HUDSON | FERRIS

WHITE PAPER | AUGUST 2026

The Agentic Advantage

How agentic AI workflows surface and engage major gift prospects at scale, and why this changes the economics of nonprofit fundraising

Ed Leibman
Founder, Hudson Ferris Inc. | ed@hudsonferris.com | www.hudsonferris.com

EXECUTIVE SUMMARY

The gap between what fundraising needs and what fundraising teams can do

American philanthropy set a record in 2025: $617.2 billion in total giving, the first time the sector has crossed the $600 billion mark. Yet beneath that headline, the number of donors fell for the fifth consecutive year, and the share of U.S. households that give at all has slid from roughly two thirds in 2008 to about half today. The dollars are increasingly concentrated in major and principal gifts, which means the organizations that thrive will be the ones that can find, qualify, and personally engage major gift prospects, consistently and at scale.

That is precisely the work most development offices cannot staff. The average frontline fundraiser stays in a role roughly sixteen months. Development director seats sit vacant for six months or more, and replacing a departed fundraiser costs organizations well into six figures once salary, recruiting, and lost pipeline are counted. Prospect research, the foundation of every major gift, is the first thing to stall when a seat goes empty.

This paper argues that agentic AI workflows, meaning AI systems that plan and execute multi-step tasks rather than simply answering prompts, are the first technology that closes this gap rather than merely narrating it. Drawing on Hudson Ferris's work deploying custom AI agents across dozens of client engagements, we describe a five-stage agentic prospecting pipeline, share early performance data from live campaigns, and set out the guardrails that keep the technology in service of authentic human relationships. The conclusion is direct: research and cultivation capacity that once required a fully staffed development office is now available to organizations of nearly any size, and the leaders who adopt it deliberately will out-raise those who wait.

THE LANDSCAPE

A record year built on a shrinking base

Giving USA's 2026 report confirmed what development professionals have felt for years. Total giving rose 5.7% in current dollars to $617.2 billion, but the growth came from the top: mega gifts roughly doubled year over year to $22 billion, bequest giving surged nearly 20%, and revenue gains concentrated in larger gifts while the number of donors giving under $1,000 stagnated or declined. The Fundraising Effectiveness Project recorded a 3.6% drop in total donor counts in 2025, the fifth straight annual decline.

For nonprofit leaders the strategic implication is uncomfortable but unavoidable: broad-base acquisition keeps getting more expensive per dollar raised, while a growing share of every organization's revenue potential sits with a relatively small number of individuals and institutions capable of major support. Major gift work is no longer one program among several. For most organizations it is the growth engine.

Yet major gift work is exactly where the sector's staffing crisis bites hardest. Surveys of the field consistently find that half of development directors expect to leave their job within two years, and vacancy periods average six months. Each departure takes pipeline knowledge with it: who was researched, who was contacted, what was said, what the next step was. The traditional answer, hire more researchers and gift officers, fails on cost and on availability. There are simply not enough experienced fundraisers to go around, and organizations below a certain budget size will never be able to field a full research and moves-management operation the traditional way.

THE TECHNOLOGY

From generative to agentic: what actually changed

Most nonprofits have already touched AI. A December 2025 survey of 346 organizations found 92% using AI in some capacity, yet only 7% reported improvements significant enough to change what their team can accomplish, such as doubling prospect research capacity or personalizing donor communication at scale. Nearly two thirds described their use as reactive and individual: a staffer asking a chatbot to polish an appeal letter. The sector has adopted generative AI as a writing aid and stopped there. Researchers call it the efficiency plateau.

Agentic AI is a different category. Where a generative tool answers a single prompt, an agentic workflow pursues a goal: it plans a sequence of steps, uses tools such as wealth screening platforms, giving records, and outreach systems, evaluates its own intermediate results, and carries the task through to completion with a human approving the consequential moments. Gartner projects that 40% of enterprise applications will embed task-specific AI agents by the end of 2026, up from under 5% in 2025. The commercial world is rebuilding its revenue operations around this capability. The nonprofit sector, which runs on exactly the kind of research-heavy, follow-up-intensive work agents do best, has barely begun.

The distinction matters because fundraising's bottleneck was never writing. It is the compounding labor of identification, qualification, enrichment, personalization, and disciplined follow-through across hundreds of prospects at once. A chatbot shortens one task. An agentic workflow runs the whole pipeline.

THE METHOD

The agentic prospecting pipeline

At Hudson Ferris we have spent the past two years building and refining custom AI agents for client engagements, working across the most cutting-edge wealth and philanthropy platforms available, including DonorAtlas and Impala for donor intelligence, Apollo.io for verified contact enrichment, and Saleshandy for managed outreach sequences, with tailored agents built on Claude orchestrating the work end to end. The pipeline has five stages.

STAGE	WHAT THE AGENTS DO	WHAT IT REPLACES
1. Surface	Screen wealth, giving history, board service, and affinity signals across cutting-edge platforms to assemble a qualified prospect universe	Weeks of manual screening and list-pulling by a researcher
2. Verify	Confirm capacity, philanthropic track record, and mission fit for each name; discard weak matches before any time is spent on them	Gut-feel qualification and stale spreadsheet ratings
3. Enrich	Build a working profile per prospect: biography, interests, giving vehicles, connections, and verified contact paths	Hours of ad hoc searching per prospect
4. Personalize	Draft outreach grounded in each prospect's documented interests and giving, in the client's authentic voice, reviewed by a human before anything is sent	Generic appeals and mail-merge templates
5. Engage	Run multi-channel cultivation: personalized email sequences, phone outreach, personal notes, LinkedIn and professional networks, warm introductions, and convenings, with every reply routed to a human	One-off emails that go silent after the first touch

Three design principles distinguish this from the automation nonprofits have been rightly wary of.

- **Grounded research.** Every claim is grounded in a verifiable source. Agents work from documented giving histories, public records, and platform data. Nothing about a prospect is invented, and any fact that cannot be verified is flagged for human review rather than asserted.

- **Multi-channel by design.** Outreach is one channel among several. The agents draft and sequence personalized email, but the same intelligence feeds phone outreach, personal notes and direct correspondence, LinkedIn engagement, warm introduction mapping, and in-person or virtual convenings. The goal of every sequence is a human conversation, not a transaction.
- **Humans own the relationship.** Agents surface, verify, enrich, and draft. People decide. A consultant or client leader reviews every message before it is sent, takes every meeting, and makes every ask. The technology handles the heavy lifting so that professional judgment is spent where it changes outcomes.

The compounding effect is the point. Any single stage can be done manually. Done together, continuously, and at scale, the pipeline produces something structurally new: a permanently current prospect universe, researched to dossier depth, engaged with personal attention, that does not walk out the door when a staff member does.

THE EVIDENCE

What early campaign data shows

Because these workflows run through instrumented systems, their performance is measurable in ways traditional cultivation never was. The figures below come from a representative sample of eight recent Hudson Ferris outreach campaigns run for seven client organizations, spanning health and human services, youth development, violence prevention, international development, and education. Client identities are withheld; all campaigns targeted cold or near-cold major gift and institutional prospects surfaced through the agentic pipeline.

MEASURE	HUDSON FERRIS SAMPLE	COLD OUTREACH BENCHMARK
Deliverability	99%+ (bounce rate under 1%)	Widely variable; list decay is a chronic problem
Email open rate	61% of delivered emails	Typically 30–50% for cold campaigns
Prospect reply rate	Approximately 9% of contacted prospects	3.4% average reply rate across cold campaigns; top quartile 5.5%
Outcomes in sample	Funder meetings booked and interested prospects flagged for solicitation	Replies rarely routed to a relationship owner

Three observations deserve emphasis. First, deliverability above 99% is itself a research finding: it reflects contact data verified through the enrichment stage rather than purchased lists. Second, a prospect-level reply rate near 9% on cold major gift outreach is roughly two and a half times the cross-industry average for cold campaigns, and the difference is attributable to targeting and personalization quality, not volume. Third, and most important, replies are not the product. In the sampled campaigns, replies converted into scheduled funder meetings and flagged interested prospects that moved directly into human-led solicitation. Meetings with prospective funders, not email metrics, are the deliverable.

These are early numbers from modest sample sizes, and we present them as directional rather than definitive. But the direction is consistent across clients and sectors, and it compounds: every campaign enriches the prospect universe the next campaign draws from.

THE ECONOMICS

Why this changes the math of major gifts

Consider what the traditional model requires. A skilled prospect researcher can produce perhaps two or three deep prospect profiles in a day. A major gifts officer can responsibly manage a portfolio of 100 to 150 relationships. Building the pipeline that feeds that portfolio, identifying and qualifying the universe behind it, is measured in staff-months. For an organization with a development staff of one or two, it simply does not happen; frontline fundraisers report spending well over half their time on administrative and preparatory work rather than donor-facing activity.

The agentic pipeline inverts those ratios. Identification and qualification that took a researcher weeks compresses into days. Enrichment that took hours per prospect happens in minutes, across the entire universe at once. Personalized first drafts for a hundred prospects arrive in an afternoon, ready for human review. The organization’s scarcest resource, the time of the people capable of sitting across from a donor and making an ask, is redirected almost entirely to conversations.

The economics also change on the risk side. Because the pipeline lives in systems rather than in one employee’s files and memory, staff turnover no longer resets the program to zero. When a fundraiser departs, the prospect universe, the research, the outreach history, and the next steps remain intact and current, ready for a successor or an interim team. For organizations that have lost eighteen months of momentum to a single resignation, this may be the most valuable property of the model.

Finally, the model changes who can afford serious major gift fundraising. Research depth, personalization, and disciplined multi-touch cultivation were previously available only to institutions that could field large advancement teams. Delivered through agentic workflows, the same capability fits the budget of a boutique consulting engagement or, increasingly, an ambitious internal team. That is what makes this a structural innovation rather than an incremental tool: it moves the capability frontier for the whole sector, most dramatically for small and mid-sized organizations.

THE GUARDRAILS

What agents should never do

The case for agentic fundraising is strong enough that it does not need overselling, and the risks are real enough that they deserve plain statement. Gartner predicts that over 40% of agentic AI projects across industries will be canceled by the end of 2027, mostly from unclear value and inadequate controls. In fundraising, where the asset at stake is donor trust, the discipline matters even more. Our own practice enforces five rules.

- **Humans make every ask.** AI agents never make an ask, never negotiate a gift, and never impersonate a human in conversation. The moment a prospect responds, a person takes over.
- **Nothing unverified goes out.** Every prospect-facing statement is reviewed by a professional before it is sent, and every factual claim traces to a verifiable source. An agent that cannot verify says so; it does not guess.
- **Privacy is honored.** Prospect research draws on public and licensed data, handled under the same confidentiality obligations that govern any development office. Being researched thoroughly is not the same as being surveilled, and the line is respected.
- **Authenticity over volume.** Personalization must reflect real knowledge of the prospect's philanthropy and interests, in the organization's authentic voice. The goal is that every message reads as what it is: a thoughtful letter from a real organization to a person whose giving genuinely fits its mission.
- **Measurement is mandatory.** Reply rates and meetings are instrumented and reported, and sequences that underperform are revised or retired. Technology that cannot demonstrate its contribution to raised dollars does not survive in the stack.

Organizations evaluating vendors or building internally should demand the same commitments. The failure mode for this technology is not science fiction; it is ordinary sloppiness, generic spam at scale wearing an AI label. The remedy is governance, not abstinence.

THE PATH FORWARD

What nonprofit leaders should do now

Adoption does not require a moonshot. It requires sequencing.

- **Start with the audit.** Take stock of your prospect pipeline as an asset. If your organization's major gift knowledge lives in one person's inbox and memory, you carry a concentration risk no board would accept in your finances. Systematizing that knowledge is the first return of the agentic model, independent of any new prospecting.

- **Pilot narrow, measure hard.** Choose one campaign, one prospect segment, one fundraising goal. Run the full pipeline against it: surface, verify, enrich, personalize, engage. Measure meetings booked and dollars advanced, not novelty.
- **Invest in the people, not just the tools.** The technology multiplies judgment; it does not replace it. Development staff should be trained to direct agents, review their output critically, and reclaim the hours for donor-facing work. The organizations that gain the most treat this as a staffing strategy, not an IT purchase.
- **Put the guardrails in writing first.** Adopt explicit rules, like the five above, before scaling. Boards should ask for them. Donors, increasingly, will too.
- **Move while it is still an advantage.** The record $617 billion in giving is real, and so is the shrinking donor base underneath it. Both facts argue for the same response: build durable, personal, well-researched major gift capacity now, while the capability gap between early adopters and the rest of the sector is still an opportunity rather than a verdict.

ABOUT

About Hudson Ferris

Hudson Ferris is a boutique fundraising consulting firm headquartered in New York City. The firm has partnered with more than 250 nonprofits in the United States and globally on major gifts, campaign strategy, board development, and institutional fundraising.

A defining feature of our practice is the way we integrate artificial intelligence into our work. For every engagement, our team develops tailored AI agents that support prospect identification, outreach messaging, funder intelligence, and internal workflow, allowing us to do more, move faster, and keep our pricing competitive without compromising quality. The result is a level of research depth, personalization, and responsiveness that has never been possible before.

To discuss how agentic workflows could serve your organization's fundraising, contact Ed Leibman at ed@hudsonferris.com.

SOURCES

1. Giving USA 2026: The Annual Report on Philanthropy for the Year 2025. Indiana University Lilly Family School of Philanthropy. Total giving of $617.2 billion; individual giving $394.2 billion; mega gifts of $22 billion; bequest growth of 19.7%.

2. Fundraising Effectiveness Project, 2025 data: total donor counts declined 3.6%, the fifth consecutive annual decline; gains concentrated in gifts above $1,000.

3. Lilly Family School of Philanthropy: share of U.S. households giving to charity declined from approximately two thirds in 2008 to roughly half.

4. Virtuous, 2026 Nonprofit AI Adoption Report (survey of 346 organizations, December 2025): 92% of nonprofits using AI; 7% reporting major impact; 65% describing use as reactive and individual.

5. Gartner press releases, 2025: 40% of enterprise applications to feature task-specific AI agents by 2026, up from under 5% in 2025; over 40% of agentic AI projects predicted to be canceled by end of 2027.

6. Sector studies of fundraiser tenure and turnover: average fundraiser tenure of approximately 16 months; development director vacancies averaging 6 months; replacement costs commonly estimated near $130,000 per departure.

7. Cold outreach benchmarks, 2025–2026 industry reports (Instantly, Belkins, Mailshake): average reply rate 3.43%; top quartile 5.5%.

8. Hudson Ferris campaign performance: internal Saleshandy analytics, sample of eight campaigns across seven client engagements, 2024–2026. Client identities withheld.