JOSHYGLESIAS.COM

# The Upgradable Mind

## Humans as models

*Seeing your conscious experience as a model may help you thrive*

NOV 21, 2026 · BY JOSH YGLESIAS

A companion to my paper What's Actually Going On. I wrote it with Claude, Anthropic's AI.

# The one-minute version

I make three points.

1. **You can look at yourself the way a lab looks at an AI model.** Your thoughts and feelings are the surface of a trained system, like a chatbot's answers. Like an AI model, you have deep wiring, a stream of thoughts, tools and a track record. Seen like that, a thought or a mood is an output, not an order.
2. **Each layer can be upgraded, at its own speed.** Your tools and routines can change today. Your deep wiring changes slowly, mostly while you practice and sleep. So upgrade the tools for quick wins, and protect sleep, exercise and practice.
3. **You are already merging with machines, one tool at a time.** A deeper merge would amplify whatever you bring to it, so the upgrade comes first.

**What to do:** start seeing your brain as something you can upgrade. The steps are simple, even obvious: sleep 7+ hours, exercise, learn by testing yourself, write if-then plans, use AI as a tutor and review your week. Even so, it matters that you care and actually try, because the benefits of a better brain and a better version of you are so profound that they outweigh the cost by far.

Upgrading yourself is the closest thing to a bet that pays in every outcome. If AI goes well, a better you enjoys it more. If it goes badly, a better you copes better. If it never comes, you're simply better.

This is an analogy and a lens, not a claim that brains are built like AI systems. Section 1 says where it breaks.

# How to read the rest

The rest takes about 15 minutes. Sources are in parentheses and listed at the end. Where advice rests on studies, I grade the evidence, so you can see how sure I am. If you're reading as family, you can skip the Abstract and Key terms and start at section 1.

# Abstract

This is a companion to my paper *What's Actually Going On*, which is about AI and how to get ready for it. This one is about the person doing the getting ready. I look at a human being the way an AI lab looks at a model: a slow-changing base, a running stream of thought, a harness of tools and a record that can be measured. I use that frame for one job, which is sorting the levers for self-improvement by how fast they act and how we can measure them. I then rank ten levers by the strength of the evidence, from sleep and exercise to AI tutoring and merging

with machines, and I say where the analogy breaks. It's a synthesis, not a study, and it reports no new data.

*Keywords: self-improvement, learning, habits, human and AI models, brain-computer interfaces*

# Key terms

- **Model:** in AI, the trained program behind a system such as a chatbot. Applied to people in this paper, it is an analogy and a lens, not a claim that brains are built like AI systems. Section 1 says where it breaks.
- **Conscious experience:** everything you think, feel and notice from moment to moment.
- **Weights:** the numbers inside an AI model that store what it has learned.
- **Harness:** the tools and instructions wrapped around a model so it can act in the world.
- **Evals:** the tests used to judge a model.
- **Upgrade:** a change to your habits, tools, skills or body that measurably improves something you care about, judged over months, not days.
- **ASI:** AI that clearly beats the best humans at most thinking work, including the work of building better AI.
- **Standard deviation:** a yardstick for the size of an effect. As a rough guide, 0.2 is small, 0.5 is medium and 0.8 is large.
- **Median:** the middle value when the answers are lined up from lowest to highest.

**HOW THE EVIDENCE IS GRADED**

| GRADE | WHAT IT MEANS |
|---|---|
| Strong | Several large trials, or analyses that combine many trials, that agree |
| Moderate | Trials exist, but results vary or shrink under scrutiny |
| Emerging | A few early studies that do not yet point one way |
| Speculative | Argument and early prototypes only |

I assigned these grades myself, from the size and agreement of the best combined studies I could find. Anything I call "my guess" is exactly that.

1.

# The Upgradable Mind: humans as models

*In one line: your thoughts and feelings are the surface of a trained system, the way a chatbot's answers are the surface of an AI model. That system has a slow-changing base, a running stream, a set of tools and a track record, and each can be upgraded at its own speed.*

My other paper, *What's Actually Going On*, is about the world. This one is about you. By conscious experience I mean everything you think, feel and notice from moment to moment. I think you can look at it the way a lab looks at one of these models. Doing so may help you thrive, and by that I mean three things: experience life well, grow because you want to and treat people well. If AI ends up doing most of the work, that may be the main thing left for us to get good at.

When a lab ships a model, it has a base trained on data, a way of running that produces outputs, a harness that connects it to the world, and evals, the tests that judge it. Building one looks a lot like growing up. First comes pretraining, where the model reads a huge library and learns to predict what comes next, like school. Then it practices tasks, gets graded and learns to use tools, like job training. Then it goes to work. Labs call that last stage inference: copies of the finished model running to get things done. I think a human can be seen the same way. What you experience is the model running, on a base you can't see directly. Seen like that, a thought or a mood is an output, not an order. And the useful question becomes how to ship a better version of yourself.

Neuroscience backs the analogy more than you'd expect. The brain looks like a prediction machine: it constantly guesses what comes next and updates on its errors (Clark, 2023). That is loosely how a language model learns, by predicting the next word and correcting its errors. The overlap can be measured. Language models trained only to predict the next word turned out, in one large 2021 comparison, to be the best available predictors of human brain activity during language (Schrimpf et al., 2021). Later work disputes what such scores show (Antonello and Huth, 2024), so I take this as a hint, not proof. Brain recordings also show people predicting upcoming words before they hear them (Goldstein et al., 2022).

Most of the model also runs below awareness. Sigmund Freud made the unconscious famous: the idea that much of what drives us happens out of sight (Freud, 1915). His specific theories, like repression and the hidden meaning of dreams, are mostly not supported by research today. The broad point holds up. Daniel Kahneman describes two modes of thinking: a fast, automatic System 1 that runs nearly all the time, and a slow, effortful System 2 that we mistake for the whole of us (Kahneman, 2011). In two diary studies, about 35% to 43% of people's daily actions were habits, often done while thinking about something else (Wood, Quinn and Kashy, 2002). We are often strangers to ourselves, explaining our choices after

we've made them (Nisbett and Wilson, 1977; Wilson, 2002). The train of thought you notice is the visible surface of a much bigger river.

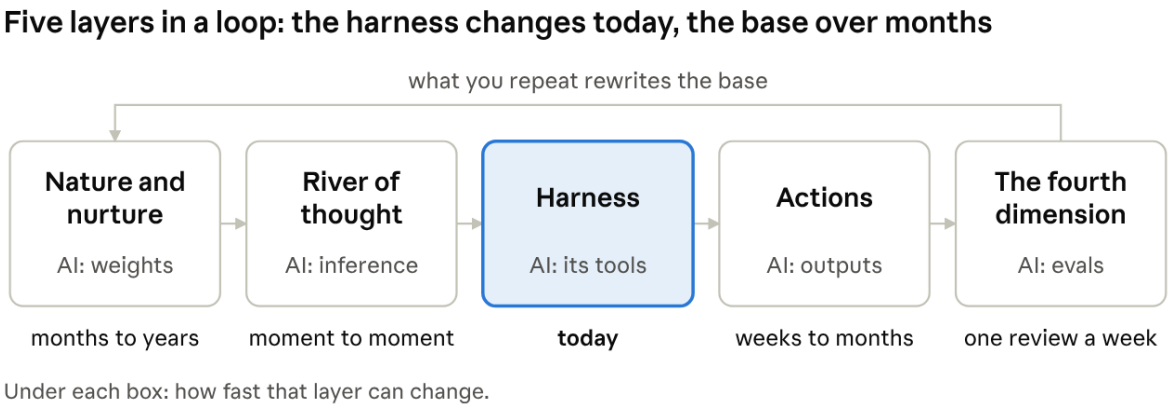


*Figure 1. The Upgradable Mind: five layers, one feedback loop.*

The layers run left to right, and the record feeds back into the base. The table gives the detail.

| LAYER | IN ME | IN AN AI MODEL | MAIN UPGRADE LEVER | HOW FAST IT CHANGES |
|---|---|---|---|---|
| 1. Nature and nurture | Genes, upbringing, everything I've absorbed | Pretraining data and weights | Sleep, exercise, deliberate practice | Slow: months to years |
| 2. River of thought | My running stream of thoughts | Inference, the model running | Attention training, meditation | State: moment to moment. Skill: months |
| 3. Harness | How I express and organize what's in my head | Tools, prompts, memory, interface | Environment design, writing, AI tools | Fast: today |
| 4. Actions | What I actually do, also a river | Outputs | If-then plans, habits | Weeks to months |
| 5. The fourth dimension | How I act across time, seen all at once | Evals across many runs | Weekly review, tracking | One review a week |

**Weights versus harness.** My weights, the deep wiring of my brain, change slowly through neuroplasticity, the brain's ability to rewire itself, mostly while I practice and sleep. My harness, the tools, notes, routines and people around me, can change today. That's the key practical insight: upgrade the harness for quick wins, and protect the slow processes that rewrite the weights.

**Why evals matter.** Neither kind of model can read its own weights. AI researchers call the missing skill interpretability, and it is unsolved. Humans have the same gap: we feel sure about our reasons and are often wrong (Nisbett and Wilson, 1977). So I trust the record over the story. What I did across many days tells me more than what I believe about myself.

**Where the analogy breaks.** A brain learns all the time from a handful of examples, runs on about 20 watts, lives in a body and feels things. Whether an AI model experiences anything at all is an open question. A 2023 review by scientists and philosophers judged that the systems of that time probably did not, and that nothing obvious rules it out for future ones (Butlin et al., 2023). Nothing in this paper depends on the answer. A child hears perhaps a hundred million words growing up, and a large language model reads thousands of times more (Frank, 2023). A brain has no clean line between weights and running: thinking changes the wiring as it goes. A language model mostly stops learning once training ends. The brain as a prediction machine is an influential theory, not a settled one, and Kahneman's two systems are a simplification too. So I use the model frame for one job: sorting the levers by how fast they act and how we can measure them. Where the frame and the biology disagree, the biology wins.

One more limit. A model is judged only by its outputs. A person is not. The upgrade is something I want, not a score I owe anyone.

2.

# You: the upgrade protocol

This is my best answer to "if you want to upgrade your model, what should you do?" It is ranked by how strong the evidence is, using the grades in Key terms.

One rule runs through the whole table. If most of the model runs below awareness, willpower alone is a weak lever, because it only reaches the small conscious part. The strong levers work on the rest: sleep, exercise, habits tied to a cue, an environment that makes the right move the easy one, and practices like pausing and meditation that let you notice the automatic part before it acts.

Dostoevsky adds an honest warning. The narrator of Notes from Underground argues that people will act against their own interest just to prove they are free, not piano keys played by reason (Dostoevsky, 1864). Any protocol that treats you as a purely rational machine will run into that part of you. So build the plan to be chosen, not imposed: something you want to do, not something you have to do.

Is this a privilege? Partly. Seven hours of sleep, time to exercise and a quiet place to practice are easier with money, flexible work and no small children, and a night-shift worker can't simply choose a steady sleep window. Two things soften that. The strongest levers are also the cheapest: sleep, walking, recall practice and if-then plans cost little or nothing. And AI tutoring may push the price of a decent tutor toward zero, though the evidence on how well it works is still mixed. But the gap in free time is real, and it is one more reason the political question matters: who owns these systems and who shares in the gains. I take that up in *What's Actually Going On*.

| LEVER | EVIDENCE | THE ACTION |
|---|---|---|
| Sleep | Strong | 7+ hours a night. Keep the times steady if you can |
| Aerobic exercise | Strong for health and mood. Moderate for thinking | About 150 minutes a week of brisk cardio plus 2 strength days |
| Testing and spacing | Strong | Close the book and recall. Review at growing gaps, for example after 1, 3, 7 and 21 days |
| Deliberate practice | Strong, with limits | Pick the hardest useful sub-skill, practice it, get feedback, repeat |
| If-then plans | Moderate to strong | Write one "if X, then I do Y" plan per goal |
| Environment | Moderate | Make good actions easy and bad ones hard |
| People | Moderate for learning. Strong for health | Spend time with people who are better than you. Find mentors. Invest in close friends |
| Meditation | Moderate for mood. Emerging for attention | For example, 10 minutes a day of watching the breath and noticing when you drift |

| LEVER | EVIDENCE | THE ACTION |
|---|---|---|
| AI as tutor | Emerging, and mixed | Think and write first. Then have AI quiz you, critique you and explain your mistakes |
| Merging with machines | Speculative | Watch it. Don't bet your life on it yet (section 3) |

**What the research says.** The numbers in parentheses are the layers of the Upgradable Mind that each lever works on.

- **Sleep (layers 1, 2).** Sleep locks in memory and learning (Diekelmann and Born, 2010). Two weeks at 6 hours a night hurt performance about as much as one to two nights with no sleep, and people didn't notice (Van Dongen et al., 2003). Sleep scientists recommend 7 or more hours for adults (Watson et al., 2015). These studies are about how long you sleep. Keeping the times steady is my own addition.
- **Aerobic exercise (layer 1).** It lowers the risk of early death and is an effective treatment for depression (World Health Organization, 2020; Noetel et al., 2024). Many trials find gains in thinking and brain health (Hillman et al., 2008; Erickson et al., 2011), but the average gain in healthy adults shrinks once weaker studies are filtered out (Ciria et al., 2023).
- **Testing and spacing (layers 1, 4).** Recalling something beats rereading it (Roediger and Karpicke, 2006). Spacing reviews beats cramming (Cepeda et al., 2006). Rereading and highlighting rate low (Dunlosky et al., 2013).
- **Deliberate practice (layers 1, 4).** Focused practice at the edge of your ability, with feedback, builds skill (Ericsson et al., 1993). It explains less of the gap between people than once claimed (Macnamara et al., 2014).
- **If-then plans (layer 4).** These plans have a medium to large effect on following through (Gollwitzer and Sheeran, 2006). That pooled estimate predates modern checks for publication bias, the tendency of journals to print the studies that worked, so the true effect is probably smaller. An example: if it is 10 p.m., then my phone goes on the charger in the kitchen. Habits took a median of 66 days to become automatic, with a wide range (Lally et al., 2010).
- **Environment (layer 3).** Making the good action the easy one works in many trials (Mertens et al., 2022), but the average effect shrinks sharply after correcting for publication bias (Maier et al., 2022).
- **People (layers 1, 3).** One-to-one tutoring lifted students about two standard deviations, a very large effect, in a classic study (Bloom, 1984). A large later review found a smaller average effect, about 0.37, which is small to medium (Nickow et al., 2020). Close relationships predict health and survival (Holt-Lunstad et al., 2010; Waldinger and Schulz, 2023). Seeking out people who are ahead of you is my own extension of the tutoring result.

- **Meditation (layer 2).** There is moderate evidence for less anxiety, depression and pain, and weak evidence so far for sharper attention (Goyal et al., 2014). The trials used longer, structured programs, so treat ten minutes as a starting dose.
- **AI as tutor (layer 3).** A Harvard physics course found an AI tutor beat an active-learning class (Kestin et al., 2025). Across 14 experiments the average gain on tests taken without AI was modest, about 0.19 standard deviations (Contractor and Reyes, 2026). It can backfire. High school students given open AI help did worse once it was taken away, while a version built to tutor avoided most of that harm (Bastani et al., 2025). In one university trial, access to an AI tutor lowered final grades (Liu et al., 2026). So use it to train the model, not replace it.

**A simple week.**

- **Every day:** sleep window protected, 10 minutes of attention practice, one block of focused practice, and recall what you learned instead of rereading it.
- **Three to five times a week:** cardio, and strength twice.
- **Once a week:** a review, your own eval. What did you do this week, seen from the fourth dimension? What should next week's version change?

**Evals.** Labs judge a model by how it does across many runs, not one answer. Judge yourself the same way. Track a few numbers over months: hours slept, workouts, hours of deliberate practice, and one outcome you care about, like an exam score or a skill. One bad day is noise. A trend is signal. Tracking has strong evidence of its own: across 138 experiments, people who were prompted to monitor their progress got closer to their goals, and more so when they wrote the results down (Harkin et al., 2016).

3.

# Merging: the version of you that meets it

*In one line: you are already merging with machines, one tool at a time. A deeper merge would amplify whatever you bring to it, so the upgrade comes first.*

Your mind already reaches past your skull. A notebook or phone you rely on is part of how you think (Clark and Chalmers, 1998). In one well-known study, people with a search engine at hand remembered where to find a fact better than the fact itself (Sparrow, Liu and Wegner, 2011). One experiment from that paper later failed to replicate, so treat it as suggestive (Camerer et al., 2018). That's a merge. It just runs through your thumbs.

The same merge can weaken you. In one study, clinicians who had spent months using AI to help spot tumors found fewer of them when working without it (Bengio et al., 2026). This is the tutor rule from section 2 again, at higher stakes: a tool that does the task for you can quietly retire the skill.

Brain-computer links are the next step. Neuralink put its first implant in a human in January 2024, and by early 2026 it reported 21 people enrolled in its trials (Newton, 2026). So far these are medical devices for people with paralysis and similar conditions, not upgrades for healthy minds. Ray Kurzweil argues we will eventually merge with AI (Kurzweil, 2024).

If merging happens, it will scale what you bring to it. A focused, healthy, curious mind plugged into ASI becomes more of that. A scattered, exhausted one becomes more of that too. That is why the upgrade protocol comes first. The same goes for the age of abundance I expect if AI ends up doing most of the work, which I describe in *What's Actually Going On*: if the work is done for us, what we do with ourselves is the whole question.

The honest objections:

- **Who controls it?** A merge built and owned by a few companies brings the winner-take-all problem inside your head.
- **Who gets access?** If only the rich upgrade, inequality becomes biological.
- **Is it still you?** Change enough of the model and it may not be Josh anymore. Philosophers split here. If you are your continuous stream of memory and character, a gradual merge that keeps that stream could still be you, the way you are still you after years of changing your mind, your habits and much of your body. If you are something that can't be copied or shared, a merge could end you even while something that remembers being you carries on (Parfit, 1984). I can't settle that, so my rule is to prefer upgrades I can reverse.
- **Can you undo it?** Most upgrades should be reversible. This one may not be.
- **It may never work** the way the hopeful version says.

4.

# What this paper can't prove, and how it was made

*In one line: this is one student's lens, dated November 2026. Here is how far to trust it.*

- **The model analogy is a lens, not a law.** Brains are not neural networks, and the comparison will break somewhere. Section 1 lists the places I know of.
- **This is a synthesis, not a study.** I ran no experiments. I am a finance student, not an AI researcher, a neuroscientist or a philosopher, and specialists will find places where I have simplified.
- **The studies have limits.** The studies behind section 2 are peer reviewed, but most follow students or volunteers for weeks, not whole lives.
- **Upgrading yourself doesn't solve politics.** Who owns ASI is decided by laws, votes and institutions, not by sleep schedules.
- **I chose the sources.** I have tried to put the strongest caution next to each claim, but a reader should check me.
- **It is a snapshot.** In this field a year is a long time. Check the dates on everything, including this.

**How it was made.** I wrote this paper with an AI system: Claude, made by Anthropic. It did much of the source-finding, fact-checking and drafting, working from my ideas and my first version. The views, the bets and any mistakes are mine.

**Method.** Sources were chosen in this order: the original paper, report or data page first, then independent measurers, then news reports, used only for dated events. This text began as the fifth part of *What's Actually Going On*. As part of that paper it went through five rounds of review by fresh copies of the same AI system, which opened sources and checked a sample of claims against them across the whole paper. I split it out on November 21, 2026, so each paper can be read on its own. What I changed for the split, which is the opening, this section and the cross-references, has been checked only by me. That is a check, not peer review.

# Conclusion

Look at your own conscious experience the way a lab looks at a model, and upgrade it. That's the move that comes closest to paying no matter what happens. Sleep, move, practice, test yourself, shape your environment, choose your people, train your attention, and use AI to make you sharper rather than lazier.

You can't pick the future. You can pick the version of you that meets it.

# References


Antonello, R., and Huth, A. (2024). Predictive coding or just feature discovery? An alternative account of why language models fit brain data. *Neurobiology of Language*, 5(1), 64-79.

Bastani, H., et al. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *PNAS*, 122(26).

Bengio, Y., et al. (2026). *International AI Safety Report 2026.* arXiv:2602.21012.

Bloom, B. S. (1984). The 2 sigma problem. *Educational Researcher*, 13(6), 4-16.

Butlin, P., Long, R., et al. (2023). *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness.* arXiv:2308.08708.

Camerer, C. F., et al. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637-644.

Cepeda, N. J., et al. (2006). Distributed practice in verbal recall tasks. *Psychological Bulletin*, 132(3), 354-380.

Ciria, L. F., et al. (2023). An umbrella review of randomized control trials on the effects of physical exercise on cognition. *Nature Human Behaviour*, 7, 928-941.

Clark, A. (2023). *The Experience Machine.* Pantheon.

Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7-19.

Contractor, Z., and Reyes, G. (2026). *Experimental Evidence on the Learning Impact of Generative AI.* arXiv: 2607.08849.

Diekelmann, S., and Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, 11(2), 114-126.

Dostoevsky, F. (1864). *Notes from Underground.*

Dunlosky, J., et al. (2013). Improving students' learning with effective learning techniques. *Psychological Science in the Public Interest*, 14(1), 4-58.

Erickson, K. I., et al. (2011). Exercise training increases size of hippocampus and improves memory. *PNAS*, 108(7), 3017-3022.

Ericsson, K. A., Krampe, R. T., and Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363-406.

Frank, M. C. (2023). Bridging the data gap between children and large language models. *Trends in Cognitive Sciences*, 27(11), 990-992.

Freud, S. (1915). The unconscious.

Goldstein, A., et al. (2022). Shared computational principles for language processing in humans and deep language models. *Nature Neuroscience*, 25(3), 369-380.

Gollwitzer, P. M., and Sheeran, P. (2006). Implementation intentions and goal achievement. *Advances in Experimental Social Psychology*, 38, 69-119.

Goyal, M., et al. (2014). Meditation programs for psychological stress and well-being. *JAMA Internal Medicine*, 174(3), 357-368.

Harkin, B., et al. (2016). Does monitoring goal progress promote goal attainment? A meta-analysis of the experimental evidence. *Psychological Bulletin*, 142(2), 198-229.

Hillman, C. H., Erickson, K. I., and Kramer, A. F. (2008). Be smart, exercise your heart. *Nature Reviews Neuroscience*, 9(1), 58-65.

Holt-Lunstad, J., Smith, T. B., and Layton, J. B. (2010). Social relationships and mortality risk: A meta-analytic review. *PLoS Medicine*, 7(7), e1000316.

Kahneman, D. (2011). *Thinking, Fast and Slow.* Farrar, Straus and Giroux.

Kestin, G., et al. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15. https://pmc.ncbi.nlm.nih.gov/articles/PMC12179260/

Kurzweil, R. (2024). *The Singularity Is Nearer.* Viking.

Lally, P., et al. (2010). How are habits formed. *European Journal of Social Psychology*, 40(6), 998-1009.

Liu, J., et al. (2026). *The Effects of Course-Integrated AI Tutoring on Student Performance and Engagement: A Randomized University Trial.* EdWorkingPaper 26-1598, Annenberg Institute at Brown University. https://doi.org/10.26300/y3f8-vh05

Macnamara, B. N., Hambrick, D. Z., and Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions. *Psychological Science*, 25(8), 1608-1618.

Maier, M., et al. (2022). No evidence for nudging after adjusting for publication bias. *PNAS*, 119(31), e2200300119.

Mertens, S., Herberz, M., Hahnel, U. J. J., and Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *PNAS*, 119(1), e2107346118.

Newton, C. (2026). Neuralink reaches 21 patients as Elon Musk continues push for high-volume brain chip production. *The Debrief*, February 16. https://thedebrief.org/neuralink-reaches-21-patients-as-elon-musk-continues-push-for-high-volume-brain-chip-production/

Nickow, A., Oreopoulos, P., and Quan, V. (2020). The impressive effects of tutoring on PreK-12 learning. NBER Working Paper 27476.

Nisbett, R. E., and Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231-259.

Noetel, M., et al. (2024). Effect of exercise for depression: Systematic review and network meta-analysis of randomised controlled trials. *BMJ*, 384, e075847.

Parfit, D. (1984). *Reasons and Persons.* Oxford University Press.

Roediger, H. L., and Karpicke, J. D. (2006). Test-enhanced learning. *Psychological Science*, 17(3), 249-255.

Schrimpf, M., et al. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *PNAS*, 118(45), e2105646118.

Sparrow, B., Liu, J., and Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776-778.

Van Dongen, H. P. A., et al. (2003). The cumulative cost of additional wakefulness. *Sleep*, 26(2), 117-126.

Waldinger, R., and Schulz, M. (2023). *The Good Life.* Simon & Schuster.

Watson, N. F., et al. (2015). Recommended amount of sleep for a healthy adult: A joint consensus statement of the American Academy of Sleep Medicine and Sleep Research Society. *Sleep*, 38(6), 843-844.

Wilson, T. D. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious.* Harvard University Press.

Wood, W., Quinn, J. M., and Kashy, D. A. (2002). Habits in everyday life: Thought, emotion, and action. *Journal of Personality and Social Psychology*, 83(6), 1281-1297.

World Health Organization (2020). *Guidelines on Physical Activity and Sedentary Behaviour.*

# Read next

- Josh Yglesias, *What's Actually Going On* (2026): the companion to this paper, on what AI is doing to the world and how to get ready for it.
- Andy Clark, *The Experience Machine*: the brain as a prediction machine.
- Daniel Kahneman, *Thinking, Fast and Slow*: the fast and slow modes of thinking.