A GUIDE FOR ORGANIZATIONS THAT RECRUIT PATIENTS INTO CLINICAL RESEARCH ONLINE

Detecting Bot Attacks on Patient Research Studies

For patient communities, academic and clinical researchers, and contract research organizations (CROs) that bring patients into studies online.

Patient communities are a digital frontline

Health care and public health drew more ransomware complaints to the FBI in 2025 than any other critical infrastructure sector.[1] Hospital attacks make headlines. Bot attacks on patient research rarely do, yet they reach the same communities through a different door: a survey link built to welcome patients into a study gets found by scripted bots and paid fraudsters, and fake answers travel into the datasets that shape care.

Two cancer organizations found bots in their patient studies, and nobody to call

Research studies from patient organizations, like LUNGevity Foundation and Cancer Support Community, have faced bot attacks. The data cleanup required a huge investment of time and resources, and when the organizations went looking for help, no playbook existed. This guide is what they learned about bots and human fraudsters, written down so the next patient organization starts further ahead.

> "Even after diverting our resources to this cause, we ran into a wall because we, as a single organization, could only do so much and we didn't know where to turn for help."
>
> Erica Fortune, Vice President of Research, Cancer Support Community[3]

How to use this guide against bots in patient research

2 **Take the bot evidence to your board**
A 2025 review of 23 online health studies shows how common fake participants have become. Use it to budget for prevention, screening and cleanup before your next study.

3 **Catch a bot wave in your patient study**
Seven warning signs drawn out, a three-question test, and eight steps in order. Keep it next to whoever reviews incoming survey data.

4 **Set up the next patient study to catch bots**
Eight steps before your patient study opens, questions for whoever holds your data, and where to turn.

5 **Get funders, IRBs and journals to help**
What a bot attack costs a patient organization, and what each institution can do about it.

If you recruit patients online, these bots are your problem too

You don't need a security background to catch bots. You need to know your patient community, your survey data and what a normal week looks like.



Registry and study managers

You hold the patient data and remove the bot records.



Academic, clinical and program researchers

You design the patient survey, set the incentive and write up the findings.



Community managers

You share the study where patients gather online, so you see the bot wave first.



CRO recruitment teams

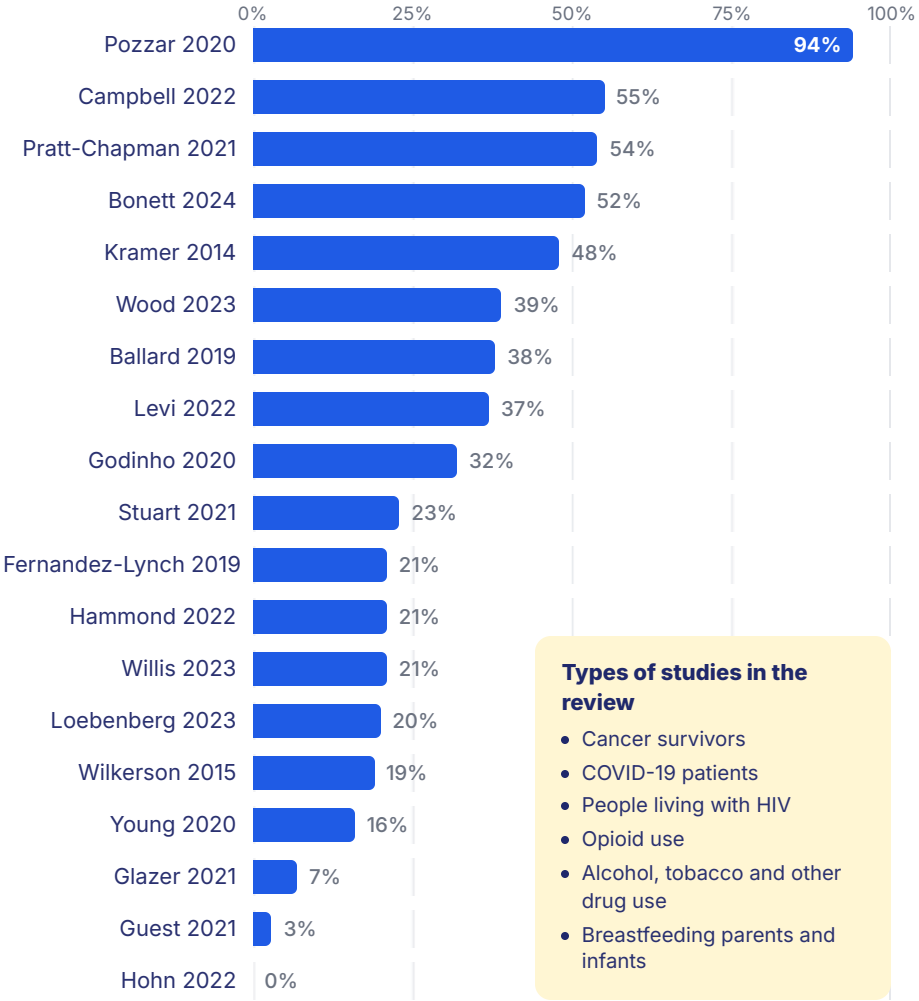
You recruit patients for sponsors at scale through the same public channels.

Nearly every online health study that looked for bots found them

A 2025 scoping review in BMJ Evidence-Based Medicine examined 23 health studies that recruited participants online. 18 reported fraudulent responses, and in some, fake answers outnumbered real patients.[2]

Share of each study's survey responses thrown out as fraudulent or suspicious

Each bar represents one of the 19 studies that reported how many responses were suspicious, and the length of the bar is the percent of that study's responses the authors threw out: 32% means about one response in three. In the study at the top, 256 of 271 responses were excluded. Four more studies in the review reported no count.



Types of studies in the review

- Cancer survivors
- COVID-19 patients
- People living with HIV
- Opioid use
- Alcohol, tobacco and other drug use
- Breastfeeding parents and infants

What drew the bots: paid incentives and public links

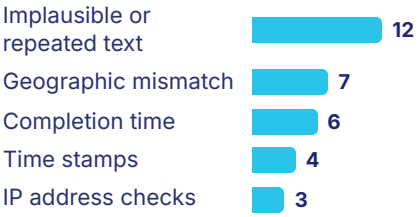
18 of 23 online health studies reported excluding fraudulent responses. One study looked and found none, and four reported no count.

17 of 23 paid participants. Incentives ran $5 to $200, and one study found that higher payments drew more fraud. Fraud showed up without pay too: a study that offered no incentive still excluded 23% of its responses.

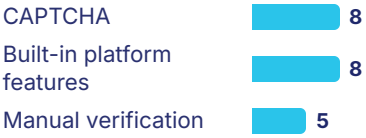
11 of 19 of those naming a channel recruited patients through social media

How studies spotted the bots

Studies using each method, of 23



What researchers tried against bots



No single method stopped fake participants. CAPTCHA "cannot stop fraudulent responses from actual people," the authors write.

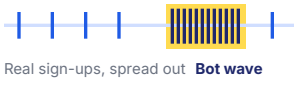
Recent bot attacks in clinical research, undercounted

22 of the 23 studies were published in the five years before the review's April 2024 search, and its authors call fake participants a growing problem: "Increasingly, the risk of fraudulent responses presents a significant challenge." Every figure here comes from a study that looked for bots and reported what it found. Studies that never checked their patient data are missing from the count.

How to tell a bot wave from ordinary noise in patient survey data

Every patient study collects a few careless answers, and a few bots that slip past cleaning are noise. A bot wave is different: it has a shape you can draw, and this page draws it.

Seven warning signs of a bot wave in patient survey data



Survey responses arrive in bursts

Bot responses land minutes or seconds apart, often overnight, over holidays, or within hours of a public post.



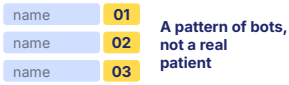
Sign-ups jump far above a normal week

Bots push sign-ups far above what your patient community produces in a typical week.



Open-ended answers repeat the same text

Bot answers reuse the same chunks in a new order, or read as fluent and say nothing a real patient would say.



Names and emails follow a bot pattern

Real patient data is never tidy. Templated names or emails, evenly spread answers, or IP addresses far from your patients all point to bots.



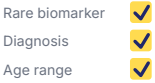
The clinical survey responses don't add up

Stage, treatment and timeline contradict each other. Surveys with few clinical questions give fewer such clues, so lean on timing and text.



Surveys finish faster than any real patient could

Bot completion times fall well under the benchmark you set in your pilot.



Answers fit the study's eligibility too well

Rare biomarkers, every criterion met, or demographics unusual for an online survey. Real patients rarely match the recruitment ad line for line.

Steps to protect your patient study from bots

1. **Pause the patient recruitment link**
   Take down the post or ad. Real patients can return later through a new link; a compromised link stays compromised, and repeat attacks follow.
2. **Audit every clinical survey entry**
   Copy flagged entries, with their time stamps, to a separate file; deleting now destroys the evidence that confirms a bot pattern. In a longitudinal study, also block flagged accounts from the platform.
3. **Confirm the bot pattern with a second signal**
   A time cluster plus repeated text, for example. One sign is a hunch; two together make a pattern.
4. **Hold incentive payments on flagged entries**
   Verify before you pay. Confirming a participant through a second channel such as phone may need IRB approval first.
5. **Alert whoever holds the patient data the same day**
   Your contractor, platform or research partner may see bot signals you can't. Ask what they see on their side.
6. **Keep a log of clinical survey results**
   Dates, counts, bot signals, decisions. You will need the log for cleaning, for partners and for anyone you report to.
7. **Report confirmed bots to your IRB**
   Tell the board that reviews your study what you found and how you are handling it, so the record of the study stays honest.
8. **Remove bot records with a rule you can rerun**
   Take out confirmed bot entries under a documented rule, keep a list of removed IDs, and apply the same rule to every later export.

A clean patient dataset, and a record of how you got there

Noise or bot attack? Three questions about your patient survey data

Clustered in time?
Did dozens arrive within minutes, overnight, or right after a public post? Real patients trickle in over days.

Patterned?
Do entries share wording, email formats, locations or clinical details that no real group of patients would?

Above your baseline?
Are there more responses than you have seen before with similar groups, in your last study or last month?

One signal: flag the records and keep watching, since a newsletter blast brings a burst of real sign-ups and one support group can describe its care in similar words. **Two or more signals:** treat it as a bot attack and start the eight steps; timing, pattern and volume together separate an attack from noise. **Unsure?** Pause the link and audit before you decide.

PREPARE YOUR NEXT PATIENT STUDY

Plan for bots before you recruit patients online

Protecting a patient study from bots costs planning, staff time and a few survey platform settings.

Eight steps before your patient study opens

1. **Budget for bot screening and cleanup**
   Put screening and cleanup time in the study plan and in the grant.

2. **Design patient incentives with bots in mind**
   Keep amounts modest, leave the amount out of public posts, pay only verified participants, and say up front that fraudulent responses go unpaid.[2]

3. **Pilot the survey to learn what real patients look like**
   Use the pilot to set a completion-time benchmark and learn what normal weekly volume looks like.

4. **Use a screener and unique links, never an open survey link**
   Screen interested patients first, then send each eligible person a unique, single-use link. A public link can be copied anywhere, and once bots find it they keep coming.

5. **Turn on the survey platform's bot protections**
   Block repeat submissions, keep the survey out of search engine indexing, and treat CAPTCHA as one check among several.[2]

6. **Decide who reads incoming patient data, and how often**
   Weekly review by someone who knows the community, even when a contractor or partner holds the data.

7. **Agree with your IRB on how you will report bots**
   Put the screening plan in the protocol, and settle in advance how a confirmed bot wave gets reported and how recruitment can pause.

8. **Check the patient studies you already published for bots**
   Rerun the warning signs on past datasets. If bot records reached a published finding or a recruitment list, tell your research partners, your IRB and the journal so they can correct it.

Open recruitment with unique links and no public survey link

Where to turn today if bots hit your patient study

For now, patient organizations handle bot attacks on their own. First calls: your survey platform's support team, your research partner's data team, and peer organizations that have been through it. Share what you learn with them, so the next team catches a bot wave sooner.

Ask whoever holds your patient survey data

- How often do you screen for fraudulent responses, and which signals do you check?
- When you find something, how and how fast will you tell us?
- Can we see raw responses every week?
- How are flagged records handled, and who decides what gets removed?
- Who pays for cleanup, and is it written into the contract?

More ways to keep your patient community safe online

The Light Collective Online Safety Toolkit
Practical steps for patients and for the people who run patient communities.
lightcollective.org/safety-toolkit/

Cyber Safety Is Patient Safety
A guide by patients, for patients, from The Light Collective and the Health Sector Coordinating Council, October 2026.

A fraud screening checklist for online recruitment
The detection, prevention and verification steps collected by the 2025 review, source 2.
ebm.bmj.com/content/30/3/173

Stay tuned for more on bots in patient research
The Light Collective will add resources for patient research teams as we find support for this work.

Patient organizations are facing bot attacks without support

Research studies from patient organizations, like LUNGevity Foundation and Cancer Support Community, have faced bot attacks. The data cleanup required a huge investment of time and resources, with nobody waiting to help.

What a bot attack costs a patient organization

A bot wave hits a patient organization three times. Incentive money goes to fake participants. Staff who were hired to run support and other programs for the community spend weeks screening survey entries and rewriting cleaning rules. Research built with and for patients slows down, and findings that already went out may need correcting. These organizations have resources. What they lack is a budget line, a tool or a partner that treats bot defense as part of the job, so every dollar and hour comes out of work meant for people with cancer.

Where support for patient research should come from, and doesn't

Survey platforms leave bot detection to the customer. Funders rarely budget for screening when they pay for recruitment. No federal agency has taken responsibility for attacks on research recruitment, so an organization that catches a bot wave has nowhere obvious to report it. Institutional review boards (IRBs), the committees that approve and oversee studies, and the journals that publish results have not yet said what they expect when bots reach patient data. Each of them can change that.

What funders, platforms, agencies, IRBs, journals and peer organizations can do about bots in patient research

FUNDERS

- Fund bot screening and cleanup in every grant that pays for patient recruitment.
- Treat a documented bot attack as a reason to extend a timeline, never as a mark against the team.

SURVEY PLATFORMS

- Build fraud detection into the tools patient organizations already use, and turn it on by default.
- Give small research teams the same bot signals large customers get: time stamps, IP checks, duplicate detection.

AGENCIES AND RESEARCH SPONSORS

- Create one place where organizations can report attacks on research recruitment and get help.
- Publish what a bot wave looks like, so patient organizations hear about new patterns before the next study launches.

INSTITUTIONAL REVIEW BOARDS

- Ask every protocol that recruits patients online for a bot screening plan.
- Set a clear, blame-free path for reporting a confirmed bot wave, with fast approval to pause recruitment or add verification.

PEER-REVIEWED JOURNALS

- Ask authors of online-recruited studies to report how they screened for bots and how many entries they excluded.
- Publish corrections when bot records turn up in published patient data, and credit the authors who disclose them.

PEER PATIENT ORGANIZATIONS

- Share bot patterns and cleanup rules with other organizations running patient research.
- Ask research partners and CROs what bot screening they run before lending your community's trust to a study.

Acknowledgements

Andrea Downing
Co-founder and Executive Director, The Light Collective

Erica Fortune
Vice President of Research, Cancer Support Community

Bellinda King-Kallimanis, PhD
Vice President of Patient-Focused Research, LUNGevity Foundation

Suggested citation: The Light Collective, LUNGevity Foundation, Cancer Support Community. Detecting Bot Attacks on Patient Research Studies. October 2026.

Citations

1. Federal Bureau of Investigation, Internet Crime Complaint Center. *Internet Crime Report 2025.* April 2026. Health care and public health ranked first among critical infrastructure sectors for ransomware complaints; complaints are reports, not a full count of attacks. ic3.gov/AnnualReport/Reports

2. Comachio J, Poulsen A, Bamgboje-Ayodele A, et al. Identifying and counteracting fraudulent responses in online recruitment for health research: a scoping review. *BMJ Evid Based Med.* 2025;30(3):173-182. ebm.bmj.com/content/30/3/173

3. Erica Fortune, Vice President of Research, Cancer Support Community. Comments to The Light Collective, October 2026.