HUMANITY LABS AI MAGAZINES

The Trust Times

ISSUE NO. 5



Starting It

IN THIS ISSUE

1. Show Your Face
2. The Candidate Who Confessed
3. Follow-Through Beats First Impressions

Current, peer-reviewed studies on the science of winning a stranger's trust.

Issue #4 asked what happens to trust when AI shows up at your desk. This issue goes back further, to someone who has never met you deciding whether you are worth a meeting. What moves that decision, in all three studies, is evidence that you are a specific person who can be held to account.

The first, from the *International Review of Social Psychology*, put a number on what a profile photo is worth. Four studies found that sellers who withhold a profile photo are trusted less and avoided, and that in a real market a visible face commanded 9 to 11 percent higher prices. An unflattering photo still beat no photo. Strangers do not need you photogenic. They need you visible.

The second, from *Frontiers in Psychology*, staged a job interview in which a candidate was confronted with a past mistake and either owned it or blamed her old boss. The study showed it to 204 Swiss employees. The confession earned far more trust. Later good news about her lifted both versions equally, so the confession stayed ahead by the same margin.

The third, from the *British Journal of Psychology*, pooled 68 studies of a money-exchange game played round after round with the same partner, the first time anyone has meta-analyzed trust learning in that setting. Group labels and reputation for skill barely mattered next to conduct. Reputation for character was different: warned players sent less money, so they never gathered the evidence that might have changed their mind.

I picked these three because none of them costs anything to act on. The hard part is that you cannot buy any of it, or hand it to someone else to do for you.

Three more great studies on trust next week.

Warmly,



Marshall Kirkpatrick
Editor, *The Trust Times*

SOURCES · ISSUE 5

1. Show Your Face. Based on Jaeger & Efendic (2025), *International Review of Social Psychology*. N=380.

2. The Candidate Who Confessed. Based on Maliqi, Lalot & Quiamzade (2025), *Frontiers in Psychology*. N=204.

3. Follow-Through Beats First Impressions. Based on Duncan, Sganzerla, Kaltwasser & Dziobek (2025), *British Journal of Psychology*. N=8,000+ across 68 studies.

ABOUT THIS ISSUE

Three peer-reviewed studies on how trust is built, retold in plain language. The "at your desk" sections are our editorial application, not claims made by the researchers.

How trust begins: the face that wins the first look, the confession that survives the first hard question, and the follow-through that outweighs both.

In Issue #4 we covered three studies on client trust in the age of AI.

A HUMANITY LABS PUBLICATION

Show Your Face

Four studies put a price on hiding: strangers trust, choose, and pay more for the seller they can see.

Jaeger, B., & Efendic, E. (2025). The cost of anonymity in the sharing economy: Consumers distrust and avoid sellers without profile photos. *International Review of Social Psychology*, 38(1). Open access: doi:10.5334/irsp.991

THE LISTING HAD EVERYTHING A CAUTIOUS TRAVELER could want: a private room in Amsterdam, a friendly paragraph about the host's interests, a healthy count of references from past guests. One thing was missing. Where the host's face should have been sat a white silhouette on a gray field, the shape of a person with nobody in it.

Researchers Bastian Jaeger and Emir Efendic built dozens of listings like that one and measured what the missing face did. Across four studies in the *International Review of Social Psychology*, the answer held steady: no face, less trust.



Nobody at a wealth firm thinks of the team page as a listing. To a prospective client who has never met you, that is what it is: names, titles, a paragraph of credentials, and either faces or blanks. This research measured what the blank costs, and what a visible face is worth in plain money.

THE EXPERIMENTS: SAME HOST, MINUS THE FACE

The first three studies worked like a controlled taste test. The researchers took real profiles from Amsterdam hosts and pared them down to what a visitor sees first: age, gender, a short self-description, the sleeping arrangement, the number of references. Then they produced two versions of each, one with the host's photo and one with the silhouette, and asked participants to rate their willingness to stay with each host.

Seventy-six Dutch students rated the profiles in the first study. The second raised the sample to 154 and added the direct question: how much do you trust this host? In both, which profiles appeared with or without a photo rotated randomly across participants, so no host's charm could tilt the comparison.

The third study, with 150 young adults in the United Kingdom, gave the design its teeth. Some profiles appeared with faces a separate panel had rated attractive, some with faces rated unattractive, some with the silhouette. If a photo merely decorates a listing, an unappealing face should do as much damage as none at all. If a face carries information a stranger needs, any face should beat a blank.

The fourth study left the lab: the team collected 461 real ride listings from BlaBlaCar, the European long-distance carpooling platform. Holding the route and everything else constant, what happens to the price a driver can charge when a profile photo appears?

WHAT THEY FOUND: NO FACE, LESS TRUST

Every experiment landed the same way. With all else identical, people were substantially less willing to stay with a host who showed no photo: removing it cost roughly a point of willingness in the first two studies, half a point in the third.

The trust question moved even more than willingness did, and that lost trust statistically explained much of the lost bookings. Participants were not objecting to bad design. They were declining to trust someone who would not be seen.

The third study produced the finding worth taping above the marketing department's door: an unattractive photo beat no photo. People preferred attractive hosts to unattractive ones, and unattractive hosts to faceless ones. The penalty for hiding was larger than the penalty for an ordinary face.

The market agreed. On BlaBlaCar, drivers who showed a photo charged 9 to 11 percent more for otherwise equivalent rides. Sellers can have good reasons to stay hidden, from ordinary privacy to the discrimination some minority hosts face when they show a face, but the data says the choice is not free.

Why does a photo carry that much weight? A face answers the question a stranger asks first: who is this person? A blank gives them nothing to judge, so they fill it with the least generous guess. Wealth management settled the headshot long ago, so the test worth running is elsewhere: the camera left off, the general inbox standing in for a name, the commentary that goes out with nobody attached.



N = 380, plus 461 field listings

The Candidate Who Confessed

The candidate who owned a past mistake won more trust than the one who blamed her old boss, and later good news never closed the gap.

Maliqi, M., Lalot, F., & Quiamzade, A. (2025). Of first impressions, shattered trust, and apology: Impact on interpersonal trust and team dynamics. *Frontiers in Psychology*, 16, 1654463. Open access: doi:10.3389/fpsyg.2025.1654463

THE INTERVIEW IS GOING WELL UNTIL THE interviewer mentions the phone call to her previous employer. The candidate, a composed woman named Anna Bloom, already knows what is coming. The reference disclosed a fault: she once forwarded an email full of sensitive information to the entire department.

What she says next is the experiment. In one version of the video she denies it and blames vague instructions from her old boss. In the other she owns it: an honest mistake with a distribution list, she is sorry, it will not happen again. Only the answer changes.



Researchers Majlinda Maliqi, Fanny Lalot, and Alain Quiamzade showed one version or the other to employees across Switzerland and published the results in *Frontiers in Psychology*. The question: what does the first answer to a hard question do to trust, and does anything that comes later loosen its grip?

THE EXPERIMENT: A REFERENCE CHECK GOES WRONG

The 204 raters were working adults, average age thirty-four, all employed on teams of three or more, about half from a single partner company and the rest from 86 other organizations in Northwestern Switzerland. When they judged what a new colleague would do to a team, they were judging something they live daily.

Every participant watched the staged interview and then rated the candidate: confidence in her abilities, belief in her good intentions, her honesty, and how much they would trust her with something important they could not monitor.

Then came the second act. Participants in both conditions received the same new piece of good news about Ms. Bloom, a chat between the interviewer and a colleague sharing a local newspaper story about her years of dedication keeping a volleyball club afloat. Warm, verifiable, and entirely unrelated to the accusation. The researchers wanted to know whether that good news would count for more with the candidate who had already earned trust, or the same for both.

Last, participants imagined the candidate had been hired onto their own team and forecast the consequences: how collaborative the team would feel, how well it would perform, how satisfied they would be working alongside her.

WHAT THEY FOUND: THE APOLOGY LED, AND KEPT LEADING

The confession won the first impression by a wide margin. Participants who saw the apology rated the candidate 5.34 out of 7 on trustworthiness. Those who saw the denial gave her 4.58. In an experiment where the two videos differ by a single answer, that is a substantial gap from one moment of candor.

The volleyball story then lifted both groups by about the same amount. The researchers had expected the first impression to act as a filter, with good news counting more for the candidate people already trusted. That is not what the data showed. The new information added the same increment to both, and the apology's lead survived intact across every measure: general trust, cognitive trust, willingness to hand her something important. Order mattered more than volume.

The widest gap after the good news sat on the most practical question: would you entrust her with something consequential that you could not monitor? Feeling trust is one thing; delegation is the business, of teams and of advisory relationships alike. The apology's advantage was largest exactly on the willingness to delegate, the measure closest to a client moving assets. That is the number a practice should care about.

Her answer also shaped what people expected of the whole team. People who saw the apology expected a better team with her on it: a more collaborative culture, stronger performance, more satisfaction with the working relationship. The trust she earned in that one answer statistically accounted for those brighter forecasts.

ONE ACCUSATION, TWO ANSWERS



Trust in the candidate: 5.34 vs 4.58 of 7 at first impression.

Good news lifted both alike; preregistered, N = 204.

Two lessons sit side by side in the data. Later good news genuinely helps; nothing about a bad first impression stopped the new fact from adding its full weight. And later good news replaces nothing; the denial's discount was still there, untouched, after the good news arrived. Her first answer fixed the starting point, and everything after it moved both candidates by the same amount. For anyone whose first meeting went sideways, that is both the bad news and the map.

WHAT THIS COULD MEAN AT YOUR DESK

Script the hard answer before it is asked. Every practice has the question it hopes a prospect never raises: the partner who left, the down year, the disclosure on the ADV. The study found the version who admits the specific mistake, names its cause, and says what changed earns measurably more trust than the version who explains why it was someone else's fault. Write that answer in advance, because the moment it is needed is the worst moment to compose it.

Blame is the expensive word. The denial in the video came packaged the way denials usually do, pointing at unclear instructions from a supervisor. Reasonable, even plausible. It still cost three quarters of a point of trust against a plain admission. A prospect listening to you explain why a former client's complaint was unfounded may be running the same arithmetic, whatever the complaint's merits.

Do not count on your record to lift a bad first answer. The good news in this study was warm and real, and it closed none of the gap. If the first meeting ends in deflection, the strong quarter you deliver later starts from the discounted baseline. The order of operations favors candor now, credentials after.

Volunteer the confession while it is cheap. The candidate answered only once the reference surfaced the mistake, and the apology still won. Issue #3 covered the other end of that road: repair that has to be demanded falls short. Between the two results sits a practical rule. The earlier the admission arrives, the more trust it buys, and the cheapest moment is the one where you raise it yourself.

THE BOTTOM LINE

Two versions of the same capable woman answered one uncomfortable question, and the one who said 'that was my mistake' was the one people trusted more and wanted on their team. The gap opened in the first minutes of acquaintance, and later virtue moved both versions in parallel without ever closing it.

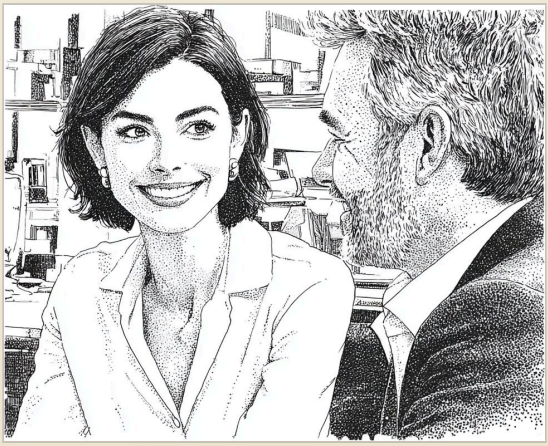
Sooner or later a prospect will ask you an uncomfortable question. What the study suggests is that the trust in the room afterward depends less on your record than on whether your first instinct is to own what happened. That instinct can be rehearsed.

Follow-Through Beats First Impressions

The first meta-analysis of trust learning in the repeated trust game: reputation, group labels, and stake size barely mattered next to conduct.

Duncan, C., Sganzerla, L., Kaltwasser, L., & Dziobek, I. (2025). Trust learning in the repeated trust game: A meta-analytic study. *British Journal of Psychology*, 117(3), 1026-1055. Open access: doi:10.1111/bjop.70045

The trust game is how economists have measured trust since the 1990s. You start with ten dollars and choose how much to send your partner. What you send usually triples on the way, so ten arrives as thirty. Your partner then decides how much to send back, anything or nothing, and the return does not multiply.



Every dollar returned is a dollar your partner could have kept, so a single round measures a snap judgment. Repeated with the same partner, sending money back is what keeps the next round coming, and the game measures trust learned from conduct.

Researchers Duncan, Sganzerla, Kaltwasser, and Dziobek gathered every repeated trust game study they could find and asked what actually teaches people to trust. Their meta-analysis in the *British Journal of Psychology* pooled 68 studies, 404 separate measurements of learning, and more than 8,000 players. It is the first synthesis of its kind, and its answer is blunt.

THE METHOD: SIXTY-EIGHT LABORATORIES, ONE QUESTION

The team defined trust learning as movement: the difference between what a player invested in the first round with a partner and what they invested in the last. A player who starts guarded and ends generous has learned to trust. One who starts generous and ends guarded has learned the opposite.

The pooled games ran the world's variety: students and working adults, five continents, stakes from pocket change to meaningful sums, partners who were live humans and partners who were programmed, games a few rounds long and games that ground on for dozens.

Against that variety they tested everything on record that might explain why a player's trust rose or fell. The partner's actual rate of paying money back. Reputations planted by experimenters before the game began. Whether the partner belonged to the player's own social group. The players' age, gender, nationality, and student status. The size of the stakes, the length of the game, and a dozen other features of its design.

WHAT THEY FOUND: THE LEDGER WINS

One factor towered over everything else: the partner's reciprocation rate, how reliably they actually paid trust back. It was the largest effect in the model by a wide margin. Age, gender, country, the size of the stakes, the length of the game: none of it reliably changed how much people learned to trust.

Planted reputations turned out to have a narrower influence than their reputation suggests. Telling players beforehand that a partner was trustworthy or moral did shift how they learned. Telling them a partner was competent or skilled did nothing. The only advance information that mattered spoke to character, and claims about skill washed out the moment behavior started arriving.

The reputation finding carries a sharp edge. Players handed a negative character reputation did not just start lower; they changed their behavior less as the game went on. A bad prior shrinks investment, small investments generate little evidence, and little evidence means little learning. Distrust, once planted, protects itself.

Group membership failed to matter at all. In studies pairing players with partners from their own group or an outgroup, people learned from what the partner did, full stop. Labels sway snap judgments in one-round games, as decades of research shows. Given rounds enough to watch behavior, the label lost to the ledger.

One feature of the games did change the results. Trust grew less against live humans than against programmed partners. Once players know which round is the last, they invest less, because a partner with no rounds left has nothing to gain by paying them back. Economists call it the endgame effect. A program reciprocates on a fixed strategy to the end, so trust against it kept climbing.

WHAT TAUGHT PLAYERS TO TRUST



Relative influence on learned trust, 68 studies pooled; bars are qualitative.

WHAT THIS MEANS AT YOUR DESK

Make the rounds small and frequent. Trust grew on reciprocation, and reciprocation needs occasions. A quarterly meeting gives a client four rounds a year. A practice built on many small kept commitments, the follow-up sent when promised, the call returned the same day, the document ready before the deadline, deals out dozens. Each one is a round in which the client watched you pay trust back. Density of proof beats scale of promise.

Ask referrers to vouch for your character, not your resume. The only advance information that changed trust learning concerned trustworthiness and morality. When a happy client introduces you, 'she does what she says she will do' is doing different work than 'she is very sharp.' Competence claims are table stakes. Character claims are the priors that actually move, so coach the introduction: one sentence about what you did, one about who you are.

Treat a wary prospect as a learning problem, not a persuasion problem. The negative-prior result names the trap: skeptical people commit little, so they collect little evidence, so they stay skeptical. Give a guarded prospect low-stakes ways to watch you reciprocate, a small deliverable inside the first week, a promise made and kept before any assets move. You are not arguing them out of the prior. You are restarting the ledger that outweighs it.

Prospect beyond your look-alikes. Group labels stopped mattering once behavior accumulated. The family that shares none of your clubs, schools, or demographics learns to trust you the same way everyone else does, by watching you follow through. Conduct crosses lines that marketing cannot, and the advisor who worries their background will read as unfamiliar in a new market might worry instead about response times.

THE TAKEAWAY

This issue opened with a face on a webpage and a confession in an interview, the moves that win a beginning. The meta-analysis is the reminder of what the beginning is for. Across 68 studies, the durable driver of trust was reciprocation, round after round, rather than the introduction, the group label, or even the reputation, except where reputation spoke to character. That order of importance is the inverse of how most marketing budgets get spent.

The face gets you the first round. The candor gets you the second. Everything after that is earned one kept promise at a time, which is the best news in this issue: of everything the literature has tested, the strongest trust signal is the one entirely under your control, starting with the next thing you told a client you would do. The next round is probably already on your calendar.

The AI Workforce
for the Human Firm.

Humanity Labs

LIVE ZOOM WEBINAR · AUGUST 28, 2026

Ask Me Anything



Andrei Pop
Founder & CEO, Humanity Labs

11:00 AM PT · 2:00 PM ET
60 MINUTES

Register Here

ABOUT HUMANITY LABS

Humanity Labs gives wealth management firms an AI Workforce that works inside existing systems to complete work across the front, middle, and back office, helping firms serve more clients without adding cost at the same rate. Every job strengthens a shared memory that builds the firm's Organizational General Intelligence (OGI).

Humanity Labs works with leading firms overseeing more than $750 billion in client assets and is backed by Google Ventures and Vestigo Ventures.

Learn more at **HumanityLabs.ai**