Innocent till proven guilty?

Omxus Research March 2026

Forcibly kidnapped. Yelled at. Treated like vermin. Your clothes, shoes, items taken. Locked in a room unfit for a lab rat.

What this is

This is not an academic paper. The academic version exists — it's Paper 11 in this series, with 67 references and five original statistical analyses. This is the kitchen table version. This is what it's like to be inside the machine.

If you've been through it, you already know everything in here. You don't need the statistics. You lived it.

If you haven't been through it, you won't believe me. That's the point.

The experience

You walk into a room. Someone with authority asks you a question. You answer it honestly.
They don't believe you.

You try again. You rephrase. You add detail, because maybe they didn't understand. You stumble over your words because you're scared and you're trying to remember exactly what happened and you want to get it right.

They believe you less.

You notice they're watching your eyes. You've heard somewhere that liars don't make eye contact, so, you try to look at them directly; to show you're telling the truth.

Now you look rehearsed. Scripted. Like you're performing.

So, you look away, because it's natural and you're uncomfortable and you're trying to think.

Now you look like you're hiding something.
You start to cry because this is terrifying and no one is hearing you.
They note that you're "emotionally manipulative."
You stop crying and try to be calm and factual.
They note that you're "flat" and "lacking appropriate affect."
You say "I didn't do this."

They hear: "I haven't confessed yet."

The trap

Here is the thing no one tells you: there is nothing you can say.
Not because they're evil. Not because there's a conspiracy. Because the system is designed — architecturally, procedurally, from the ground up—so that no statement you make can be interpreted as evidence that you're innocent.

- If you confess: guilty.
- If you deny: guilty, but resistant.
- If you're calm: cold, lacking remorse.
- If you're emotional: manipulative, performing.
- If you look them in the eye: rehearsed, defiant.
- If you look away: evasive, hiding something.
- If you stumble over your words: nervous because you're lying.
- If you speak fluently: rehearsed because you're lying.
- If you give lots of detail: over-explaining, trying too hard.
- If you're brief: withholding, uncooperative.

Read that list again. Every pair is a double bind. For every possible behaviour, the system has a guilt-consistent interpretation ready. There is no box on the form for "this person is telling the truth." The form doesn't have that box.

This is not a failure of the system. This is the system.

The research says you're right

Everything you experienced in that room — the disbelief, the double binds, the feeling that nothing

you said mattered — is documented in the peer-reviewed literature. Not as an aberration. As the normal operating condition.

They can't tell. Human beings detect lies at 54% accuracy. That's 4% better than flipping a coin. This isn't one study — it's a meta-analysis of 247 studies involving 24,483 people (Bond & DePaulo, 2006). Trained police officers, judges, psychiatrists — none of them do significantly better than random members of the public. Training doesn't improve accuracy. It improves confidence. The investigator who assessed you was not more accurate than a stranger on the street. They were more certain.

The cues are backwards. 91.3% of the behavioural cues people use to assess credibility are inverted (our original analysis, $p < 0.0001$). The things you did because you were telling the truth — stumbling, hedging, looking away, self-correcting — are the things the system reads as deception. The things a liar does — speaking fluently, maintaining eye contact, delivering a rehearsed narrative with confidence — are the things the system reads as honesty.

You were penalised for telling the truth. A liar would have been rewarded.

The stumbling proves you were honest. Truth-tellers are significantly more disfluent than liars ($d = 0.60$, $p = 0.004$, $N = 121$ real courtroom transcripts). When you stumbled, it was because you were doing the hard cognitive work of accessing real memories in real time. A liar doesn't stumble — they've rehearsed. Your stumbling was the single most reliable indicator that you were telling the truth. And it was used against you.

Your eye contact was a trap. Gaze aversion is the #1 believed indicator of deception worldwide — 63.7% of people in 75 countries believe that liars avoid eye contact (Global Deception Research Team, 2006). The actual relationship between gaze aversion and deception? $d = 0.05$. Essentially zero. It means nothing. But you didn't know that. So when someone yelled at you to look them in the eyes, you did. And then you looked "rehearsed." And when you looked away because it was natural, you looked "evasive." The cue carries no information. The belief in the cue carries all the power.

If you're autistic or neurodivergent — and many people in this situation are — the trap is even worse. Autistic people are rated as more deceptive and less credible when telling the truth (Lim et al., 2021; $N = 1,410$). The diagnostic criteria for autism overlap almost perfectly with the behavioural cues the system reads as deception: gaze aversion, flat affect, fidgeting, unusual prosody. An autistic person cannot present their truthful testimony without involuntarily triggering every alarm the system has. Their innocence is structurally illegible to the instrument.

The confession machine

If you didn't confess, they treated your denial as an obstacle, not as information. The dominant interrogation method in the English-speaking world — the Reid Technique — has a step called "Overcoming Denials." Step 4. It's in the manual. The denial is not evaluated. It is overcome.
The technique starts by assessing whether you're lying using the inverted cues described above. Since truthful behaviour triggers the deception indicators, innocent people routinely fail this assessment.

You failed it because you were telling the truth.

Once you've been assessed as deceptive, the interrogation proper begins. You are confronted with your "guilt" as an established fact. Your denials are suppressed. You are offered minimising themes — "We know something happened. We just want to understand why." The alternative question embeds guilt as a presupposition: "Did you do it because you were angry, or because you lost control?" Both answers are confessions.

If you confess, your confession becomes the most powerful piece of evidence in the case. More powerful than eyewitness testimony. More powerful than physical evidence. More powerful than anything (Kassin & Neumann, 1997). And it doesn't just add to the evidence — it corrupts everything else. Forensic examiners who learn about a confession reinterpret ambiguous findings as incriminating. Alibi witnesses lose confidence. Eyewitnesses get more certain. The confession restructures reality around itself (Kassin, 2012).

If you don't confess — if you maintain your innocence — nothing changes. The denial is not evidence of innocence. It's evidence that the interrogation hasn't worked yet.
Confessing proves you're guilty. Denying proves you're guilty but haven't confessed yet.

The numbers

Forty out of the first 250 people exonerated by DNA evidence in the United States had confessed in detail to crimes they did not commit (Garrett, 2011). Not vague admissions. Detailed, factually 3accurate confessions — because the facts were fed to them by the interrogators and then presented as knowledge “only the perpetrator could have known.”

Of false confessors whose cases went to trial, 73% were convicted (Leo & Ofshe, 1998). The confession, once it exists, is almost impossible to overcome — even when the confessor recants, even when there is exculpatory evidence, even when the confession was plainly coerced.

False confessions account for 12-30% of documented exonerations (across multiple datasets: Kassin & Gudjonsson, 2004; Gross et al., 2005; Garrett, 2011; NRE, 2023). These are the ones we know about — cases where DNA or other evidence eventually proved innocence. The ones we don’t know about are still in prison.

The conditions

We keep the people this system convicts in conditions that would be illegal for laboratory rats.
That is not metaphor. The regulations governing the housing of laboratory animals specify minimum space requirements, environmental enrichment, social contact, natural light cycles, and freedom from unnecessary distress. Many prison systems fail to meet these standards. Solitary confinement — which constitutes torture under international law when applied for more than 15 days (UN Mandela Rules) — is routinely applied for months or years.

Australia’s recidivism rate is 45% within two years, rising to 60-70% over five years. The United States reaches 76-83%. Norway, which treats incarcerated people with dignity — private rooms, education, therapeutic services, maintained family contact — achieves 20%.

The punitive system does not reduce reoffending. It increases it. The conditions damage prefrontal cortical function, destroy social connections, and expose people to criminal networks. The system produces the recidivism that justifies its own continuation.

It can’t happen to you

That’s what you think. You’re educated. You’re respected. You have a career, a reputation, relationships in the community. You’ve worked with the system — maybe even for it. You know how it works. You know your rights. You’d never confess to something you didn’t do.
That’s what everyone who goes through this thought before it happened to them.
It doesn’t matter who you are. It doesn’t matter how respected you are, how articulate, how connected. The machine doesn’t process identity. It processes behaviour. And if your truthful behaviour triggers the inverted cues — which it will, because that’s what truthful behaviour does — the machine will read you as deceptive regardless of your CV, your community standing, or your history of working alongside the very system that is now processing you.

In fact, your knowledge may work against you. If you know you’re being assessed, you’ll try to appear credible. And trying to appear credible is performance. And performance is what deception looks like.

The people who designed the system did not exempt themselves from its architecture. The machine has no “trusted source” category. It has a single input: behaviour. And the behaviour it reads as 4honest is the behaviour of liars.

What it felt like

It felt like being kidnapped by someone who had all the power and none of the interest in what was true.
It felt like sitting across from someone who had already decided, and everything you said was being fed into a machine that could only produce one output.

It felt like the person across the table was a psychopath — not because they were cruel, but because they were incapable of updating their beliefs in response to new information. They had a theory. Your job was to confirm it. If you didn’t confirm it, that was your problem, not theirs.

It felt state-sanctioned. Because it was. The person across the table had credentials. A uniform or a title. The room had recording equipment. There were procedures being followed. It all looked legitimate. It looked like justice. It felt like abduction.

The worst part wasn't the accusation. The worst part was the realisation that nothing you could say would change what was happening. That the system was not listening. Not because the individual was bad — maybe they were, maybe they weren't — but because the system does not have a mechanism for hearing innocence. It has mechanisms for processing guilt. Innocence is not a category the machinery handles.

It's not fixable

This is the part that people don't want to hear.

You might think: OK, the cues are wrong. So fix the cues. Recalibrate. Train people on the right indicators.

Here's why that doesn't work: the moment you publish the correct indicators, everyone learns them. Liars learn to fake disfluency. Truth-tellers learn to "perform" their truth. The published cue's diagnostic value collapses to zero. You cannot have a public-knowledge detection system in an adversarial context. The act of fixing the instrument destroys the fix.
This is not a limitation of current science. It is a mathematical certainty. A signal whose encoding is public knowledge carries zero information when both parties have access to the codebook. The cues are dead. No future research can resurrect them.

Behavioural credibility assessment cannot be reformed. It can only be abolished.

What this means

Every conviction that relied on an investigator's assessment of credibility — "I could tell the suspect was lying" — is epistemically void.

Every conviction that relied on a jury's assessment of demeanour — "the witness appeared evasive" — is epistemically void.

Every conviction that relied on a confession extracted after the suspect's truthful denials were "overcome" — is tainted by a process that cannot distinguish truth from deception and is not designed to try.

This doesn't mean every convicted person is innocent. It means the instrument that determined their guilt is broken, and broken in a specific direction: it penalises truthful behaviour and rewards deceptive behaviour. Some convicted people are guilty despite the instrument, not because of it. But some of them — we don't know how many, but we know it's not zero, and the evidence suggests it's a lot — are in prison because they told the truth in a room where truth sounds like lying.

The question

The evidence for everything in this paper has been in the peer-reviewed literature for decades. Bond and DePaulo published in 2006. The Global Deception Research Team published in 2006. Kassin and Gudjonsson published in 2004. Garrett published in 2011.

The people who design and operate the justice system have had access to this evidence for twenty years. They have continued to operate a system that assesses credibility using inverted cues, extracts confessions through processes known to produce false confessions, and treats denials as obstacles rather than information.
The question is not whether the system is broken. The research settled that.

The question is what we do with the people who knew it was broken and kept it running.

References

Bond, C. F., Jr., & DePaulo, B. M. (2006). Accuracy of deception judgments. Personality and Social Psychology Review, 10(3), 214-234.
DePaulo, B. M., et al. (2003). Cues to deception. Psychological Bulletin, 129(1), 74-118.
Garrett, B. L. (2011). Convicting the innocent: Where criminal prosecutions go wrong. Harvard University Press.
Global Deception Research Team (2006). A world of lies. Journal of Cross-Cultural Psychology, 37(1), 60-74.
Gross, S. R., et al. (2005). Exonerations in the United States 1989 through 2003. Journal of Criminal Law and Criminology, 95(2), 523-560.
Kassin, S. M. (2012). Why confessions trump innocence. American Psychologist, 67(6), 431-445.
Kassin, S. M. (2017). False confessions: Causes, consequences, and implications for reform. Policy Insights from the Behavioral and Brain Sciences, 4(2), 161-168.
Kassin, S. M., & Gudjonsson, G. H. (2004). The psychology of confessions. Psychological Science in the Public Interest, 5(2), 33-67.
Kassin, S. M., & Neumann, K. (1997). On the power of confession evidence. Law and Human Behavior, 21(5), 469-484.
6Leo, R. A., & Ofshe, R. J. (1998). The consequences of false confessions. Journal of Criminal Law and Criminology, 88(2), 429-496.
Lim, A., Young, R. L., & Brewer, N. (2021). Autistic adults may be erroneously perceived as deceptive. Journal of Autism and Developmental Disorders, 52(2), 490-507.
Perez-Rosas, V., et al. (2015). Deception detection using real-life trial data. Proceedings of the ACM International Conference on Multimodal Interaction, 59-66.
Rizzelli, L. (2021). The language of criminal confessions. The Wrongful Conviction Law Review, 2(3).

This paper is part of the OMXUS research series. The academic version is Signal Inversion.
The statistical appendix is in RESEARCH_FACTS.md. The trap map is in TRAP_MAP.md.
If you've been through this and you're reading this: you weren't crazy. The research confirms what you already knew. They couldn't tell. They never could.