Part III — Where We Are Now: Life at 4 p.m. (2025)

Chapter 13–18 Overview

Defining “Now”: Why 2025 Feels Like 4 p.m.

If you’re going to use the clock as a way of thinking, you have to be willing to say where the hand is right now. Not in a hand-wavy, “something is happening with AI” way, but with enough precision that it changes how you act.

The claim in this part of the book is simple:

In 2025, we are no longer at smartphone Noon.
We are not yet at full AI 6 p.m. either.
We are somewhere around 4 p.m. on the way there.

This isn’t prophecy. It’s pattern-matching. You’ve already seen this movie twice: PC → Internet, Internet → Smartphone. Each time, there was a period that felt exactly like this—exciting, uneven, a little chaotic, with big claims and small practical realities living side by side. That’s the texture of 4 p.m.

This chapter is about making that concrete:

- what 4 p.m. actually means on the clock,
- how we know we’re there for AI in 2025,
- what that implies about smartphones behind us and robots in front of us,
- and why getting this “time of day” roughly right matters to anyone trying to build, invest, or plan.

The Meaning of 4 p.m.

On the innovation clock, Noon and 6 p.m. are the two anchor points.

At 12 p.m. we're in a concentrated state: a thing you can point at, buy, hold, and call yours has become the dominant way intelligence and capability show up in daily life. The PC at its height was that thing. The smartphone at its height is that thing. In the future, robots—embodied AI you own—will sit there as well.

At 6 p.m. we're in a diffused state: the action is in networks, platforms, and shared services. The Internet in its mature form, with billions of people participating through browsers, apps, and APIs, is the canonical 6 p.m. example. AI as a networked, subscription, "user of our platform" phenomenon is on its way to that position.

What does 4 p.m. mean between those two?

By 4 p.m.:

- The previous Noon technology is still everywhere, but it's no longer the frontier.
- The new 6 p.m. technology is clearly real and increasingly unavoidable, but not yet ubiquitous.
- The infrastructure, culture, and business models for the new thing exist, but they are still fluid.
- You can still opt out for a while, but the cost of doing so is obviously rising.

Historically, the late 1990s were "4 p.m. web": the browser was clearly important, many enterprises still lived largely in PC-centric thinking, and e-commerce and web apps were just beginning to prove they weren't toys. Around 2010–2011 you had "4 p.m. mobile": smartphones were everywhere in the hands of early adopters, but large swaths of government, enterprise, and even consumer brands still behaved as though the desktop web was the real thing and mobile was a secondary channel.

Each of those periods shared a particular personality: everyone knew something irreversible was happening, but institutions hadn't fully reorganized around it. That is exactly where AI sits in 2025.

We're Not at AI 6 p.m. Yet

It's tempting to talk as though "AI is already everywhere." The marketing decks certainly do. But 6 p.m. has a very specific meaning in this framework: ubiquity at the level where taking it away would break the week of billions of people.

For AI, we can phrase that as a working threshold: something like five billion humans having at least one AI-mediated interaction every week that materially changes what they can do. Not "they saw an AI-generated ad" or "a ranking algorithm touched them somewhere in the background," but that they personally are relying on AI as part of a task: to understand, to decide, to create, to transact.

We're not there yet.

You can see that by looking at where AI is truly indispensable versus where it's an impressive option:

- In software development, design, and content creation, a growing number of practitioners will tell you they feel slower and less capable without modern AI tools. In those niches, you can argue we're already past 4 p.m. and approaching 5 p.m.
- In white-collar knowledge work—strategy, marketing, operations, finance—you see serious usage, but it's a gradient. Some teams are deeply AI-first, others barely touch it beyond a few experiments. Within the same company, you'll meet a power user and, two doors down, someone who has never meaningfully used an AI system.
- In frontline and service work, in most of the global south, in small local businesses, AI shows up more sporadically: surfaced through a better camera app, a slightly smarter messaging app, an automated translation, a payment system that quietly uses models behind the scenes. It is there, but not in a way that people would name as "my AI."

If you turned off the major AI platforms globally tomorrow, the headlines would be big, but large sections of daily life would look surprisingly similar for billions of people. Important workflows would get worse or slower, but the whole fabric wouldn't collapse. That is not what 6 p.m. looks like.

By contrast, imagine turning off the Internet in its mature form. You don't just lose a convenience; critical infrastructure fails. Payments, logistics, basic communications, governance, education, energy management—large parts of modern life stall or go dark. That was the Internet at 6 p.m. We're not yet there with AI.

So yes, AI is powerful. It is changing how early adopters work. It is starting to reshape industries. But we're still in that phase where you can find two people in the same role, in the same city, one of whom uses AI dozens of times a day and one of whom barely touches it. That's the hallmark of 4 p.m.: strong local intensity, uneven global penetration.

Smartphone Noon Is Behind Us

At the same time, it's equally clear that we are no longer living at smartphone Noon as the edge of possibility. Smartphones are still the dominant physical thing in people's lives—the primary 12 p.m. object of this era—but they are no longer the object we argue about.

You can tell a technology has passed Noon when nobody finds it interesting anymore as a technology question. Phones are assumed:

- Every serious consumer product is already "mobile-first" or has adapted to that expectation.
- Nobody sits in a strategy meeting in 2025 bragging that they have an app; that's like bragging you have electricity.
- There is very little philosophical debate about whether we'll "adopt smartphones." We already did.

That is exactly what PCs felt like by the late 2000s. They were still everywhere, still essential, but no longer where the frontier was. The arguments had moved elsewhere.

For the last decade, the serious debates—at least in the tech and business world—have shifted to questions like:

What happens when you can translate anything into anything instantly?
What happens when any image, any video, any document can be synthesized?
What happens when you can talk to your tools instead of learning their menus?

Those are AI questions, not smartphone questions. The phone is now one of the places where AI shows up, not the thing that defines the era on its own.

This doesn't mean the smartphone disappears from the cycle; it stays as one of the anchors of the current generation, just like the PC stayed. It simply means that if you are trying to understand *change*, looking at phones directly won't tell you much anymore. They've become part of the scenery.

The Texture of 4 p.m. in 2025

Once you accept that we're between smartphone Noon behind us and AI 6 p.m. ahead of us, the particular feel of 2025 snaps into focus.

Several things tend to coexist at 4 p.m.:

First, hype and uncertainty live side by side. You can sit in a boardroom, a government hearing, or a startup pitch and hear incredibly ambitious claims about AI—changing productivity curves, transforming education, collapsing the time to build new products—often in the same week that you see a major system hallucinate, misbehave, or run into an obvious social or legal constraint. The story oscillates between "this will change everything" and "this is not reliable enough" depending on which anecdote is on screen.

Second, adoption is highly uneven. Within a given sector, you'll see islands of the future surrounded by large pools of the past. A design team might be deeply AI-native, while the legal department still works primarily in Word attachments. A single teacher might quietly use generative tools to redesign everything in their classroom, while the institution as a whole has no coherent AI policy. Small startups might have AI woven into every process while large

incumbents are still publishing “exploratory task forces” and running endless pilots.

Third, the product landscape is fragmented. There are half a dozen major foundational platforms, dozens of high-quality specialized tools, and an almost comical long tail of “AI-powered” everything. Some of these products are clearly durable; others will vanish as quickly as early dot-com sites or novelty mobile apps. You can’t yet tell which categories will harden into infrastructure and which are transitional experiments.

Fourth, regulation is reactive, but real. Governments and institutions are no longer ignoring AI, but they are not yet synchronized around a stable regulatory regime. You see hearings, draft laws, voluntary commitments, informal norms, and a steady stream of public debates about safety, copyright, labor displacement, misinformation, and more. That’s exactly what the late 1990s and early 2010s looked like in their respective transitions: the law lagged, scrambled, overreacted in some places and underreacted in others, then eventually converged on a set of norms.

Fifth, talent has already moved. The most ambitious builders, thinkers, and operators are not asking “should I pay attention to AI?” They are already building in and around it. Students choose their projects and majors with AI in mind. Researchers frame their work assuming AI will be a baseline capability. Founders pitch ideas assuming that powerful models are just there, like the Internet or GPS. You don’t see that degree of talent migration toward something that isn’t real; you see it at 4 p.m., when early proofs of value have convinced people that the new thing is worth their careers.

Finally, the old world is still intact enough to feel normal. You can absolutely still run a company, a school, or even a government office with only minimal AI integration. It will be slower and less adaptive than it could be, but it will not be untenable yet. That was true of PC-centric businesses in the late 90s and desktop-web-centric businesses in the early 2010s: they could still function while the world shifted under them. The break only becomes obvious in hindsight.

All of those characteristics apply to AI in 2025. It is everywhere in the news, uneven in practice, powerful in pockets, immature in others, and already attracting most of the energy from the next generation of builders. That is the signature of 4 p.m.

Why Getting the Time Right Matters

Calling this moment “4 p.m.” is not academic. It changes how you think about risk, opportunity, and timing.

If you misread 4 p.m. as Noon—“this is done, mature, and we can treat the current dominant interfaces and platforms as permanent”—you’ll over-optimize for the current layer: the exact chat UI, the current roster of model providers, today’s pricing and licensing schemes. You’ll build brittle strategies that assume what you see now is what you’ll see in ten years.

If you misread 4 p.m. as 1 p.m.—“this is still early, we can wait”—you’ll underinvest in learning, in tooling, and in organizational change. You’ll be the PC-era incumbent in 1998 explaining why the web isn’t quite ready for your “serious” business, or the desktop-web brand in 2011 insisting that mobile browsing is a second-class experience and you’ll get to it “once things settle down.”

The truth in 2025 is in between. It’s late enough that you can’t afford to ignore AI without falling behind, and early enough that you should be careful about anchoring everything to the current user interface and market structure.

In practical terms, a 4 p.m. mindset for an organization looks like this:

You accept that AI will be woven into almost every workflow you care about. That’s no longer a question; it’s a time horizon issue. You begin the work of making your processes, data, and culture AI-compatible. You identify roles where AI augmentation is already compelling and start there, rather than trying to boil the ocean. You build internal literacy so that your people can distinguish between hype and real capability.

At the same time, you don’t bet your entire future on owning “your bot” as the final interface. You design APIs, connectors, and clean capabilities that can be called by many different agents, human or robotic. You recognize that the path from AI 6 p.m. to Robot Noon will involve a shift in interface power away from platforms toward things users own, and you adjust your architecture accordingly. You invest in being good at something concrete—logistics, safety, curation, transformation, manufacturing, whatever your domain is—rather than in being

the only place where users will ever want to type prompts.

You also start to notice early signals of the next Noon. Wearables, AI glasses, desktop devices, and “pucks” begin to look less like gadgets and more like early robots. Tools for orchestrating complex, multi-step tasks on behalf of users start to appear. When you see these, you don’t dismiss them as toys. You file them mentally as prototypes of Robot Noon and ask what part of your business they might eventually mediate.

In that sense, 4 p.m. is a planning gift. It gives you enough clarity to say:

We are leaving the smartphone-as-frontier era.
We are not yet in fully diffused AI.
We have a window during which we can learn, position, and adapt.

But it also gives you a warning: that window is not indefinite. The hand moves whether you acknowledge it or not. The organizations that treated late-90s web as a phase missed the chance to become default Internet properties. The companies that treated early 2010s mobile as a cute add-on discovered a few years later that half their traffic had quietly gone elsewhere.

The same will be true for AI. In twenty years, the question won’t be “were you excited about AI in 2025?” Everyone will claim they were. The question will be: “did you behave as though it was 4 p.m.—late enough to get serious, early enough to play the long game?”

The rest of Part III takes that as a working assumption. We will describe what full AI 6 p.m. likely looks like, what the early signals of Robot Noon are, how to classify technologies by their “time of day,” and how to build a clock for your own industry. All of that starts here, with the decision to stop talking about AI as an abstraction and start recognizing 2025 as a very particular, very consequential moment on the cycle.

Chapter 13 — Defining “Now”: Why 2025 Feels Like 4 p.m.

The clock model is only useful if we can confidently say *where we are on it*.

For our current moment, the claim is:

> In 2025, we're not at AI Noon and we're not at AI Night.
> We're around **4 p.m.** — on the way from Smartphone 12 p.m. to AI 6 p.m.

This chapter is about justifying that statement. Why 4 p.m., not 2 p.m. or 5:55 p.m.? What is "now," structurally?

What 4 p.m. Means on the Clock

On our cycle:

- 12 p.m. → **concentrated, owned thing** (PC, Smartphone, eventually Robot)
- 6 p.m. → **diffused, shared network** (Internet, AI-as-platform)

4 p.m. is:

- closer to 6 than to 12,
- not yet ubiquitous,
- high hype + real utility,
- uneven distribution,
- lots of experimentation,
- business models still in flux.

Historically, 4 p.m. looks like:

- the web around 1997–1999,
- smartphones around 2010–2011.

The pattern:

- Everyone *knows* something big is happening.

- No one has fully rearranged their institutions around it yet.

Evidence That AI Is Not Yet at 6 p.m.

For AI to be at full 6 p.m., we've defined a practical threshold:

> ~5 billion weekly active participants in some form of AI-mediated interaction.

Rough checks for 2025:

- Most people **have heard** of generative AI.
- A non-trivial fraction **have tried** ChatGPT, Gemini, Claude, etc.
- Regular, weekly dependence is still skewed:
  - knowledge workers, developers, students, creators, founders, etc.
  - early adopter geographies and industries.

In many domains:

- A lot of key workflows are **AI-optional**:
  - code can be written without AI (though it's slower),
  - documents can be drafted without AI,
  - customer support can run without AI,
  - design can still be done with classical tools.

That's 4 p.m. energy: very real, but not yet **inescapable**.

We are not in a world where:

- the majority of humanity touches AI every week,
- and notices its absence when it's gone.

We're heading there, but we're not there.

Evidence That Smartphone Noon Is Behind Us

At the same time, it's clear we are no longer in smartphone 12 p.m. as the *frontier*.

You can tell because:

- People are **bored** by phones.
    - There's no existential debate about "mobile-first" anymore.
    - Everyone assumes everything will work on a phone.
- Product pitches rarely center on:
    - "We've built an amazing mobile app."
    - That's table stakes, not differentiation.
- The interesting work has moved:
    - to AI models,
    - to multimodal experiences,
    - to agentic workflows,
    - to "what can this system decide and do on its own?"

Smartphones are still **critical infrastructure**—they're the main 12 p.m. thing we own—but they are not where the frontier questions live.

That's exactly how PCs felt by the mid-2000s: essential, but not the edge.

Signals of 4 p.m.

The "feel" of 4 p.m. has some recognizable features:

- **Hype + confusion coexist**
    - Board decks and government hearings are saturated with AI.

    - At the same time, many organizations have no coherent AI strategy; they're in pilot purgatory.
- **Uneven adoption**
    - Some teams are deeply dependent on AI already.
    - Others barely use it beyond PR and experiments.
- **Tool fragmentation**
    - Many overlapping products, unclear winners.
    - People maintain multiple accounts across models and platforms.
- **Regulation lagging but catching up**
    - Serious policy conversations have begun.
    - Frameworks are incomplete and often reactive.
- **Talent migration**
    - Builders and ambitious people are flowing into AI-adjacent work,
    - even if their host industries haven't fully reorganized yet.

If you look at historical 4 p.m. moments, the texture is similar:

- everyone's talking,
- no one's finished re-wiring,
- the old world is still mostly intact,
- but you can feel it loosening.

Why "Now" Matters for Decision-Making

Calling 2025 "4 p.m." is not a poetic flourish; it's a way to shape how you act:

- At 4 p.m., it is too late to ignore the shift.
- At 4 p.m., it is too early to optimize *only* for the currently-dominant platforms.

Strategic implications:

- You **must** be learning, experimenting, and building in AI now.
- You **must not** bet everything on the exact shape of today's interfaces and business models.

Part III will sit in this tension:

> we are clearly leaving smartphone Noon behind,
> we have not yet arrived at full AI 6 p.m.,
> and the robots are already casting a shadow from the next 12 p.m.

Chapter 14 — What Full 6 p.m. AI Looks Like: 5 Billion Weekly Participants

If 2025 is 4 p.m., where are we heading before the hand swings toward robots?
This chapter sketches a plausible picture of **full AI 6 p.m.**:

> a world where AI is as ubiquitous and backgrounded as the Internet is today.

The shorthand:
5 billion weekly participants in AI-mediated activity.

Defining "5 Billion Weekly Participants"

We're not talking about:

- 5 billion people using a specific app, or
- 5 billion people writing prompts into a chat box.

We mean:

> ~5 billion humans whose week involves at least one meaningful interaction mediated by AI, whether or not they *label* it "AI."

Examples:

- A farmer using a cheap smartphone app that:
  - diagnoses plant diseases via camera,
  - recommends fertilizer schedules.
- A worker in a small business using:
  - AI-augmented messaging,
  - AI-assisted accounting,
  - AI booking and logistics.
- A teenager:
  - editing video with AI tools,
  - getting homework help,
  - using translation and creation tools constantly.
- An elder:
  - interacting with AI-enhanced interfaces for health, banking, communication.

The important property:

- If you removed AI, **they feel the loss**. Something practical in their week breaks or becomes noticeably worse.

Characteristics of Full 6 p.m. AI

At full 6 p.m., you'd expect to see:

- **Invisible AI**
  - Most AI isn't a separate destination (like "go to ChatGPT").

    - It’s inside workflows, devices, and services:
        - search, messaging, editing, commerce, navigation, support.
- **AI-default workflows**
    - The default way to do many tasks:
        - summarize, decide, draft, translate, optimize
          is AI-first, with manual-only as a fallback.
- **Broad global reach**
    - Not just wealthy countries and tech sectors.
    - Penetration into:
        - low-cost phones,
        - offline-ish use cases,
        - local languages,
        - small businesses and public services.
- **Stable platform layer**
    - A relatively small number of infrastructural model providers and tool ecosystems,
    - with clear interfaces, pricing norms, and safety regimes.
- **Cultural normalization**
    - Arguments shift from “Should we use AI?” to:
      “How should we regulate and constrain AI we obviously use?”
    - Moral panics cool into ongoing policy debates.

In other words:
AI becomes more like “electricity” or “the web” than “a fancy optional tool some people know about.”

Where Humans Show Up in a 6 p.m. World

Even at full AI diffusion, humans don’t disappear. They just move up a level:

- Humans set goals, values, constraints.
- AI automates:
    - pattern recognition,
    - search,
    - transformation,
    - many micro-decisions.

Typical patterns:

- You rarely:
    - start from a blank page,
    - search without summaries,
    - analyze raw data manually.
- You frequently:
    - edit drafts,
    - inspect suggested decisions,
    - tune behaviors via preferences and rules.

This is the “human as editor / curator / principal” pattern.

What Stays Platform-Like at 6 p.m.

Even as robots start to emerge, a full 6 p.m. AI world still has:

- **platform control points**:
    - clouds that host models,
    - services that manage identity and billing,
    - ecosystems of tools.
- **subscription and usage models**:

    - tokens, seats, tiers, per-feature pricing.
- **“users”**:
    - individuals and organizations treated as accounts in someone else’s environment.

AI is still primarily experienced as a **network phenomenon**, not as a thing you own.
It’s ubiquitous, but not yet embodied as your personal physical artifact.

Why 6 p.m. Matters Before Robot Noon

Why dwell on 6 p.m. if we care about robots?

Because the path to Robot Noon passes *through* full AI diffusion:

- Robots depend on cheap, ubiquitous cognition.
- The cultural comfort with “AI in everything” has to be established.
- The economic and regulatory stack around AI needs to exist before robots plug into it.

Full 6 p.m. AI is the **substrate** on which Robot 12 p.m. will be built.

If you’re building for robots, you’re implicitly assuming:

- that something like this 6 p.m. world arrives and stabilizes first.

Chapter 15 — Early Signals of Robot Noon: Wearables, Glasses, Pucks, and Pods

Even while AI is marching toward 6 p.m., the next 12 p.m. is already whispering.

This chapter focuses on **early embodied AI form factors**:

- wearables,
- glasses,

- “pucks” and “pods,”
- home units and quasi-robots,

and why they matter as *prototypes of Robot Noon*.

The Key Idea: Proto-Robots Before Robots

You don’t go from:

- “AI is a website”
  to
- “Here is your full humanoid household robot”

in one step.

Just as:

- the iPod, BlackBerry, and early PDAs preceded the fully-formed “smartphone” era,

we will see (and are seeing):

- many intermediate embodiments of AI:
  - glasses, pendants, badges, desktop companions, speaker-like units, mobile devices.

These are **proto-robots**:
embodied, owned, persistent AI experiences that don’t yet feel like “robots” in the sci-fi sense.

What Makes Something a Proto-Robot?

Regardless of the marketing label, you can ask:

- Does the thing live with you (vs you “visiting” it)?
- Is it owned (bought) rather than purely subscribed to?

- Does it know you and your context over time?
- Is it the primary surface you speak to when you want something done?
- Does it orchestrate other services on your behalf?

If the answers are mostly “yes,” that’s proto-robot territory.

Glasses and pucks are interesting not because of their shape, but because they:

- sit at the edge,
- interface with your life,
- increasingly *feel* like “mine.”

Wearables and Glasses: Embodiment on Your Face

AI glasses and other wearables are early Robot Noon candidates because:

- They are **continuous**:
  - worn for long stretches, not opened like an app.
- They are **contextual**:
  - see what you see,
  - hear what you hear,
  - can overlay information in place.
- They are **identity-linked**:
  - it’s clear who they belong to,
  - they are rarely shared.

They also push design toward:

- low-friction interaction:
  - subtle voice, gaze, gesture, taps.

- proactive suggestions:
    - “you’ve parked here before, should I mark this spot?”
    - “this is the same product you looked at last week at home.”

This is what early Robot Noon feels like:
my embodied AI noticing, remembering, and helping *in the moment*.

Pucks, Pods, and Home Units

Small desktop or countertop devices — little cylinders, spheres, pods — also serve as proto-robots when:

- they become the **default interface** for getting digital things done in a space:
    - “Hey, order more detergent.”
    - “What’s on the schedule?”
    - “Pay this bill.”
- they integrate:
    - voice,
    - screens or projectors,
    - presence detection,
    - sometimes cameras.

If:

- the family interacts with the pod more than with individual websites for routine tasks,
- it mediates commerce, media, and information flow,

then functionally, you have a **room robot**, even if it doesn’t move.

Why These Experiments Are Strategically Important

You don’t care about the specific hardware revision as much as you care about:

- **interaction patterns**:
  - do people talk to these things naturally?
  - do they trust them with money and decisions?
- **ownership patterns**:
  - do people name them? personalize them? feel attachment?
- **delegation patterns**:
  - what do people actually hand off to them?
  - where do they still insist on doing things manually?

These experiments teach:

- what kinds of embodiment feel comfortable,
- where the creepiness lines are,
- how much initiative people will tolerate at the edge.

How to Read These Signals

When you see a new wearable or home device, don’t ask:

- “Is this *the* robot?”

Ask:

- What part of the eventual robot experience is this normalizing?
- What new form of delegation is being tested?
- What new failure modes (privacy, security, agency) are being surfaced?

By 4 p.m., these early forms are already here.
They’re the **training wheels** of Robot Noon.

Chapter 16 — Classifying Technologies: Is This 6 p.m. or 12 p.m.?

One of the most practical tools in this book is a simple classifier:

> "Is this technology behaving more like a 6 p.m. network or a 12 p.m. thing?"

This chapter gives you a checklist so you can apply the clock lens to real products, platforms, and trends.

The 6 p.m. Profile: Diffused Networks

A classic 6 p.m. technology has these traits:

- **Access model**: sign up, log in, subscribe.
- **Ownership**: you're a user, not an owner.
- **Location of intelligence**: mostly in the cloud/platform.
- **Identity**: your "self" lives as an account in their database.
- **Economics**: recurring revenue, usage fees, ads, sponsorship.
- **Language**: "users," "participants," "MAUs," "DAUs."

Examples:

- Web platforms, SaaS products, cloud-hosted AI services, social networks.

In 6 p.m. land, you *visit* them; they don't live with you as objects.

The 12 p.m. Profile: Concentrated Things

A 12 p.m. technology has different energy:

- **Access model**: you buy or receive a specific thing.
- **Ownership**: you are the owner, with control expectations.
- **Location of intelligence**: at the edge (device/agent), even if it calls back home.
- **Identity**: your "self" is attached to the device/agent; accounts are supporting details.
- **Economics**: hardware sales plus optional services.
- **Language**: "my phone," "my laptop," "my robot," "my glasses."

You don't log in to "use" a 12 p.m. thing in the same way.
It's already yours, always on, in your space.

The Classification Checklist

When you encounter a technology, run it through questions like:

- Do people **own** it, or do they **use** it?
- Is the primary **product** a *service* or a *device/agent*?
- Where does the **canonical profile** of the human live:
  - on the provider's servers,
  - or on/with the device?
- Does the business model mostly depend on:
  - ongoing **usage and engagement**,
  - or upfront purchase + optional add-ons?
- When something goes wrong, who does the user feel betrayed by:
  - "their system" (platform),
  - or "my thing" (device/agent)?

You're deciding:
is this primarily a platform-era creature or a thing-era creature?

Hybrids and Transitional Forms

Many real-world systems blend both energies. That's fine — in fact, that's what 4 p.m. looks like.

Examples:

- Smartphones:
  - physical, owned devices (12 p.m.),
  - heavily dependent on cloud services and app ecosystems (6 p.m.).
- AI wearables:
  - bought hardware (12 p.m.),
  - cognitive layer mostly in the cloud (6 p.m.).

For hybrids, you can ask:

- Which side is likely to become dominant over time?
- If you had to design the future version, would you:
  - push more intelligence to the edge,
  - or pull more control into the cloud?

This tells you whether something is likely *migrating toward* 6 p.m. or *toward* 12 p.m.

Why This Classification Matters

You classify not for taxonomy's sake, but because:

- 6 p.m. and 12 p.m. technologies:
  - obey different economics,
  - invite different UX approaches,
  - require different safety and permission models,
  - create different kinds of customer relationships.

If you misclassify, you misdesign:

- Treat a 12 p.m. robot like a 6 p.m. platform → you break loyalty and ownership expectations.
- Treat a 6 p.m. platform like a 12 p.m. device → you underutilize network effects and shared intelligence.

The clock lens is forcing you to align the **shape of your design** with the **shape of the era** you're really in.

Chapter 17 – Psychological Drift: From Owning Smartphones to Using AI

If you want to understand what comes next with robots, you first have to understand what already happened inside people's heads as we moved from the smartphone era into the AI era.

The underlying technology story is one thing: models got better, chips got cheaper, bandwidth increased, smartphones saturated the market, and cloud platforms industrialized machine learning. That's the external arc.

But there's an internal arc too, and it matters just as much for predicting behavior:

- We went from **owning a thing** that felt like part of us
- to **using systems** that live somewhere else.

This chapter is about that psychological drift. It traces the emotional and cognitive move from:

> "This is my phone; it's where my digital life lives"

to

> “I use their AI; sometimes it helps me.”

Once you can see that clearly, you’re in a better position to understand why robots are so psychologically attractive as the next 12 p.m. object—and what will be required for them to actually earn that place.

Smartphone Ownership as a Mental Model

Start with the thing everyone already understands: the way people relate to their phones.

Smartphones are not just gadgets. They sit in the same psychological bucket as:

- house keys,
- wallets,
- glasses,
- maybe even wedding rings.

You can see this in how people talk and act:

- Losing your phone feels like losing a part of yourself. People use that language without irony.
- You don’t say, “I use a smartphone” in any abstract sense. You say, “Where’s *my* phone?”
- The setup process—choosing wallpapers, installing apps, arranging the home screen—feels like decorating a room you’re going to live in.

Three elements reinforce that feeling of ownership.

First, **physical possession**.
The phone is literally in your hand, pocket, or within arm’s reach most of the time. You feel its weight, see its wear, know the cracks and scratches like scars.

Second, **persistent personalization**.
The device isn't just generically capable; it is increasingly filled with your:

- photos, messages, and documents,
- social graphs,
- apps and passwords,
- preferences, shortcuts, and automations.

Two identical phones, out of the box, become irreversibly different after a few weeks of use. They accrete a history with their owner.

Third, **identity linkage**.
For most people, "my phone" is the primary anchor for:

- two-factor authentication,
- messaging and social accounts,
- payment tokens,
- contact lists and calendars.

If your phone is gone, you're temporarily cut off from a lot of things that define your digital identity. It's not the only key, but it's a major one.

All of this adds up to a specific psychological stance:

> This is mine.
> It lives with me.
> If you mess with it, you're messing with me.

That's classic 12 p.m. psychology: concentrated, owned, proximal.

AI as Something You Visit, Not Something You Are

Now contrast that with how people experience today's AI systems.

The dominant AI experiences in 2025 are:

- model-front ends (ChatGPT, Claude, Gemini, Grok, etc.),
- copilots embedded in tools (Office, design suites, IDEs, CRMs),
- APIs that developers call from their own apps.

People talk about those systems very differently:

- "I use ChatGPT sometimes."
- "Our company has access to Gemini."
- "We turned on the AI features in our CRM."

It's the language of **visitation and use**, not ownership.

You rarely hear someone say:

- "Where is my AI?"
- "I lost my AI."
- "This AI feels like a part of me."

Even people who use AI heavily tend to describe it as:

- "a tool I reach for,"
- "a site I log into,"
- "a feature built into my software."

There are exceptions—people who name their favorite system, talk to it daily, and treat it as a quasi-companion—but culturally, those are still edge cases. The mainstream relationship is closer to the way people talked about search engines in the early 2000s:

- “I Googled it.”
- “I checked that site.”

You don’t own Google. You use it.

That’s the 6 p.m. psychological stance: you are a participant in a network, a “user” of someone else’s platform, a subscriber to a service that lives elsewhere.

The Drift from “Mine” to “Theirs”

When you zoom out, the move from smartphone to AI is a move from:

- a **thing** that feels like “me”
  to
- a **system** that feels like “theirs.”

It’s not that smartphones disappeared; they’re still there, and they still feel owned. But the locus of *intelligence*—the thing that surprises you, helps you, builds for you—has drifted toward the cloud.

To get a feel for this:

Think about where “decisions” live.

- In the smartphone era, most of the intelligence you experienced was wrapped in apps that ran on your device or in thin clients calling a server you didn’t think much about. You chose the app, you tapped the button; it felt like the phone was doing the work for you, even when it wasn’t.
- In the AI era, you’re far more aware that the cognitive heavy lifting is happening in models far away. You send text, images, or audio into a box; it responds with something beyond what any local app could do. The “brain” is clearly not on your phone.

Cognitively, that creates a subtle distance:

- The phone is “mine,” but the intelligence is “theirs.”
- I carry the access point, but not the agent.

That’s the psychological drift: the center of gravity for power and agency moves out of your pocket and into other people’s data centers.

How Cloud-Native AI Reshapes the Sense of Self

This drift doesn’t just change how you talk about tools; it changes how you think about yourself in relation to them.

In the smartphone era, your “digital self” is anchored in devices:

- If you’re asked, “Where does your stuff live?” you point to your phone, your laptop, your backups.
- You restore from a device backup and feel like “you” again.
- Even cloud services feel like attachments to your devices, not the other way around.

In the AI era, your “digital self” is increasingly scattered across services:

- recommendation histories,
- embeddings and profiles generated by models,
- traces of your behavior feeding many different systems.

Most people don’t see or control these representations directly. They’re inferred, opaque, and distributed:

- One model builds an implicit profile of your preferences for shopping.
- Another builds a latent representation of your writing style.
- Another models your social behavior.

Each of these “you”s lives inside someone else’s infrastructure. None of them live *with* you in the way the data on your phone does.

Psychologically, this creates three feelings, even if they’re rarely articulated explicitly.

First, a sense of abstraction.
It’s harder to point to “where I am” in the network. You know there are many systems modeling you, but you don’t have a unified picture of that, and you don’t think of it as something you own.

Second, a sense of dependence without possession.
You may grow materially dependent on AI systems—for translation, coding, planning, writing—while fully aware that you don’t control them and can’t keep them if access patterns change.

Third, a quiet erosion of territoriality.
With a phone, you have a sense of digital territory: this is my device, my lock screen, my home screen, my apps. With AI systems embedded everywhere, your presence is less bounded. You interact as a user among many, through interfaces that are the same for everyone.

That’s what 6 p.m. feels like internally: powerful and somewhat placeless.

Emotional Side Effects: Excitement, Anxiety, and Mild Alienation

The drift from owning smartphones to using AI brings a cluster of emotional side effects that often show up mixed together:

Excitement.
People feel real awe at what AI can do: generating fluent text, images, audio, and code; summarizing complex material; translating across languages and modalities. This feels new compared to anything a phone could do on its own.

Convenience.
AI reduces friction:

- drafts instead of blank pages,
- instant explanations instead of searching and filtering manually,
- semantic search instead of keyword guessing.

This is the same “it just works” feeling people had when they first experienced Google or GPS.

Anxiety.
Alongside excitement, there’s a nagging unease:

- Who is seeing what I’m sharing with this system?
- How much of me is being learned, stored, recombined?
- What if access changes? What if the system is biased? What if it fails at the wrong moment?

People feel this explicitly around jobs (“Will I be replaced?”) but also implicitly around autonomy and control.

Mild alienation.
As more of your interactions are mediated by systems that are not yours, you can feel slightly dislocated:

- Support experiences that are obviously driven by bots with platform interests.
- Recommendation systems that seem to nudge you more than serve you.
- AI-driven interfaces that update and change with no regard for your past habits.

At smartphone Noon, the “mine-ness” of the device compensated for some of that alienation. You could say, “At least this piece of the system is mine.” In the pure AI platform world, that anchor weakens.

All of this sets the stage for why the idea of a personal robot is so psychologically resonant even before it fully exists. It’s an answer to a question people haven’t quite found words for yet:

“Can I get something that is as smart as these systems,
but feels as much like mine as my phone does?”

Early Attempts to Re-anchor Identity in AI

You can already see people trying to create islands of ownership inside the sea of platform AI.

Some of that is language:

- People naming their favorite AI system or a specific custom agent,
- talking about “my assistant” or “my bot,”
- saving and reusing conversation threads as if they were personal spaces.

Some of it is behavior:

- customizing system prompts,
- building personal knowledge bases,
- creating bespoke workflows and integrations.

These moves are attempts to pull AI closer to the phone-psychology of “this is mine, tuned to me,” even though the underlying systems are still fundamentally 6 p.m. platforms.

They’re also limited. No matter how much you customize an account inside a platform AI, you know a few things:

- It exists at the provider’s pleasure.
- It has siblings identical in structure for millions of other people.
- It is not a physical artifact that lives with you; it’s an account in somebody else’s house.

That gap—between the desire for a personal, persistent agent and the reality of cloud accounts—is exactly the space robots are going to try to fill.

Robots as Psychological Counterweight

From a technical perspective, you can define robots in many ways: embodied AI, autonomous agents, mechatronic systems, devices with sensors and actuators.

Psychologically, you can define them much more simply:

> A robot is a home for AI that feels like a thing you own.

It doesn't have to be humanoid. It might be:

- a pair of AI glasses,
- a home device on your counter,
- a pod in your car,
- a companion on your desk.

What matters for the clock—and for this chapter—is that it satisfies the same three conditions that smartphones did:

1. **Physical presence** – It lives with you. You can point at it.
2. **Persistent personalization** – It accumulates a history with you.
3. **Identity linkage** – Your "self" with respect to AI is anchored there.

Robots are psychologically attractive because they promise to bring these 12 p.m. qualities back into a world where the intelligence has drifted away from owned objects.

You can think of them as a rebalancing move:

- AI pulled cognition into the cloud.
- Robots pull at least one instance of that cognition back to the edge and attach it to a body that is yours.

For many people, that's not a luxury. It's how they will resolve the tension between wanting the power of AI and wanting the comfort of "mine."

Design Implications: How Not to Fight the Drift

Understanding this psychological drift is not just an interesting narrative; it should change how you design and build.

First, if you build AI purely as a platform service, recognize that you are asking people to give up some of the ownership psychology they had with smartphones. For some use cases, that's fine—people are used to cloud services. But for others, especially long-lived, intimate, high-trust relationships, the lack of a personal anchor will limit how far people will go.

Second, if you're building embodied AI—wearables, devices, robots—you are stepping into the role the smartphone plays today in people's minds. That comes with expectations:

- The device will feel like "mine," not like a remote terminal for your brand.
- Personalization will not be cosmetic; it will define the experience.
- People will expect continuity over years, not months.

You can't ignore the smartphone inheritance. Whatever robot-like thing you build is competing not just with other AI systems, but with the emotional bar set by the phone in people's pockets.

Third, you should assume that people will try to pull AI back toward 12 p.m. on their own:

- creating their own scripts, agents, and wrappers,
- putting AI into local devices where possible,
- favoring tools that allow export, portability, and local memory when they can.

If you fight that—by locking everything into opaque cloud profiles, preventing personalization, and refusing to expose capabilities in a way edge agents can use—you're pushing against a psychological current as well as a technical one.

Fourth, you can strategically decide where on the clock you want to live:

- Some businesses are naturally 6 p.m. businesses: shared infrastructure, central intelligence, participation at scale. Their main job is to be good platforms.
- Some businesses will be primarily 12 p.m. businesses in the robot era: building the things individuals or households own and live with, anchoring identity and loyalty at the edge.

Many organizations will do both: platform in the back, embodiment at the edge. But you can only do that well if you are conscious of the different psychological promises each side makes.

Looking Ahead: From Using AI to Owning a Robot

The main work of this chapter has been to make a simple point precise:

We have already drifted from a world where “my phone is my digital self” to a world where “I use their AI to get things done.” That drift brings power and unease. It creates appetite for something that restores a feeling of ownership without giving up the new capabilities.

Robots—broadly defined—are the likely answer to that appetite. They are the next 12 p.m. thing waiting at the far side of AI 6 p.m.:

- bought, not just subscribed to,
- persistent, not just session-based,
- personalized, not just parameter-tuned,
- loyal to the owner, not to a platform.

If you are building in this space, your job over the next decade is not just to wire up models and sensors. It is to manage this psychological transition responsibly:

- to use AI in ways that don’t erase the sense of “mine,”
- to design embodied agents that actually deserve people’s trust,
- to make sure that as cognition gets more powerful, it doesn’t get more distant.

The rest of the book will talk a lot about architectures, tools, connectors, and ecosystems. Underneath all of that is this quieter story: the shift from owning smartphones to using AI, and the longing to own something again that is both smart and unambiguously ours.

Chapter 18 — Building Your Own Clock for a Specific Industry

The clock is only useful if you can apply it locally.

This chapter is about:

> building a **4 p.m. clock** for your own industry, company, or domain.

You'll answer:

- Where is 12 p.m. for us historically?
- What was our 6 p.m. moment?
- Where does AI sit right now?
- What would Robot Noon look like *here*, not in the abstract?

Step 1 — Identify Your Past 12 p.m. and 6 p.m.

Pick your industry (examples: healthcare, finance, education, logistics, media).

Then identify:

- **Past 12 p.m. (Thing)**
  - What was the dominant *thing* people owned that anchored your domain?
    - examples:
      - paper records, physical branch offices, TV sets, on-prem software, printed textbooks.

- **Past 6 p.m. (Network)**
    - What was the big **network phase**?
        - examples:
            - online banking, EHR systems, streaming platforms, learning management systems, marketplaces.

Ask:

- How did behavior change from 12 to 6 for your customers?
- How did your business model change?
- Who won and who lost in that transition?

You're just retelling your industry's PC → Internet → Smartphone story in its own language.

Step 2 — Place Today on Your Industry Clock

Now:

- Where does **AI** currently sit in your domain?

Questions:

- Is AI still at "pilot project" stage?
- Is it already embedded in core workflows?
- Is it mostly being used:
    - internally (decision support),
    - externally (customer-facing),
    - or both?

Your goal is to roughly say:

> “For us, AI feels like [2 p.m., 3 p.m., 4 p.m.].”

If you’re reading this around 2025, there’s a good chance it feels somewhere between 3 and 5.

Step 3 — Sketch Full 6 p.m. AI for Your Industry

Next, imagine:

- What does **full 6 p.m. AI** look like in *this* domain?

Concretely:

- How many of your key processes become AI-default?
- How many of your customers rely on AI weekly (directly or indirectly)?
- What becomes *weird* to do without AI?

Examples:

- In healthcare:
  - AI triage, AI reading of images, AI documentation, AI scheduling, AI patient education.
- In finance:
  - AI-driven risk scoring, AI personal financial planning, AI fraud monitoring, AI compliance.

Write a one-page sketch:

- “When our industry is at full 6 p.m. AI, here is how a normal week looks for a typical participant.”

Step 4 — Imagine Robot Noon Locally

Now bring in the robots — but very specifically.

Ask:

- What would a 12 p.m. **thing** look like in our domain?
  - a physical robot?
  - a dedicated device?
  - an embodied agent in a specific environment (clinic, warehouse, classroom, branch)?
- Who would own it?
  - individuals, households, professionals, institutions?
- What would people say:
  - “my ___” about, in your domain?

Examples:

- Healthcare:
  - “my health companion,”
  - a device or robot that:
    - monitors, nudges, explains, schedules, advocates, tracks adherence.
- Education:
  - “my tutor,”
  - a persistent agent that:
    - knows the student deeply,
    - interfaces with all learning platforms,
    - lives on their desk or in their bag.

You’re defining Robot Noon **in context**, not as a generic humanoid fantasy.

Step 5 — Identify Your 4 p.m. Moves

Once your local clock is sketched, ask:

- If we are at 4 p.m. now, what are the **right moves**?

These often look like:

- Build **AI-default** workflows without waiting for robots.
- Expose **tools and connectors** that will later serve robots.
- Start experimenting with **embodied or edge experiences**:
    - devices, wearables, in-context agents.
- Design **ownership- and loyalty-aware** patterns:
    - even if your “robot” is currently just a phone app with a voice.

Your aim isn’t to jump straight to 12 p.m.; it’s to **aim the 4 p.m. work** so it survives into the next quarter turn of the hand.

Step 6 — Turn It into a Living Document

Finally:

- Put your industry clock into a single-page document:
    - past 12 p.m.,
    - past 6 p.m.,
    - current time,
    - future 6 p.m.,
    - local Robot Noon vision.

Update it:

- yearly if you’re conservative,
- quarterly if you’re in a fast-moving segment.

The point is to prevent two classic errors:

- acting like it's 1 p.m. when it's really 4 p.m.,
- acting like it's 6 p.m. or 12 p.m. when we're not there yet.

Your own clock is a sanity check against both panic and complacency.