ONE-PAGER

# Should Fraud Detection Settings Be Adjustable?

## TL;DR

A visitor is either fraudulent or they are real. There is no in-between. Yet some fraud and bot detection tools let you adjust their sensitivity or thresholds, a setting that shifts the error rate in one direction or the other rather than removing it. When a vendor hands you that dial, they are handing you the job you paid them to do.

## Fraud Doesn't Come in Degrees

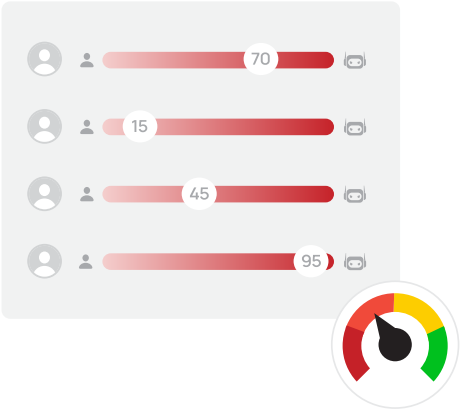


At the moment a click happens, a form gets submitted, or a login is attempted, there is exactly one true answer underneath it. Either the traffic is coming from a real, present human who intends to interact with your business, or it isn't.

The industry doesn't always treat it that way. A number of detection tools score traffic on a scale, then hand you a dial and ask you to decide where the line between "acceptable" and "block this" should sit.

> **Fraud isn't a spectrum. It's a binary condition: real or fake.**

## What "Adjustable Sensitivity" Actually Does

Look closely at how these adjustable settings are documented, and a pattern shows up. Vendors will let you raise or lower strictness levels, click thresholds, or risk cutoffs, and their own public documentation is upfront about the tradeoff. One widely used IP risk scoring tool offers four strictness levels for its fraud scoring and states plainly that the higher levels carry a greater risk of false positives, recommending that customers start at the lowest setting and only raise it if needed. A separate click fraud tool documents that setting its threshold to block after a single click will flag every click that IP generates, genuine or fraudulent alike, and cautions that this setting isn't recommended for most businesses.

These aren't buried admissions. They're written directly into public help docs. The vendors know that turning the dial in either direction creates the opposite problem. That's the tell. Adjustable sensitivity doesn't make fraud detection more accurate. It just moves the error from one column to the other and asks you to pick which column you'd rather live with.

## Whose Job Is It Really?

If you bought a smoke detector, you wouldn't want a dial on the side that reads "turn this up to catch more smoke, but expect more false alarms; turn it down to cut false alarms, but you might miss a real fire." You'd want a device that reliably tells you: fire, or no fire. That's the entire point of paying for the device.

Fraud detection should work the same way. A vendor's whole value proposition is that they have the data, the models, and the expertise to make the call so you don't have to. When they hand you a lever instead of a verdict, the outcome, and the responsibility for getting it wrong, quietly shifts back onto you.

**How can you make confident decisions when the answer can be adjusted?**

Anura.io • (888) 337-0641 • 114 Sandhill Drive, Suite 104, Middletown, DE 19709

Anura

That's not a customization feature. It's a built-in excuse. If the tool blocks a paying customer, the vendor can point to the setting you chose. If fraud slips through, same thing. Either way, the tool is never wrong, because you made the call. You're paying for a decision and getting a dial instead.

## Industry and Traffic Source Shouldn't Change the Answer

Some vendors argue sensitivity needs to flex by industry: e-commerce sees more repeat visitors, lead generation runs on urgency, on-demand services can't afford a single wasted click. But the industry a fraudster targets doesn't change what they are. A bot is a bot whether it's clicking a shoe ad or a mortgage lead. A real person is a real person whether they're buying a coffee or filling out a contact form. Tuning a threshold by industry doesn't make the visitor underneath it any more or less fraudulent. It only changes how much of each type of traffic your tool has been told to let through.

> **Fraud doesn't change.**
>
> **The standard for detecting it doesn't either.**

## What a Real Verdict Looks Like

Fraud detection that actually does its job doesn't ask you to guess. It validates the environment behind the visitor, the hardware, software, IP, and device characteristics that a real fraudster has to hide or manipulate to look legitimate, and returns a clear answer. No slider. No strictness level. No industry preset to second guess. A verdict of fraudulent should only ever be issued when the detection is certain, because a tool that isn't certain has no business making you split the difference.

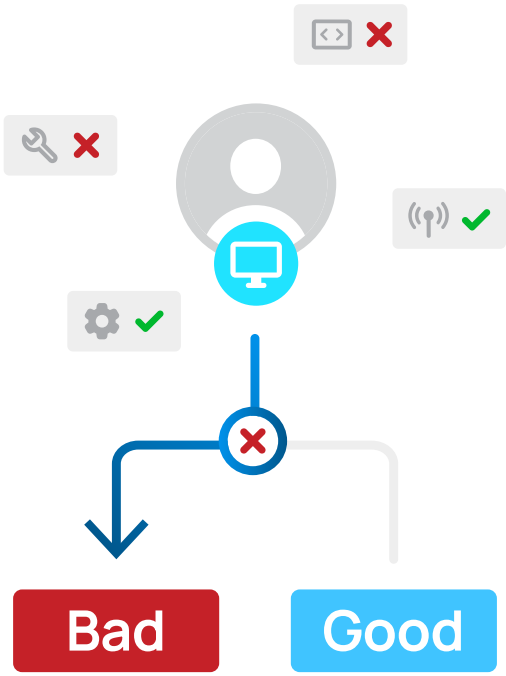


## You're Trusting This Solution With Your Marketing Budget

Step back from the mechanics for a second and think about what's actually at stake. Every dollar of ad spend, every lead form, every login assumes the traffic behind it is real. That's not a small thing to hand off to a vendor. It's your budget, your pipeline, and your customer trust riding on the answer.

So ask the question plainly: should you trust a solution that hands you a dial and lets you play guessing games with that budget, or one that does the work and gives you a straight answer? A tool that asks you to keep tuning strictness levels and thresholds isn't managing your risk. It's asking you to manage its uncertainty for it, with your own marketing dollars as the test case. That's a guessing game dressed up as a settings page.

## The Bottom Line

You don't pay a fraud detection vendor for a dashboard full of knobs. You pay them for confidence. If a solution needs you to decide how many real customers you're willing to sacrifice, or how much fraud you're willing to tolerate, it hasn't actually solved the problem you hired it to solve. It's handed the hardest part of the job, and the risk to your budget, back to you and called it flexibility.

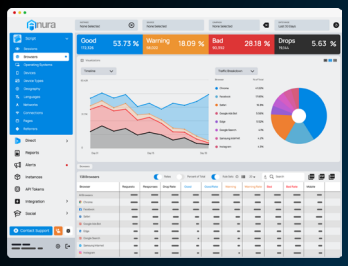

## Get a Definitive Answer on If your traffic is Real or Fraud

*Get a clear, binary answer on how much of your traffic is actually invalid and how much it is impacting your campaigns.*

*No obligation. No estimates. Just real data from your real traffic.*

GET A FREE TRAFFIC QUALITY AUDIT