

The psychological ceiling of AI adoption: Why the AI bubble is a social phenomenon, not a technical one

Author: Adrian STAN (Independent author, Pitesti, Romania)

Date: November 27, 2025

Motto: "Loading ..." OR "Thinking ..."? This is the question!

Published in memory of **Leta Stetter Hollingworth (1886-1939)**, who first understood the loneliness of intelligence.

1. Introduction

Comparisons between the current AI frenzy and the dot-com bubble of the late 1990s have become commonplace. With stratospheric valuations for still-unprofitable companies and gigantic investments in infrastructure, the sense of *irrational exuberance* is palpable. But such a parallel, while convenient, hides a fundamental difference and misses a defining **paradox of the present era: private capital is financing**, perhaps for the first time in history - on such a scale - not a product with a predictable profit cycle, but **a cognitive upgrade of the human species**. The return on this colossal investment is measured not in "quarterly profits" but in social progress measurable on a generational scale, rendering classical financial analysis "blind" to the true nature of the phenomenon - and investors paying for a vision that will never exist in the form in which it is presented today.

If the long-term social value is undeniable, why does the sense of an "imminent bubble" persist? What will be corrected if the benefit to humanity is real? The answer is that the market will correct a fundamentally flawed assumption, tacitly shared by investors, AI specialists, and the general public: that technological progress can continue exponentially indefinitely, and society will naturally absorb it. This article argues that **the true limit to artificial intelligence**, the "brake" that will flatten its growth curve, will not be found in hardware architecture or the volume of available data, but in **the architecture of the human psyche**.

Our central thesis is that AI development is hitting a "psychological ceiling", governed by a robust principle from social psychology: **the cognitive gap**. Grounded in the observations of pioneer **Leta Hollingworth**, this theory postulates that meaningful and comfortable human communication breaks down when the intelligence gap between interlocutors exceeds about 20-30 IQ points. Applied to human-machine interaction, our thesis becomes simple and perhaps brutal: humanity, in its vast majority, is not cognitively prepared to support direct and continuous collaboration with a superior intelligence. Current models are already at an **IQ equivalent** of 100-120, at the upper limit of the comfort zone for most of the population. Each additional IQ point from now on does not expand the market, but contracts it, generating cognitive discomfort - a phenomenon supported by modern research on perceived competence as a threat (Fiske et al., 2007) and cognitive status protection (Anderson et al., 2012) - that will cause the average user to seek simpler alternatives.

In the following sections, we will explore in detail the psychological underpinnings of this **cognitive gap** and show why the natural user behavior will not be to abandon the technology, but to migrate to simplified, unidirectional tools. We will mathematically

model how this social pressure forces the AI development curve to shift from exponential to linear and, ultimately, to a socially imposed asymptote. We will analyze the economic implications of a two-tiered market segmentation—a cognitive elite partnering with advanced AIs and a majority consuming derivative products—and we will argue why the realistic future of AI is not one of ubiquitous superintelligence, but one of calibrated and deeply human integration.

2. Theoretical foundations of the Psychological Ceiling

The psychological ceiling hypothesis is not speculation, but an extrapolation of well-documented principles in social psychology and intelligence. Its foundation is the idea that human interaction is governed by a need for balance and cognitive comfort, and that a pronounced intellectual asymmetry destroys this balance. To understand why this principle applies with particular force to human-AI interaction, we need to explore three key psychological mechanisms: the IQ gap as a communication barrier, perceived competence as a threat, and the "cognitive uncanny valley".

2.1 The IQ gap and the cognitive comfort zone

The origin of the concept is found in the works of pioneer **Leta Stetter Hollingworth** in the 1930s, who, studying the social dynamics of gifted children, noticed a critical phenomenon: to be an effective leader, a child had to be smarter than the group, but not *excessively* intelligent. When the gap exceeded about 30 IQ points, communication eroded, shared frames of reference disappeared, and the potential leader became socially isolated. This observation was later formalized in the concept of the "cognitive comfort zone" or "communication range", which postulates that deep and effective human communication is optimal within a range of approximately ± 15 -20 IQ points (one standard deviation). Beyond this threshold, frustration, misunderstandings, and ultimately avoidance of interaction occur.

Our central hypothesis is that this principle, validated in interhuman dynamics (Dyrenforth et al., 2010), applies with **even greater force in human-AI interaction**. If humans avoid other humans perceived as too intelligent, the reasons to avoid a non-human, infinitely patient, and factually superior entity are even stronger. Current AI systems already place themselves at the upper limit of this range for most of the population, making discomfort not a possibility but a statistical certainty.

Conceptual clarification note: This article uses terms like "operational IQ" and numbers like " ± 20 IQ points" not as exact measurements, but as discussion tools for a real phenomenon: when someone perceives you as being much more intelligent than them, communication becomes uncomfortable and people avoid you.

This is well documented - from Leta Hollingworth (1942) who observed how gifted children become socially isolated, to Gross 's (1993) research on the extreme isolation of adults with IQs above 160, to the modern theory of perceived competence as threat (Fiske, Cuddy, Glick, 2007). What we don't know exactly are the precise numbers - at what IQ point difference does the problem arise?

The figures in this article (± 20 IQ points, estimating that LLMs are at $100 \div 120$ IQ) are estimates meant to make the discussion concrete, not rigorous psychometric measurements. What really matters is not whether we have estimated the figure accurately, but whether the anticipated behavioral patterns are actually observed:

people use AI in two extreme ways (very simple or very complex, with no "middle ground"), those with less education avoid complex tasks more than those with higher education, most prefer simplified answers even when complex ones are better (details in Section 5.4).

When we talk about the "operational IQ" of AIs, we are strictly referring to how people feel when interacting with them: "Does the AI write better, analyze faster, structure more clearly than me? Then I feel intimidated". We are not talking about a scientific measurement of general intelligence, but about the perception of competence in the areas where people themselves "work" with the AI : writing, analysis, problem solving.

2.2 Competence as a threat: the warmth-competence matrix

The psychological mechanism that best explains *why* this gap is uncomfortable is described by the universal model of social cognition proposed by Fiske, Cuddy, and Glick (2007). According to this model, our perception of others (or, in this case, "the other") is structured along two fundamental axes: **warmth** (intention, friendliness, empathy) and **competence** (skill, efficiency, intelligence). An advanced AI is "maximum competence " but is perceived as having "warmth towards zero". This combination of "high competence, low warmth" places the AI in the quadrant of cold, envied, but unpleasant stereotypes (e.g., "the cold and calculating robot"). Such a perception does not trigger admiration, but an instinctive sense of threat, suspicion, and social distance. The user does not perceive that he has a partner, but that he is being evaluated by a superior tool, which turns collaboration into an examination. This type of perception amplifies what the literature calls "*perceived competence threat*", a documented mechanism by which people avoid entities perceived as being much more competent than them.

2.3 Cognitive status and ego defense mechanisms

Interaction with a superior AI is not just an exchange of information, but also an implicit negotiation of status. Research shows that people go to great lengths to protect their perceived social and intellectual status, avoiding situations that might expose their incompetence (Anderson et al., 2012). When a user receives an answer from an LLM that is not only correct, but also more nuanced, better structured, and deeper than they could have formulated themselves, they experience **a direct threat to their self-esteem**. This phenomenon explains why many users underuse AI systems, limiting themselves to simple, non-challenging tasks: it is an ego defense mechanism. The migration to simplified, unidirectional tools (such as advanced semantic search) is a rational strategy to obtain the benefits of technology without paying the psychological cost of confronting their own cognitive inferiority.

2.4 The Cognitive "Uncanny Valley"

Finally, the discomfort is amplified by a phenomenon we can call "cognitive uncanny valley", an analogy to Mori's (1970/2012) concept in robotics. Just as a robot that is nearly but not perfectly human evokes a sense of repulsion, an intelligence that perfectly mimics *the outputs* of human thought (language, logic, creativity) but is fundamentally devoid of the underlying *processes* (consciousness, subjectivity, lived experience) is deeply disturbing. This intelligence is simultaneously familiar in form and completely

alien in nature. This existential dissonance, conscious or not, contributes to the psychological ceiling, making prolonged and deep interaction exhausting and ultimately avoidable for most people.

3. Current state of LLM intelligence and First Signs of Ceiling

The psychological ceiling argument would remain a mere theoretical curiosity if the technology had not already reached the critical threshold. This section demonstrates that top-level language models (*frontier models*) of 2024-2025 is right at the upper limit of the cognitive comfort zone for most of the population. Moreover, we are already seeing the first economic and strategic indications that suggest that the AI industry is starting to react, consciously or not, to this social barrier.

3.1 Measuring Artificial Intelligence: the 120 IQ threshold

Assigning an IQ score to an artificial system is a methodologically complex and controversial task. Standard IQ tests assess a range of human abilities (spatial reasoning, temporal processing) where the architecture of LLMs has fundamental limitations. However, by triangulating performance on tasks measuring logical reasoning, pattern recognition, and general knowledge, we can obtain a reasonable estimate. Various approaches converge towards a consistent range:

- **Mensa - type puzzles** : Abstract reasoning tests indicate equivalent scores of 100-120 for top models.
- **Academic benchmark performance** : Results on standardized tests such as the MMLU (Massive Multitask Language Understanding) places the models at a level equivalent to that of a college graduate, a performance achieved by only 15% to 30% of the general population.
- **Proficiency in professional tasks**: The ability to pass bar, medical, or accounting exams suggests competence that falls within the same range of IQ 115-125.

The critical point is not the exact value, but the order of magnitude. All of these estimates place the current generation of top AIs **right at the upper limit of the cognitive comfort zone** for an individual with average intelligence (IQ 100). These are no longer simple tools; they are entities perceived as being "smarter" than most of their users, thus triggering the psychological mechanisms described earlier.

Additional methodological clarification: It is essential to clarify that attributing an "IQ equivalent of 100-120" to AI systems does not suggest that they possess general human intelligence or that there is a global equivalence of cognitive capabilities. Current LLMs have peak (sometimes super-human) performance on language domains, verbal reasoning, and knowledge retrieval, but are significantly sub-human on spatial reasoning, causal understanding in physical systems, temporal processing, and sensory-motor integration.

The "IQ 100÷120" range refers strictly to performance on tasks where the average user actually interacts with the system: complex verbal communication, logical problem solving, information synthesis and content generation. It is this perceived competence on the relevant interaction domain - not a global measure of intelligence - that triggers the psychological mechanisms of discomfort described in Section 2. For example, a user with an IQ of 95 is not comparing the AI's ability to model quantum physics (a

domain where he does not interact), but its ability to produce more sophisticated analytical texts than he himself could formulate - and this direct comparison generates the threat to cognitive status.

3.2 The first signs of Social Ceiling

The apparent slowdown in progress from one generation of models to another is often attributed to technical limitations: reaching the limits of scaling laws, depletion of high-quality training data, or architectural constraints. The psychological ceiling hypothesis, however, offers an alternative or, more likely, complementary explanation: **there is less and less economic pressure to exceed the 120 IQ threshold for the mass market.** We already see several signals that support this idea:

- **Product strategy shift:** AI companies are visibly shifting their focus from increasing raw capabilities to **optimizing efficiency, personalization, and, most importantly, improving the user interface (UI/UX).** This effort to make AI friendlier, more accessible, and less intimidating is an implicit recognition that raw intellectual power is not the primary driver of mass adoption.
- **The emergence of "tamed" models:** The dominant strategy has become to offer a range of models. Typically, a high-end, extremely capable (and expensive) model is available to a niche of developers and enterprise customers, while the mass market is offered smaller, faster, and "friendlier" versions. This is an early form of the two-tier structure we anticipate: cutting-edge technology is developed, but it is "tamed" and simplified so as not to alienate the majority of users.
- **The Utilization Gap:** Numerous market studies indicate that most users use these extremely powerful tools for relatively simple tasks (summaries, emails, translations), systematically avoiding the complex analytical tasks for which they were designed. This behavior does not reflect a lack of AI capability, but a reluctance of the user to enter into a cognitive "partnership" that could expose its limits.

These trends converge towards an inescapable economic conclusion: investing billions of dollars to develop an AI with an IQ equivalent of 150, which would be comfortable for less than 2% of the population, is a bankrupt business strategy. The market itself is starting to impose a natural brake, not because the technology cannot advance, but because its advance becomes economically irrational and socially counterproductive.

4. The S-curve of AI development: from exponential to social asymptote

The history of technology adoption is often described by a logistic curve, or **S-curve** (Rogers, 2003). Traditionally, this curve models the diffusion of an innovation in society, flattening as the market reaches saturation. We propose a reinterpretation of this model for the development of AI capabilities: the flattening of the curve will not be determined by market saturation, but by **its psychological resistance.** AI development does not follow an infinite exponential trajectory, but an S-curve whose **inflection point** and asymptote are **imposed by social, not technical, constraints.**

4.1 Mathematical formulation of the social ceiling

The trajectory of AI capabilities can be divided into three distinct phases, which correspond to the classic segments of an S-curve:

- **Phase 1 (2017-2024): Exponential Growth.** This period was characterized by an exponential growth in capabilities. Each new generation brought shocking qualitative improvements, and the market enthusiasm reflected this dynamic. In this phase, AI was still in the comfort zone of most advanced users ("early adopters").



- **Phase 2 (2025-2027): Linear Deceleration.** We estimate that we are at the beginning of this phase. As AI capabilities reach and exceed the psychological threshold of ~120 IQ, social pressure begins to act as a brake. Improvements continue, but are perceived as incremental, not revolutionary. Market forces begin to favor optimization and accessibility over raw intelligence growth, as each added IQ point brings diminishing returns by alienating a growing segment of users.

- **Phase 3 (2028+): Social Asymptote.** In this phase, the development of AI capabilities for the mass market will reach a plateau. The curve will flatten out not because technical progress is impossible, but because it becomes economically irrational (Baumol, 2002). The asymptote (C_{max}) does not represent the maximum intelligence theoretically possible, but **the maximum intelligence socially acceptable and economically viable for mass adoption**, likely in the range of 120-130 IQ. Development beyond this point will continue, but will be limited to niche products aimed at the cognitive elite. This range does not represent an intrinsic cognitive limit, but is chosen for consistency with the literature on the human cognitive comfort zone and reflects the point at which perceived differences generate discomfort and avoidance of interaction.

The logistic equation,

$$C(t) = \frac{C_{max}}{1 + e^{-k(t-t_0)}}$$

where t_0 is the inflection point (approx. 2024-2025) and C_{max} is the social ceiling, provides a mathematical description of this process : it starts fast (exponentially), slows down around 2025, then flattens out to a maximum level imposed not by technology, but by human psychological tolerance. C_{max} is not the maximum possible intelligence - but the maximum accepted intelligence.



4.2 Comparison with the “Dot-Com Bubble”: technology vs. timing

This perspective on the S-curve highlights the fundamental difference between the “AI bubble” and the “dot-com bubble”.

Dot-Com Bubble: It was a case of **immature technology**. The vision was right, but the technology, infrastructure, and user base were not ready. The bubble burst because **the**

promises could not be technically delivered. The value was projected into a future that had not yet arrived.

The AI Bubble: It is a case of **technology being "too mature" for society.** The technology is already capable of exceeding expectations, but the user base is not psychologically prepared to absorb it. The bubble (the correction of expectations) will burst not because the technology fails, but because its success **exceeds the market's ability to adapt.** It is a problem of **synchronization** between two curves with fundamentally different speeds: the exponential development of technology and the linear, generational adaptation of the human psyche.

4.3 Investment implications

If the social S-curve hypothesis is correct, current valuations of AI companies are built on a false premise: that of continued exponential growth and universal adoption. Investors who extrapolate Phase 1 successes over the next decade will face a harsh reality as the market enters a Phase 2 deceleration. The correction will come when the broader market understands that the addressable market for cutting-edge AI is, and will remain, niche for a long time. This does not negate the value of the technology, but it does require a radical reassessment of the time horizon for investment recovery—from years to decades or even generations.

5. The emergence of a two-tier market: the cognitive elite and the mass consumer

The psychological ceiling pressure will not stop the development of artificial intelligence, but will channel it into two distinct trajectories, leading to the natural segmentation of the market into a two-tier structure. This structure is not a consequence of the marketing strategy of companies, but a direct reflection of the normal distribution of intelligence in the human population. Advanced technologies, throughout history, have tended to amplify existing differences, creating hierarchies based on skills, and AI will dramatically accelerate this process.

5.1 Level 1: Symbiotic partnership of the cognitive elite (~2-5% of the market)

This level will consist of users who are at the upper extremes of the Gaussian curve - individuals with an IQ above 130, professionals in highly specialized fields (research, engineering, complex strategy), and "digital natives" with exceptional cognitive plasticity. For this segment, the psychological ceiling is not a barrier, but a starting point.

- **Mode of Interaction:** These users will not simply seek answers, but will enter into a **symbiotic partnership** with cutting-edge AIs (IQ 140+). The interaction will be bidirectional, iterative, and deeply complex, with the AI acting as an "exocortex" that amplifies and accelerates the processes of analysis and synthesis.
- **Products:** They will use frontier models, paying premium prices for access to maximal, unfettered capabilities. The focus will not be on "friendliness" or affordability, but on raw power, speed, and flexibility.
- **Value generated:** This small group will generate disproportionately large amounts of value. They will be the ones who will produce the innovations, scientific discoveries, and business strategies that will shape society. They are not just using AI, they **are thinking with it.**

5.2 Level 2: Simplified Mass Market Consumption (over 95% of the market)

This level encompasses the majority of the population, which, according to our hypothesis, feels uncomfortable interacting directly with a higher intelligence. For this segment, the value of AI lies not in cognitive partnership, but in its utility as an efficient and non-intrusive tool.

- **Interaction Mode:** Users will prefer one-way and simplified interactions. They will look for tools that provide quick responses, automate repetitive tasks, and integrate seamlessly into existing workflows without causing cognitive discomfort (*estimated: 85-92% one-way, 7-12% semi-two-way, ~1-3% deep partnership interaction*). It is important to emphasize that this phenomenon is not the result of general technological anxiety, but of the specific cognitive discomfort generated by interacting with an agent perceived as intellectually superior.
- **Products:** They will use "tamed" and specialized versions of AI technology, embedded in applications they already know. The focus will be on user experience, "personality", and psychological safety. These products will be optimized to be helpers, not mentors; assistants, not superiors.
- **Value consumed:** This segment will largely be a **consumer of products and innovations generated by Level 1**. For example, a drug discovered by a researcher (Level 1) using cutting-edge AI will be prescribed by doctors (Level 2) using an AI assistant that simplifies protocols. The economic value here is huge in volume, but small per user.

5.3 Strategic and economic implications

Recognizing this two-tier structure has radical implications for AI companies and investors:

1. **The End of the "One-Size-Fits-All" Model:** The strategy of developing a single frontier model for the entire market is doomed to failure. Successful companies will have to develop radically different product lines for the two tiers, with distinct R&D and marketing teams.
2. **The value is in the Interface, not just the Model:** For the mass market, the battle will not be over who has the smartest model, but over who creates the most **comfortable and useful interface** to deliver that intelligence.
3. **Reassessing the Addressable Market:** Investors need to understand that the market for cutting-edge AI, sold as a premium service, is a niche. The true mass market is for derivative, simplified products, where profit margins are likely to be lower. This economic model ("winner-takes-most" for Level 1, fragmented market for Level 2) is much more complex than market valuations assume (Brynjolfsson & McAfee, 2014).

In conclusion, the future is not a democratization of superintelligence, but a **stratification of access to cognitive partnership**, a reality that has profound consequences not only economic, but also social and ethical.

5.4 Testable predictions and research agenda

If the psychological ceiling hypothesis is correct, we should observe certain specific phenomena in the period **2026-2028**.

P1. People will use AI in two extreme ways, with no "middle ground"

What we anticipate: Most conversations will be either very short (1-3 questions: "Write me an email", "Summarize this text") or very long (10+ questions: deep collaboration, complex analysis). Almost no one will use AI in "medium" mode (5-7 questions).

How we check: We analyze usage statistics - how many questions a user asks per session. If the graph shows two peaks (one at 1-3, another at 10+) and a "valley" in the middle, the hypothesis is confirmed.

Additional prediction: People with higher education or technical professions will have more long sessions than the rest.

P2. Smarter AIs will be used less by ordinary people

What we predict: When a new, more capable model comes out, highly educated users will use it more. But average users will use it as little as before, or even less.

How we check: We compare how many active users the product has in the month before the launch of the new model and after, separated by groups: those with higher education vs. those without.

Quantitative prediction: If the "higher education" group increases by +20% usage, the "middle education" group will increase by <5% or decrease.

P3. People prefer simpler answers, even when complex ones are better

What we predict: If you show the same person two answers to the same question - one simple (high school language), one complex (academic language) - most will prefer the simple answer, even if both are correct and free.

How we check: Simple experiment: 500 people are asked the same question, half get a simple answer, half get a complex answer. We ask: "Which answer do you prefer?"

Quantitative prediction:

- People with a bachelor's degree: >65% prefer simple
- People with a master's/doctorate degree: <30% prefer simple

P4. "Too smart" AI makes you feel stupid, then you avoid it

What we anticipate: Users who feel that the AI is much smarter than them will experience anxiety, feelings of inferiority, and will avoid using it for tasks where they would feel directly compared.

How we verify: 6-month study with 500 people:

- we measure their IQ (approximately)
- we make them use AI
- we ask them monthly questions: "How anxious do you feel when using AI?", "Do you feel like AI makes you seem incompetent?"

Prediction: The greater the perceived difference between one's intelligence and the AI's, the greater the anxiety and the lower the usage.

P5. People will not pay for smarter AI, but they will pay for simpler AI

What we predict: If you cut the price of the "super-smart" version by 50%, sales don't double. But if you release a much cheaper "simplified" version, everyone buys it.

How we check: AI companies are doing pricing experiments:

- I'm dropping the price to "tier" premium" (smartest model) by 50% - how many new people subscribe?
- Launch a "basic" tier (simpler model) at a low price - how many new people subscribe?

Prediction: The "basic" tier will have 3-5 times more new users than the premium tier. reduced premium, although premium is more capable.

How to test everything at once: a complete experiment

It takes ~500 people, ~2 months, budget ~25,000-30,000 Euros:

Experiment:

1. People are divided into 3 groups (randomly):
 - Group A: use complex AI (academic, detailed answers)
 - Group B: use simple AI (clear, short answers)
 - Group C: use Google normally (control)
2. Before and after 2 months, the following is measured:
 - How often do I use it?
 - How happy they are
 - How anxious they feel
 - What types of questions do I ask?

If the hypothesis is correct:

- People with lower IQs avoid Group A (complex AI) and use Group B (simple AI)
- People with higher IQs love Group A and are bored in Group B
- Anxiety is higher in Group A for people with secondary education

If the predictions are confirmed: The psychological ceiling is real and AI companies need to take it seriously when making product strategies and when estimating how many people will use the product.

6. Conclusions: The Future of AI in a "Human World"

I have argued throughout this article that the anticipated AI market correction, the so-called "bubble burst", will not be triggered by a failure of the technology, but by its success—a success so rapid and profound that it **exceeds the human psyche's capacity to absorb it**. I have shifted the analysis from financial metrics to psychological factors, proposing that the real brake on the universal adoption of artificial intelligence is a **psychological ceiling**, a barrier of cognitive comfort imposed by the natural limits of human interaction.

6.1 Summary of arguments

Our thesis started from a simple but fundamental premise: as in human interactions, human-machine cognitive partnerships are productive only in a comfort zone where the intelligence gap is not overwhelming. Once AI has crossed the critical threshold of ~120 IQ, it has entered a territory where, for the majority of the population, it is no longer an assistant, but a source of cognitive discomfort, a threat to self-esteem and intellectual status.

This social ceiling does not stop progress, but reshapes it. We have shown how market pressure forces the exponential development curve of AI to slow down and flatten into a social asymptote. The direct consequence is the inevitable segmentation of the market

into a two-tier structure: a cognitive elite (<5%) engaged in *symbiotic partnerships* with frontier AIs, and a mass market (>95%) that naturally migrates to simplified, one-way tools that offer the benefits of technology without the psychological cost.

6.2 A calibrated, not exponential, future

This perspective forces us to radically rethink the future. The vision of a universal superintelligence, accessed equally by all, is a technocratic utopia that **ignores the reality of human nature**. The far more likely - and perhaps more sustainable - future is one of **calibrated integration**. It involves:

- **for investors:** A reconciliation between short-term profit expectations and the reality of a cognitive leap in society that will mature on a generational scale, not financial quarters.
- **for developers:** A paradigm shift, from maximizing raw capabilities to optimizing psychological comfort and real utility for the end user. The battle for the mass market will be fought on interface and experience, not IQ points.
- **for society:** An awareness that this technology, more than any other, will create **new cognitive hierarchies**. Managing these gaps to avoid a deep social fracture becomes one of the greatest ethical and governance challenges of the 21st century.

7. Speculative Epilogue: Homo Symbioticus - a possible direction of evolution

Note: This section goes beyond the analytical framework of the article and enters speculative territory, not being a *prediction*, but a theoretical exploration of a possible direction of social evolution in the long term (2-3 generations).

This section represents a prospective theoretical extension, not an empirical conclusion, and is intended to illustrate a possible direction of evolution of social structures resulting from the phenomenon under analysis.

While this article has focused on the limits of the present - what most people can't do today - there is a natural question: what happens to those who can? That minority of 2-5% who don't feel the psychological ceiling, but see AI as a real cognitive partner - could it constitute the germ of a new **form of human organization** ?

A variant of this concept also appears in Slot (2025) under the name *Homo Symboticus*, but the present work adopts the form **Homo Symbioticus**, which more faithfully reflects the idea of cognitive symbiosis and evolutionary partnership between humans and AI. From another perspective, that of this article, this idea can be reinterpreted **sociologically**: as the emergence of a **symbiotic caste**, a class of users capable of overcoming the psychological ceiling and entering into an authentic form of cognitive coevolution with artificial intelligences.

These individuals could form what we call here **Homo Symbioticus**: people who don't just use AI, but think with AI. For them, AI is no longer an external tool (like a computer or Google), but becomes something closer to a mental extension - as if you had a permanent cognitive co-pilot, with whom you reflexively consult on any complex problem - an *exocortex* in a practical sense.

Such a group would inevitably have massive advantages in economics, science, and strategic decision-making. If for the past 10,000 years power has come from controlling

resources (land, oil, money), in the future power could come from the ability to think in symbiosis with AI. It no longer matters how rich you are, but how well you can cognitively collaborate with an artificial intelligence - **economic capital becoming subsidiary to augmented cognitive capital**. In this sense, the symbiotic class would not be a genetic elite, but a **cognitive-technological one**.

This idea also appears in other authors - for example Slot (2025) speaks of "*Homo Symbioticus*" in neuropsychological terms. Our perspective is **sociological**: we are not talking about biological changes in the brain, but about the emergence of a new type of "social class" - defined not by resources or origin, but by the ability to work in tandem with AI. If the psychological ceiling is ever to be exceeded at a broad social level, it will probably happen gradually: first by this minority that is now overcoming it, then - perhaps in a few generations - through normalization and education.

In other works, I have argued that **human-AI symbiosis** may be the most likely path for the future evolution of the human species (Stan, 2025). While this is a much broader view than the subject of this article, it provides a context in which the phenomenon observed here-the emergence of an **augmented cognitive elite** - becomes not an anomaly but a **natural first step**.

From this perspective, the emergence of *Homo Symbioticus* it is not a utopian or technological scenario, but a **possible mutation of the social order**. Instead of traditional economic hierarchies, the future could be marked by a cognitive stratification: between those who can think in symbiosis with AI and those who prefer simplified interactions. It would be, perhaps, **the most profound social transformation** since the agricultural revolution.

In this scenario, **Homo Symbioticus** is therefore a possible social reality: a new form of human stratification. We no longer divide society into "rich vs. poor", but into those who can think in partnership with AI vs. those who can't (or won't).

It would be perhaps the most profound social change since humans invented agriculture 10,000 years ago. Agriculture created social classes based on land. Industrialization moved them to capital. Symbiosis with AI would move them to augmented cognitive capacity.

Final considerations

The AI bubble will correct, and the technology will remain. The difference from the dot-com bubble : then the technology had to catch up with the vision. Now the vision (and the investments) has to catch up with **the human psychological reality**.

The survivors of this correction will be the companies that understand: you don't build artificial intelligence in a vacuum, you integrate it into the (complex, fragile, and often irrational) human psyche.

This analysis has limitations: it describes trends, not absolute laws. Better interfaces, education, and habituation could alter the trajectories described here. But the general direction-that there is a real psychological ceiling-remains plausible.

The successful future of AI is not one that surpasses us, but one that learns to walk alongside us, at our pace.