



A decorative orange line graphic in the top left corner, consisting of a large arc and a smaller loop.

Smarter, faster, *kinder*.
Ethical AI tools every gift
planner should be using now

Cherian Koshy, CFRE, CAP[®], ACFRE

Kindsight


Two curves. One axis.



Capability: measured task-completion horizon (METR, 2019–2026). Donor trust: share of donors comfortable with AI in fundraising. Curve shapes illustrative.



Only one of these is moving
on *our* schedule.



This is not about automation. It is about *amplification*.

AUTOMATION

Replaces the person.


It takes the task and the judgment together, and the donor meets the system instead of you.

AMPLIFICATION

Returns the person.

It takes the friction and gives you back the hours, the preparation, and the conversation.

Why this one is
different



AI capability is measured in task length now.

Not accuracy. Not benchmark scores. Length of work.


Researchers measure how long a task takes a human expert, then ask how long a model can go before reliability breaks down. That single number tracks the whole curve.



TIME HORIZON

The length of task, in human-expert hours, a model completes with 50% reliability.

Doubling every seven months. Recently, every four.



Logarithmic scale. 50% reliability horizon. Plotted from published METR measurements; intermediate points illustrative.

Under an hour to twelve hours, in twelve months.

EARLY 2025

< 1

hr

Frontier models held together for about the length of one meeting.

EARLY 2026

12

hrs

About the length of a full working day, and a half.

Measurements above sixteen hours are flagged unreliable on the current task suite.



What it does half the time today, it
does *reliably* in about a year.

The 80% reliability horizon runs about two doublings behind the 50% horizon.

Capability is jagged, not uniform.


WHERE IT'S STRONG

Well-specified tasks
Self-contained work
Clear success criteria
Anything scoreable

WHERE IT DROPS

Messy, partial inputs
Work judged holistically
Ambiguous goals
Anything relational

It performs worst on exactly our kind of work.



Messy inputs

Fifteen years of notes in four systems and one notebook.



Holistic judgment

No answer key. Only a read of the whole situation.



Relational stakes

Family dynamics, timing, grief, and trust built over decades.




Irreversible gifts

A bequest conversation happens once. There is no second draft.

Gift planning is messy work
judged *holistically*.

The *authority* gap



We ran the study because the sector was arguing without data.

Blinded

No one knew which recommendation came from a person.


340

Nonprofit professionals rated real fundraising decisions.

1,360


Recommendation comparisons across the full public report.

AI's answers came back nearly on par with human experts.



Mean quality rating, five-point scale. Values illustrative of the study's directional finding: near parity in many scenarios.

98%



98%

would reject AI acting without human oversight.

The same professionals who rated its answers nearly as good as an expert's.

Humans led here. AI led there.

HUMANS WERE FAVORED



Ambiguity

Major and planned gifts

Key relationship moments

Donor escalations

AI LED




Performance insights

Predictive pattern-finding


Less donor-proximate work

Portfolio-level triage

Quality and authority are two
separate judgments.



92% of donors say transparency matters.




Donor Perceptions of AI, Koshy & Chappell, 2025 edition.



34%

name deception — AI portraying itself as human —
as the line not to cross.

Drawing the *line*



Prepare. Produce. Protect.

PREPARE • AI CAN LEAD

Before the room

Readiness research with sources

Conversation rehearsal

Plain-language explainers

PRODUCE • AI CAN DRAFT

After the room

Contact reports from notes

Follow-ups and next steps

Stewardship drafts

PROTECT • ONLY YOU

In the room

The ask and the timing

Family and legacy decisions

Grief, conflict, and trust

Two questions, and you have your answer.

1 Is it
reversible?

2 Is a donor in the room at the moment of
decision?


REVERSIBLE • NO DONOR PRESENT

AI can lead.

IRREVERSIBLE • AMBIGUOUS • RELATIONAL

AI informs. You decide.

PREPARE
Smarter



The alternative to rehearsal isn't no rehearsal. It's rehearsing on the *donor*.

We've done that for forty years and never once called it an ethics problem.

The rehearsal prompt, before I run it.

1 Play this composite donor — mid-seventies, fifteen-year annual donor, no children, mortgage-free home, just asked for estate planning materials.

2 Stay in character. Don't be easy on me.

3 At the end, drop character and coach me.

The third instruction matters more than the first two.

LIVE DEMONSTRATION

Voice rehearsal with a composite donor.

A gift officer opens the conversation. Somewhere in the middle, the donor asks how a charitable remainder trust actually works.



Rehearse the conversation, not the person.

ARCHETYPE IN

The composite is built from life stage and role. Never from a donor record.

THE DISCLOSURE TEST

If you'd have to disclose it to the donor, you did it wrong. You don't disclose that you practiced.

ONE DIRECTION ONLY

The synthetic voice may talk to you. It never talks to your donor.

So AI's best contribution to the sacred moments is preparing the person who will be in them.

The explanation, fixed.

DONOR-FRIENDLY EXPLAINER

A charitable remainder trust, in plain words

You move an asset into a trust.

The trust pays you income for life.

What remains comes to the cause you chose.

Plus: five questions to ask your own advisor.




THE GUARDRAIL

Numbers come from your software or the ACGA table.

Never from the chat window.

PRODUCE
Faster



LIVE DEMONSTRATION

Messy notes in. Contact report out.

Fragments, no punctuation, a misspelled name — into a filed report, a follow-up email, and three next steps with dates.



Same output. No donor data.

WHAT YOU TYPED

“Margaret, 76, widow, 1247 Oak Street, gave \$50k in 2019, husband's memorial fund...”


Identifiers in a consumer tool.

WHAT YOU SHOULD TYPE

“Donor A, mid-seventies, widowed, long-tenured, prior six-figure memorial gift...”

Context without data. Same quality of draft.

PROTECT AND PERSONALIZE
Kinder



LIVE DEMONSTRATION

Gratitude in the donor's own words.

Generic thanks first. Then the version that echoes what she told you about why she gives.



Emotion. Identity. Trust. Logic.

Emotion

Why this matters to her at all.

Identity

Who she is when she gives. The step the echo hits.

Trust

That the promise will be kept after her.

Logic

The vehicle, the tax treatment, the paperwork.

Identity is the step we skip. It's also the one AI can help you hit consistently.



LIVE DEMONSTRATION

Watch it be confidently wrong.

One gift annuity rate. One tax threshold. Checked against the real source, on screen.



Three failures you will actually hit.

01

Stale figures, stated confidently

Last year's rate, delivered in this year's voice.

02

Fabricated citations

A plausible source, a real-looking link, nothing behind it.

03

Advice that is subtly wrong

Right vehicle, wrong treatment. The hardest kind to catch.

AI drafts language. It never
sources numbers, law, or rates.

Four red lines. One named human.

01 No AI representing itself as human to a donor.

02 No donor identifiers in consumer tools.

03 No numbers, law, or rates sourced from a model.

04 No AI-only message at a moment of grief, escalation, or solicitation.



One accountable human, named, per workflow.


One sentence on your website, written to the 92%.

“We use AI to prepare and draft. A person on our team reads, edits, and signs everything you receive from us.”

Donors aren't asking for prohibition. They're asking to know.




Automation replaces the
person. Amplification *returns*
the person.



The curve isn't yours to set. The
line is.

Kindsight

Thank *you!*



Get Neurogiving Lab Notes

Brain-based fundraising, in 9 minutes a week.

One research-backed idea each week on why donors give, remember, trust, and return. If I gave you one useful idea today, I send one more every week. Totally free!

www.labnotes.neurogivingbook.com

Cherian Koshy, CFRE, CAP[®], ACFRE

Vice President
Kindsight

 [linkedin.com/in/cheriankoshy](https://www.linkedin.com/in/cheriankoshy)
www.neurogivingbook.com