

AI vs.Human

AI vs.Human:The risks of AI adoption for emerging markets



Venugopal Garre
+6563267643
venugopal.garre@bernsteinsg.com



Nikhil Arela
+912268421482
nikhil.arela@bernsteinsg.com

AI adoption is not just a developed market story. Emerging markets are increasingly becoming large-scale users, and as we argued in our previous report, per capita metrics often understate the extent of AI adoption in these economies. In this report, we explore what that means for one of the countries most exposed to both the opportunities and risks of AI adoption: India.

EMs except for China are adopters of technology and rarely originate their own tech platforms. In the various phases of the technology percolation - the internet era offered an interesting adoption path. India for example boasts of the world's largest digital user bases. Whether it is WhatsApp, YouTube, Instagram, Google Pay or Amazon, India consistently ranks among the largest markets globally by users. The reason for this is that, most of these services were offered for free, allowing rapid adoption in a highly price-sensitive market, while platform companies monetized through advertising, transactions and ecosystem expansion and free data to power their ML/AI training. The first phase of the digital revolution, however, was largely additive rather than disruptive as digital platforms rarely replaced existing jobs. Instead, they made commerce, communication, entertainment and payments more efficient, convenient and engaging.

AI adoption-paradigm shift from other tech products: In AI as well, India has higher users of various US based LLM offerings than in the US driven by subsidized pricing. However, AI represents a fundamentally different shift. Large language models are expensive to build, train and operate, the variable costs per user are high - in sharp contrast to traditional games and apps where initial higher costs in R&D are followed by lower upkeep costs over a large user base. The industry has embraced subscription-based monetization from day one: be it ChatGPT, Gemini, Claude, Microsoft Copilot or AI-native enterprise applications, users have been expected to pay for access and productivity. In a market like India, where paid software still remains lowly-penetrated, does this spell a disaster for adoption after LLM subsidies moderate?

Subscriptions may become a norm: Beyond the obvious dependency on the exact territory that LLM costs chart, the usage in India as per both Open AI and Anthropic stats suggests that stickiness could indeed be high and token usage will expand. A higher usage for technical tasks (6.5% vs 4% global), high usage in computers and mathematical, more productive usage than advanced economies and a greater share of users under 35 years (78% vs 66% global) are all pointing towards higher lifetime value and more propensity to pay. Initially, business use cases can especially expand, partly funded through lower headcounts. It's thus, a mistake to assume AI will not impact jobs as the adoption remains fringe and at low cost in most organizations currently.

From petrodollar to LLM the new crude?: Our broad calculations indicate that local AI usage may lead to as high as \$20-30 billion of annual outflows for India as LLM fees in a decade, excluding any spending with export offsets. The application layer will charge additional fees which could be in tens of billions, with most of those flows moving to US tech companies. This will make AI equivalent to a third of current oil imports. Any development leading to massive LLM cost reduction and locally owned AI applications can change this scenario somewhat.

DETAILS

Few weeks back, we had written on how countries around the world are using AI—which drew a sharp contrast between emerging and advanced economies. We had covered quite a few sectoral implications back then, but a look at data from OpenAI cements a few trends concretely—and something deeper, more macro level change is in store for us. Here, we take a look at what the usage stats tell us about the future.


TOWARDS BEING A POWER USER

India crossed the threshold of 100 million active weekly ChatGPT users in February, and is the second-biggest market for OpenAI after the US. The per capita usage remains low, where India ranks 109 out of 150 economies. This is virtually the same picture seen in the Anthropic index, where India at 7.1% absolute usage was second to the US, but ranked 103 out of 121 economies when looking at per capita usage. We have already determined why per capita metrics do not reflect the extent of benefits reaped out of AI—especially for countries like India where large workforce is in industries relatively lesser exposed to AI.

With the relatively narrow base using AI in India, the use case still amounts to immense proportions. Despite lagging global averages severely when it comes to per capita usage—both for Anthropic and for OpenAI, India's usage of data analysis tools from ChatGPT is four times the global average when it comes to paid subscribers. Similarly, the usage of ChatGPT's coding tools is three times the global average.

Overall, India's work-related usage in overall ChatGPT comes out at 35%, much ahead of the 31% global average. It is also higher than many advanced economies like the US, UK and other European economies. We see emerging economies significantly ahead of advanced economies here—32.4% work-related use vs 29.3%. Does that remind of the takeaways from our extensive study on Anthropic data a while back?

EXHIBIT 1: A far larger share of Emerging economies' AI usage is for work related tasks, compared to advanced economies




Source: OpenAI Signals Data, Bernstein Analysis

COMPLEX TASKS, REAPING MORE BENEFITS-TREND, NOT AN ANOMALY

Another thing that came up in our study was a high usage of India and other emerging economies in fields like computers and engineering. OpenAI segments its data differently—clubbing the technical fields under "Technical Help". Following directly from those trends, here too we see technical usage especially high among emerging economies, with India being the second highest in the 28 key economies we tracked. This is despite this use case falling rapidly over the past two years—back in 2024 almost


20%(or one fifth)of the overall queries from Indian users on ChatGPT were clubbed under "technical help",which has dropped to 6.3%today.Much of the gains have been taken over by multimedia(1.4%to 9.8%),self-expression(2%to 6.3%),practical guidance(26%to 33%),or seeking information(15%to 17%).However,when it comes to self-expression or information seeking,advanced economies are still placed comfortably ahead in usage.

EXHIBIT 2: India's usage in technical help witnessed a sharp decline(similar to global trends),though it still sees one of the highest usages globally in this segment.Writing,Seeking information and practical guidance witnessed an increase

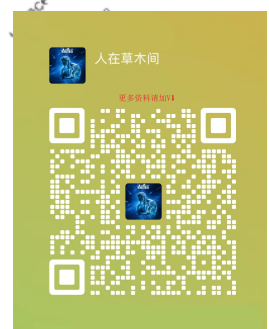


Source:OpenAI Signals Data,Bernstein Analysis

EXHIBIT 3: Even with reduced usage,technical help is one of the most prominent use cases of AI in India globally




Source:OpenAI Signals Data,Bernstein Analysis



AN EERILY SIMILAR RESULT AND VIEWING THE AUTOMATION VS AUGMENTATION DEBATE

This gives India's usage stats a profile remarkably similar to what we observed with Anthropic's Economic Index data. While emerging economies lead when it comes to technical usage and writing tasks given to AI, the advanced economies use it far more for information gathering, self-expression and practical guidance. Quite simply, emerging economies are using AI for more work related tasks or to attain objectives linked to automation, writing work emails or similar writing work, while advanced economies are using AI as an augmentation tool. While doing the work largely themselves, they use AI to gather more information and to help them express better. This, perhaps, also explains the lower per capita usage in developing economies: since the usage remains more centred around work, the professions that are relatively lesser exposed to AI haven't become material users.

EXHIBIT 4: India's writing usage remains relatively low, comparable to advanced economies rather than emerging economies which are near 25% or more




Source: OpenAI Signals Data, Bernstein Analysis

Despite the strikingly similar results we observe, we couldn't help notice some sharp divergences when it comes to ChatGPT vs Anthropic usage. Computer and Mathematical accounted for over a quarter of the usage (>25%) across all emerging economies, while even advanced economies shared a lower but healthy usage profile. When it comes to ChatGPT, the usage in technical aspects has shrunk very fast: around 13% (2024) to 4% now globally. Clearly, the tech users have migrated elsewhere, making the platform lesser suitable for development and coding problems compared to newer entrants and newer tools targeted at coding problems.


A similar market share loss has happened with writing tasks, and instead we see practical guidance and multimedia shoot up in usage stats, demonstrating arising ChatGPT use in seeking personal help as well as for video/image generations or editing.

EXHIBIT 5: Developed economies are using AI more for information gathering, conforming to Augmentation vs Automation debate that we saw in Anthropic usage



Source: OpenAI Signals Data, Bernstein Analysis

EXHIBIT 6: While Anthropic is more used for technical help and writing, ChatGPT is immensely popular for practical guidance and information seeking



The use categories are as per reported by OpenAI signals data and the 37 generalized work activities from Anthropic are split into differing types (each activity can be of 3 different categories with differing mix) as per our estimates to reach these numbers.


Source: Anthropic Economic Index, OpenAI Signals Data, Bernstein Analysis and Estimates

SUMMING UP THE USAGE DATA-THE KEY TAKEAWAYS


Some very interesting points arise when we look at the data from Anthropic and OpenAI together and this is what we've attempted to do throughout this note rather than making it an isolated note on AI usage.

The first thing, which we did see traces of when analyzing Anthropic data, but never had the concrete data, is the age profile of users. Globally, nearly 70% of the users of ChatGPT back in 2024 were under 35 years of age. Today, that number is down to 66% (Exhibit 7), which indicates a steady rise in older people accessing AI with their share at 34%. This is spread rather unevenly among emerging and advanced economies. Nearly 40% of users in advanced economies are over 35 years of age, with 18% above 45 years of age (Exhibit 9). In emerging economies, the corresponding numbers are 28% (above 35) and 11% (above 45, Exhibit 10). India is the most skewed towards a younger population. 92% of users are under 45 years of age, indicating a marginal share of users in the older age brackets. For a country where we have nearly 30% population above the age of 45 and over a billion internet users, this is clearly a metric of a skewed usage. While the Anthropic data did give this hint given the higher usage in technical fields and coursework, OpenAI's data gives a conclusive proof.


EXHIBIT 7: Globally, the mix is more even - with ~35% AI users above 35 years of age



Source: OpenAI Signals Data, Bernstein Analysis


EXHIBIT 8: Nearly 80% AI users in India are under 35 years of age

Source: OpenAI Signals Data, Bernstein Analysis

EXHIBIT 9: Advanced economies have the most even mix-with some countries having nearly half of AI users above 35 years of age

Source: OpenAI Signals Data, Bernstein Analysis

EXHIBIT 10: India's profile is a more extreme case of emerging economies where a young workforce is more prominent-EMS appear to be more balanced across age groups than India




Source: OpenAI Signals Data, Bernstein Analysis

Another thing we noticed is the secular progression of AI tasks, which leads to an interesting hypothesis/conclusion. At its core, the advent of Generative AI back in 2022 set a few things in motion—the transition from learning to experimenting to scaling to evolving (which we discussed in one of our earlier notes). If one goes by this evolution, he would expect a high degree of tasks at "information gathering" in the initial stages, and a gradual progression to multimedia and technical tasks with more automation in the later stages. However, we notice the share of information gathering has stayed broadly similar in emerging economies (a slight decrease), and has actually increased in advanced economies. Automation is higher in India (and emerging economies), along with usage for work-related tasks. It seems that emerging economies are graduating to more advanced use cases of AI faster than the advanced economies, putting them at a more matured scale of AI usage lifecycle.

Let's reflect back on the usage data by age group, and club it with the usage for tasks like technical help, writing or practical guidance—which we now club under "productive use", as it gives some end result that's likely to be used at work or in a similar setting. Now, the age group usage gives an indication of how long someone is likely to remain associated with an AI tool, the productive use gives some indication of "willingness to pay". People who are using AI for work or their own entrepreneurial set up are more likely to pay for a subscription. Using these two, we reach the matrix described in [Exhibit 12](#). The results are quite unlike what one would expect. The total lifetime value appears way higher for emerging economies compared to the advanced ones—given a higher share of young users plus a greater use towards productive tasks. Thus, we have both lifetime value plus propensity to pay. The only challenge here is ability to pay—given the difference in income levels between emerging and advanced economies. Hence, if the pricing of subscription plans were to take care of income levels, one would see far more paid users emerging in the developing countries.


EXHIBIT 11: Emerging economies lead in using AI for productive and work related tasks, while advanced economies lead in information gathering or more experimental tasks



Productive usage is considered to be usage spent in technical help, writing or practical guidance. Knowledge/Work linked usage is only writing and technical help. Discovery/Experimentation is usage for information seeking or self-expression.

Source: OpenAI Signals Data, Bernstein Analysis

EXHIBIT 12: Emerging economies (red dots) present some of the best cases for lifetime economic value through subscriptions of LLMs



LTV refers to lifetime value of a consumer, assumed higher for younger users. P2P refers to propensity to pay for a subscription, which has been assumed to be higher for users which are using AI for productive tasks

Source: OpenAI Signals Data, Bernstein Analysis

THE EMERGING MARKETS' AI SUBSCRIPTION-AN OVERAMBITIOUS STORY?

The above conclusion—where emerging economies, with India among them, become a more plausible subscription market for LLMs, may not sit down well with many. India's adoption of paid software remains low, along with several other emerging economies. So what makes us believe that it makes an excellent market for subscription of AI models? Several factors.

Before we delve into them, it's important to note the fundamental differences between the software age and the AI age. A good deal of software were either offered free or had free alternatives present in market. The costs mostly comprised high initial development/R&D with a very low ongoing costs, especially when spread over a large user base. The biggest factor was piracy—where India ranks third in the world according to Revenera data.

AI is different, the variable costs per user are high, and more usage increases the costs proportionally. There is higher power and compute footprint that is not localized on the user's system or resources, but with global impact. Moreover, sharing credentials for a single paid LLM subscription may not work with usage and token limits fully set in, plus reasons of privacy: sharing Gemini for instance means one is sharing his Google credentials with someone, giving access to emails, drive, and much more. Thus, it's not as simple as sharing a Netflix password, for instance.

Thus, there's a high chance that adoption of paid models for AI will eventually pick up in emerging economies. The consumer adoption (token use not just no. of users) may prove more gradual. Since AI will not be universally free unlike social media or digital payments, consumers choosing annual AI subscriptions will effectively need to reallocate spending from other discretionary categories. In that sense, India's status as one of the world's largest users of AI may not automatically translate into being one of its largest economic beneficiaries. The key question is whether India can build AI-native businesses and capture value, or whether it primarily becomes a large paying user of AI developed elsewhere.

INDIA'S AI INVESTMENTS-A CONCENTRATED BET

When it comes to India's AI investment story, it becomes an overwhelming infrastructure and manufacturing story rather than a frontier LLM or applied AI products story. Most of the investments that emerged in the 2026 AI Impact summit in New Delhi were centred around infrastructure related commitments-as much to the tune of \$250 billion as cited by some government officials. Reliance (not covered) alone committed \$110 billion for sovereign AI infrastructure over 7 years, while Microsoft (covered by Mark Moerdler) put aside \$50 billion by 2030 to extend AI to "lower-income" countries via India. Broadly, three domains of investments exist:

1. Data Centers and Compute capacity remains the most ambitious one, with both global hyperscalers and Indian conglomerates pouring resources into it. India's current data center capacity stands at around 1500 MW. 30 projects with additional 3.5 GW of capacity were announced between March 2025 and April 2026 alone. Some estimates put the capacity to reach 12 GW by 2030 based on announcements.
2. Semiconductors is India's next most popular leg of foray into the AI value chain. A large part of this bet is into Outsourced Assembly and Testing rather than fabrication (though some fab capacity is indeed planned in India-especially by Tata Electronics (Private)).
3. The two above leave out sovereign AI compute and foundational models-a gap which government is trying to address by its India AI mission. There's a shared compute facility with over 45,000 GPUs funding 15 domestic large/small language models, alongside 20 AI solutions across 12 sectors including agriculture and health.

HEADING TOWARDS THE FUTURE OF AI ADOPTION

Based on the usage trends and data seen throughout this report, we see the following future trends shaping up:

1. We see some sort of niche developing in LLMs. Certain use cases are dwindling fast in particular LLMs while other use cases retain their share. Anthropic (private) is clearly emerging the preferred one for coding, while OpenAI retains its supremacy in information seeking or personal guidance tasks.
2. High usage of AI in performing software tasks for countries like India, Vietnam and other EMs further fortifies our belief that IT services will be reshored within same borders: the AI alternatives to IT and agentic boom will occur within the same countries that are providing IT workforce to the world. Of course, this will mean some of the largest organizations might take a hit as more equitable distribution occurs with the help of AI.
3. While the per capita ranking remains low, India remains the second most active geography overall behind US when it comes to ChatGPT (owned by OpenAI-private) usage. We saw similar trends with Anthropic. Clearly, when it comes to usage, India isn't holding back. Per capita usage will remain low by design-but the overall usage could very well start nearing the top spot or even overtake it in the future.
4. With young users and more productive use cases, India will emerge as a high spender when it comes to LLMs. Our case, assuming 0.75%-1.25% conversion of retail users given monthly active users of each of the three major LLMs (Gemini (owned by Alphabet which is covered by Mark Shmulik), ChatGPT and Claude); and c1.4% enterprise users across all three, gives an estimate of nearly \$950 million dollars of annual spending on LLMs, just through the subscription model. This is a realistic case of \$1 billion outflows with the current adoption levels. With higher penetration-especially in enterprise and more premiumization in plans-a \$4-\$5 billion outflows in near future are very much likely. While ChatGPT remains a distant hit with retail users (~330 million monthly active users), the **revenue share will not be as skewed**, with Claude trailing with a very small margin.
5. In a decade, with application layer charging a premium on top of existing LLM subscriptions, we foresee tens of billions of dollars of outflows linked to LLM subscriptions alone. Most of these will flow out of India in the scenario of lack of any major homegrown foundational model or widely used application layer software domestically. Thus, AI linked outflows could very well become similar to oil import bill-with macro consequences.
6. The next point may seem like a counter to what we mentioned in the above two points-but has a solid standing. With around 1.5 GW capacity at present, India's projected 12 GW data center capacity by 2030 forecasts a rather aggressive build out. Even if we move over power availability, water, contractor capacity and GPU access, demand is an additional inhibitor. Our

high level calculations yield that a 1GW data center at 70%utilization in a day is able to address around 3 billion queries a day in inference,which would stand at around 55-60 million users each day.12GW data centers,thus,may be able to address queries from over 700 million users each day.This might mean a material excess capacity i the entire 12GW data center plans are realized.We believe this euphoric phase of announcements will gradually pave way for some misses and shelved plans down the road unless Indian data centers are also used for inference on AI usage across borders-something which is thwarted by latency-sensitive usage and data localization issues.We will,in the future,discuss this separately where we further refine our calculations in this area to reach a more conclusive deduction.