Prediction, Control, and the Regulating Mind

Executive Summary

Dr. Greg Moody

August 3, 2026

Executive summary

What this is, and what it isn't

This is not an alternative to cognitive behavioral therapy. CBT has decades of outcome data. A control-theoretic account of its mechanism has none. Nothing here asks a clinician to stop doing what they're doing or to add a technique.

What it offers is a layer underneath. CBT says to test the thought, and CBT is right. But "test the thought" covers at least three mechanically different situations, and nothing in the standard model separates them.

The belief can be inaccurate. The belief can be accurate and acted on with four times the force the situation calls for. Or the belief can be roughly right and held with certainty it hasn't earned.

Three different problems, three different levers. CBT already contains all three, which is most of why it works. What it doesn't supply is a mechanical read on which one you're reaching for and why it lands.

That read comes from a gap between two literatures that have run in parallel for fifty years without citing each other.

The two halves

The mind as a control system. Targets, error signals, corrective action. Emotions are the error signals. Hunger is what sitting below a set point feels like from the inside. Fear is what it feels like to have estimated more danger than you're built to tolerate. This is cybernetics, it goes back to Wiener and Powers, and it has been the losing side of an argument since the cognitive revolution.

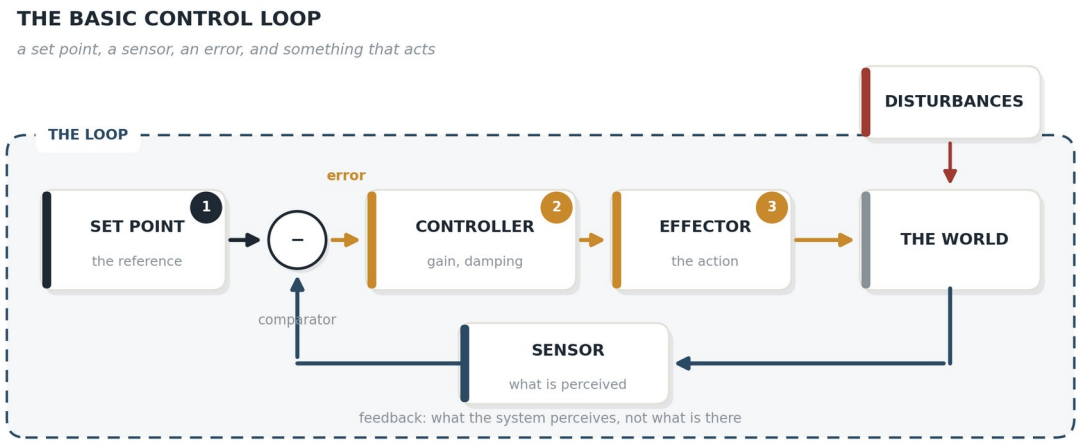

Figure 1. The basic control loop.

On the positive side of the ledger, the same account holds that **happiness is not an emotion**, because nothing works to drive it to zero. It's what error correction feels like, which is why a large or fast correction produces more of it than a slow one. The best meal of your life was eaten hungry. Amusement, awe and curiosity are left unaccounted for, and this document says so rather than stretching to cover them.

The mind as a prediction machine. The brain guesses at the world continuously and corrects the guess only where the evidence insists. Perception is inference. This is predictive processing, and it currently dominates computational psychiatry.

> *Aside… if you have read Plato's The Allegory of the Cave, this is kind of congruent with that… but enough intellectualizing here….*

Figure 2. The allegory of the cave. The prisoners are chained facing a wall, a fire burns behind them, and carried objects throw shadows onto the wall in front of them. They see the shadows, not the objects casting them, and take the shadows for the world.

The prisoners aren't reasoning badly. Their inference is right given what reaches them, and what reaches them is a projection rather than the thing itself. That's the predictive account of perception, arrived at about twenty-four centuries early. The shadows are the sensory input, the prior is what a prisoner already expects a shadow to mean, and precision is how much authority the shadow gets against that expectation.

Neither is complete. The shape of what each one is missing is the other one.

The control account has a hole where its sensor should be. Real biological systems can't read the variable they regulate. Plasma osmolality is not thirst. A racing heart is not danger. The system has to infer.

The prediction account has a hole where its purpose should be. It explains how an estimate gets built and says almost nothing about why the organism cares or how hard it acts once it knows.

The combination

Put a prediction engine on the front of a control loop and you get a system with two dials instead of one. This is not a new idea. It's the architecture of every serious control system built since 1960, and it's where I spent my first career before I spent this one.

Two features of the wiring matter clinically. **The estimate becomes the target:** the estimator's output is also your updated prediction of what should be happening, and that prediction is what the controller defends. There's no separate goal sitting off to the side.

And **the sensory return does double duty:** the same sensation your action produces feeds the estimator as evidence and feeds the controller as feedback against the set point.

Observer gain (K) governs how much your senses move your estimate of the situation, relative to what you already expected. Predictive processing calls this *precision*. Because the estimate becomes the target, K determines where the target sits.

Controller gain (L) governs how hard you act on the gap between that estimate and the incoming sensation. It has nothing to do with belief.

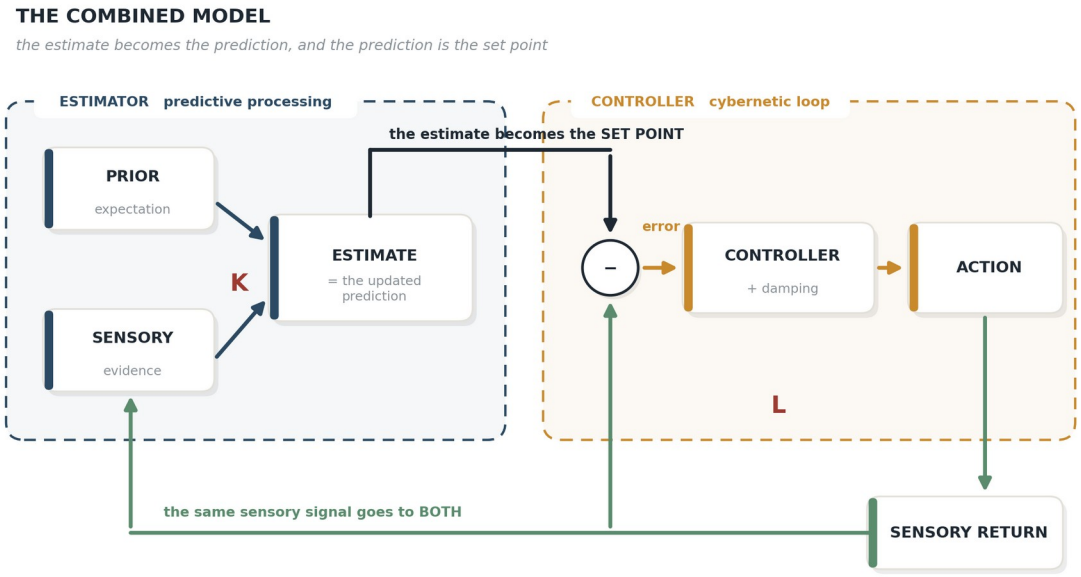

Figure 3. The combined model. The estimate becomes the prediction, which becomes the set point.

The two are in series, in that K's output is what L works on, and simultaneous, in that both run on every cycle of the loop. Change K and you move the target. Change L and you change how hard the system drives toward a target that hasn't moved.

Two more parameters matter. **Damping** decides whether corrections settle or overshoot. **Delay** decides how stale the information driving each correction already is.

All four are learned, from the statistics of the environment the person actually lived in. Volatility teaches observer gain. Consequence teaches controller gain. Overshoot teaches damping.

Which produces the claim worth carrying out of this document:

Most of what presents as dysregulation is a correctly learned parameter running in the wrong environment.

The settings aren't defects. They're accurate solutions to a previous problem, carried into a situation with different statistics. That's a rational-adaptation account of psychopathology

rather than a deficit account, it matches what clients tell you about their histories, and it's considerably easier to say out loud than most alternatives.

What it gives a clinician

Four questions that separate presentations which look identical from across the room. None of them replaces a technique. Each one tells you which technique you already have is the right one to reach for.

How this fits CBT. The four terms map one to one. The event is sensory input, the cognitive assessment is the estimator running on schemas as priors, the emotion is the error signal, and the behavior is the controller's output. What the redraw adds is that the behavior produces a result, the result is the next event, and the loop closes — which is why a maladaptive core belief survives being wrong. It runs a loop that manufactures its own evidence. Not all core beliefs are maladaptive; it's the maladaptive ones that get in the way.

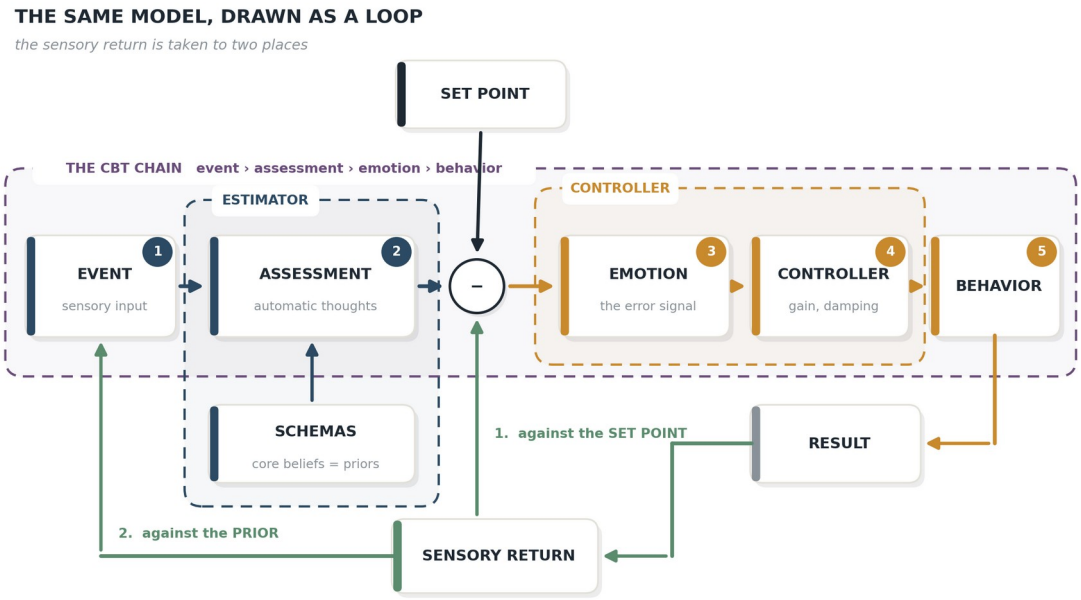

Figure 4. The CBT chain drawn as a control loop.

Is this belief inaccurate, or is it accurate and being acted on too hard? Cognitive restructuring addresses the second and does very little for the first. The first needs work on response intensity.

Is the belief wrong, or is it held with too much confidence? Content and confidence come apart. A moderately accurate belief held with total certainty behaves worse than a somewhat inaccurate belief held loosely, because certainty blocks updating. Behavioral experiments work on confidence. Arguments work on content. That's why the experiment wins.

Is the person failing to update, or being prevented from updating? Safety behavior and avoidance stop the error signal from ever being generated. The prior isn't stubborn. It's starved.

What environment taught this setting, and does it still exist? Usually answerable, usually informative, and it changes the conversation from correction to recognition.

What it says about autism

The autism literature has spent fourteen years arguing about observer gain. Whether autistic perception involves attenuated priors or over-precise sensory evidence is a question about K. Controller gain has never been modeled, measured, or discussed.

That matters because of a piece of elementary engineering. A loop with high gain and insufficient damping doesn't just respond too strongly.

It oscillates.

And repetitive behavior is what an oscillating loop produces.

Four findings already fit, from independent groups, each reported as a separate deficit and none interpreted this way:

- Excess low-frequency force oscillation during **dynamic** tracking only, specific to autism against ADHD and fetal alcohol controls (Lidstone et al., 2020, *Autism Research*, 13(12), 2177–2189)
- Increased postural trembling that gets **worse with eyes open** than closed (Bloomer et al., 2025, *Molecular Autism*, 16(1), 44)
- Increased regularity of center-of-pressure trajectories (Pettinato et al., 2024, *J NeuroEng Rehabil*, 21(1), 225)
- Force irregularity worsening disproportionately as visual display gain rises (Mosconi et al., 2015, *J Neurosci*, 35(5), 2015–2025)

Worse with more feedback is not a sensory-overload signature. It's a stability signature.

The practical implication doesn't wait on any of that being confirmed: **predictability and sensory intensity are different variables and should be intervened on separately.** A quiet room that's unpredictable can be worse than a noisy one that isn't.

What the evidence actually supports

The honest accounting, because the citation counts in this area may overstate the evidence.

The largest synthesis of Bayesian accounts of autism, 83 studies, reports that a slight majority find **no group difference** at all (Angeletos Chrysaitis & Seriès, 2023). The second largest, 47 studies, reaches the opposite conclusion (Cannon et al., 2021) — the two biggest reviews of this literature disagree about what it shows. The first meta-analysis of the simple model found a small effect, $g = 0.37$, that no moderator could explain, concluding the evidence offers "only limited support for a universal 'simple Bayesian model'" (Cui et al.,

2026). The flagship computational finding failed a preregistered test using the identical model (Plank et al., 2026). The mismatch negativity pools to a **near-null effect that reverses direction between children and adults** (Schwartz et al., 2018; Sapey-Triomphe et al., 2025).

The predictive account of stimming has been asserted in high-profile theory papers since 2014 and never tested. No study has measured autonomic state time-locked to spontaneous stimming in autistic people. The nearest attempt is from 1998.

The endorphin explanation is worse than untested. The historical opioid theory ran the opposite direction, proposing chronically elevated opioid tone, which is why the drug tested was an opioid *antagonist*. Naltrexone then got three decades of randomized trials and failed for core symptoms and stereotypy, with one well-powered trial finding stereotypy got **worse**.

Status, stated plainly

This is a mechanistic reframe. It has no outcome data. Nobody has tested whether thinking in these terms helps a single client. The autism application reinterprets data collected to answer other questions, which is legitimate and weaker than predicting a result in advance.

The model is also flexible enough to accommodate a great deal, which is the same objection the full document levels at the literature it reviews. The only defense is naming what it forbids.

It forbids repetitive behavior with no frequency structure. It forbids damping interventions failing while sensory reduction succeeds. It forbids observer and controller differences that never dissociate. It forbids gain manipulations leaving repetitive behavior untouched.

If those results come in, it's wrong and should be thrown out rather than patched.

What's worth doing now

Three things don't require waiting on any of it.

Ask the two questions separately. How confident is this belief, and how hard is this person acting on it. Different answers, different interventions, and asking costs nothing.

Treat predictability and intensity as distinct variables in sensory distress, instead of assuming less stimulation is always the answer.

And when a regulatory pattern looks maladaptive, ask what environment would have made it correct.

Full treatment, including the first-person accounts, the treatment evidence, the falsification criteria, and the eight ways this fails, is in Prediction, Control, and the Regulating Mind: A Clinician's Framework for Autism, Stimming, and Emotional Regulation. *The underlying evidence audit is in the companion literature review.*