NEURO-RIGHTS · FROM PRINCIPLE TO STATUTE

After the Frontier

Writing the Law for the Mind — A Sequel

The first essay argued that the mind had become, for the first time, a place that could be measured from the outside — and that this made it a problem for law. This one takes the next step. If we accept that the inner forum must be protected, what exactly do we write down, at which level, against which threats, and starting on which Monday morning? This is a guide to the answers: the eight routes by which the mind is reached today, the model statutes that close them, the geopolitics that will decide whether any of it holds, and the new capabilities — from neural inference to synthetic telepathy — that the next two decades will force onto the statute book.

A long read · policy & ethics · written for legislators, policymakers and concerned citizens · reads as a standalone

CONTENTS

1 Where We Left Off — And Why Principles Are Not Enough

2 The Eight Doors — How the Mind Is Actually Reached in 2026

3 The Inference Problem, Restated as the Central Drafting Choice

4 The Legislative Stack — Six Tiers, One Architecture

5 What Goes in the Constitution

6 What Goes in the Statute — Prohibitions, Remedies, a Regulator

7 The Inalienable Core — Rights You Cannot Sign Away

8 The Black Box for the Brain — Making Abuse Provable

9 Whistleblowers — The People Who Make Oversight Real

10 The Map of Power — Six Blocs and Four Horizons

11 Synthetic Telepathy — Legislating Before Deployment

12 The Cases That Are Coming — Where Today's Law Already Bites, and Where It Breaks

13 Victims — A Standard of Care That Tells the Difference

14 The Citizen's Toolkit — What You Can Do Before the Law Catches Up

15 Three Futures, and the Choice Between Them

- In Closing — The Same Window, Now Narrower

Published as part of the TI One Voice initiative. TI One Voice (one.witysk.org) is an initiative founded by Stephane van der Aa to give a unified voice to the community of so-called "Targeted Individuals" (TIs) — people who report being subjected to organised harassment, surveillance, and what they describe as remote forms of physical and psychological abuse. The platform brings together their testimonies and supporting documentation in a single, structured place, so that experiences which are too often dismissed or scattered can be recorded, preserved, and examined seriously. Its goal is to document these accounts rigorously, to support the individuals affected, and to advocate before Belgian, European, and international institutions for recognition, protection, and proper investigation.

1. Where We Left Off — And Why Principles Are Not Enough

A companion essay, *The Last Private Place*, made a single argument at length: that neurotechnology and artificial intelligence together have made the interior of a human being legible from the outside in a way no prior technology approached, and that the rights we have always assumed protected the mind were written for a world in which this was impossible. It ended where most good arguments end — with a principle. The mind has been private for as long as there have been minds; whether it stays that way is now, for the first time, a policy choice.

This essay begins where that one stopped, because a principle is not a law. "The inner forum must be protected" is a statement nobody on a parliamentary committee will vote against and nobody on a product team will be bound by. Between the principle and the protection lies the unglamorous work that actually decides outcomes: the definition of a term, the choice of which body enforces it, the penalty schedule, the carve-out an industry lobby asks for at the last reading. The history of technology regulation is a history of arriving late — we wrote the rules for the car after the deaths, for finance after the crashes, for social media after the damage was done. Neurotechnology offers the rare chance to legislate before the harm becomes structural. Taking that chance means getting specific.

So this is a drafting guide as much as an essay. It assumes no memory of the first one. It restates what a reader needs — what neurotechnology is, why now, what the five neuro-rights are — in passing, and then turns to the harder questions: *which* law, at *which* level, against *which* threat, enforced by *whom*. Where the evidence is solid, it says so and cites it. Where it is contested or absent — the exotic weapons, the classified capabilities, the strongest claims of covert control — it says that too, and resists resolving the uncertainty in either direction for effect. On a subject that touches the deepest fears and hopes we have about our own minds, that discipline is not a hedge. It is the whole point.

> **The five-second refresher.** Neurotechnology is any tool that measures, interprets, modulates or stimulates the nervous system — from a cochlear implant to a meditation headband to an experimental cortical implant. The "neuro-rights" most legislation now turns on are five: **mental privacy** (keep neural data from disclosure), **personal identity** (an intact, continuous sense of self), **free will / agency** (decisions that are genuinely one's own), **fair access to augmentation** (enhancement without a new caste), and **protection from neural-data discrimination**. In November 2025 the world's governments adopted the first global standard on the subject — UNESCO's Recommendation on the Ethics of Neurotechnology — but it is "soft law": urged, not binding. Turning it into something with teeth is what the rest of this essay is about.

2. The Eight Doors — How the Mind Is Actually Reached in 2026

"Mind control" is a loaded phrase, and the cinematic version — a remote operator seizing irresistible command of a person's behaviour — is also the version least supported by evidence. It is more useful, and more honest, to define the risk broadly: *any technological capability that materially impairs a person's ability to form, hold or act on their own thoughts, beliefs and decisions, free of coercion or of manipulation that bypasses their conscious will.* By that definition the risk is not a future possibility. It is a present condition, and it arrives through eight distinct doors.

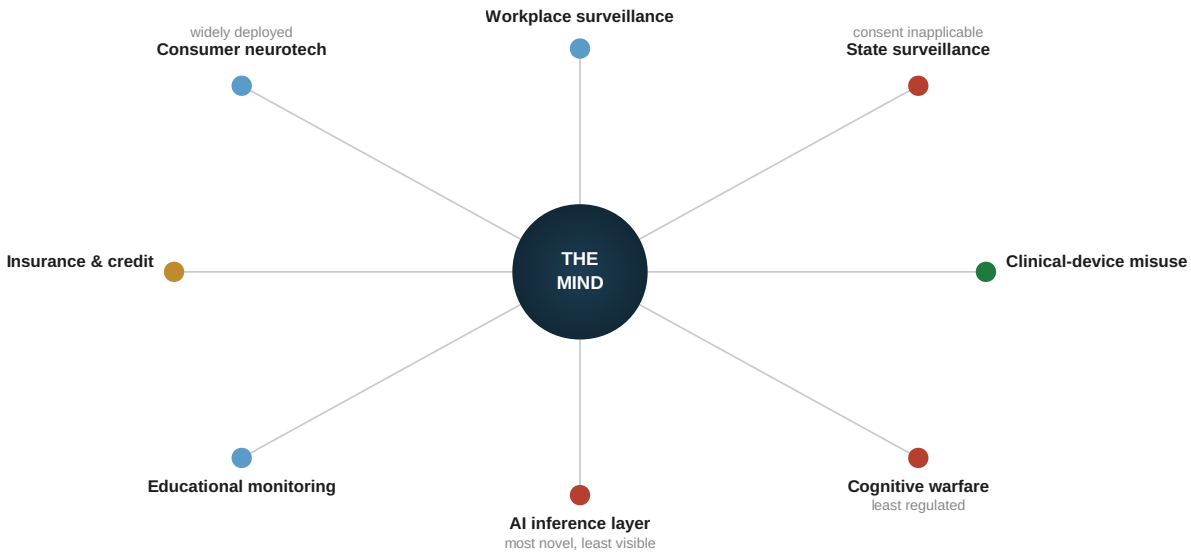


Figure 1. The eight principal routes by which the mind is reached in 2026. Consumer neurotech and workplace surveillance are the most widely deployed; the AI-inference layer and cognitive-warfare vectors are the most novel and the least regulated.

1. Consumer neurotechnology. The most prosaic door is the wearable on sale today. A 2024 study of thirty consumer neurotech companies found that twenty-nine reserved broad rights over the brain data their devices collect, with essentially no meaningful limits and the right to share with third parties. No single headband performs covert mind control; the risk is that millions of them generate a continuous stream of data which, aggregated and fed to ever-better models, becomes a window the user never agreed to open. Chile's Emotiv case — a meditation headset whose data a constitutional court ordered deleted — is the proof of concept.

2. Workplace neurosurveillance. Here the danger is not exotic but structural: a manager with a power asymmetry, equipped with a tool that claims to measure attention, fatigue and stress in real time. Consent is coerced by economic necessity; data gathered for "wellness" migrates to performance scoring and termination; and the slow-acting harm is the worker who suppresses idle thought because thought is being scored. The EU AI Act's ban on workplace emotion-recognition is the only consequential rule anywhere on this — and even it has safety-critical exemptions large enough to drive trucks through.

3. State surveillance. The most consequential door, because consent does not even arguably apply. Credibly imminent or already-extant uses include neural "deception detection" at borders, neural data folded into social-credit systems, EEG in prisons under the banner of rehabilitation, and interrogation augmented by brain-state monitoring. The leading edge is visible; analogous capability is being built by other states under varying transparency.

4. Clinical-device misuse. Medicine is the most regulated tier and still leaks: off-label use of deep-brain stimulation for behaviour rather than the approved indication; data exfiltrated from implant-makers' servers; the device that remains in a patient who has withdrawn consent but cannot easily have it explanted; closed-loop systems that adjust a patient's state in real time on opaque inferences.

5. Cognitive warfare. A soft end and a hard end. The soft end — disinformation and algorithmic manipulation of attention and emotion at scale — is unquestionably happening and touches every networked person. The hard end — hypothesised directed-energy or acoustic devices, of which the Havana Syndrome cluster is the most-discussed example — is genuinely contested. The responsible posture, developed in Section 10, is to treat the

doctrine as real (because NATO and adversary documents demonstrably make it so) without overcommitting to claims about specific deployed weapons.

6. The AI inference layer. The most novel door, and the one invisible to users: foundation models applied to neural and behavioural data already collected. Someone who shared sleep-tracking data in 2022 did not consent to what a 2026 model can extract from it. Inferences leap ahead of consent. This is why protecting only the raw signal protects the lock and not the valuables — the subject of the next section.

7. Educational monitoring. Children in classrooms cannot meaningfully consent. Attention-tracking headbands — most famously deployed in some Chinese classrooms but spreading — chill developing minds, disadvantage neurodiverse children whose attention patterns differ, and establish a rolling baseline of monitoring that runs from school into university into work. These harms are the slowest-acting and may be the most consequential, because they shape who a generation becomes.

8. Insurance and credit. The most economically structural door. Insurers and lenders have strong incentive to fold neural-derived inferences — sleep quality, attention patterns, mood — into risk scoring once the actuarial value is shown. Most jurisdictions prohibit the obvious forms, but the prohibitions are leaky and wellness programmes are an established workaround. The result is a new and largely invisible redlining: people whose neural patterns deviate from a normative profile pay more, often without knowing why.

> **Different doors, different adversaries.** The eight vectors imply a portfolio of adversaries, not a single villain. **Commercial actors** present a risk of incentive misalignment — they intend to monetise, not to harm — and the remedy is regulatory: defining what may not be monetised or repurposed. **States** present capability concentration; the remedy is constitutional entrenchment and international constraint. **Criminal and abusive private actors** — stalkers, ex-partners, lawless employers — present low-cost capability dispersion as devices cheapen, and the remedy is criminal-law innovation. **Foreign and non-state adversaries** present attribution difficulty by design, and the remedy is collective defence. A single statute cannot answer all four. That is why the architecture has to be layered.

3. The Inference Problem, Restated as the Central Drafting Choice

If a legislator remembers one sentence from this essay, it should be this: **the single most consequential clause in any neuro-rights bill is the definition of "neural data," and the question that clause must answer is whether the term includes the algorithmic inferences drawn from the signal, or only the raw signal itself.**

The reason is that the inference is almost always the part that can hurt you. The raw electrical trace from an EEG headband is, on its own, a meaningless smear. It becomes a window into your mood, your attention, your political reaction, your likely solvency only when a model interprets it. A law that protects the signal but not the inference has locked the door and left the valuables on the porch. The US state laws already disagree on exactly this point — California excludes data an algorithm derives from neural signals; Colorado includes it; Montana excludes "downstream physical effects" — and these are not pedantic differences. They decide whether the law reaches the harm at all.

The inference problem also dissolves the comfort of consent. Consent is necessarily consent to what the data could reveal *at the time it was given*. "Foundation models for the brain," trained on large pooled neural datasets, promise to decode new individuals with little calibration — which means each person’s data, pooled with

millions of others, makes everyone easier to read, including people who never consented at all. The migration of the risk is therefore from what is *measured* to what is *inferred*, precisely the category most existing law fails to protect. A neuro-rights regime that does not govern the AI applied to neural data is guarding a locked door beside an open window.

The drafting fix is blunt and should be copied verbatim into every tier: *"Neural data" means (a) any data obtained by direct or indirect measurement of the nervous system of a natural person, by any modality; and (b) any algorithmic inference derived from such data.* Everything downstream — consent rules, deletion rights, the ban on commercial transfer — should attach to both limbs of that definition or it attaches to nothing that matters.

4. The Legislative Stack — Six Tiers, One Architecture

Neural data crosses borders at the speed of a cloud upload: a Chilean's brain data was processed on servers in San Francisco. Protection at any single level is therefore necessary and never sufficient. The realistic answer is to legislate at every level at once, with each tier doing the job only it can do. Six tiers form the stack.

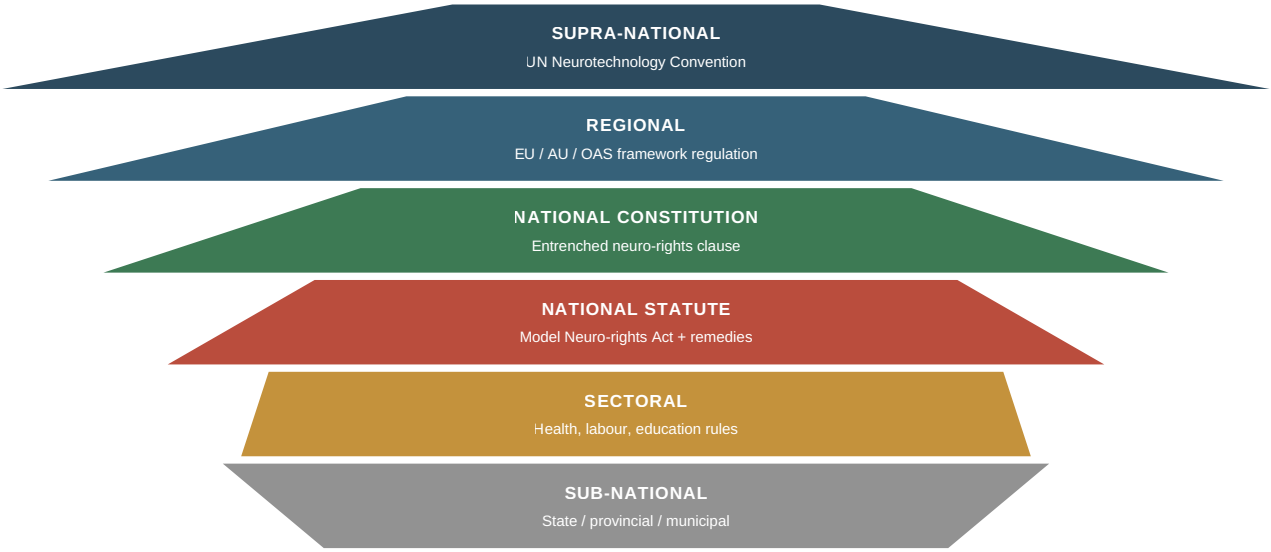


Figure 2. The ideal legal stack — six tiers, one architecture. Each tier does the job only it can do; a gap in one must be compensated by the others. The Chilean Emotiv remedy worked because the constitutional clause and the data-protection statute delivered it *together*; either alone might have failed.

The **supra-national** tier sets the prohibitions no state can violate and supplies the enforcement layer for cross-border and weapons questions no nation can solve alone. The **regional** tier harmonises within a bloc and operationalises the global instrument. The **national constitutional** tier entrenches the foundational right beyond ordinary political bargaining — the level only a constitution can reach. The **national statute** supplies the operative prohibitions, the remedies and the regulator. The **sectoral** tier translates the rules into the granular specifications that actually bind a doctor, an employer or a school. The **sub-national** tier captures local context and, in many jurisdictions, is where class actions and aggressive protection are most achievable.

The crucial property of the stack is that gaps in one tier can be compensated by the others. A jurisdiction without a constitutional clause can still achieve serious protection through aggressive sectoral regulation, as the EU has shown for data protection. A jurisdiction with a constitutional clause but weak statutory implementation will watch the constitutional protection erode into uncertainty. No tier is a substitute for the rest; each is load-bearing. The following four sections walk up the stack from constitution to enforcement.

Tier	Instrument	Force	Unique job
Supra-national	UN Neurotechnology Convention (proposed)	to build	Prohibit categories no state can police alone; verification; cross-border data & weapons
Regional	EU Neurotechnology Regulation; AU / OAS protocols	emerging	Harmonise within the bloc; operationalise the global instrument
National constitution	Neuro-rights clause (Chile model)	binding	Entrench the foundational right beyond ordinary politics
National statute	Model Neuro-rights Act	binding	Operative prohibitions, remedies, the regulator
Sectoral	Health / labour / education rules	binding	Translate rules into product- and practice-level specifications
Sub-national	State / provincial / municipal codes	binding	Local context; procurement; class-action venue

5. What Goes in the Constitution

The foundational protection belongs at the constitutional level, because only a constitution can make a right resistant to the ordinary political bargaining that erodes statutes. Chile proved it can be done: its 2021 amendment to Article 19 gave brain activity explicit constitutional protection, and its Supreme Court applied that protection two years later to order a US company to delete a citizen's brain data. But constitutional text alone is too abstract to bind a product manager, so the clause must be drawn tightly enough to anchor the statute beneath it.

A model clause, building on the Chilean text and the lessons of the Brazilian and Spanish processes, has five working parts. First, a **right to the integrity of the mind**, expressly including the inviolability of mental content, mental privacy, cognitive autonomy, personal identity, and freedom from discrimination on the basis of neural characteristics. Second, a statement that neurotechnology which reads or writes to the mind **shall be subject to law**. Third — and this is the part most constitutions omit — a short list of practices **prohibited absolutely, not authorisable by any law or contract**: involuntary reading of neural activity; involuntary modulation of the nervous system; neural monitoring of children in education; neural monitoring of workers as a condition of employment; neural interrogation of people deprived of liberty; and neural influence intended to alter electoral choice. Fourth, a declaration that **neural data is inalienable** — donatable for research or medicine, never commercialised. Fifth, an **effective-remedy guarantee**: compensation, deletion, and protection of the violated right going forward. Rights without remedies are decoration; the Chilean case mattered precisely because a court ordered data deleted.

6. What Goes in the Statute — Prohibitions, Remedies, a Regulator

Below the constitutional clause, the national statute does the operative work. A model Neuro-rights Act has four moving parts, and a legislator can check a draft against each.

Definitions that close the inference gap. "Neural data" must include algorithmic inferences (Section 3). "Cognitive autonomy" should be defined as the capacity to form, hold and act on thoughts and decisions free of

coercion or technological manipulation that bypasses the conscious will — so that the harm is named, not merely implied.

Substantive prohibitions with teeth. No reading of neural activity without free, specific, informed and revocable consent; no modulation beyond an approved therapeutic purpose without such consent; no employer use of neurotech as a condition of employment; no school monitoring of a child's attention, emotion or cognition; no neural influence in the electoral process; no neural interrogation of detainees. Each should carry a real criminal penalty, doubled where the victims are children, detainees or other vulnerable populations, and escalated again where the violation forms a pattern affecting a hundred or more people.

Remedies people can actually use. A private right of action with actual *and* statutory damages (so a plaintiff need not prove a precise dollar figure to recover), attorney's fees, injunctive relief, mandatory deletion of unlawfully obtained neural data, and availability of class actions. Corporate-officer liability — the principle that an officer who knew or should have known of the violation is personally on the hook — is what converts a fine that a company treats as a cost of doing business into a deterrent that changes behaviour.

A regulator that can say no. A National Neurotechnology Commission with technical staff, the power to license high-risk neural products *before* market, the power to audit and to fine, a national neural-data register, and a statutory duty to publish an annual report on use, compliance and incidents. The single most common failure mode in technology regulation is the under-resourced regulator issuing guidance it cannot enforce; the statute should fund the Commission to staffing levels comparable to the national medical-device regulator, partly from fees on the entities it oversees.

> **A note on which model of ownership.** Three approaches contend for how to treat neural data. **Property models** let you own and license your neural data — intuitive, but anything you can sell you can be pressured or impoverished into selling, and a market in thoughts finds the desperate first. **Data-protection models** (the EU and US-state path) treat it as a sensitive category — the path of least resistance, but it inherits every weakness of click-to-accept consent. **Inalienability models** treat neural data like a donated organ: it may be given for research or medicine but never commercialised, regardless of consent. The strongest protection, and the hardest to enact, because it forecloses a business model investors are counting on. The honest question for a legislator is not which is most elegant but which you choose — and whether you choose it before or after the markets are built around the weakest one.

7. The Inalienable Core — Rights You Cannot Sign Away

The deepest design decision in the whole architecture is which protections survive consent. Most privacy law is consent-based: you may agree to almost anything, and the law's job is to make sure your agreement was informed. But consent-based privacy has failed comprehensively everywhere it has been tried — people click "agree" without reading, and the protection collapses. For the most sensitive category of data humanity has ever generated, repeating that mistake would be a choice, not an accident.

The alternative is to make a small core of neural rights **inalienable** — non-waivable by consent, contract, terms of service or any other instrument — on the same logic by which a person cannot validly consent to be enslaved. The defensible core is narrow and should stay narrow: freedom from involuntary neural reading or writing; the freedom of children from neural monitoring in education; the freedom of workers from neural monitoring as a condition of employment; the freedom of detainees from neural interrogation; and freedom from neural influence aimed at electoral choice. These are the cases where the consent a checkbox purports to capture is structurally

inadequate — because the consenter is a child, or economically coerced, or deprived of liberty, or being manipulated below the threshold of awareness. Naming them as inalienable is not paternalism. It is the recognition that some interventions erode the very conditions for valid consent, so consent to them cannot be the thing that authorises them.

8. The Black Box for the Brain — Making Abuse Provable

There is a hard, unsentimental fact at the centre of this field, and pretending otherwise corrodes trust: with today's science, there is **no validated forensic test** that can examine a person and establish that their thoughts or actions were externally directed. This is not because the question is forbidden. It is because the capability such a test would have to detect has not been demonstrated to exist, and the brain's ordinary variability is so large that distinguishing "influenced" from "normal" is, at present, not possible to a forensic standard. What forensic neuroscience *can* do is narrower: document physical injury sometimes consistent with energy exposure (the Havana Syndrome imaging debate showed how inconclusively); establish whether an implanted device was present and what it logged, *if* it kept honest logs; and document coercion or surveillance through ordinary evidence.

The provability gap is not a reason to dismiss the concern. It is a reason to build the infrastructure that would close it — and that infrastructure is achievable now. If neurotech devices were legally required to maintain **tamper-evident, independently auditable logs** of every signal read and every stimulation delivered — the neural equivalent of an aircraft's black box — then abuse by a *known* device would become provable after the fact. The log should record each measurement and stimulation event, the parameters used, the firmware version, and the destinations to which data was transmitted; and it should be accessible to the subject, to regulators, and to independent forensic experts in a dispute. This is the single most useful thing a lawmaker can do for anyone who fears their mind has been tampered with. We cannot yet detect the exotic. We can absolutely make the deployed accountable.

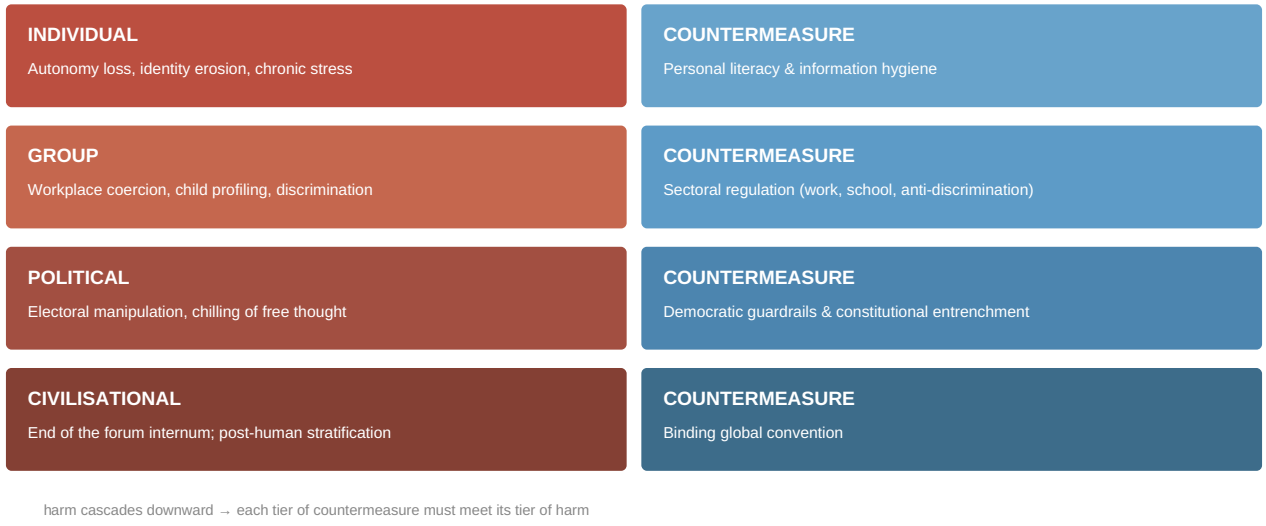


Figure 3. The consequence cascade of unaddressed risk, and the layered countermeasures that arrest it. Each tier of countermeasure must meet its corresponding tier of harm; black-box logging sits within the sectoral and statutory layers as the mechanism that makes group- and individual-level abuse provable.

9. Whistleblowers — The People Who Make Oversight Real

Every oversight regime in this essay depends, in the end, on insiders willing to report wrongdoing. Engineers, scientists, executives and officials who know that a neural decoder has been deployed against civilians, or that a stimulation profile was quietly modified, are the people who turn a paper regime into a real one. Neurotech whistleblowing is also distinctively hard. The field is technical enough that a true account can be incomprehensible to a non-expert audience — the whistleblower is "technically right but rhetorically unconvincing." Much of the most concerning research is classified, so disclosure can carry criminal exposure. And the credibility gap created by a dismissed victim population means a potential whistleblower must reckon with being labelled a conspiracy theorist for a well-documented disclosure.

Existing frameworks — the EU Whistleblower Directive, the US Whistleblower Protection Act — were not built for this and share predictable weaknesses: narrow subject-matter scope, channel requirements that can route a report straight back to the wrongdoer, national-security exclusions that exempt the deployments of greatest concern, and remedies that do not come close to compensating career destruction. A framework adequate to neurotech needs five reforms. **Explicit subject-matter coverage** for neurotech wrongdoing, so the protection is not left to interpretation. **Multiple authorised channels** — independent regulators, parliamentary oversight, accredited civil-society bodies, and, where those fail, journalists — so no single channel can be captured. **Modified national-security treatment**, protecting disclosure to cleared oversight bodies and mitigating penalties where the public interest substantially outweighs the harm. **Burden-shifting** so that adverse action soon after a protected disclosure puts the onus on the employer to prove a non-retaliatory cause. And **substantive remedy plus affirmative support** — not only back pay but compensation for career damage, secure channels, legal-representation funds, and, where safety requires it, protected relocation. The cases of the last decade show that retaliation reaches across borders; the protection must too.

> **Why this is in a neuro-rights essay at all.** It is tempting to file whistleblower law under "administrative detail." It is closer to a structural prerequisite. A black-box logging mandate (Section 8) makes abuse provable *if* someone preserves and surfaces the log. A licensing regime (Section 6) catches misconduct *if* someone inside the licensee is willing to say what the audit missed. Take the people who can speak out of the equation, and the rest of the architecture quietly becomes ceremonial.

10. The Map of Power — Six Blocs and Four Horizons

None of this is decided inside any one parliament. Neurotechnology is reshaping the relations among states, and the choices made in the near term will determine which longer-term trajectory the world follows. In 2026 the landscape organises into six recognisable blocs.

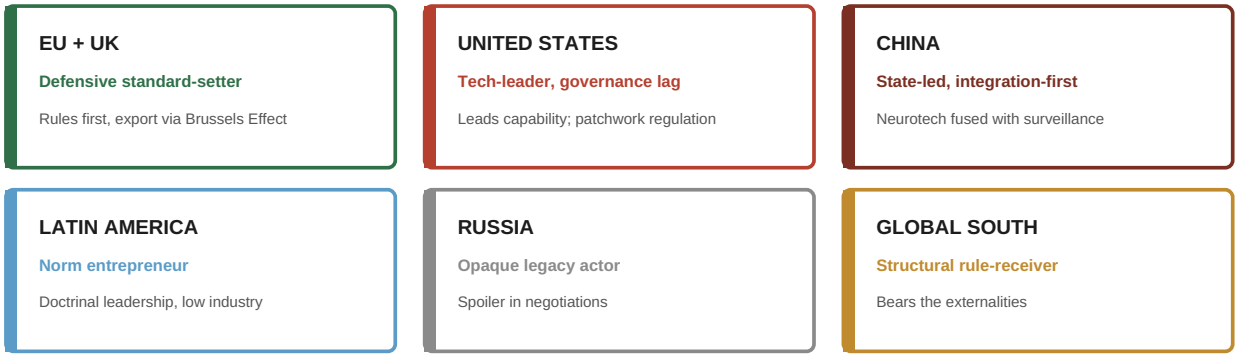


Figure 4. Geopolitical blocs and their neurotech postures in 2026. The EU is the defensive standard-setter; the US the tech-leader with a governance lag; China the state-led, integration-first actor. Latin America is the principal norm entrepreneur; Russia an opaque legacy actor; the Global South a structurally vulnerable rule-receiver.

The pattern is familiar from the data-protection era. The EU and UK are the principal **defensive standard-setter** — the GDPR, the AI Act and an emerging dedicated Neurotechnology Regulation give the bloc the most articulated rights-based framework, exported through the "Brussels Effect" even without a constitutional clause. The United States is the **tech-leader with a governance lag**: its companies lead across most invasiveness tiers and its defence research leads in dual-use capability, but federal regulation trails badly and the MIND Act remains pending. China is the **state-led, integration-first actor**, fusing neurotech with surveillance, military use and population management in ways no Western state would attempt. Latin America — Chile, Brazil, Spain as a European partner, Mexico, Colombia, Argentina, Uruguay — is the **norm entrepreneur**, with a louder voice on this topic than its combined economic weight suggests. Russia is the **opaque legacy actor**, more spoiler than primary mover. And the Global South is structurally a **rule-receiver**, bearing the externalities of technology designed and regulated for other markets, whose principal leverage is collective bargaining in multilateral forums.

Across four horizons, the trajectory is one of consolidation followed by a fork. The **near horizon (~2029)** is regulatory consolidation: an EU Neurotechnology Regulation, a US MIND Act with — if drafted well — a substantive inference clause, Chinese industrial consolidation, and the first trade disputes over cross-border neural-data flows. The **medium horizon (~2031)** is international-framework construction: a first UN draft convention circulating, NATO declaring the cognitive domain operational, more Latin American states adopting Chile-style clauses, an African Union protocol. The **long horizon (~2036)** forks on the choices made earlier — toward a binding convention with verification, or toward a fragmented order in which neuroweapons proliferate among middle powers with no agreed rules. The **generational horizon (~2046)** is where forecasting becomes conditional, and where the plausible structural development is two parallel digital ecosystems, one for the neurally augmented and one for the unaugmented, with "neural sovereignty" beginning to replace "data sovereignty" as the frame for state authority over information flows.

11. Synthetic Telepathy — Legislating Before Deployment

Of the capabilities on the horizon, none commands more fascination or more dread than synthetic telepathy: the technological transfer of meaningful content — words, images, concepts, intentions — from one brain to another, in real time, without speech, writing, gesture or display. The capability does not yet exist. Most of its constituent technologies do. It decomposes into three branches, and an honest forecast treats each separately rather than collapsing them into a single science-fiction image.

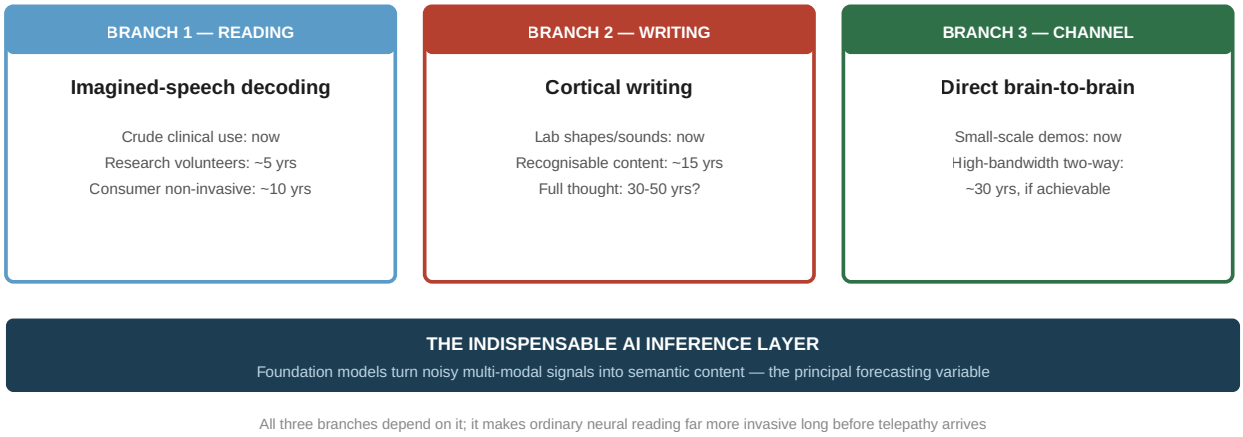


Figure 5. Synthetic telepathy decomposed into three capability branches, with the AI inference layer on which all three depend. The same models that would one day make telepathy possible will, long before, make ordinary neural reading far more invasive — which is why the AI layer is the principal forecasting variable.

Reading — decoding imagined or attempted speech — has advanced fast: implanted electrode arrays have decoded attempted speech in paralysed patients at sixty-plus words per minute, and word-level decoding from imagined speech is now a published result. Crude clinical decoding is real now; useful decoding in research volunteers is plausibly five years out; mass-market non-invasive decoding perhaps ten. **Writing** — delivering content directly to a recipient's cortex — is harder by orders of magnitude; today's state of the art produces perceived flashes of light or shifts of mood, not content. Recognisable semantic content is plausibly fifteen years away; a whole thought delivered as the speaker conceived it, thirty to fifty years, if achievable at all. The **brain-to-brain channel** integrates the two with an AI bridge between them, and consumer-grade two-way telepathy between healthy users is, realistically, several decades off if it arrives.

The policy point is not the timeline but its implication: this is the rare case where law can precede deployment rather than chase it. The constraints should be drafted now, embedded in the proposed convention and national statutes, so the technology develops *within* them. Eight are essential. **Symmetry of consent** — both sender and receiver must consent freely, specifically and revocably to every session, because sender-only consent is psychic intrusion and receiver-only consent is theft of thought. **No pre-configuration** — a recipient's interface must not be set up so that future writing can be delivered without fresh per-session consent. An **inalienable content filter** — political coercion, commercial advertising and state-actor content must be prohibited regardless of consent. **Mandatory black-box logging** by an independent authority, with both parties receiving a verifiable record. A **receiver override** — the receiver can halt the channel and expunge content at any time. A **capability-asymmetry prohibition** — channels in which one party has far more capability must be specifically licensed, because the powerful-sender-to-weak-receiver configuration is the one most prone to exploitation. **Absolute protection** for

children, the cognitively vulnerable, and people in coercive contexts. And a **cross-border restriction** placing inter-jurisdictional sessions under the global authority.

Synthetic telepathy poses the deepest constitutional question in the whole field, because it ends the historical inviolability of the forum internum — at least for the receiver. Once content can be delivered directly to a cortex, thought is no longer technically beyond reach. The right answer is not to ban the technology (which would foreclose real benefits for the locked-in and the paralysed) but to reinterpret the absolute prohibition as a prohibition on *unauthorised reception*. A person retains sovereignty over the cognitive stream that constitutes their inner life; what they choose to let in is theirs to choose, and what is delivered without their choice is a violation regardless of the sender's intent.

12. The Cases That Are Coming

Abstract law is hard to evaluate. The most useful test of a regime is to run plausible near-future disputes through it and ask, of each, two questions: what can existing law already do, and where does it break? The following composites — deliberately invented, none describing real people or events — are set across the next decade to show the trajectory.

The schoolchild's EEG (a 2028 composite)

A school district adopts an "attention enhancement" programme: every child wears a low-cost EEG headband during certain lessons, scores display on the teacher's screen, and data flows to a private start-up. A dyslexic ten-year-old is flagged for "sustained attention anomalies," the characterisation enters his permanent record, and his parents later discover the data was retained for a decade and shared with universities — none of it in language the consent form made comprehensible. *Existing law does a great deal:* children's-data rules require genuine informed parental consent that this plainly lacked; a constitutional mental-integrity clause lets a court read those rules stringently; a public prosecutor can intervene; the programme is suspended, the start-up fined, the data deleted. *Where it breaks:* the inference has already shaped how teachers see the child, and deletion does not erase that; models trained on the data before deletion continue to exist; and without a national prohibition, parents in the next district must litigate the whole thing again. *What is needed:* a categorical ban on neural monitoring of minors in primary and secondary education — an inalienable right, not a consent question.

Insurance redlining (a 2030 composite)

A mortgage applicant is approved at a rate above the median for her profile. The reason, surfaced only through a regulator's investigation, is a credit-score adjustment derived from sleep-pattern variability in data licensed from a wearable she has worn for years, under terms she never read. *Existing law does something:* automated-decision and credit-explanation rules give her a remedy, a re-underwriting and damages. *Where it breaks:* millions of similar borrowers were disadvantaged without ever complaining; the original data transfer through a chain of brokers may have been lawful when it happened; and the inference may be discriminatory even if statistically valid. *What is needed:* a categorical, non-waivable prohibition on neural-derived inferences in credit and insurance underwriting, modelled on existing bans on the use of genetic information.

The modified sleep device (a 2030 composite)

An honestly-marketed consumer device delivers gentle stimulation to improve sleep. Internally, a product team — without leadership's or the regulator's knowledge — alters the stimulation profile to nudge users' suggestibility and lift ad-revenue. The change is discovered three years and 1.4 million devices later. *Existing law*

does a great deal: the modification arguably converts a wellness device into an unapproved medical device (a crime), officers face liability, the consumer-protection regulator can pursue disgorgement, class actions follow. *Where it breaks:* whether the parameters actually changed behaviour is contested and may take years; meanwhile users made real, unrecoverable decisions during the exposure. *What is needed:* a categorical prohibition — inalienable, strict-liability — on cortical writing intended to influence behaviour beyond a device's disclosed purpose, backed by the black-box logging mandate so the firmware change would have left an auditable trace.

The election nudge (a 2032 composite)

Two months before a parliamentary vote, a foreign-affiliated entity runs a microtargeted campaign — no visible disinformation, just subtle shaping of the attentional and emotional patterns of several hundred thousand swing voters, with a measurable effect in a handful of constituencies. *Existing law does something:* platform and AI-manipulation rules, broker liability under data law, an electoral-commission investigation. *Where it breaks:* the voters were not deceived in the conventional sense; they were nudged below the threshold of conscious deliberation, which electoral law was not built for; the foreign entity is largely beyond domestic reach; and a re-run would not undo the cognitive effect. *What is needed:* treatment of cognitive operations against elections as a category subject to specific prohibition and collective international response — a normative red line, drawn before the next election rather than after it.

> **The pattern across the cases.** In each, today's law reaches the individual remedy and misses the structural harm. It deletes one child's record but not the model trained on it; it re-underwrites one mortgage but not the million; it punishes one firmware change but not the diffuse, unrecoverable effect on users' lives. The missing piece is always the same: a prohibition that is *categorical* (it forecloses the practice rather than litigating it case by case), *inalienable* (consent cannot cure it), and *cross-border* (it reaches the procurement chain and the foreign actor). That is exactly what the six-tier stack is built to supply.

13. Victims — A Standard of Care That Tells the Difference

"Victim of neurotechnology" needs a careful definition — both to help real victims and to avoid a category so loose it helps no one. The defensible definition centres on demonstrable harm from the use or misuse of a neurotechnology, and it falls into recognisable classes: **data-rights victims** (the largest and most provable class, whose neural data was collected, inferred from or shared without valid consent — the Chilean case is the template for redress); **clinical and research harm** (patients injured by a device, an explantation, a malfunction, or inadequate consent); **coercion and surveillance victims** (pressured or covertly monitored by an employer, institution or state); and **discrimination victims** (disadvantaged in hiring, insurance, credit or justice on the basis of neural data or inferences). A model statute should establish recognition mechanisms, redress including compensation, and rehabilitation for each — with a trust fund for cases national mechanisms fail to reach.

A growing number of people also report believing they are victims of covert, remote neural influence. This requires a response written with care, because both reflexive dismissal and uncritical credence cause harm. The honest position holds two things at once. The distress is real and deserves a serious, humane response. And, as Section 8 established, there is at present no validated way to confirm covert remote influence, and for the symptoms most often described, thorough clinical assessment finds a range of explanations, many of them treatable. The compassionate and truthful response is the one good medicine already offers: take the suffering seriously, pursue thorough evaluation, document everything, and press for the detection science and the device-accountability rules — above all the black box — that would let any future claim be tested rather than merely believed or denied. The field’s present inability to detect is itself a policy failure worth fixing. It is not evidence that anyone’s distress is unreal, and it is not, by itself, evidence of the mechanism they suspect. Helping victims means building systems that can tell the difference — not assuming the answer in either direction.

14. The Citizen’s Toolkit — What You Can Do Before the Law Catches Up

The work of defending the inner forum is not only the work of legislators, scientists and activists. It is also the work of ordinary citizens at the scale of daily life, and several of the most effective protections are available before any statute passes. None requires technical expertise.

- **Read before you grant.** Do not accept neurotech terms of service you have not read. The brain data you license away under a meditation app’s default terms is the same data a future model can re-interpret. The single most common point of failure is the unread checkbox.
- **Keep sensitive spaces device-free.** Do not bring consumer EEG or biosensing wearables — including some earbuds — into sensitive workplaces, negotiations or political meetings. The point is not paranoia; it is that you cannot control what is inferred from a signal you did not need to emit.
- **Document anything unusual, plainly and contemporaneously.** Dates, times, devices present, what happened, who was there. Documentation that is specific and timely is worth far more than documentation assembled in retrospect — both for clinical care and for any future legal claim.
- **Build community.** Isolation compounds every harm in this field. Local networks of mutual aid — for sharing information, accompanying each other to appointments, and pooling the effort of civic engagement — are among the most protective structures available, and the only ones individuals can build without waiting for permission.

- **Engage your representatives with specifics.** "Protect mental privacy" is easy to nod along to and impossible to act on. "Define neural data to include inferences; make neural monitoring of children and workers inalienable; fund a regulator that can license before market" is a request a legislator can actually take to a drafting table.

Institution	The action it can take now, without waiting for the others
Legislative branch	Define "neural data" to include inferences; classify it as a special protected category; make a core of neural rights inalienable; fund a regulator; constitutionalise the foundational right where possible.
Executive branch	Adopt the UNESCO Recommendation domestically; convene a national strategy across health, labour, justice and defence; fund detection and forensic science; require neural-device black-box logging in procurement.
Judiciary	Recognise neural data within existing privacy, bodily-integrity and dignity jurisprudence; set evidentiary standards that exclude neural "lie detection" until validated; build capacity to hear redress claims.
Law enforcement	Treat covert neural-data collection as a serious offence; build capacity to investigate device-based abuse; refuse neural-surveillance tools absent statutory authority and oversight.
Medical sector	Tighten informed-consent standards; insist donated research data is never commercialised; develop standardised forensic protocols for anomalous neurological incidents; lead rather than follow on regulation.

15. Three Futures, and the Choice Between Them

Which trajectory the world follows is not yet determined; it will be settled by the policy choices made in the near and medium horizons. Three scenarios for the generational horizon, offered as scenario-planning rather than prediction, mark out the range. Their value is in the contrast.

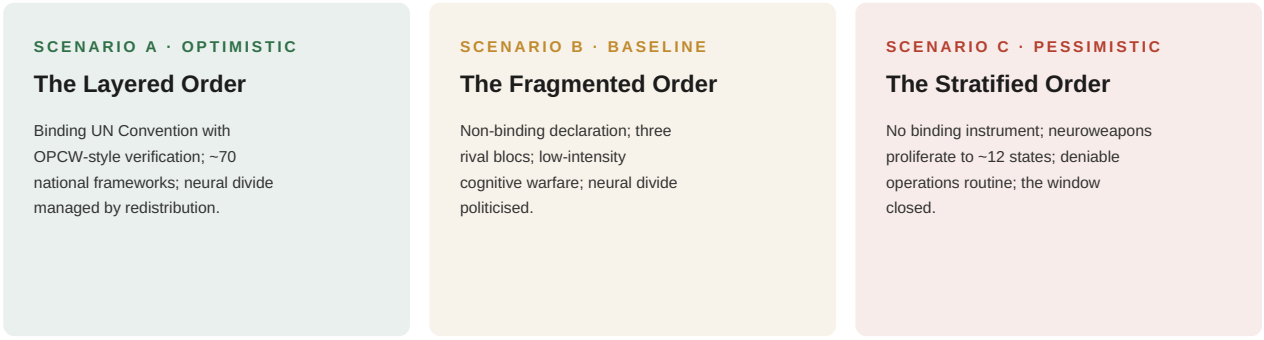


Figure 6. Three scenarios for 2046. The contrast among them — not the precision of any one — is the point. Which unfolds depends on choices that are not yet made.

In the **Layered Order**, a binding convention with real verification has been operational for a decade; national frameworks are in force in dozens of states; augmentation is broadly available through public pathways and the neural divide is managed by redistributive policy; cognitive warfare is constrained by mutual deterrence and verification. In the **Fragmented Order** — the baseline, the path of least resistance — the convention process produced a non-binding declaration, three rival regulatory blocs operate with significant friction between them, low-intensity cognitive warfare is endemic and restrained only by case-by-case calculation, and the neural divide

is significant and politicised. In the **Stratified Order**, no binding instrument exists, neuroweapons capability has proliferated to roughly a dozen states, deniable operations against adversary civilians are routine, and within wealthy states the neural divide has produced a stratified society in which the augmented carry separate political weight — the window described in these essays has closed.

The argument of this essay, and of the book behind it, is that the layered order is technically achievable and morally required — and that the effort to reach it must begin now, because the baseline is not neutral. The default outcome, the one that requires no political will and no investment, is the fragmented order sliding toward the stratified one: consumer neurotech becomes ubiquitous, the inference layer becomes dominant, workplace and school deployments normalise, state capability expands, and cognitive-warfare techniques diffuse from intelligence services to political consultants to private parties — until the assumption of inner privacy has quietly become obsolete. Most of that is already happening. The question is whether the trajectory is bent by deliberate choices or simply allowed to run.

In Closing — The Same Window, Now Narrower

The first essay argued that the mind had become a policy choice for the first time in history. This one has tried to show what choosing actually looks like: a definition that reaches the inference and not just the signal; a constitutional clause that entrenches the right; a statute that prohibits, remedies and funds a regulator; a small core of protections that no checkbox can waive; a black box that makes the deployed accountable even where the exotic stays undetectable; whistleblower law that keeps the whole apparatus from going ceremonial; a global instrument for the threats no nation can police alone; and constraints on synthetic telepathy written before the capability arrives rather than after it has been abused.

None of this requires resolving the questions that cannot yet be resolved — what exists behind classification, whether the most exotic claims of covert control are true. Good policy does not require knowing the secret. It requires assuming that states keep some, that markets will exploit every gap left open, and that the harms of this technology are structural, slow to manifest and very hard to reverse once entrenched — and then binding all of it anyway. UNESCO has supplied a shared vocabulary. Chile has shown that constitutional protection is possible and enforceable. A handful of states and provinces have shown that even partial protection can pass quickly. What is missing is not the idea but the will: the decision, by enough legislators in enough places, to treat the inside of a human head as the one place that markets and states may not enter without permission. The window is the same one the first essay described. It is simply a few years narrower now. This essay is an argument for choosing — deliberately, specifically, and soon — to keep the last private place private.

A note on method. *Where the evidence is solid — the laws, the rulings, the technology that demonstrably exists — this essay states it plainly. Where it is contested or absent — the exotic weapons, the classified capabilities, the strongest claims of covert influence — it says so, and resists resolving the uncertainty in either direction for effect. The fictional cases in Section 12 are composites: none describes any real person, company or event, and they are offered only to test how law would apply. On a subject that touches the deepest fears and hopes we have about our own minds, that discipline is not a hedge. It is the point.*