

CASE STUDY | OCTOBER 2026

AI disclosure and voter engagement

*Findings from 166,552 AI-assisted text
conversations with voters and donors,
across three disclosure modes.*

Vector Political runs voter and donor texting programs on the Convos conversation engine. At a volume few programs reach, we asked one question of every conversation: when the AI is disclosed, do people reply less, opt out more, or stop talking? The answer turned on honesty, not disclosure.

Start reading →

PROGRAMS RUN BY
Vector Political

CONVERSATIONS POWERED BY
Convos

EVIDENCE BASE
166,552 threads, 3 disclosure modes, 201 identity questions
read by hand



166,552

Delivered text threads analyzed across every Vector program on Convo

+0.04

Replies per 100 texts with the state-mandated AI footer vs without (no difference)

54%

Keep talking after an honest "yes, I'm AI" answer, against 9% evasive and 0% silent

0

Disclosure modes that produced an opt-out exodus

THE QUESTION

Every campaign is about to be asked this

States are starting to require that AI-generated political texts say so. Operators have feared disclosure since a 2019 sales call study found that admitting to AI upfront tanked purchases.

Does that fear hold up in a real texting program?

Vector had the data to answer it: 166,552 delivered threads across voter, donor and advocacy programs, every one of them running on Convo. We looked at the three ways disclosure actually happens in the field.

THE FINDINGS

Three modes of disclosure, three answers

MODE 1 | MANDATED FOOTER

The state compliance line appended to texts going to California and Oregon phone numbers. No human promise, no softening.

No effect

6.01 vs 5.97 replies per 100 delivered (n ≈ 123k). Opt-outs up at most 0.55 points, driven by one warm donor program.

MODE 2 | UPFRONT, IN THE MESSAGE

"I'm a digital assistant, but a real person reads every reply." Used by two programs on 45,699 threads.

Higher replies

+4.1 points vs a matched cold voter program; +2.6 points vs matched warm donors. Confounded by client and copy, never A/B tested.

MODE 3 | ONLY WHEN ASKED

About one in ten real conversations reaches "is this a person or AI?" What the agent says next decides everything.


Depends on the answer

Honest answers hold engagement at baseline. Evasion and silence destroy it.

THE LEVER

Answer the question. Every time.

We hand-read all 201 conversations where a voter or donor asked whether they were talking to AI. Asked as a first reply ("who is this?"), continuation matched baseline: 48% vs 49%. The gap opened only in how the agent answered.



Share who kept talking after asking if they were texting with AI. Honest answers held opt-outs to 10%. All 9 hostile threads in the corpus came from a question that went unanswered.

Candor converted, too. **A monthly donor in Vermont set up a recurring gift after the agent confirmed it was AI.** A donor in Maryland wrote "I will donate thru ActBlue" in the same breath as learning he'd been texting a machine. Both gifts verify in the contribution data.

The critique people actually make is not "you're AI." It is "you should have told me sooner." Upfront disclosure does not pre-empt the question (identity mentions ran 2.7 vs 0.6 per thousand on disclosed sends), but it turns the question into confirmation rather than a reveal.

ONE CONVERSATION

The whole study in three texts

JEAN, MICHIGAN

Is a real person typing these questions or is it AI?

CONVOS AGENT, FOR THE CAMPAIGN

It's AI. A real person on the team reads every reply, and I'll be honest with you about anything you ask.

JEAN, MICHIGAN

You should have identified yourself upfront. I still would have responded had I known. You professed honesty.

Jean kept talking. She closed the conversation with "You're welcome Ms Machine" and never opted out. Engagement survived the disclosure. What she punished was its lateness.

"They provide AI to talk to them. That's the kindness of our Democrats."

Terri Rae, donor, on learning she was texting with an AI agent. The deepest conversations in the corpus belong to people who knew it was AI from the start.

WHAT TO DO WITH IT

Four rules for your next program

- 1 Comply without fear.** The mandated footer costs nothing measurable in replies and does not trigger opt-outs. Turn it on everywhere it's required and stop treating it as a tax on performance.
- 2 Disclose upfront, and pair it with the human promise.** The programs that said "I'm a digital assistant, but a real person reads every reply" outperformed matched peers. The promise alone did not move replies; the pairing did.
- 3 Never leave "is this AI?" unanswered.** Every Convos agent brief now carries the rule: honest on ask, never evasive, never silent. Silence is the one agent behavior that manufactures hostility.
- 4 Track identity questions per thousand threads as a quality metric.** The corpus average is 1.0; one AI-skeptic donor base ran 5.0. When the number climbs, the copy or the audience needs attention before the volume does.

ABOUT VECTOR POLITICAL

Vector helps campaigns and causes find and talk with their audience at any scale, turning what people say into conversational intelligence about why they care, give, vote and act. Vector's texting programs run on the Convos conversation engine.

ABOUT CONVOS

Convos is a controlled AI texting platform that turns every conversation into clear, actionable campaign intelligence. Real-time tagging, alignment and sentiment tracking, and AI replies kept inside campaign-approved messaging. Clients run the platform themselves; the founders advise when asked.

convos with **Vector Political**

Tom Carroll, Co-founder, Convos

tom@pubsent.com | convos.ai

Marty Santalucia, Vector Political

marty@vectorpolitical.com | vectorpolitical.com

Built for dialogue.

Method: three independent analyses of one conversation corpus (state footer, in-body, reactive), cross-checked on thread universes and detectors. All results observational. Bump threads (82% of raw replies) excluded. Reply rates per 100 delivered texts, within program. Program and client names withheld.