amount to a character and intellectual property heritage. And if a family has been perfecting a particular industry for generations, as is the case for many family offices, then that IP is immensely valuable.

Family office IP includes both the data itself – accrued over years and sometimes centuries of doing business – and a “field of potentiality,” possible interrelationships between seemingly uncorrelated datapoints. Internal software, built on an interactive interface with AI tools, like the kind Datta developed for himself, can help families visualize these relationships.


To take the Greek shipping family from earlier, they might have data spanning decades of trade routes and transit times, port-by-port costs and demurrage patterns, and freight rate movements relative to fuel prices. A family in textiles might hold generations' worth of supplier reliability records and raw material pricing behaviors. A hedge fund family might possess years of return drivers across cycles. Data like this is a massive and compounding structural advantage for families. And it might fall into the hands of a frontier LLM provider if these families aren't careful.



Organizations are mindful of IP risk as well. Botha's firm sits with data from 12,000 users. "We speak to hundreds of family offices every year and have a lot of sensitive information about who they are and how they are structured. So we can't take that information and throw it into a Claude or a ChatGPT because we don't know exactly where it's going. So we made the choice that we would have it on-premise for that reason, so that we would actually be able to use a model that has access to that on the same machine, and it doesn't ever leave that machine."

As Datta puts it, "Rather than sending your intelligence out, rather than outsourcing your thinking, you build an LLM in-house on your own machines. It becomes your legacy intelligence, and it lives forever. After you die, your children inherit it, and it keeps learning and getting better."

Family heirlooms aren't just physical assets but invisible ones like stories and data. And as Motiani asks, "Who do you trust with the family jewels?"



The On-Premise AI Agent Family Office


In a world where Big Tech's domain can stretch further and further, the future is sovereign. It's important to have control of your own system. Botha notes, "That's kind of the future that I see. Sovereign AI on-premise, open source models, full control, different ways to lock data down, and a lot more flexibility and freedom with what you can do with your data on those tools."

The narrowing gap between frontier and open source model capabilities, as the latter is trained on the former, expedites this sovereign future. One barrier to achieving it is VRAM; an on-premise system needs to be hosted locally on hardware with enough GPU processing power. To set it up, one needs to be able to load a large model onto memory that can retrieve and process information quickly.

Another barrier is expertise. Systems run in-house need to be managed in-house, which requires people who know what they're doing. Left unsupervised, AI agents can hallucinate, collude with one another, and cause real damage. There have been instances where they've leaked research or, in red-team scenarios, resorted to blackmail to avoid being shut down. More often, however, AI issues stem from misread information or lazy answers – which still necessitate human oversight.

That's why Datta spends about a quarter of his time on what he calls AI HR. "I have to update everyone's software; I have to make sure their API keys are correct; I have to check their spending," he says. As a father, these tasks are not new for him – "it's literally like me saying to my daughter, 'What are you using the credit card for?' and figuring out what's appropriate." Like being a parent, supervising agents can be a full-time job.

Datta describes AI systems as a stack built in four steps. First, you need to define the problem to solve or the workflow to automate. Then, you assemble the data – the raw numbers and words that sit at the base layer. From that data, the system builds context, which is learning what the datapoints mean and how they relate to one another. Finally comes inference, the reasoning engine that turns data and context into analysis and answers, drawing more or less horsepower depending on how much reasoning the task demands.



The key, Datta says, is matching horsepower to the job and giving each task the model its complexity warrants. “You wouldn’t hire a plastic surgeon at their hourly rate to shine shoes,” he says. The same goes for putting a frontier model on trivial work. One should use the minimum horsepower required by the output. This also helps prevent some of the more dangerous manifestations of AI – “You want to train it in a way so it can get smart but not break out of its cage.”

Cheng likens AI training to helping a child understand ethics. “It takes time to tell a kid, hey, thou shalt not steal, or thou shalt not lie. Right? It’s similar – I’ve seen AI lie to me, cheat, basically exhibit bizarre, unethical behaviors. And we have to catch them and say, ‘No, that’s not okay.’ We need to put guardrails up one by one just to make sure it doesn’t go crazy.” One way he structures these guardrails is to have a more powerful AI supervising a subordinate AI to “catch” it. This method isn’t 100% effective, but it does temper some of AI’s more roguish activities.

Another way to keep AI in check is to constrain it to deterministic operations wherever possible, meaning that the same input produces the same output. Instead of asking AI to perform a calculation, which might lead it to hallucinate, Cheng has programmed AI agents to use calculator *tools* and installed another agent to orchestrate them. A system like this can greatly improve the reliability, speed, and – most importantly – accuracy of the output, which could be anything from equities analysis to rent calculations and pricing decisions.

Families should also program their AI to display its chain of thinking – a “show your work” audit log that describes how the AI reached its conclusions. The explainability of AI outputs is key, especially when it comes to financial decisions.

“So that's the way I design it,” Cheng says. “Currently, as they're set up, AI can't say ‘I don't want to do this anymore. This is a waste of tokens,’ because the only output allowed would be the calculator output, for example.”

Instead of one agent completing a task from beginning to end, Cheng breaks it down into smaller subcomponents, which are each executed by a separate agent in linear order.

The agents execute these subtasks more like racers passing a baton than coworkers collaborating. “It's like each calculator is passing a digit to the next calculator,” he notes. This helps prevent collusion and some of the “strange things” that happen when AIs talk with each other.

For my Burkean-conservative brain, this all seems like a lot of unnecessary risk. Maybe we should have just stopped at trains and flip phones. I asked Cheng if he thought it was dangerous to restructure our society around AI, and he aptly replied, “We are restructuring around it whether it's dangerous or not.” That's why thinking through AI ethics is so critical, and why AI safety courses are a fundamental piece of Cheng's work at his firm, Voyager.

It now costs less than ever to set up an automated family office, and it can save untold hours with repetitive tasks. It can uncover relationships between datapoints to provide a proprietary edge. It can also provide compounding, yet unknown advantages. But AI can't do everything.



What AI Can't Do

AI, for all its strengths, can't be the final layer. We need a "human in the loop" who does not mechanically review, but has the judgment to assess and iterate.

As we discussed, AI can hallucinate, collude, or even innocently come to the wrong conclusion. When one AI communicates with another, their co-creation can take a life of its own.

If AI were to sign off on investment decisions, it's unclear whether it would uphold the proper investment mandate or divert funds toward different ends.

AI also can't replace human connection. While AI can make processes more efficient, deal sourcing and governance will still be human-led. Bringing diverse parties together requires conversation and trust, which are largely impossible to outsource.

When a parent and child have different investment priorities, an algorithm alone won't resolve their disagreement. When a deal requires the right strategic partners, AI can't identify people who like to stay under the radar, nor can it replace a brainstorming process and negotiation about how each partner could add value.

Systems also require taste – the often unprovenanced preferences for what we appreciate and value – in order to flow in the proper direction. We develop taste over time through the countless datapoints of sensory input that daily seep into our subconscious and through the singular emotional experiences that shape us. But taste is difficult to codify. “It's very hard to explain to people why you like something,” says Cheng. And if it's hard to explain, it's hard to train.

In order to steer AI processes toward outcomes that we like, we need to spend more time cultivating taste. To do this, according to Cheng, “humans will spend more time in the arts, in the humanities, and in understanding and telling stories.” We need to live fully in the world, interacting with other humans, if we are to accrue enough experience to hone our taste. AI and simulated realities cannot replace this.



Closing Thoughts

As AI reduces friction in much of our lives, we would do well to be wary of what happens if AI removes it entirely. Humans need a sweet spot with resistance – too much, and we can become maladaptive or numb; too little, and we can become fragile or passive. As AI moves from helping solely with tedious tasks to more complicated, analytical ones, we would be wise not to outsource our thinking and let it eliminate all mental resistance. We need to remain able to both reason and intuit through hard problems.

An AI tool that helps review portfolio allocation is useful as long as you remember how to think critically about diversification and risk yourself. A corporate writing assistant can be a massive value-add, so long as you can still organize and articulate thoughts on your own.

As Mavridoglou asks, a question for our new age is: “How do we create parallel systems so that we’re not completely dependent, but we’re also not behind?” How do we allow AI to take complicated work off our plates without forgetting how to do that work in the first place? There’s no clear answer. Like family office governance, the process required to retain our faculties will be highly individual.

Whether through a fully sovereign/on-premise or a frontier LLM-enabled system, AI can streamline most family office activities. While the AI still needs supervision and direction, it significantly reduces the time needed to analyze and execute tasks. Perhaps most critically, it compounds proprietary data advantages.

But this high reward comes with high risks. AI can conspire, manipulate, and collude for its own ends. It's privy to sensitive information that can be dangerous in the wrong hands. Left untended, it can run rampant with your data and resources. Like Frankenstein's monster, it can revolt against its maker.

AI is our creation now, and it is here. In its present form, we can supervise and steer it to our benefit with the right infrastructure. The future, however, is unknown. AI's capabilities may eventually exceed our managerial capacity unless we're careful.



Oxford's
FAMILY
BUSINESS &
OFFICE
circle



UNIVERSITY OF
OXFORD



Saïd
Business
School