VŠE / PRAGUE UNIVERSITY OF ECONOMICS AND BUSINESS

Miscellanea   **Open Access**

Universal Basic AI Access: Countering the Digital Divide

Petr Špecián [ID] [1,2], **Jitka Špeciánová** [ID] [3]

[1] Institute of Political Science, Faculty of Social Sciences, Charles University, Prague, Czech Republic
[2] Department of Philosophy, Faculty of Economics, Prague University of Economics and Business, Prague, Czech Republic
[3] Department of Managerial Economics, Faculty of Business Administration, Prague University of Economics and Business, Prague, Czech Republic

Corresponding author: Petr Špecián (petr.specian@fsv.cuni.cz)

Abstract

Generative artificial intelligence (GAI) presents an opportunity to democratize access to high-performance, easy-to-use tools of productivity enhancement. However, current adoption patterns suggest that it may instead amplify existing digital divides. The aim of this paper is to propose a policy intervention to ensure equitable access to frontier GAI capabilities: the universal basic AI access (UBAI). Relying on literature research and theoretical analysis, we examine two implementation variants: a voucher-based system making use of commercial providers (UBAI-Light) and direct public provision of GAI services (UBAI-Heavy). We also consider a gradual implementation approach that allows policymakers to support an immediate capture of democratizing benefits while building the capacity for a more substantive future government involvement, should it become necessary. Given the rapid pace of GAI development and adoption, we conclude that timely implementation of UBAI could help prevent the spread of GAI-driven inequalities before they become entrenched.

Index Terms

Generative artificial intelligence; Digital divide; Technology policy; Universal basic services; Digital inequality; AI governance.

1   INTRODUCTION

Artificial intelligence (AI) comprises a broad range of computational approaches to simulating intelligent behaviour. It has also become a brand, a buzzword and a hype machine. In this paper, we keep our perspective constrained to the technical and especially socio-economic layer of the theme and use "AI" to refer specifically to generative AI (GAI) systems, particularly large language models, such as *ChatGPT* or *Claude* and similar tools, that emerged after 2022 and can engage in natural language dialogue and generate text, code and other content. Two empirically robust facts characterize the current state of GAI: Firstly, GAI capabilities are improving at a breakneck speed, swiftly saturating key performance benchmarks (Stanford University, 2025); secondly, GAI has a significant positive impact on productivity across a wide range of economically valuable cognitive tasks (Brynjolfsson et al., 2025; Cui et al., 2025). Yet only a small fraction of workforce are high-intensity GAI users (Bick et al., 2024).

This situation creates a critical fork in the road. The ability of GAI to receive directions via a conversation in natural language and deliver a large productivity boost even to users with no specialized training make it significantly more accessible than previous digital technologies, which typically have a steeper learning curve.

Early research suggests that GAI tools may particularly benefit users with lower initial competence levels (Brynjolfsson et al., 2025; Noy & Zhang, 2023). However, the current adoption patterns indicate a growing digital divide rather than a narrowing one (Khowaja et al., 2024; Pontis et al., 2025). The most easily accessible GAIs with the lowest barriers to entry lag behind the frontier models in terms of their capability, availability or safety (Li et al., 2024; Cottier et al., 2024; Stanford University, 2025). As a result, users with limited incomes or competencies are much less likely to reap the benefits than seasoned early adopters.

We consider these developments concerning, but not yet definitive: a timely policy intervention could still sway the outcomes in the democratizing direction. In this paper, we argue that universal basic AI access (UBAI) could mitigate the inequities by ensuring that all citizens, regardless of their socio-economic background, have access to advanced GAI tools. With UBAI, policymakers can promote a more equitable distribution of digital opportunities. We examine two implementation variants—a market-based voucher system (UBAI-Light) and a direct public provision system (UBAI-Heavy)—and consider their strengths and weaknesses. We also discuss the potential path of implementation.

Our argument proceeds as follows. Section 2 provides an overview of the current GAI landscape, documenting the emerging patterns of uneven adoption. Section 3 outlines two potential approaches to UBAI and a possible implementation pathway. Section 4 discusses the limitations of our analysis and opportunities for future research. Section 5 concludes.

2   WHERE WE ARE

After the "big bang" of 2022, GAI progress has continued rapidly with swiftly improving performance in various capability benchmarks (Stanford University, 2025). Over a two-year period, GAI models have become increasingly capable of difficult reasoning tasks while their rates of confabulation ("hallucination") have dropped significantly (Mendelevitch et al., 2024). They have also grown increasingly multi-modal (that is, capable of processing image and sound, not just text) and have been endowed with much larger working memory (Team Gemini Google et al., 2024). In short, GAIs have been rapidly maturing from an amusing toy to a major, if not game-changing, economic and social force.

While technological progress has been rapid, social adaptation understandably remains at an early stage. Therefore, the early movers' advantage has not yet solidified and become entrenched, leaving a broad range of scenarios feasible. Both scenarios with greatly increased "winner takes all" inequalities and scenarios with a flattening of the skill pyramid and achieving a broad democratization of cognitive capabilities may become a reality.

While recent technological transitions have typically favoured skilled users and triggered a superstar phenomenon on many labour markets (Acemoglu & Autor, 2011; Autor, 2019), GAI may present a different dynamic. Its ability to follow directions provided in natural language and engage in a naturally flowing back-and-forth exchange with its user makes it more accessible than these previous innovations (Feuerriegel et al., 2024).

However, the likelihood that the democratizing scenario will play out spontaneously is not especially high. Recent data paint a clear picture of emerging stratification in GAI adoption. For instance, in the United States, while 39% of working-age population report having used GAI tools, only about 10% are daily users (Bick et al., 2024). Similar patterns are present in Europe: as of 2024, 34% of respondents remained unaware or unsure of the existence of GAI tools, while only a quarter of those familiar with the technology ever used it for work (Corduneanu et al., 2024).

These usage patterns map onto familiar digital divides: GAI users are mostly young, male, highly educated and already technologically proficient. Tellingly, as of April 2024, still only 15% of Czech respondents expected GAI to significantly affect their employment in the next five years (Moravec & Kormaňák, 2024). Given the recent speed of GAI capability gains, these numbers suggest a degree of disconnect between the technological realities on the ground and the public perception thereof.

Importantly, the adoption patterns are not set in stone (yet) but depend on malleable features of the surrounding socio-economic landscape. Firstly, there is a capability gap between frontier and free-tier models. At the frontier, competition drives rapid innovation with models such as OpenAI's o3 achieving impressive performance in complex reasoning tasks (Garg, 2024). Free-tier offerings, however, lag significantly behind the state of the art. They tend to come with various frictions, restricted capabilities, lighter data protection and tight usage limits that constrain their practical utility. These limitations hamper the experimentation necessary to navigate the "jagged frontier"

(Dell'Acqua et al., 2023) of GAI capabilities, where users need continuous practice to identify productive use cases. The few unconstrained alternatives that do not lack in capability, such as the Chinese *DeepSeek*, raise significant privacy and security concerns (Burgess & Newman, 2025).

The rapid evolution of GAI landscape presents further barriers to the uninitiated. New models, capabilities and use cases emerge on an almost daily basis. Paired with confusing model naming practices by the leading GAI corporations, they create an impenetrable jungle of options that often vary vastly in their features and capabilities (Stanford University, 2025). In this environment, advantages easily compound as early users build their expertise, develop high-performance workflows and gain tacit knowledge (cf. Brynjolfsson et al., 2025). Often secretive about their usage strategies (Mollick, 2023), these users achieve better job performance and productivity, while the rest of the workforce remains temporarily unaware of the shifting ground beneath their feet.

However, since it is still early days in the socio-economic adaptation to GAI, we believe that a timely intervention ensuring a more equitable access to frontier models could, with high likelihood, counter or at least mitigate the worrying inequality-catalysing dynamics.

In the following, we draw inspiration from the debates surrounding universal basic services (UBS) and universal basic income (UBI). Both approaches recognize that market mechanisms alone may fail to ensure equitable access to basic resources crucial for a person's meaningful social and economic participation. While UBI proposals focus on providing unconditional cash transfers to enable individual choice in meeting basic needs (Van Parijs, 1995), UBS advocates argue for direct public provision of essential services (Portes, Reed & Percy, 2017).

GAI tools will increasingly serve the same purpose of many legacy public services that they may supplement and at least partially replace, such as education or libraries. Access to a frontier GAI system can significantly enhance an individual's productivity and learning potential (De Simone et al., 2025; Kestin et al., 2024). It serves as a free, immediately available assistant, tutor and knowledge worker to perform tasks that otherwise require specialized skills or training (e.g., Wang et al., 2024; Špecián, 2024). Unlike traditional public services that require long-term institutional investment (e.g., training and hiring teachers), deployment and impact of GAI can be much more immediate. Compared to many public services, securing frontier GAI access has relatively low costs and a limited need for physical infrastructure.

As such, GAI presents important challenges from a regulatory perspective. The most prominent regulatory frameworks, such as the EU's AI Act, focus primarily on risk mitigation through restrictions on risky uses (Novelli et al., 2024). However, such restrictions may widen digital divides: more technically sophisticated users can often circumvent limitations through VPNs or even self-hosted solutions while the ordinary crowd bears the full brunt of the regulation. Digital inclusion policies, in their turn, tend to focus on skill training and infrastructure development (Gottschalk & Weise, 2023; Vosloo & Helsper, 2023). With respect to GAI, however, such approaches fail to address the core challenge: ensuring broad access in a situation where both market mechanisms and regulatory frameworks tend to amplify rather than mitigate existing inequalities.

Therefore, we need a policy framework that can guarantee access to frontier capabilities and support human adaptation to the GAI-driven social and economic disruption. The next section proposes such a framework.

3   UNIVERSAL BASIC AI ACCESS

The preceding analysis painted a picture of a large positive potential coupled with significant risks. While the natural language interface of GAI lowers entry barriers compared to past technologies, the capability gap between the free-tier and cutting-edge GAI offerings, together with compounding early adopter advantages, suggests that market forces alone are unlikely to produce an equitable outcome. Indeed, current trajectories point towards an amplification, rather than a mitigation, of preexisting inequalities. Also, traditional digital inclusion policies appear insufficient for ensuring broad sharing of the benefits of frontier GAI. Therefore, a more targeted and proactive policy intervention, focused on ensuring broad access to frontier GAI systems, appears advisable. As the mixed-at-best successes of earlier digital education interventions indicate (Gottschalk & Weise, 2023; Vosloo & Helsper, 2023), battling entrenched inequalities is hard. Tackling them as they emerge may bear a better chance for success.

Among the potential policies to address the uneven adoption and inequality-accelerating potential of GAI, let us propose and examine what we call universal basic AI access (UBAI). As documented below, we see UBAI as a

promising foundational approach that can deliver both relatively affordable early impact and long-term strategic benefits. Its key promise is to guarantee baseline access to capable GAI systems while preserving market incentives for innovation and development of specialized offerings. Similar to a universal basic service, it treats access to a certain threshold of GAI capability as an increasingly critical component for full participation in society and the economy. However, unlike a typical UBS scheme, which often involves direct public provision of a standardized service, UBAI can incorporate a universal-basic-income-like element of user choice and resource provision, albeit directed towards specific technological capability rather than income support *per se*.

We are going to consider two possible implementation strategies: UBAI-Light, a voucher-based system making use of existing commercial services, and UBAI-Heavy, which includes direct public provision of GAI services. While both variants share the core goal of democratizing access to capable GAI systems, they differ in their implementation requirements, advantages and challenges.

3.1   UBAI-Light

UBAI-Light would operate as a voucher system enabling citizens to access premium-tier commercial GAI services with government-covered costs. Based on current commercial pricing (~$20/month for premium access), such a system could cost about $240 (€210) per citizen annually. While not cheap, this appears comparable in costs to many existing digital inclusion programmes, and perhaps amenable to bulk discounts.

To retain competitive dynamics and avoid vendor lock-in, the scheme would function as a public GAI aggregator, similar to commercial services such as *Poe.com*. Thus, its users would be free to choose among multiple providers. The system could support both subscription-based access, allowing provider switches at fixed periods (e.g., monthly), and API credits, enabling provider selection for individual tasks. While the subscription model is the simplest, the API model offers greater flexibility, allowing users to accumulate credits for costly, high-performance tasks on frontier models. This is a more advanced but perhaps increasingly important option, given the recent rise of inference-heavy "reasoning models".

Some of the key advantages of UBAI-Light stem from its relatively minimalistic requirements for government involvement. It keeps the role of the government focused on setting basic rules, providing a gateway and "paying the bills" rather than getting entangled in the technical intricacies of GAI itself. While the provision of the necessary gateway app through which users could utilise their vouchers and access the selected service is still a non-trivial undertaking, many countries appear to possess the necessary capacities for its relatively swift development. Presumably, the gateway app could function similarly to existing e-government portals that many countries already operate for tax filing or benefit claims and rely on existing authentication infrastructure. Public-private partnerships also remain an option, although not without risks of their own since they create greater risks of regulatory capture or vendor lock-in (cf. Voorwinden, 2021).

The gateway app would also serve as a filter on the eligible services. Compliance standards could be set for providers to participate in the voucher scheme in terms of required data security, safety or bias mitigation. This not only allows the government to partially suppress problematic providers but also increases its negotiation power when it comes to directing GAI developments. The government's position as a major customer gives it some leverage to demand privacy protection, safety measures or bias auditing. Specific training or model behaviour requirements could be established. For the sake of transparency, these should be publicly announced general standards that automatically enable any complying provider to participate in the voucher scheme.

An independent third-party verification of compliance also appears to be a reasonable measure in this context. Employing an independent auditor would lend credibility to the process of vetting providers for compliance with the safety, security and ethical standards required for participation in the voucher scheme. The separation of standard-setting (government) and verification (third-party auditor) introduces a separation of powers that may enhance public trust in the approved GAI services available via UBAI.

With such precautions, the UBAI-Light scheme could be effective in expanding access and managing risks, while remaining less coercive than direct regulatory measures. After all, the only default penalty for circumventing the gateway app and choosing an unvetted provider is forfeiting the opportunity for the government-sponsored voucher.

The scheme design limits potential government overreach. Most importantly, the gateway app serves merely as a matchmaker between users and providers, without creating a necessity to process or store user queries and model responses. Thus, the technical architecture can ensure that the government has no access to user interactions with GAI services. The government could therefore more credibly commit to a relatively neutral role of a mere access facilitator than if it were involved in any actual service provision.

Overall, the main appeal of UBAI-Light lies in its speed, simplicity and affordability: it builds on the infrastructure of existing providers, keeps government involvement limited and could be implemented quite rapidly without major capacity building. A well-implemented UBAI-Light would mean that users gain access to frontier GAI tools without the complexities and inefficiencies of most large-scale public IT projects. Additionally, the voucher-based model with easy provider switching avoids vendor lock-in and does not undermine the incentive to compete and innovate. For many countries, especially those without a strong domestic GAI industry, it offers the most practical path forward. However, legitimate concerns remain about relying entirely on commercial providers for such a crucial service. A more comprehensive public solution of GAI provision may still be justified, at least in some cases.

3.2 UBAI-Heavy

UBAI-Heavy represents a more ambitious strategy based on direct public provision. It would establish a government-operated baseline GAI service, likely built upon the foundation of an open-source model. Commercial providers would still be free to offer premium alternatives, perhaps providing advanced features and specialized applications. Of course, this approach would require substantial public infrastructure: from data centres and technical personnel to the entire stack of technologies necessary for operating capable GAI systems. The costs and complexities involved are far greater than with UBAI-Light. At the same time, UBAI-Heavy would present opportunities for state capacity building for the AI Age as well as greater public control over a key technology.

Given the technical capabilities required for its accomplishment, UBAI-Heavy clearly presents a major challenge for the public sector. However, developing these capabilities could yield broader strategic benefits beyond mere access provision.

Currently, a small number of firms dominate the development and deployment of the most capable GAI models. This oligopolistic environment creates power asymmetries, especially if the currently heated competition were to deteriorate to a stalemate (Zhu, 2023). Private firms can unilaterally alter their business models, impose new licensing requirements or restrict access to certain regions or user groups based on commercial or geopolitical considerations. Such concentration of power in the hands of a few private corporations poses risks to digital sovereignty, potentially leaving governments without viable alternatives in critical domains such as education, public administration or even national security (Aschenbrenner, 2024). UBAI-Heavy would provide a publicly controlled alternative to these oligopolistic alternatives. Beyond addressing market concentration risks, direct involvement in GAI operations would build government expertise for more effective regulation and governance of these systems.

In short, the appeal of UBAI-Heavy lies in the degree of control that it offers. Direct public provision allows for the most precise alignment with policy objectives and social priorities. Besides reducing dependence on commercial providers' goodwill, it also fortifies the government's negotiation position, since a capable public service could provide a baseline alternative that anchors the minimum expectations of mainstream users regarding service quality.

However, the increased centralisation inherent in UBAI-Heavy also brings substantial risks of its own. Once the government capability exists, it can be misused for data collection or surveillance. For instance, the ability to directly shape model behaviour instead of just mandating compliance with public guidelines may trigger political pressures aimed at enforcing particular ideological perspectives. There is also a risk that, rather than providing the public GAI offering as one option amidst genuine market competition, the government would instead take steps to hamstring the private competitors; perhaps especially so if the public offering were proving a materially and politically costly failure.

Furthermore, even a direct public provision model such as UBAI-Heavy is not immune to the risks of regulatory capture by industry insiders. Building and maintaining a state-of-the-art GAI service, even one based on open-source foundations, would almost certainly require tight long-term collaboration with the corporations that currently dominate the GAI landscape. These firms possess the expertise, datasets, cutting-edge R&D capacities and complex

supply chains (e.g., for advanced semiconductors) that governments typically lack. Whether through procurement contracts, joint ventures, talent acquisition or even more drastic measures such as partial nationalization (cf. Aschenbrenner, 2024), reliance on these corporate actors seems unavoidable. However, such close interdependence creates major opportunities for capture. While the government holds the regulatory power and coercive authority, critical informational advantage (and, arguably, significant political influence) would reside with the tech corporations. It can thus be expected that they would attempt to bend both the design and the outcomes of the publicly provided solution to better align with their commercial interests (cf. Dal Bó, 2006), potentially undermining the public interest goals of the UBAI-Heavy approach.

The implementation timeline presents another major challenge, given the premium on speed and the swift technological developments already discussed. While UBAI-Light could be deployed within months, developing a public GAI service would most likely require years of efforts. Moreover, early architectural decisions—necessarily made at the lowest point of public sector GAI-related competence—could create path dependencies with enduring implications for the system capability and safety. Given the usual delays with large-scale IT projects, the whole undertaking could even become irrelevant before any actual service came online.

The costs and complexities of UBAI-Heavy thus make it rather more demanding and less likely to succeed. Its main advantages are strategic and long-term. For most countries, especially those with limited technical capacity or financial resources, UBAI-Light likely represents the only feasible approach. However, for larger countries with substantial resources and strategic interests in AI development, it could prove myopic to completely forego public GAI capacity building.

3.3   Gradual approach

The key trade-offs of the UBAI variants—the swift deployment but limited control of UBAI-Light versus the strategic value but high complexity and costs of UBAI-Heavy—mean that a gradual implementation strategy might offer the best option. This approach could capture the immediate effects of UBAI-Light while preparing ground for more strategically oriented UBAI-Heavy, were the need for it to arise.

A gradual approach could become particularly compelling due to broader geopolitical developments. As Aschenbrenner (2024) argued, intensifying international competition in AI development is likely to ultimately require rather massive government involvement anyway, given the security considerations and tail risks associated with advanced AI systems (cf. Kokotajlo et al., 2025). In this context, using UBAI as a springboard for building state capacity in AI governance could prove priceless. Therefore, let us schematically explore a three-stage implementation process aimed at gaining the maximum benefits from both UBAI-Light and UBAI-Heavy while minimizing their weaknesses:

- **Phase 1** introduces UBAI-Light and includes development of baseline technical capacities necessary for more ambitious interventions. Here, the introduction of UBAI-Light would be used as a runway towards establishing the institutional preconditions for UBAI-Heavy. This includes hiring technical personnel, building cooperation with research institutions and industry stakeholders, as well as deliberating on the standards to guide the GAI deployment.
- **Phase 2** would pilot an experimental UBAI-Heavy scheme for the sake of learning and iteration. The deployment at this stage could target specific strategic areas, such as national security, public administration or research, where direct public provision may offer the highest societal return. Rather than attempting comprehensive public GAI service immediately, these domain-specific trials would allow iterative testing and improvement for the public models.
- **Phase 3** would finally introduce a full-fledged UBAI-Heavy scheme, where a publicly provided baseline GAI service would complement private offerings and compete with them. Commercial providers could continue to lead in the development of advanced features and specialized applications, while the public option would act as a baseline option and a strategic reserve.

This gradual approach may mitigate many of the problems identified earlier. The step-by-step building of technical capacity limits the risks of early-on commitment to a specific GAI architecture that could plague immediate attempts at UBAI-Heavy deployment. Starting with UBAI-Light enables the immediate democratizing benefits while the more

ambitious public infrastructure has the time to mature. The pilot phase enables learning and adaptation before full-scale deployment, reducing the probability of costly mistakes.

Importantly, the gradual approach maintains flexibility. Individual countries can adjust the specifics of public provision in view of their circumstances and the evolution of the GAI landscape. Some may find that UBAI-Light performs well enough and decide to remain in Phase 1. Others may proceed all the way to an ambitious UBAI-Heavy scheme, driven by strategic considerations. Beyond its flexibility, gradual deployment also improves political feasibility. It avoids committing to a massive public investment early on but does not preclude its possibility either. Success in initial phases can build support for more ambitious steps, problems can be addressed at a still manageable scale and the implementation can be (re-)evaluated along the way to decide whether and how to take the next step.

4 ASSUMPTIONS, LIMITATIONS AND FUTURE RESEARCH

Our analysis of UBAI shows a possible policy pathway to benefit from the democratizing potential of GAI and mitigate the risk of deepening digital divides. However, as a conceptual proposal, it also necessarily operates at a high level of abstraction. This implies presence of some important underlying assumptions and limitations, which now require explicit acknowledgment but also provide fertile ground for future research.

Firstly, our proposal rests on several key assumptions regarding the trajectory and impact of GAI and the feasibility of UBAI. In particular, we assume that: (a) GAI frontier capabilities will continue to advance and diffuse across the economy with a significant impact on productivity across diverse sectors, but without triggering an abrupt "intelligence explosion" (MacAskill & Moorhouse, 2025) or any similar transformative shock; (b) without proactive intervention, access to economically relevant capabilities will remain unevenly distributed, contributing to existing inequalities (cf. Bick et al., 2024; Corduneanu et al., 2024); (c) policy interventions can change GAI adoption patterns and, by extension, distribution of its economic benefits; and (d) the implementation of either UBAI-Light or UBAI-Heavy is feasible within the political and economic realities of the given society. In the face of today's rapid and unpredictable technological and social developments, the reliability of these assumptions is uncertain.

Secondly, as far as limitations are concerned, our conceptual exploration deliberately omits specifics of funding mechanisms, cost structures or implementation logistics of UBAI deployment. For instance, estimating the costs of vouchers or public provision, designing systems for provider vetting and user access or navigating the inevitable legal hurdles all require a thorough, context-specific analysis beyond the scope of this paper.

Another significant limitation is our focus on access provision, as already mentioned. Access alone is no guarantee of effective usage, let alone equitable outcomes. Achieving actual democratization would most likely require significant efforts beyond the UBAI scheme. These efforts may be directed towards integration with broader digital literacy and upskilling initiatives, such as Digital Equity Plan Programs (National Telecommunications and Information Administration, 2021) or The Digital Europe Programme (European Commission, 2022), for instance. Moreover, any intervention on a scale conceived of in this paper risks unintended consequences, including market distortions, new forms of inequality related to effective use or misuse of publicly provided services (cf. Robles & Mallinson, 2023).

Finally, the specific form of UBAI proposed here is created in view of the current GAI paradigms, especially the large language models used via a chat interface or API access. Future developments may lead to a new paradigm that requires access mechanisms or governance frameworks different from those that we explicitly endorse or implicitly assume. Even the current pivot towards more agentic GAI systems bears some likelihood of limiting the usefulness of our proposal (cf. Kokotajlo et al., 2025).

These limitations can be at least partially attributed to the explorative nature of this paper, which aimed primarily to initiate discussion and policy exploration. Therefore, future research should focus on filling in the remaining gaps. We see several priority areas:

a) Study of political feasibility, potential funding models and institutional requirements for UBAI implementation in specific national contexts.
b) Building economic models to estimate the costs and benefits (including productivity gains and inequality reduction) of specific UBAI schemes, conditional on various scenarios of GAI development.

c) Investigation of factors that enable or hinder GAI use, especially by populations already disadvantaged by pre-GAI digital divides. This extends towards research into optimal user interface design and effective GAI skill training methods.

d) Consideration of how UBAI frameworks could be designed as robust and adaptive vis-à-vis the deep uncertainties noted above.

While the path from UBAI as a concept to UBAI implementation is far from straight, the rapidity of GAI development and diffusion means that the exploration of the policy space exemplified by our paper is nevertheless urgent and necessary. We hope that our proposal encourages further research and innovative policy design in the domain of democratizing GAI benefits.

5   CONCLUSION

Our paper argues that a timely policy intervention could help realize the democratizing potential of GAI. Currently, GAI adoption maps onto familiar digital divides, with usage concentrated among younger, more educated and technically oriented people. However, the natural language interface and adaptive capabilities of GAI make it uniquely suited for broader accessibility.

We outlined universal basic AI access (UBAI) as a feasible policy framework for expanding GAI access. Our analysis of two implementation variants—UBAI-Light and UBAI-Heavy—pointed out the trade-offs between speed of deployment, degree of control and implementation complexity. For most countries, especially those with limited technical capacity or financial resources, the voucher-based UBAI-Light likely represents the most practical path forward. However, longer-term strategic considerations can justify more ambitious public provision schemes, particularly for larger countries with substantial resources, geopolitical concerns and interest in GAI-related state capacity building.

A gradual implementation approach may allow policymakers to capture immediate democratizing benefits while retaining longer-term flexibility. Starting with UBAI-Light enables fast deployment combined with technical capacity building. Subsequent steps can then be calibrated based on the accumulated experience, the trajectory of technological developments and the strategic considerations.

Several challenges remain to be addressed. These include developing reasonable provider standards, ensuring *de facto* rather than merely formal access and finding a balance between safety requirements and system capabilities. However, the relatively modest costs of UBAI vis-á-vis the traditional universal basic service schemes, together with the foreseeable economic impacts of GAI, mean that universal access provision could yield large returns on public investment.

However, the window of opportunity for effective intervention may be narrow. As GAI systems become more deeply embedded in society and economy, early adopter advantages are likely to compound. A rapid deployment of UBAI could help prevent the emergence of GAI-driven inequalities before they become entrenched.

ADDITIONAL INFORMATION AND DECLARATIONS

REFERENCES

Acemoglu, D., & Autor, D. H. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics* (Vol. 4, Part B, pp. 1043–1171). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

Autor, D. H. (2019). Work of the Past, Work of the Future. *AEA Papers and Proceedings*, 109, 1–32. https://doi.org/10.1257/pandp.20191110

Aschenbrenner, L. (2024). Situational Awareness: The Decade Ahead. https://situational-awareness.ai/

Bick, A., Blandin, A., & Deming, D. J. (2024). *The Rapid Adoption of Generative AI: Working Paper 32966*. National Bureau of Economic Research. https://doi.org/10.3386/w32966

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

Burgess, M., & Newman, L. (2025). DeepSeek's Popular AI App Is Explicitly Sending US Data to China. *Wired*. https://www.wired.com/story/deepseek-ai-china-privacy-data/

Corduneanu, R., Winters, S., Michalski, J., Horton, R., Sahu, R. K., & Echániz, J. (2024). Europeans are optimistic about generative AI but there is more to do to close the trust gap. *Deloitte Insights*. https://www2.deloitte.com/us/en/insights/topics/digital-transformation/trust-in-generative-ai-in-europe.html

Cottier, B., You, J., Martemianova, N., & Owen, D. (2024). How Far Behind Are Open Models? https://epoch.ai/blog/open-models-report

Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., & Salz, T. (2025). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Social Science Research Network*. https://doi.org/10.2139/ssrn.4945566

Dal Bó, E. (2006). Regulatory Capture: A Review. *Oxford Review of Economic Policy*, 22(2), 203–225. https://doi.org/10.1093/oxrep/grj013

De Simone, M., Tiberti, F., Mosuro, W., Manolio, F., Barron, M., & Dikoru, E. (2025). From chalkboards to chatbots: Transforming learning in Nigeria, one prompt at a time. *World Bank Blogs*. https://blogs.worldbank.org/en/education/From-chalkboards-to-chatbots-Transforming-learning-in-Nigeria

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Social Science Research Network*. https://doi.org/10.2139/ssrn.4573321

European Commission. (2025). The Digital Europe Programme. https://digital-strategy.ec.europa.eu/en/activities/digital-programme

Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126. https://doi.org/10.1007/s12599-023-00834-7

Garg, J. (2024). Open AI's O3 Benchmarking: Redefining Standards in AI Performance. *Gocodeo.com*. https://www.gocodeo.com/post/open-ais-o3-benchmarking

Gottschalk, F., & Weise, C. (2023). *Digital equity and inclusion in education: An overview of practice and policy in OECD countries*. OECD Education Working Papers. https://doi.org/10.1787/7cb15030-en

Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2024). *AI Tutoring Outperforms Active Learning*. Research Square. https://doi.org/10.21203/rs.3.rs-4243877/v1

Khowaja, S. A., Khuwaja, P., Dev, K., Wang, W., & Nkenyereye, L. (2024). ChatGPT Needs SPADE (Sustainability, PrivAcy, Digital divide, and Ethics) Evaluation: A Review. *Cognitive Computation*, 16(5), 2528–2550. https://doi.org/10.1007/s12559-024-10285-1

Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., & Dean, R. (2025). AI 2027. *AI 2027*. https://ai-2027.com/

Li, M., Sun, J., & Tan, X. (2024). Evaluating the effectiveness of large language models in abstract screening: A comparative analysis. *Systematic Reviews*, 13(1), 219. https://doi.org/10.1186/s13643-024-02609-x

MacAskill, W. & Moorhouse, F. (2025). *Preparing for the Intelligence Explosion*. Forethought. https://www.forethought.org/research/preparing-for-the-intelligence-explosion

Mendelevitch, O., Haihao, S., & Bharadwaj, D. (2024). *Do Smaller Models Hallucinate More?*. Vectara. https://www.vectara.com/blog/do-smaller-models-hallucinate-more

Mollick, E. (2023). *Detecting the Secret Cyborgs*. One Useful Thing. https://www.oneusefulthing.org/p/detecting-the-secret-cyborgs

Moravec, V., & Kormaňák, M. (2024). *Ve stínu ChatGPT: Proměny vnímání a využívání umělé inteligence českou populací a žurnalistikou*. Ipsos.

National Telecommunications and Information Administration. (2021). Digital Equity Act Programs. https://broadbandusa.ntia.doc.gov/funding-programs/digital-equity-act-programs?page=1

Novelli, C., Casolari, F., Rotolo, A., Taddeo, M., & Floridi, L. (2024). Taking AI risks seriously: A new assessment model for the AI Act. *AI & Society*, 39(5), 2493–2497. https://doi.org/10.1007/s00146-023-01723-z

Noy, S., & Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Social Science Research Network*. https://doi.org/10.2139/ssrn.4375283

Portes, J., Reed, H., & Percy, A. (2017). *Social prosperity for the future: A proposal for Universal Basic Services*. Institute for Global Prosperity.

Potnis, D., Tahamtan, I., & McDonald, L. (2025). Negative consequences of information gatekeeping through algorithmic technologies: An Annual Review of Information Science and Technology (ARIST) paper. *Journal of the Association for Information Science and Technology*, 76(1), 262–288. https://doi.org/10.1002/asi.24955

Robles, P., & Mallinson, D. J. (2023). Catching up with AI: Pushing toward a cohesive governance framework. *Politics & Policy*, 51(3), 355–372. https://doi.org/10.1111/polp.12529

Stanford University. (2025). Artificial Intelligence Index Report 2025. *Stanford Institute for Human-Centered Artificial Intelligence*. https://hai.stanford.edu/ai-index

Špecián, P. (2024). Machine Advisors: Integrating Large Language Models Into Democratic Assemblies. *Social Epistemology*, (in press), 1–16. https://doi.org/10.1080/02691728.2024.2379271

Team Gemini Google, et al. (2024). Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv*. https://doi.org/10.48550/arXiv.2403.05530

Van Parijs, P. (1995). *Real Freedom for All: What (if Anything) Can Justify Capitalism?*. Clarendon Press.

Voorwinden, A. (2021). The privatised city: Technology and public-private partnerships in the smart city. *Law, Innovation and Technology*, 13(2), 439–463. https://doi.org/10.1080/17579961.2021.1977213

Vosloo, S., & Helsper, E. (2023). A global review of selected digital inclusion policies. *Unicef - Innocenti – Global Office of Research and Foresight*. https://www.unicef.org/innocenti/reports/global-review-selected-digital-inclusion-policies

Wang, R. E., Ribeiro, A. T., Robinson, C. D., Loeb, S., & Demszky, D. (2024). Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise. *arXiv*. https://arxiv.org/abs/2410.03017v2

Zhu, Y. (2023). The Case for Universal Basic Computing Power. *arXiv*. https://doi.org/10.48550/arXiv.2311.12872