Naming the Error Without Naming the Person

A Humanistic Foundation for Surfacing Cognitive Distortions, Dissonance, and Reasoning Biases in Mental Health Care

Jeffrey Lourie, PMHNP, FNP
2026-06-02

Abstract

A recurring clinical intuition holds that if patients could only *see* the faulty machinery of their own thinking — the distortions, the dissonance, the leaps of inference — they would be freed from some part of their suffering. This whitepaper takes that intuition seriously and then takes it apart. Four clinical moves are routinely collapsed into one: cognitive restructuring, dissonance/discrepancy work, metacognitive bias training, and formal fallacy instruction. The first three rest on distinct and substantial literatures; the fourth is mostly an open research question. The evidence supports surfacing in specific forms — strongest when it is collaborative rather than didactic, self-generated rather than imposed, and depersonalized rather than confrontational — and gives little support for declarative correction from authority. The paper then makes the humanistic case that the *manner* of surfacing matters more than the *content* surfaced, because telling a person their reasoning is wrong carries real risks: invalidation, alliance rupture, a deflation of the person's standing as a knower, and the mislabeling of cognition that is actually accurate. It closes with design principles and required failure-mode gates for a clinical tool built on a single organizing commitment: **surface the structure of thinking to support a human relationship, never to win an argument with the patient.**

In this paper, "naming the error" means rendering a pattern visible in plain, tentative, human language — not attaching a formal debate label to the patient's thought. That distinction is load-bearing and recurs throughout.

Clinical Intuition and the Trap Inside It

Anyone who has sat across from a person in distress has felt the pull. The patient describes a chain of reasoning — *I failed the interview, so I am worthless, so things will never improve, so why try* — and the structure of the error is almost luminous to the clinician. The temptation is to point at it. To name it. To say, in effect: *look, here is the broken step, and once you see it you can step around it*.

This intuition is not naive. It is, in compressed form, the founding premise of cognitive therapy, and there is real evidence behind it. But the intuition contains a trap, and the trap is the word

"pointing." Pointing implies that the therapeutic work is *epistemic* — a matter of getting the patient to hold true beliefs instead of false ones. A great deal of what actually helps people is not epistemic at all. And in the cases where epistemic correction does help, *how* the correction is delivered routinely matters more than *whether* the correction is right.

The purpose of this document is to hold both truths at once. Surfacing the architecture of a person's thinking can be genuinely useful — the literature supports this in several distinct forms. And surfacing it badly can be inert, or worse, harmful. A tool built in this space inherits both possibilities. It cannot be designed responsibly without understanding both.

Four Moves, Often Collapsed Into One

The phrase "cognitive dissonance, distortions, and fallacies" feels like a single family of phenomena. It is not. It collapses four distinct clinical moves drawn from separate traditions. The first three have real literatures; the fourth is mostly an open question. Conflating them produces a further error: borrowing the credibility of the well-evidenced moves to license the unevidenced one.

Cognitive distortions and the Socratic method

The cognitive distortion — overgeneralization, catastrophizing, all-or-nothing thinking, mind-reading — is the most clinically familiar of the four, and the best supported. What is easy to miss is that the supporting evidence is not for *labeling* distortions but for a particular *process* of surfacing them: Socratic questioning, also called guided discovery, in which the clinician asks questions that lead the patient to examine the evidence for a belief rather than being told the belief is wrong.

The distinction is not cosmetic, and it has been examined. In a study of cognitive therapy for depression, within-patient increases in a therapist's use of Socratic questioning predicted a drop in depressive symptoms in the *following* session, and the effect held after accounting for the therapeutic alliance (Braun, Strunk, Sasso, & Cooper, 2015). This is not definitive causal proof — it is a process-outcome association in 55 depressed adults, drawn from observer ratings of early sessions. But the within-patient design rules out stable patient traits as an explanation, which makes it unusually relevant: collaborative questioning tracked subsequent symptom improvement more closely than a didactic "correction" model would predict. The active ingredient looks like *collaborative examination*, not *declarative correction*. A clinician who simply announces "that is a cognitive distortion" is using the vocabulary of the evidence base while skipping the part the evidence base actually credits.

A second, sobering caveat sits beside this. Cognitive therapy contains both cognitive restructuring and behavioral activation, and a landmark component analysis found that behavioral activation alone produced outcomes comparable to full cognitive therapy at termination and six-month follow-up — and, strikingly, altered negative thinking as much as the cognitive components did (Jacobson et al., 1996). The correction of distorted thoughts is real, but it is not the whole engine, and changing how a person *thinks* may not require directly

disputing their thoughts at all. A tool that treats distortion-spotting as the center of gravity risks over-weighting a component the broader literature has struggled to isolate as necessary.

Cognitive dissonance and discrepancy as a lever

Cognitive dissonance enters the clinic through a different door. In its most direct form, the goal is to have a person voluntarily argue against an attitude they hold, so that the gap between their words and their belief becomes an engine of change.

The clearest demonstration is the Body Project, a dissonance-based eating-disorder prevention program in which participants actively argue against the thin ideal. Randomized trials show durable reductions in body dissatisfaction, disordered eating, and related risk factors. A Brazilian replication found reductions across body dissatisfaction, disordered eating, and negative affect, largely sustained at six months (Hudson, Amaral, Stice, Gau, & Ferreira, 2021), and a separate randomized trial found the same approach improved intuitive eating, body appreciation, and self-esteem with large between-condition effects (Resende, Almeida, dos Santos Alvarenga, Brown, & de Carvalho, 2022). The important limiter: this demonstrates that dissonance-based methods can be effective in *specific, structured, prevention-oriented* contexts. It is not general evidence that dissonance induction is broadly therapeutic across diagnoses or levels of clinical acuity.

A gentler cousin of the discrepancy idea underwrites motivational interviewing, widely used in substance-use and health-behavior change. MI can be understood partly as *eliciting* discrepancy between a person's values and their behavior — but it does so in an autonomy-supportive way, and the discrepancy is invited, not imposed; confrontation is explicitly avoided. The honest characterization is that MI works but modestly, and that discrepancy is a *theorized* contributor rather than a cleanly isolated mechanism: a meta-analysis of MI for adolescent health behavior found a small but durable aggregate effect that did not differ meaningfully from MI applied to substance use (Cushing, Jensen, Miller, & Leffingwell, 2014).

The design implication across this tradition is consistent. The clinician does not point out the patient's error. The clinician arranges conditions under which the patient *generates and examines their own* discrepancy. Authorship of the insight stays with the patient — and that appears to be a feature, not an accident.

Reasoning biases and metacognitive training

The closest existing analog to "showing patients the flaws in their reasoning" is Metacognitive Training (MCT) for psychosis. MCT targets specific reasoning biases — jumping to conclusions on thin evidence, overconfidence in memory, a bias against dis-confirmatory evidence — thought to contribute to the formation and maintenance of delusions.

The evidence is real but moderate, and it has been stress-tested. A meta-analysis found significant effects on positive symptoms and delusions, with a large effect on patients' acceptance of the intervention; under sensitivity analysis restricted to low-bias studies, the effect on delusions shrank and lost significance while the effect on positive symptoms held at a smaller magnitude (Eichner & Berna, 2016). A more recent meta-analysis of 37 studies reproduced significant reductions in positive symptoms and delusions and a gain in cognitive

insight, with demographic factors exerting limited influence (Gelner et al., 2025). There is neuroimaging work tying the jumping-to-conclusions bias to specific connectivity patterns that shift after training, and longitudinal evidence that improvement in this bias is linked to delusion improvement specifically in the metacognitive condition — a content-specific mechanism rather than a generic cognitive boost (Andreou et al., 2018; Andreou et al., 2015).

The design of MCT is as instructive as its outcomes. MCT does not sit a patient down and dispute their delusion. It works through *general, depersonalized examples* — exercises about how anyone can leap to a wrong conclusion — deliberately keeping a step's distance from the patient's own fixed belief. The program's own metaphor for its mechanism is the planting of seeds of doubt, not the felling of the tree. This indirection reflects a long-standing clinical principle, shared with cognitive therapy for psychosis: fixed delusional beliefs are not productively addressed by direct confrontation, which tends to invite defense rather than reconsideration. (This principle is convention and design rationale across CBT-for-psychosis and MCT; it is more firmly established as clinical practice than as a single controlled finding.)

Formal fallacy instruction: sparse and preliminary

The fourth move has the most intellectual appeal and the thinnest clinical evidence: explicitly teaching the *named, formal taxonomy* of logical fallacies and reasoning errors — ad hominem, base-rate neglect, false dichotomy, the gambler's fallacy — as objects a patient learns to recognize and avoid.

This approach has no mature clinical evidence base comparable to CBT, dissonance-based prevention, or MCT. A small adjacent literature does exist: Logic-Based Therapy (LBT), a philosophical-counseling modality derived from rational-emotive behavior therapy, in which the clinician identifies the "fallacy" driving a self-defeating conclusion and matches it to a guiding virtue. A preliminary randomized controlled study compared a single one-hour LBT session to a one-hour mindfulness activity in 20 family caregivers and reported a promising signal for reduced *state* anxiety relative to the control, while explicitly framing the result as preliminary and pending a larger study (Chaukar, 2024). That is genuine adjacent evidence, but at n=20, a single session, and a narrow outcome, it does not establish that teaching formal fallacy taxonomies is clinically useful — let alone safe — in ordinary mental health care.

The pattern across the mature traditions is telling: cognitive therapy and MCT both teach the *shape* of an error (overgeneralization, jumping to conclusions) while studiously avoiding the *formal vocabulary* of logic, and they have done so for decades. For a builder who has invested real effort in cataloguing biases, fallacies, and distortions, this is the most consequential finding in the document, and it cuts two ways. It marks genuinely open territory — the formal-taxonomy approach is underexplored and could, in principle, be developed and tested. It also marks a *caution*, because the disciplines that have studied this longest have, by their own design choices, routed *around* the formal taxonomy rather than through it. The next section explains why.

The Humanistic Core: What It Means to Tell Someone Their Thinking Is Wrong

Everything above is a question of efficacy. This section is a question of what we owe a person. The two are not separable, because in mental health care the relationship is not merely the delivery vehicle for the treatment — much of the time, the relationship *is* the treatment.

The patient as a knower

When a clinician tells a patient that their reasoning is distorted, fallacious, or biased, something happens that has no analog in fixing a broken bone. The clinician is not only proposing a different conclusion; they are, however gently, lowering the patient's standing as a reliable interpreter of their own life. Philosophers call the systematic deflation of a person's credibility *epistemic injustice* (Fricker, 2007), and there is a direct argument that psychiatric patients are *especially* exposed to it — more so than patients with purely somatic illness — because a person whose perceptions have already been labeled symptomatic learns quickly that their account of reality is treated as provisional, subject to correction by someone with more authority in the room (Crichton, Carel, & Kidd, 2017).

A tool that surfaces "errors" in a patient's thinking is, by default, an engine of exactly this deflation, scaled and automated. That is not a reason not to build it. It is a reason to treat the deflation risk as a primary design constraint rather than an afterthought.

The cost of getting it wrong is measurable

The humanistic concern is not soft. When the surfacing of an error strains the bond between patient and clinician, the literature has a name for it — an alliance rupture — and it has measured what follows. Across studies, the *repair* of such ruptures is moderately associated with better treatment outcomes (Eubanks, Muran, & Safran, 2018). Read in the other direction: a clinician — or a tool — that creates ruptures and does not repair them is working against outcome. Pointing out a distortion is precisely the kind of move that can produce a rupture: a disagreement about what is true, a strain in the bond. So the same gesture the cognitive literature credits with symptom change is the gesture the alliance literature flags as a hazard. Both are correct. The reconciliation is in the *how*.

Insight is not change

There is a persistent, seductive belief that insight precedes and produces change — that once a person *sees* the distortion, the distortion loses its grip. The evidence is more humbling. People routinely understand, with great clarity, the irrationality of a fear they cannot stop feeling. And the component-analysis finding above cuts directly here: behavioral activation changed patients' negative thinking *without* targeting it through cognitive disputation (Jacobson et al., 1996). *Doing* — approaching a feared situation, acting before believing — can drive cognitive change that direct correction does not. A tool optimized to deliver accurate insight may be optimizing for an endpoint that is neither necessary nor sufficient.

When the "distortion" is the truth

This is the sharpest edge. Catastrophizing is a distortion when the catastrophe is unlikely. It is *accurate forecasting* when the person is genuinely poor, genuinely sick, genuinely unsafe, or genuinely being mistreated. Hypervigilance is a distortion in a safe environment and a survival skill in a dangerous one. The cognitive model's vocabulary — "distortion," "error," "irrational belief" — carries a built-in assumption that the patient's appraisal is the broken part, when often the appraisal is the only part of the situation working correctly.

A system that flags accurate cognition as distorted does not merely fail to help; it functions as gaslighting, telling a person who is reading their situation correctly that the problem is their reading. This is a reasoned design concern rather than a measured rate, but it is not idle: for a practice serving rural patients facing real material hardship, any distortion-detecting logic carries a standing risk of misfiring on exactly the populations whose appraisals are most likely to be accurate and whose grievances are most legitimate. The cost of that error is borne by the people least able to absorb it.

The moralizing trap

Finally, the formal-fallacy framing carries a tonal hazard the plain-language frameworks avoid. To tell someone they have committed an "ad hominem" or a "false dichotomy" is to import the register of *debate* — of winning and losing, of being caught out. It positions the clinician as the holder of correct reasoning and the patient as the one *caught reasoning badly*. Even when accurate, this is experienced as a judgment about the person's mind rather than an observation about a moment. MCT's depersonalized design is, in large part, an engineered escape from exactly this trap. Any tool that names fallacies must solve the same problem, or it will reproduce the harm.

A Reasoned Synthesis: The Conditions Under Which Surfacing Helps

Pulling the threads together, surfacing the structure of a person's thinking appears to help under a recognizable set of conditions, and to be inert or harmful outside them.

It tends to help when the patient retains authorship of the insight — when the move is a question rather than a verdict (the Socratic evidence) or a self-generated discrepancy rather than an imposed one (the dissonance evidence). It tends to help when the error is named in plain, depersonalized, non-technical language that describes a *moment of thinking* rather than indicting a *kind of thinker* (the MCT design). It tends to help when there is enough alliance to absorb the friction, and when ruptures that do occur are noticed and repaired (the rupture-repair evidence). And it tends to help when the cognition in question is genuinely inaccurate rather than an accurate appraisal of a hard reality.

It tends to fail, or to harm, when delivered as declarative correction from authority; when it uses a technical taxonomy that signals debate; when it outpaces the alliance; when insight is treated as sufficient for change; and when the "distortion" is in fact the truth.

The following table summarizes the evidence grade behind each move, to keep the design work from over-reading any one line.

Domain	Evidence strength	What is supported	What is *not* supported
Socratic questioning / cognitive restructuring	Moderate to strong	Collaborative examination of beliefs	Didactic labeling of distortions
Dissonance-based prevention	Strong in specific prevention contexts	Structured, self-generated counter-attitudinal exercises	Generalized "discomfort induction" across diagnoses
Motivational interviewing	Modest to moderate	Autonomy-supportive discrepancy work	Dissonance as the isolated mechanism
Metacognitive training (psychosis)	Moderate	Depersonalized reasoning-bias training	Direct confrontation of fixed delusions
Formal fallacy taxonomy	Sparse / preliminary	An open research area	Routine patient-facing fallacy labels

This synthesis is the spine of everything that follows. A tool in this space is not, at bottom, a classifier of errors. It is an instrument for arranging the *conditions* under which a human relationship can do epistemic work safely.

Implications for a Future Clinical Tool

What follows is a set of principles, not a specification. They are written to constrain the design space before any feature is chosen, because the failure modes here are not bugs to be patched later — they are direct consequences of the framing one starts from.

Surface to the clinician by default; surface to the patient only under gates. The strongest, safest version of this tool makes the structure of a patient's reasoning *visible to the clinician* as a scaffold for a Socratic conversation, and treats any patient-facing naming of an "error" as a separate, gated action conditioned on alliance, belief fixity, acuity, and readiness. The central design decision of this entire project is *who the tool talks to.* The evidence points towards mostly the clinician, sometimes the patient, never the patient unguarded.

Name the shape, not the person, and never in the register of debate. If the tool ever surfaces a pattern to a patient, it must do so in plain language describing a moment of thinking, not a class of thinker, and it must abandon formal-logic vocabulary entirely. "It sounds like one hard moment is standing in for the whole future" is therapeutic. "That is a hasty generalization" is a debate score. The formal taxonomy can live in the clinician-facing layer as an organizing schema; it should be presumptively barred from patient-facing output.

Validate distress while withholding premature correction; fail toward epistemic humility, not factual endorsement. When context could justify the patient's appraisal, the default is to withhold the distortion frame and surface the situation instead. This is not a rule to believe every claim — it explicitly is not factual endorsement of paranoia, of allegations that require careful assessment, or of hopelessness. It is a rule to *not assert distortion under uncertainty*: when the picture is unclear, the tool withholds the label and prompts the clinician to assess reality-testing, safety, social context, and available evidence. A system that cannot distinguish catastrophizing from accurate forecasting must not claim that it can.

Make the reasoning inspectable and evidence-linked. Any system analyzing language will *infer* candidate patterns — pretending otherwise is dishonest engineering. The discipline that matters is preventing opaque, trait-level, diagnostic, or motive-level inference, and making every surfaced pattern auditable. For each candidate, the tool should show the utterance, the rule or heuristic involved, a confidence level, and the reason a clinician might override it. This refines rather than abandons the determinism commitment: determinism governs *how candidates are presented and justified*, not a pretense that classification is inference-free. Transparency here is what lets a clinician override the tool when it misfires on a real grievance.

Treat the alliance as a first-class signal, not an externality. Because surfacing can rupture, and because rupture repair tracks outcome, the tool should lower rather than raise the temperature of the room — giving the clinician material for a collaborative question, not a confrontation. A feature that makes the clinician more likely to declare the patient wrong is a regression even if its classification accuracy is perfect.

Reject the useless-guardrail failure mode. The opposite error is equally real: a tool so hedged, so caveated, so frightened of harm that it tells the clinician nothing and changes no decision. A system that avoids harm by becoming inert has not succeeded; it has only failed quietly. The goal is proportionate, plain, genuinely useful surfacing — bounded by the principles above, but not strangled by them.

6. Failure Modes and Required Gates

A clinical-tool foundation document is incomplete without an explicit account of how it goes wrong. These gates should be treated as requirements, not aspirations.

Acuity gate. No patient-facing reasoning challenge during active suicidality, mania, acute psychosis, intoxication, severe agitation, or acute trauma response, unless explicitly clinician-directed. In these states, surfacing "errors" is at best inert and at worst destabilizing.

Context gate. Before any thought is treated as distorted, the tool must prompt assessment of whether the appraisal may reflect poverty, abuse, discrimination, unsafe housing, medical illness, chronic pain, or a realistic threat. The default under ambiguity is to surface the *situation*, not the *distortion*.

Diagnosis gate. The same surfacing strategy should not be applied uniformly across depression, OCD, PTSD, psychosis, personality pathology, and substance use. An identical reasoning label can function very differently across conditions — reassurance-seeking in OCD, hopelessness in depression, and a fixed belief in psychosis are not the same target and do not respond to the same move.

Documentation risk. Clinician-facing output can enter the record and later harm the patient if written as "patient demonstrates fallacious reasoning." The tool must generate careful, non-pejorative language and should never emit a phrase that would read as a credibility indictment in a chart, a disability determination, or a custody proceeding.

Outcome humility. Classification accuracy is not the endpoint. More appropriate endpoints include alliance preservation, patient-perceived validation, clinician-rated usefulness, reduced distress, improved functioning, and rupture/repair tracking. A tool that classifies perfectly and helps no one has failed.

Open Questions and a Research Agenda

The honest gaps are also the opportunities. Five questions stand out.

First, the sparse quadrant: can the *formal* taxonomy of fallacies be made therapeutic beyond the preliminary signal in the logic-based-therapy literature, or do the plain-language traditions route around it for good reason? This is testable and, so far, barely tested.

Second, the audience question: does a reasoning-structure tool produce better outcomes when its output is directed to the clinician versus the patient? The synthesis predicts clinician-facing wins, but it is a prediction, not a finding.

Third, the accuracy-detection problem: can a tool reliably distinguish a distortion from an accurate appraisal of a hard reality, and what is the cost of its false positives in populations with legitimate grievances? This is the failure mode most likely to harm a rural, resource-constrained patient population, and it should be characterized first.

Fourth, the insight-versus-action question: if surfacing reliably produces insight but action drives change, what is the tool actually for, and how should that reframe its success metrics?

Fifth, the rupture question: does surfacing structure raise or lower rupture rates in practice, and can the tool be instrumented to detect and prompt repair?

Conclusion

The intuition we started with — that seeing the machinery of one's own thinking is liberating — is half right, and the half that is right is conditional on a manner of delivery that the strongest evidence keeps quietly insisting on. People are not freed by being told they are wrong. They are sometimes freed by being asked a question they had not thought to ask themselves, by arguing their way into a discrepancy they then cannot unsee, by acting before they believe. The role of a tool in this space is not to be right about the patient. It is to make it easier for one human being to help another think — without, in the process, telling them they are the kind of person who thinks badly.

That is the line the title points at. We can name the error. We must not name the person.

References

Andreou, C., Schneider, B. C., Balzan, R., Luedecke, D., Roesch-Ely, D., & Moritz, S. (2015). Neurocognitive deficits are relevant for the jumping-to-conclusions bias, but not for delusions: A longitudinal study. *Schizophrenia Research: Cognition, 2*(1), 8–11. https://doi.org/10.1016/j.scog.2015.02.001

Andreou, C., Steinmann, S., Leicht, G., Kolbeck, K., Moritz, S., & Mulert, C. (2018). fMRI correlates of jumping-to-conclusions in patients with delusions: Connectivity patterns and effects of metacognitive training. *NeuroImage: Clinical, 20*, 119–127. https://doi.org/10.1016/j.nicl.2018.07.004

Braun, J. D., Strunk, D. R., Sasso, K. E., & Cooper, A. A. (2015). Therapist use of Socratic questioning predicts session-to-session symptom change in cognitive therapy for depression. *Behaviour Research and Therapy, 70*, 32–37. https://doi.org/10.1016/j.brat.2015.05.004

Chaukar, H. (2024). A randomized, controlled, preliminary study to assess the efficacy of logic-based therapy in reducing anxiety and/or depression in family caregivers. *Journal of Rational-Emotive & Cognitive-Behavior Therapy, 42*(3), 582–609. https://doi.org/10.1007/s10942-023-00532-z

Crichton, P., Carel, H., & Kidd, I. J. (2017). Epistemic injustice in psychiatry. *BJPsych Bulletin, 41*(2), 65–70. https://doi.org/10.1192/pb.bp.115.050682

Cushing, C. C., Jensen, C. D., Miller, M. B., & Leffingwell, T. R. (2014). Meta-analysis of motivational interviewing for adolescent health behavior: Efficacy beyond substance use. *Journal of Consulting and Clinical Psychology, 82*(6), 1212–1218. https://doi.org/10.1037/a0036912

Eichner, C., & Berna, F. (2016). Acceptance and efficacy of metacognitive training (MCT) on positive symptoms and delusions in patients with schizophrenia: A meta-analysis taking into account important moderators. *Schizophrenia Bulletin, 42*(4), 952–962. https://doi.org/10.1093/schbul/sbv225

Eubanks, C. F., Muran, J. C., & Safran, J. D. (2018). Alliance rupture repair: A meta-analysis. *Psychotherapy, 55*(4), 508–519. https://doi.org/10.1037/pst0000185

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing.* Oxford University Press.

Gelner, H., Krężołek, M., Aleksandrowicz, A., Piwińska, J., Lamarca, M., Berna, F., Acuña, V., König, C., Moritz, S., Ochoa, S., & Gawęda, Ł. (2025). The role of demographic variables in the effectiveness of metacognitive training for psychosis: A meta-analysis. *Schizophrenia Research, 283*, 85–97. https://doi.org/10.1016/j.schres.2025.06.017

Hudson, T. A., Amaral, A. C. S., Stice, E., Gau, J., & Ferreira, M. E. C. (2021). Dissonance-based eating disorder prevention among Brazilian young women: A randomized efficacy trial of the Body Project. *Body Image, 38*, 1–9. https://doi.org/10.1016/j.bodyim.2021.03.008

Jacobson, N. S., Dobson, K. S., Truax, P. A., Addis, M. E., Koerner, K., Gollan, J. K., Gortner, E., & Prince, S. E. (1996). A component analysis of cognitive-behavioral treatment for depression. *Journal of Consulting and Clinical Psychology, 64*(2), 295–304. https://doi.org/10.1037/0022-006X.64.2.295

Resende, T. R. O., Almeida, M., dos Santos Alvarenga, M., Brown, T. A., & de Carvalho, P. H. B. (2022). Dissonance-based eating disorder prevention improves intuitive eating: A randomized controlled trial for Brazilian women with body dissatisfaction. *Eating and Weight Disorders, 27*(3), 1099–1112. https://doi.org/10.1007/s40519-021-01245-x